

1 Probabilistic inference of multi-Gaussian fields from
2 indirect hydrological data using circulant embedding
3 and dimensionality reduction

Eric Laloy,^{1,3} Niklas Linde,², Diederik Jacques¹ and Jasper A. Vrugt^{3,4}

Eric Laloy, Institute for Environment, Health and Safety, Belgian Nuclear Research Centre, Mol, Belgium. (elaloy@sckcen.be)

Niklas Linde, Applied and Environmental Geophysics Group, Institute of Earth Sciences, University of Lausanne, Lausanne, Switzerland.

Diederik Jacques, Institute for Environment, Health and Safety, Belgian Nuclear Research Centre, Mol, Belgium.

Jasper A. Vrugt, Department of Civil and Environmental Engineering, University of California Irvine, Irvine, CA, USA.

¹Institute for Environment, Health and

4 **Abstract.** We present a Bayesian inversion method for the joint infer-
5 ence of high-dimensional multi-Gaussian hydraulic conductivity fields and
6 associated geostatistical parameters from indirect hydrological data. We com-
7 bine Gaussian process generation via circulant embedding to decouple the
8 variogram from grid cell specific values, with dimensionality reduction by in-
9 terpolation to enable Markov chain Monte Carlo (MCMC) simulation. Us-
10 ing the Matérn variogram model, this formulation allows inferring the con-
11 ductivity values simultaneously with the field smoothness (also called Matérn
12 shape parameter) and other geostatistical parameters such as the mean, sill,
13 integral scales and anisotropy direction(s) and ratio(s). The proposed dimen-
14 sionality reduction method systematically honors the underlying variogram

Safety, Belgian Nuclear Research Centre,
Mol, Belgium.

²Applied and Environmental Geophysics
Group, Institute of Earth Sciences,
University of Lausanne, Lausanne,
Switzerland.

³Department of Civil and Environmental
Engineering, University of California Irvine,
Irvine, USA.

⁴Department of Earth Systems Science,
University of California Irvine, Irvine, USA.

15 and is demonstrated to achieve better performance than the Karhunen-Loève
16 expansion. We illustrate our inversion approach using synthetic (error cor-
17 rupted) data from a tracer experiment in a fairly heterogeneous 10,000-dimensional
18 2D conductivity field. A 40-times reduction of the size of the parameter space
19 did not prevent the posterior simulations to appropriately fit the measure-
20 ment data and the posterior parameter distributions to include the true geo-
21 statistical parameter values. Overall, the posterior field realizations covered
22 a wide range of geostatistical models, questioning the common practice of
23 assuming a fixed variogram prior to inference of the hydraulic conductivity
24 values. Our method is shown to be more efficient than sequential Gibbs sam-
25 pling (SGS) for the considered case study, particularly when implemented
26 on a distributed computing cluster. It is also found to outperform the method
27 of anchored distributions (MAD) for the same computational budget.

1. Introduction

28 High-parameter dimensionality poses considerable challenges for the inversion of ground-
29 water flow and transport data [e.g., *Kitanidis, 1995; Hendricks-Franssen et al., 2009;*
30 *Laloy et al., 2013; Zhou et al., 2014*, and references therein]. What is more, conceptual
31 and structural inadequacies of the subsurface model and measurement errors of the model
32 input (boundary conditions) and output (calibration) data introduce uncertainty in the
33 estimated parameters and model simulations. Another important source of uncertainty
34 originates from sparse data coverages that rarely contain sufficient information to uniquely
35 characterize the subsurface at a spatial resolution deemed necessary for accurate modeling.
36 This results in an ill-posed inverse problem with many different sets of model parameter
37 values that fit the data acceptably well. Inversion methods should consider this inherent
38 uncertainty and provide an ensemble of model realizations that accurately span the range
39 of possible models that honor the available calibration data and prior information.

40 Hydraulic conductivity (K) fields are typically assumed to be stationary and log-
41 normally distributed with a spatial structure determined by a two-point geostatistical
42 model or variogram [e.g., *Rubin, 2003*]. Unfortunately, a lack of (sufficient) point K mea-
43 surements (if any) makes it difficult to estimate directly the geostatistical parameters
44 (mean, sill, variogram model, integral scales and anisotropy factors) from variographic
45 analysis [*Ortiz-Deutsch, 2002; Nowak et al., 2010*]. Simultaneous (inverse) inference of
46 conductivity values and associated geostatistical parameters is therefore attractive yet
47 computationally challenging. Indeed, only a few studies can be found in the literature
48 that have attempted simultaneous estimation using global and probabilistic search meth-

ods. For example, *Jafarpour and Tarrahi* [2011] used the Ensemble Kalman Filter [EnKF, *Evensen*, 2003] to estimate jointly the conductivity field and associated geostatistical properties. However, their attempt was not particularly successful, and they attributed this lack of success to a complex and nonunique relationship between the parameters of interest and available flow data. On the contrary, *Jardani et al.* [2012] used Bayesian inversion and backed out successfully the transmissivities at 72 pilot-points together with 45 leakage coefficients, and the sill, and correlation range of a 2D spherical isotropic variogram. However, their case study was made relatively simple by assuming a fairly smooth multi-Gaussian field with narrow ranges for the two unknown variogram parameters. *Fu and Gómez-Hernández* [2009] introduced a Markov chain Monte Carlo (MCMC) simulation method that iteratively produces local perturbations of a multi-Gaussian parameter field by resimulation of only a fraction of the field at each realization. This approach cannot directly cope with variogram uncertainty, but *Hansen et al.* [2012] and *Hansen et al.* [2013a, b] proposed methodological extensions to enable joint inference of the values of a multi-Gaussian field and its associated variogram properties. Their method, known as sequential Gibbs sampling (SGS), was applied to the joint estimation of a 2D velocity field and corresponding correlation lengths using crosshole ground penetrating radar tomography data [*Hansen et al.*, 2013b]. However, the variogram inference only concerned two parameters (i.e., the range in two different directions), and the inverse solution was stabilized by using a Gaussian prior. Furthermore, variogram inference with SGS simply involves sampling of the prior variogram distribution, which may not be efficient if the prior variogram uncertainty range is large and/or the information content of the calibration data is high. Lastly, the method of anchored distributions (MAD) introduced by

72 *Zhang and Rubin* [2008] and *Rubin et al.* [2010] simultaneously derives variogram param-
73 eters and hydraulic conductivity/transmissivity values at selected locations (the so called
74 “anchors”) within the hydrogeologic domain of interest. In the examples presented by
75 *Rubin et al.* [2010] and *Murakami et al.* [2010], the estimated geostatistical properties are
76 restricted to the mean, sill and range of a 2D exponential isotropic model. Perhaps more
77 importantly, the MAD framework relies on plain Monte Carlo (MC) simulation and is thus
78 computationally very demanding, particularly in high-dimensional parameter spaces [e.g.,
79 *Murakami et al.*, 2010]. Indeed, MC simulation is very inefficient if the prior distribution
80 is large with respect to the size of the posterior distribution.

81 Dimensionality reduction of the parameter space can help solving high-dimensional in-
82 verse problems. This includes methods such as the MAD technique described above or
83 the Karhunen-Loève (KL) transform [*Loève*, 1977]. The latter is widely used in sub-
84 surface hydrology to represent multi-Gaussian parameter fields. The KL transform uses
85 the covariance function to describe a spatially Gaussian process in a reduced basis [e.g.,
86 *Zhang and Zu*, 2004; *Li and Cirpka*, 2006; *Laloy et al.*, 2013]. The base functions are
87 the eigenfunctions of the covariance function multiplied by the square root of the associ-
88 ated eigenvalues. The sorted eigenvalues and corresponding eigenfunctions can then be
89 truncated, thereby leading to a reduced parameter space if the number of dominant eigen-
90 values is smaller than the number of simulation grid points. Overall, the smoother the
91 covariance kernel the larger the parameter reduction. Hence, the KL expansion cannot
92 cope efficiently with rough random fields and/or short integral scales. In such cases, the
93 number of base functions needed for accurate field reconstruction may approach the size
94 of the original discretized domain. Another difficulty arises from the fact that numerical

95 estimation of the required eigenfunctions and eigenvalues of the considered covariance
96 kernel can be CPU-demanding for large grids [though different kernel-specific solutions
97 exist for speeding up efficiency, see, e.g, *Zhang and Zu, 2004; Li and Cirpka, 2006*].

98 Here we present a novel Bayesian inversion approach for the simultaneous estimation
99 of high-dimensional hydraulic conductivity fields and associated two-point geostatistical
100 properties from indirect hydrologic data. Our method uses Gaussian process generation
101 via circulant embedding [*Dietrich and Newsam, 1997*] to decouple the variogram from grid
102 cell specific values, and implements dimensionality reduction by interpolation to enable
103 MCMC simulation with the DREAM_(ZS) algorithm [*Vrugt et al., 2009; Laloy and Vrugt,*
104 *2012*]. We use the Matérn function to infer the conductivity values jointly with the field
105 smoothness and other geostatistical parameters (mean, sill, integral scales, anisotropy
106 direction(s) and anisotropy ratio(s)). Moreover, conditioning on direct point conductivity
107 measurements (if any) is straightforward. We illustrate our method using synthetic, error
108 corrupted, data from a tracer simulation experiment involving a 10,000 dimensional 2D
109 conductivity field.

110 This paper is organized as follows. Section 2 presents the different elements of our
111 inversion approach, and demonstrates the merits of the proposed dimensionality reduction
112 method by comparison against the widely used KL transform. This is followed in section 3
113 with numerical experiments involving a fairly heterogeneous reference field. Our numerical
114 experiments involve benchmark tests against standard implementations of the state-of-the-
115 art SGS and MAD techniques. Section 4 takes a closer look at the required CPU-time and
116 provides further analysis of the performance of our method. Finally, section 5 concludes
117 this paper with a summary of the most important findings.

2. Methods

118 For ease of understanding, we proceed first by describing the general circulant em-
 119 bedding technique in section 2.1. Next, section 2.2 details our proposed dimensionality
 120 reduction approach. This step is of utmost importance as it reduces significantly the di-
 121 mensionality of the parameter space used by the circulant embedding, thereby enabling
 122 Bayesian inference via MCMC simulation. Section 2.3 then compares for different lev-
 123 els of dimensionality reduction the proposed approach against the KL transform, before
 124 sections 2.4, 2.5 and 2.6 present the other ingredients of our inversion approach.

2.1. Stationary Gaussian Process Generation via Circulant Embedding

125 We use the stationary Gaussian process generation through circulant embedding of the
 126 covariance matrix proposed by *Dietrich and Newsam* [1997] [see also the excellent review
 127 by *Kroese and Botev*, 2013]. We provide herein a short description of this methodology
 128 for a 2D domain but extension to a 3D grid is straightforward. Further details on the
 129 method can be found in the cited references.

130 Multi-Gaussian field generation over a regularly meshed 2D grid can be implemented as
 131 follows. If \mathbf{Y}_0 , the zero-mean stationary Gaussian field to be generated, is of size $m \times n$
 132 then:

- 133 1. Build a $(2m - 1) \times (2n - 1)$ symmetric covariance matrix \mathbf{S} for a given covariance
 134 kernel \mathbf{C} . The entries of \mathbf{S} are the 2-point covariances between any point in the $(2m - 1) \times$
 135 $(2n - 1)$ domain and the domain center point. The dimensions $(2m - 1) \times (2n - 1)$ rep-
 136 resent the minimal length for which a symmetric nonnegative definite circulant matrix \mathbf{S}
 137 can be found.

138 2. Compute the $\mathbf{\Omega}$ matrix of eigenvalues as

$$139 \quad \mathbf{\Omega} = \frac{\text{real}[\text{FFT2}(\text{FFTSHIFT}\{\mathbf{S}\})]}{(2m-1) \times (2n-1)}, \quad (1)$$

140 where $\text{real}[\]$ denotes the real part, FFT2 signifies the two-dimensional fast Fourier trans-
 141 form, and FFTSHIFT is a function that swaps the first quadrant of a matrix with the
 142 third and the second quadrant with the fourth.

143 3. Make sure that all elements $\Omega_{i,j}$ where $i = 1 \cdots 2m-1$ and $j = 1 \cdots 2n-1$ are greater
 144 than zero (nonnegative embedding). Negative eigenvalues might appear when the integral
 145 scale of \mathbf{C} becomes large compared to the domain size. For instance, *Dietrich and Newsam*
 146 [1997] showed that for a two-dimensional domain of size $m \times m$ the minimum ratios of
 147 the domain side length to the integral scale for which the embedding is non-negative are
 148 about 5.6 and 4.5 respectively, for Gaussian and exponential covariance models. In this
 149 work, we simply assume that the search range of the integral scale is bounded such that
 150 the maximum number of negative eigenvalues in $\mathbf{\Omega}$ is kept reasonably small, and negative
 151 values are set to zero if they occur.

152 4. Generate two $(2m-1) \times (2n-1)$ matrices with standard normal variates, say \mathbf{Z}_1
 153 and \mathbf{Z}_2 , and construct the complex Gaussian matrix

$$154 \quad \underline{\mathbf{Z}} = \mathbf{Z}_1 + i\mathbf{Z}_2, \quad (2)$$

155 with $i = \sqrt{-1}$. Finally, two independent zero-mean stationary Gaussian realizations \mathbf{Y}_0
 156 with prescribed covariance kernel \mathbf{C} are given by the first $m \times n$ elements of $\text{real}[\mathbf{f}]$ and
 157 $\text{imag}[\mathbf{f}]$ where $\text{imag}[\]$ signifies the imaginary part and

$$158 \quad \mathbf{f} = \text{FFT2}\left(\sqrt{\mathbf{\Omega}} \odot \underline{\mathbf{Z}}\right), \quad (3)$$

159 where \odot denotes component-wise multiplication. The complex standard normal matrix
 160 $\underline{\mathbf{Z}}$ represents the uncorrelated noise component of the Gaussian field, decoupled from its
 161 covariance structure embedded into $\mathbf{\Omega}$.

162 The use of circulant embedding for generating multi-Gaussian fields (section 2.1)
 163 has several attractive features, one of which is that it enables joint inference by fast
 164 (de)coupling of the grid cell random numbers from their geostatistical properties. The
 165 computational complexity (cost) of circulant embedding is similar to that of the FFT
 166 method. The computational complexity of the FFT method is of order $\mathcal{O}(n \log(n))$ for a
 167 symmetric positive definite $n \times n$ covariance matrix, which compares very favorably to the
 168 computational complexity of $\mathcal{O}(n^3)$ for Cholesky decomposition. Yet statistical inference
 169 becomes more and more difficult with increasing dimensionality of the search space. This
 170 justifies the dimensionality reduction approach proposed in section 2.2.

171 Lastly, it is worth noting that an alternative FFT-based decoupling method was pro-
 172 posed by *Le Ravalec et al.* [2000]. Their FFT moving average (or FFT-MA) approach
 173 has been developed independently from the standard technique by *Dietrich and Newsam*
 174 [1997], though it also makes use of FFT and the circulant embedding property of covari-
 175 ance matrices. The FFT-MA generator is used as a basic building block by *Hansen et*
 176 *al.* [2013a, b] for joint inversion of field properties and variogram parameters. Yet our
 177 experience with FFT-MA suggests the need for a larger embedding domain than for the
 178 standard approach to produce consistent multi-Gaussian fields with long integral scales.
 179 Coupling FFT-MA with dimensionality reduction also led to reconstructed models that
 180 are less accurate than those derived from coupling of dimensionality reduction with the
 181 classical circulant embedding approach (see also section 2.3).

2.2. Dimensionality Reduction by One-dimensional FFT Interpolation

182 We take advantage of one-dimensional interpolation to significantly reduce the dimen-
 183 sionality of the parameter space. More specifically, we employ one-dimensional interpola-
 184 tion by the FFT method to resample two lower-dimensional vectors of standard normal
 185 variates, say \mathbf{r}_1 and \mathbf{r}_2 , to two sets of $(2m - 1)(2n - 1)$ equally spaced points, hereafter
 186 referred to as \mathbf{z}_1 and \mathbf{z}_2 . The reconstructed vectors \mathbf{z}_1 and \mathbf{z}_2 can then be reshaped into
 187 the matrices \mathbf{Z}_1 and \mathbf{Z}_2 (see equation (2)). In this work, \mathbf{r}_1 and \mathbf{r}_2 have a dimensionality
 188 that is one to two orders of magnitude lower than that of \mathbf{z}_1 and \mathbf{z}_2 .

189 The original low-dimensional \mathbf{r}_1 and \mathbf{r}_2 vectors are thus transformed to the Fourier
 190 domain using FFT and then transformed back with $(2m - 1)(2n - 1)$ points to produce
 191 \mathbf{z}_1 and \mathbf{z}_2 , respectively. We deliberately chose FFT over other methods such as linear
 192 interpolation, as this approach was found to better preserve, during resampling, the unit
 193 variance of the standard normal distribution. In other words, if \mathbf{r}_1 and \mathbf{r}_2 are $\sim \mathcal{N}(\mathbf{0}, \mathbf{1})$,
 194 then the variances of \mathbf{z}_1 and \mathbf{z}_2 were found to be closer to unity when one-dimensional
 195 interpolation is performed by FFT. In contrast, linear interpolation generally led to a
 196 variance reduction.

197 To eliminate short lag autocorrelation, the elements of \mathbf{z}_1 and \mathbf{z}_2 are permuted ran-
 198 domly after interpolation from \mathbf{r}_1 and \mathbf{r}_2 , respectively. This permutation step is necessary
 199 as the circulant embedding method breaks down if neighboring values of \mathbf{z}_1 and \mathbf{z}_2 are
 200 correlated. Therefore, we use pre-selected permutation schemes to independently permute
 201 the elements of \mathbf{z}_1 and \mathbf{z}_2 . Of the two field realizations $\text{real}[\mathbf{f}]$ and $\text{imag}[\mathbf{f}]$ produced by
 202 equation (3), we only use $\text{real}[\mathbf{f}]$ herein. This is because, for inference, each $\mathbf{r} = [\mathbf{r}_1, \mathbf{r}_2]$

203 vector must be associated with a single multi-Gaussian field and corresponding simulated
 204 data set.

205 Our dimensionality reduction approach can thus be briefly summarized as follows:

206 1. Perform circulant (periodic) embedding of the covariance function at the desired
 207 resolution. This gives the circulant matrix \mathbf{S} .

208 2. Fourier transform \mathbf{S} to obtain the matrix of eigenvalues $\mathbf{\Omega}$.

209 3. Generate two low-dimensional vector of real-valued (standard normal) random num-
 210 bers, \mathbf{r}_1 and \mathbf{r}_2 . The larger the dimensionality reduction, the smaller the selected sizes of
 211 \mathbf{r}_1 and \mathbf{r}_2 .

212 4. Using one-dimensional FFT interpolation, resample \mathbf{r}_1 and \mathbf{r}_2 to two vectors of real-
 213 valued random numbers of the right size, \mathbf{z}_1 and \mathbf{z}_2 .

214 5. Randomly permute the elements of \mathbf{z}_1 and \mathbf{z}_2 to eliminate the short-lag autocorre-
 215 lation caused by the interpolation.

216 6. Reshape the \mathbf{z}_1 and \mathbf{z}_2 vectors into the two matrices \mathbf{Z}_1 and \mathbf{Z}_2 , respectively.

217 7. Fourier transform the product of the complex Gaussian matrix $\underline{\mathbf{Z}} = \mathbf{Z}_1 + i\mathbf{Z}_2$ with
 218 the square-root of $\mathbf{\Omega}$ to obtain two fields in the spatial domain (a real-valued and an
 219 imaginary one).

220 Lastly, we would like to stress that our dimensionality reduction approach is funda-
 221 mentally different from the KL transform. This latter method describes multi-Gaussian
 222 fields in a reduced basis that reproduces the large scale variations only. The proposed
 223 approach, on the contrary, does not favor one length scale over another, and will lead to
 224 reconstructed fields that consistently honor the selected variogram independently of the

225 number of “super parameters” or dimensionality reduction variables (e.g., elements of \mathbf{r})
226 considered. This is demonstrated in the next section.

2.3. Effects of the Dimensionality Reduction

227 We first investigate the trade-off between dimensionality reduction and the accuracy of
228 reconstruction, that is, the degree to which the statistical properties of the reconstructed
229 field match those derived from direct generation of the original field. To highlight the
230 essential differences between our approach and the KL expansion, the latter is included
231 in our analysis.

232 An anisotropic exponential variogram with short integral scales is considered for re-
233 construction of a 100×100 field (that is, 10,000 grid points). This variogram model
234 characterizes the log conductivity of the reference field used in our inversions (Table 1
235 and Figure 1a). The grid point mean and variance distributions and the average ex-
236 perimental variograms calculated from 1000 field realizations are analyzed to assess the
237 performance of the dimensionality reductions. The number of variables of each dimen-
238 sionality reduction, hereafter referred as DR variables, corresponds to the length of the
239 \mathbf{r} -vector, $\mathbf{r} = [\mathbf{r}_1, \mathbf{r}_2]$. For the KL transform, the dimensionality reduction variables are the
240 coefficients that multiply the base functions [see, e.g., *Zhang and Zu, 2004; Li and Cirpka,*
241 *2006; Laloy et al., 2013*, for details] and we refer to these coefficients as KL variables.

242 Figures 2 and 3 depict the corresponding results for 100, 250 and 1000 DR (Figure 2)
243 and KL (Figure 3) variables. The mean of the reconstructed field (not shown) is not
244 affected by dimensionality reduction, yet the grid point variances clearly are (Figures
245 2a-2c and 3a-3c). As the number of DR variables increases and dimensionality reduction
246 becomes less important, the distribution of the grid point variance gets narrower and closer

247 to the statistical fluctuations derived from direct simulation of 1000 standard normal
248 fields (Figures 2a-c). A similar trend is observed for the KL transform (Figures 3a-
249 c), though with much more irregular and over-dispersed variance distributions. Indeed,
250 the proposed approach appears to honor the prescribed variogram independently of the
251 selected number of DR variables (panels 2d-2f). The spurious correlations introduced by
252 dimensionality reduction do not noticeably affect the 2-point correlation structure of the
253 reconstructed field. This is explained by the fact that the fixed permutation scheme used
254 herein causes the (artificial) additional correlations to be distributed independently from
255 the lag (separation) distance between two points. This permutation scheme also has a
256 desired byproduct which is that it simplifies the model reduction error to random noise
257 during inversion. As a consequence of the above, the associated (randomly chosen) field
258 realizations are visually similar to their counterparts derived from direct field generation
259 (compare Figure 1a with Figures 2g-i). Perhaps not surprisingly, the KL is unable to
260 honor the variogram of the reference field even when 1000 KL variables are used (Figures
261 3d-f), and the generated fields are overly smooth (Figures 3g-i).

262 We repeated the above analysis for the same geostatistical model except for the integral
263 scale along the major axis of anisotropy that we fixed to a five times larger value, that is,
264 3.33 m. The main results of this analysis (not shown) are that our proposed approach still
265 outperforms the KL. Even when considering 1,000 KL variables, the KL transform was
266 found to produce over-smoothed fields given the selected exponential variogram model
267 whereas the associated correlation lengths remain slightly overestimated (not shown).

268 We conclude from this analysis that the proposed dimensionality reduction approach
269 (1) is well suited to reconstruction of multi-Gaussian fields, and (2) outperforms the KL

transform in cases of short and moderate-lag correlation(s). For computational tractability, we use 250 DR variables in our first MCMC trial. Though larger values would ensure less bias in the grid point variance, we consider the deviations of Figure 2b to be acceptable. In a second step, a MCMC trial with 1000 DR variables is performed, and the posterior distributions resulting from using 250 and 1000 DR variables are compared.

2.4. Conditioning to Point Conductivity Measurements

The unconditional (approximately) multi-Gaussian field realizations generated by our method can easily be conditioned on point measurements via kriging [e.g., *Chilès and Delfiner, 1999*]. This reproduces the actual point measurements and preserves the prescribed variogram. Details of this procedure can be found in Appendix A. For the sake of brevity, however, we do not condition on point conductivity measurements in the present paper.

2.5. The Matérn Variogram

We use the *Matérn* [1960] variogram model to describe the geostatistical properties of the field. This function is given by

$$\gamma(|\mathbf{h}|) = v \left[1 - \frac{1}{\Gamma(\nu) 2^{\nu-1}} \left(\frac{|\mathbf{h}|}{\alpha} \right)^\nu K_\nu \left(\frac{|\mathbf{h}|}{\alpha} \right) \right], \quad (4)$$

where $\alpha > 0$ is the scale (or range) parameter, $\Gamma(\cdot)$ represents the gamma function, $K_\nu(\cdot)$ denotes the modified Bessel function of the second kind and order ν , and $|\mathbf{h}|$ signifies the norm of the lag distance vector \mathbf{h} . The Matérn function is equivalent to the exponential model for $\nu = 0.5$, the *Whittle* [1954] model for $\nu = 1$ and approaches the Gaussian model for $\nu \rightarrow \infty$. The integral scale, I , measures spatial persistence and is defined as [e.g.,

289 *Rubin, 2003]*

$$I = \frac{1}{v} \int_0^\infty C(|\mathbf{h}|) d|\mathbf{h}|, \quad (5)$$

291 with covariance function $C(|\mathbf{h}|) = -\gamma(|\mathbf{h}|) + v$. The integral scale of the Matérn model
 292 depends on the values of α and the shape parameter ν . For a fixed value of α , I becomes
 293 larger if ν increases [e.g., *Pardo-Iguzquiza and Chica-Olmo, 2008*]. This dependence might
 294 complicate inference of I . Fortunately, numerical simulations with different values of α
 295 and ν demonstrates that the ratio of I to α is constant for a given value of ν . For any
 296 value of ν , the following fitted polynomial function can be used to derive α from I

$$\frac{I}{\alpha} = -0.0014\nu^6 + 0.0242\nu^5 - 0.1745\nu^4 + 0.6558\nu^3 - 1.4377\nu^2 + 2.4506\nu + 0.0586. \quad (6)$$

298 This allows us to simultaneously infer I and ν . The coefficient of determination (squared
 299 correlation coefficient) associated with equation (6) is 0.999986.

2.6. Joint Inference of Conductivity Fields and Variogram Parameters

300 In the Bayesian paradigm, the unknown model parameters, $\boldsymbol{\theta}$, are viewed as random
 301 variables with a posterior probability density function (pdf), $p(\boldsymbol{\theta}|\mathbf{d})$, given by

$$p(\boldsymbol{\theta}|\mathbf{d}) = \frac{p(\boldsymbol{\theta})p(\mathbf{d}|\boldsymbol{\theta})}{p(\mathbf{d})} \propto p(\boldsymbol{\theta})L(\boldsymbol{\theta}|\mathbf{d}), \quad (7)$$

303 where $L(\boldsymbol{\theta}|\mathbf{d}) \equiv p(\mathbf{d}|\boldsymbol{\theta})$ signifies the likelihood function of $\boldsymbol{\theta}$. The normalization factor
 304 $p(\mathbf{d}) = \int p(\boldsymbol{\theta})p(\mathbf{d}|\boldsymbol{\theta})d\boldsymbol{\theta}$ is obtained from numerical integration over the parameter space
 305 so that $p(\boldsymbol{\theta}|\mathbf{d})$ is a proper probability density function and integrates to unity. The quan-
 306 tity $p(\mathbf{d})$ is generally difficult to estimate in practice but is not required for parameter
 307 inference. In the remainder of this paper, we will thus focus on the unnormalized density
 308 $p(\boldsymbol{\theta}|\mathbf{d}) \propto p(\boldsymbol{\theta})L(\boldsymbol{\theta}|\mathbf{d})$. As an exact analytical solution of $p(\boldsymbol{\theta}|\mathbf{d})$ is not available in most
 309 practical cases, we resort to MCMC simulation to generate samples from the posterior

pdf [see, e.g., *Robert and Casella, 2004*]. The state-of-the-art DREAM_(ZS) [*ter Braak and*
Vrugt, 2008; Vrugt et al., 2009; Laloy and Vrugt, 2012] algorithm is used to approximate
the posterior distribution. A detailed description of this sampling scheme including a
proof of ergodicity and detailed balance can be found in the cited references. Various
contributions in hydrology and geophysics (amongst others) have demonstrated the abil-
ity of DREAM_(ZS) to successfully recover high-dimensional target distributions [*Laloy et*
al., 2012, 2013; Linde and Vrugt, 2013; Rosas-Carbajal et al., 2014; Laloy et al., 2014;
Lochbühler et al., 2014, 2015].

Under Gaussian and stationarity assumptions, the field geostatistical properties and
pixel/voxel random number values that jointly define the (base ten) log conductivity field,
 $\log_{10}(\mathbf{K})$, can be inferred simultaneously using MCMC simulation. We use the Matérn
function to infer field smoothness jointly with the standard normal variates and other
geostatistical parameters. The following geostatistical parameters are sampled together
with \mathbf{r}_1 and \mathbf{r}_2 : (I) m , the mean, (II) v , the variance, (III) I_M , the integral scale along
the major axis of anisotropy, (IV) R_I , the ratio of the integral scale along the minor axis
of anisotropy (I_m) to the integral scale along the major axis of anisotropy, (V) A , the
anisotropy direction or angle (rotation anti-clockwise from the z -axis), and (VI) ν , the
shape parameter of the Matérn function. To build the covariance kernel, \mathbf{C} , we used the
mGstat geostatistical toolbox in MATLAB (<http://mgstat.sourceforge.net/>).

If we assume the N -vector of residual errors (differences between the measured and
simulated data), \mathbf{e} , to be Gaussian distributed, uncorrelated and with constant variance,
 σ_e^2 , the likelihood function of $\boldsymbol{\theta}$ can be written as

$$L(\boldsymbol{\theta}|\mathbf{d}) = \left(\frac{1}{\sqrt{2\pi\sigma_e^2}} \right)^N \exp \left(-\frac{1}{2}\sigma_e^{-2} \sum_{i=1}^N [d_i - F_i(\boldsymbol{\theta})]^2 \right), \quad (8)$$

333 where $\mathbf{d} = (d_1, \dots, d_N)$ is a set of measurements, and $F(\boldsymbol{\theta})$ is a deterministic “forward”
 334 model. The standard deviation of the residuals, σ_e (kg m^{-3}), is jointly inferred with
 335 the other unknown variables, and thus $\boldsymbol{\theta} = [\sigma_e, m, \nu, I_M, R_I, A, \nu, \mathbf{r}_1, \mathbf{r}_2]$ (see Table 1).
 336 The number of parameters sampled with MCMC is thus equivalent to the number of DR
 337 variables plus seven. This equates to a total of 257 parameters for the first MCMC trial
 338 with 250 DR variables, and 1007 parameters for the second trial with 1000 DR variables.

339 The standard normal distribution of \mathbf{z}_1 and \mathbf{z}_2 (and thus \mathbf{r}_1 and \mathbf{r}_2) can be enforced by
 340 the use of a standard normal prior

$$341 \quad p(\mathbf{r}) = \frac{\exp\left(-\frac{1}{2}\mathbf{r}^T\mathbf{r}\right)}{\sqrt{2\pi}^N}, \quad (9)$$

342 in which the superscript T signifies transpose and $\mathbf{r} = [\mathbf{r}_1, \mathbf{r}_2]$. The variogram is assumed
 343 to be largely unknown and thus characterized by a wide prior, details of which will follow
 344 in section 3.

3. Case Studies

3.1. Model Setup

345 The 100×100 modeling domain lies in the $x - z$ plane with a grid cell size of 0.2 m
 346 (Figure 1a). Steady state groundwater flow is simulated using MaFloT [*Künze and Lunati,*
 347 2012] assuming no flow boundaries at the top and bottom and fixed head boundaries on
 348 the left and right sides of the domain so that a lateral head gradient of 0.025 (-) is
 349 imposed, with water flow in the x -direction. For the tracer experiment, we consider
 350 two different boreholes that are 20 m apart. A conservative tracer with concentration
 351 of 1 kg m^{-3} is applied into the fully screened left borehole using a step function. The
 352 background solute concentration is assumed to be 0.01 kg m^{-3} . Ignoring density effects,

353 conservative transport of the tracer through the subsurface is simulated with MaFloT
354 using open boundaries on all sides, and longitudinal and transverse dispersivities of 0.1
355 and 0.01 m, respectively. Solute transport was monitored during a period of 10 days
356 with concentration measurements made every day at nine different depths (2, 4, 6, 8,
357 10, 12, 14, 16, and 18 meters) in the borehole at the right hand side. The total number
358 of observations is therefore 90. These measurement data were then corrupted with a
359 Gaussian white noise using a standard deviation equivalent to 5% of the mean observed
360 concentration. This led to a root-mean-square-error (RMSE) of 0.039 kg m^{-3} between
361 error-free and noisy data (Figure 1b).

3.2. Inference of an Heterogeneous Random Field with Short Integral Scales

362 Our case study considers a reference log conductivity field with an exponential variogram
363 model and fairly short integral scales compared to the domain size of $20 \times 20 \text{ m}$ (Figure
364 1a). The values of the geostatistical parameters are: $\nu = 0.5$, $I_M = 0.67 \text{ m}$ and $R_I = 0.25$
365 ($I_m = 0.17 \text{ m}$). Furthermore, we assume value of $m = -3$ and $v = 1$ for the mean and
366 variance of the log-conductivity field, whereas the anisotropy angle, A , is set to 75 degrees.
367 In the absence of prior information about the geostatistical parameters (with exception
368 of the ranges of the search space), we assumed either uniform or *Jeffreys* [1946] (that is,
369 log-uniform) truncated individual priors that span a wide range of values. We selected a
370 bounded uniform prior for m and a bounded Jeffreys prior for v . This is a common choice
371 in the inference of multi-Gaussian fields [e.g., *Box and Tiao, 1973; Rubin et al., 2010*].
372 We chose bounded uniform priors for I_M , A , and R_I , and a bounded Jeffreys prior for
373 ν . Also, a Jeffreys prior is selected for σ_e . Table 1 summarizes the prior distribution and

374 corresponding ranges of each parameter. For completeness, we also list the true values of
375 the geostatistical parameters used to generate the reference field.

376 We estimate the posterior distribution of the parameters using MCMC simulation with
377 DREAM_(ZS). Default values of the algorithmic variables are used. Yet the number of
378 Markov chains was increased to eight and the number of crossover values (geometric series)
379 set to 25 to enhance the MCMC search capabilities for this high-dimensional parameter
380 space. To further increase the acceptance rate of proposals, we decreased the default jump
381 rate of DREAM_(ZS) by a factor of four. Of course, we could have tuned the jump rate
382 automatically in DREAM_(ZS) to achieve a certain desired acceptance rate of proposals but
383 choose this simpler approach. Convergence of the sampled Markov chains was monitored
384 using the potential scale reduction factor, \hat{R} [Gelman and Rubin, 1992]. This statistic
385 compares for each parameter of interest the average within-chain variance to the variance
386 of all the chains mixed together. The closer the values of these two variances, the closer
387 to one the value of the \hat{R} diagnostic. Values of \hat{R} smaller than 1.2 are commonly deemed
388 to indicate convergence to a limiting distribution. Our simulation results indicate that
389 convergence is achieved after about 400,000 forward model evaluations (FEs) (shown
390 later). Visual inspection of the sampled likelihood values suggests however, that far fewer
391 model evaluations are needed for every chain to locate the posterior distribution.

392 Figure 4a presents the evolution of the \hat{R} statistic calculated from the last 90% of the
393 samples in each chain, and Figure 4b depicts a trace plot of the sampled RMSE values
394 for each of the eight Markov chains. The average acceptance rate (AR) is about 33.3 %
395 (not shown). All chains appear to sample stable RMSE (and thus likelihood) values after
396 approximately 40,000 FEs (Figure 4b). However, another 360,000 FEs are required before

397 all \hat{R} values are smaller than 1.2 and official convergence can be declared (Figure 4a). It
398 is not surprising that the sampled RMSE values stabilize much faster than the associated
399 values of the \hat{R} diagnostic. The chains converge rapidly to a point in the posterior but
400 many more function evaluations are required to fully explore this distribution and satisfy
401 requirements for convergence.

402 Marginal distributions of σ_e , the geostatistical parameters and r_1 and r_{250} are depicted
403 in Figure 5 using kernel density smoothing. The prior distribution is also shown. The
404 standard deviation of the residuals, σ_e , and the field mean of the log conductivity, m ,
405 appear very well resolved. Despite a 40-times dimensionality reduction, the posterior
406 distributions of the geostatistical parameters contain their true values used to create the
407 reference conductivity field. The posterior modes are somewhat removed from the true
408 values, especially for A and to a lesser extent ν . This is due to measurement errors and the
409 use of a reduced parameter space which inevitably introduces some bias in the sampled
410 posterior distribution.

411 Figure 6 displays the reference conductivity field and eight randomly chosen samples
412 from the posterior distribution. The posterior conductivity fields differ substantially from
413 each other, yet all of them produce simulation results that are in (statistical) agreement
414 with the observed data. The geostatistical properties of realization V are relatively similar
415 to those of the reference field. The other posterior fields include fairly different spatial
416 statistics (e.g. realizations III, IV and VI).

417 To investigate the bias introduced by the dimensionality reduction in the posterior es-
418 timates we would need to compare our results for the DREAM_(ZS) trial with 250 DR vari-
419 ables against those of DREAM_(ZS) for the original parameter space. However, such a sam-

420 pling run is computationally intractable. Instead, we performed a trial with $\text{DREAM}_{(\text{ZS})}$
421 using 1000 DR variables and thus 1007 parameters.

422 The results of this more complex run are fairly similar to those of our trial with 250
423 DR variables. Again, about 30,000 - 40,000 FEs are required to reach stable values of
424 the RMSE (not shown), yet a larger computational budget of approximately 1 million
425 FEs is required for this 1007 dimensional search space before convergence to a limiting
426 distribution can be officially declared (not shown). This is more than twice the number of
427 FEs needed for the previous trial with 250 DR variables, but arguably rather efficient con-
428 sidering the about fourfold increase in parameter dimensionality. The AR of $\text{DREAM}_{(\text{ZS})}$
429 is rather large (45%) but results in a good mixing of the individual chains. Perhaps most
430 importantly, the marginal distributions of the geostatistical parameters plotted in Figures
431 5b-g (dashed red line) are in good agreement with their counterparts derived from the
432 $\text{DREAM}_{(\text{ZS})}$ trial with 250 DR variables (solid blue line). For the trial with 1000 DR
433 variables as well, the marginal distributions include the true values used to generate the
434 reference log-conductivity field. Note though that the distributions of most of the geo-
435 statistical parameters have become somewhat more peaky, most noticeably for A . The
436 distributions of m , v , I_M and ν are also more centered on their true values.

437 Altogether we conclude that the posterior distribution of the geostatistical parameters
438 is only weakly affected by dimensionality reduction, let alone the maximum a posteriori
439 (MAP) values which are better resolved as dimensionality reduction decreases.

440 The field realizations resulting from the 1007-dimensional posterior distribution are
441 in strong qualitative agreement with the counterparts derived from the 257-dimensional

442 posterior distribution, but with less variation in the anisotropy angle, and a slight tendency
443 towards smoother fields (Figure 7).

3.3. Comparison With Other Posterior Sampling Methods

444 3.3.1. Comparison Against Sequential Gibbs Sampling

445 Now that we have discussed the main elements of our inversion methodology we are
446 left with a comparison against state-of-the-art methods in the literature. As a first test,
447 we consider the SGS method for variogram estimation as implemented in the SIPPI 0.94
448 toolbox [*Hansen et al.*, 2013a, b]. This open source MATLAB package is described in
449 detail in the cited references, and interested readers are referred to these publications.
450 We used default settings for the algorithmic variables, and the same prior distribution
451 for σ_e and the geostatistical parameters as used in our numerical experiments described
452 previously. Furthermore, we added to the SIPPI toolbox the Matérn variogram as this
453 function was not yet incorporated in the toolbox.

454 Before proceeding with our results, we would like to emphasize that SGS (or the very
455 similar but independently developed iterative spatial resampling (ISR) scheme by *Mari-*
456 *ethoz et al.* [2010]) is a powerful MCMC algorithm for sampling from complex geologic
457 prior models [e.g, *Mariethoz et al.*, 2010; *Hansen et al.*, 2012]. This method creates
458 candidate points by conditioning a field realization drawn from the prior to a randomly
459 chosen set of points from the current state (and hence model/field) of the Markov chain.
460 Nonetheless, SGS with variogram inference suffers one important drawback and that is
461 that it relies completely on sampling from the prior distribution of the variogram param-
462 eters [see *Hansen et al.*, 2013b, for details]. Moreover, SGS uses a single Markov chain in
463 pursuit of the posterior distribution. This not only makes it difficult to rapidly explore

464 multi-dimensional parameter spaces, but also complicates assessment of convergence, and
465 effective use of multi-processor resources. Pre-fetching [Brockwell, 2006] and multi-try
466 Metropolis sampling [Liu et al., 2000] offer some options for distributed, multi-core im-
467 plementation of single chains. What is more, the use of a single chain increases chances
468 of premature convergence. To mitigate this risk, it is generally recommended to perform
469 several independent trials and verify whether the different chains have converged to the
470 approximate same limiting distribution. In contrast to SGS, DREAM_(ZS) is embarrass-
471 ingly parallel and thus readily amenable to multi-processor distributed computation which
472 should drastically reduce the required CPU-time for posterior exploration. In the case
473 studies presented herein, we ran each of the eight different Markov chains of DREAM_(ZS)
474 on a different processor. This significantly reduced the CPU-time required for posterior
475 exploration, details of which will be presented in section 4 of this paper.

476 For proper convergence assessment, we performed three (independent) SGS trials using
477 starting points drawn randomly from the prior distribution. Because of the associated
478 computational costs, we terminated the calculation after a total of 500,000 FEs (that is,
479 166,667 FEs per Markov chain) and the comparison with our method is made on the
480 basis of the same computational budget of 500,000 FEs. We used default settings of SGS
481 and individually sampled, with equal probability, the different geostatistical parameters
482 and vector of (standard normal) field values. Each iteration produces a candidate log
483 conductivity field as follows. With probability 1/8, either a new 10,000-dimensional vector
484 of standard normal variates, say \mathbf{g} , is produced, or the current geostatistical model is
485 updated by replacing one of the six geostatistical parameters with a random draw from
486 its prior, or a new value of σ_e is sampled from $p(\sigma_e)$. When \mathbf{g} is updated, the candidate

487 model (proposal) is obtained by conditioning on a fraction, ϕ , of locations randomly
488 chosen from the current model. The value of ϕ is adapted during burn-in to achieve a
489 targeted acceptance rate, which we set to 20%. The upper bound of ϕ was set to 1 (i.e.,
490 totally different proposal) whereas its lower bound was fixed to 0.001, that is, only 10 of
491 the 10,000 $\log_{10}(K)$ values are perturbed per iteration.

492 With an average (adapted) value of ϕ reaching its lower bound of 0.001, the mean AR
493 values for each of the eight individual sampling steps are 20.7, 1.0, 7.9, 5.1, 9.4, 4.9, 13.4,
494 and 7.0%, for \mathbf{g} , m , v , I_M , A , R_I , ν , and σ_e , respectively. Both SGS and the proposed
495 sampling method successfully fit the data to the prescribed error level (Figures 8a and 8b),
496 but the SGS method needs somewhat fewer function evaluations to do so. Nevertheless,
497 our proposed inversion approach is more effective and efficient in exploring the posterior
498 target, as shown by the respective Markov chain trajectories of DREAM_(ZS) and SGS for
499 the field variance (Figures 8c and 8d) and integral scale along the major anisotropy axis
500 (Figures 8e and 8f). Even after a total of 500,000 FEs or 166,667 FEs per chain, the
501 three SGS trials do not converge appropriately to the true value of I_M . The DREAM_(ZS)
502 algorithm on the contrary needs about 200,000 FEs (that is, 25,000 FEs per chain) to
503 converge to the reference value of 0.67. Moreover, SGS has difficulty in sampling the
504 correct value of v as well. Two of the three chains are somewhat stuck near its lower
505 bound.

506 A convergence check of the three chains sampled by SGS is provided in Figure 9a. For
507 convenience we plot only the evolution of the \hat{R} statistic of σ_e and the six variogram
508 parameters. In practice, SGS samples 10,007 parameters, and hence convergence can
509 only be formally declared if all sampled parameters fall below the threshold value of

510 1.2 (horizontal black line). Nevertheless, it is evident that SGS is unable to converge
 511 adequately within the allowed computational budget. Even after 500,000 FEs several of
 512 the plotted trajectories remain well above the \hat{R} -threshold value of 1.2. As a consequence,
 513 the corresponding posterior distribution is quite inaccurate. The posterior mode of v is
 514 close to its lower bound, well removed from the reference value (Figure 9b). What is
 515 more, a multimodal posterior distribution is observed for I_M with true value that falls in
 516 a region with lower posterior probability (Figure 9c). Finally, the posterior fields sampled
 517 by SGS exhibit considerably more correlation at different spatial lags than its counterpart
 518 derived from our proposed approach (Figure 9d).

519 3.3.2. Comparison Against the Method of Anchored Distributions

520 The MAD method [see *Rubin et al.*, 2010, for an extensive description] is especially de-
 521 signed for inference of (multi-)Gaussian parameter fields. It differs from classical Bayesian
 522 inference methods in the treatment of the likelihood function, $L(\boldsymbol{\theta}|\mathbf{d}) \equiv p(\mathbf{d}|\boldsymbol{\theta})$. Whereas
 523 SGS and our proposed inversion method describe $p(\mathbf{d}|\boldsymbol{\theta})$ as a parametric probability dis-
 524 tribution of the residuals (e.g, equation (8)) that is specified a-priori, MAD takes $p(\mathbf{d}|\boldsymbol{\theta})$
 525 as the conditional probability density of the simulated data given a parameter set $\boldsymbol{\theta}$. This
 526 is done by approximating $p(\mathbf{d}|\boldsymbol{\theta})$ from an ensemble of conditional simulations, using a
 527 nonparametric approach whenever computationally tractable. This method has several
 528 advantages, one of which is that it avoids making strong and sometimes unjustified as-
 529 sumptions about the properties of the residual errors. In practice, however, and because
 530 of computational constraints, at least some parametric (Gaussian) assumptions often need
 531 to be made about $p(\mathbf{d}|\boldsymbol{\theta})$ a-priori [e.g., *Murakami et al.*, 2010; *Over et al.*, 2015]. Another
 532 feature of MAD is that it uses basic Monte Carlo simulation to solve for the posterior

parameter distribution. Once the anchor locations have been defined, the inferred parameters are drawn randomly from their (marginal) prior distribution for a pre-specified number of times. This is not very efficient, especially if the posterior distribution constitutes only a small part of the prior distribution. Hence, even with the assumption of a Gaussian distribution for $p(\mathbf{d}|\boldsymbol{\theta})$, the total number of forward model evaluations required by MAD will typically be on the order of several millions [e.g., *Murakami et al.*, 2010; *Over et al.*, 2015]. This requires the use of many processors on a distributed computing network [on the order of several thousands, e.g., *Murakami et al.*, 2010].

MAD distinguishes between two types of inferred variables: variogram parameters, $\boldsymbol{\theta}_V$ and conductivity values at selected locations or anchor sets, $\boldsymbol{\theta}_K$. If no direct point conductivity measurements are considered in the inference, the posterior distribution $p(\boldsymbol{\theta}_V, \boldsymbol{\theta}_K|\mathbf{d})$ reduces to

$$p(\boldsymbol{\theta}_V, \boldsymbol{\theta}_K|\mathbf{d}) \propto p(\boldsymbol{\theta}_V) p(\boldsymbol{\theta}_K|\boldsymbol{\theta}_V) p(\mathbf{d}|\boldsymbol{\theta}_V, \boldsymbol{\theta}_K), \quad (10)$$

where $p(\boldsymbol{\theta}_V)$ denotes the prior distribution of the variogram parameters, $p(\boldsymbol{\theta}_K|\boldsymbol{\theta}_V)$ signifies the prior anchor distribution given a variogram parameter vector $\boldsymbol{\theta}_V$, and $p(\mathbf{d}|\boldsymbol{\theta}_V, \boldsymbol{\theta}_K)$ is the likelihood function of $\{\boldsymbol{\theta}_V, \boldsymbol{\theta}_K\}$. While $p(\boldsymbol{\theta}_V)$ and $p(\boldsymbol{\theta}_K|\boldsymbol{\theta}_V)$ can be derived analytically, numerical estimation of $p(\mathbf{d}|\boldsymbol{\theta}_V, \boldsymbol{\theta}_K)$ is a complicated task. MAD proceeds as follows. First, define the anchor locations. Then sample $p(\boldsymbol{\theta}_V)$ n_V times and for each of the resulting $\boldsymbol{\theta}_V$ vectors, sample $p(\boldsymbol{\theta}_K|\boldsymbol{\theta}_V)$ n_K times. Third, create an ensemble of n_f random fields, $\boldsymbol{\Psi}$, from each of the $n_V \times n_K$ $\{\boldsymbol{\theta}_V, \boldsymbol{\theta}_K\}$ parameter sets using a direct conditioning method such as the one described in Appendix A. The n_f fields of $\boldsymbol{\Psi}$ thus exactly honor $\boldsymbol{\theta}_K$ and are distributed according to $\boldsymbol{\theta}_V$, and result into n_f simulated data vectors, $\boldsymbol{\Delta}$. Finally, the likelihood, $p(\mathbf{d}|\boldsymbol{\theta}_V, \boldsymbol{\theta}_K)$, of the parameter set under evaluation,

556 $\{\boldsymbol{\theta}_V, \boldsymbol{\theta}_K\}$, is approximated by fitting a multivariate density distribution to the multi-
 557 variate frequency distribution of the simulated data stored in $\mathbf{\Delta}$. For this, one can use
 558 non-parametric kernel density estimation (possibly after reduction of the data vector) or
 559 assume a Gaussian parametric model for \mathbf{d} . The total number of forward model calls is
 560 therefore equal to $n_V \times n_K \times n_f$.

561 We used 49 anchors on a regular grid (Figure 10a) and set n_V to 500, n_K to 12 and n_f to
 562 100, leading to a total computational budget of 600,000 FEs. The 500 samples from $p(\boldsymbol{\theta}_V)$
 563 were drawn randomly using Latin hypercube sampling. Our values for n_V , n_K and n_f are
 564 based on the work of *Murakami et al.* [2010] who used 44 anchor locations, $n_V = 3000$,
 565 $n_K = 12$ and $n_f = 250$ for a grid of approximatively similar dimension as in our numerical
 566 experiments herein, but with a much smaller prior variogram uncertainty. Obviously, the
 567 best choice for the MAD algorithmic parameters is problem-dependent and our settings
 568 might not be optimal. We would like to stress, however, that a total of 600,000 FEs for
 569 MAD is justified. Indeed, a computational budget of only 500,000 FEs was assigned to
 570 our proposed inversion approach (section 3.2) and the SGS method (section 3.3.1).

571 Figure 10b presents a histogram of the RMSE values derived from the 600,000 forward
 572 model calls. The minimum RMSE sampled by MAD is 0.051 kg m^{-3} (vertical black
 573 line) which is not only significantly larger than the true value of 0.039 kg m^{-3} for the
 574 reference field but also outside the posterior distribution of RMSE values sampled by
 575 our approach and SGS (see Figures 4b and 8a-b). No matter how the likelihoods of the
 576 $i = 1, \dots, n_V \times n_K \{\boldsymbol{\theta}_V, \boldsymbol{\theta}_K\}_i$ parameter sets are estimated, inference from these 600,000
 577 forward runs can thus only be flawed. Obviously, appropriate (random) sampling of the
 578 prior parameter space would require much larger values of n_V and n_K .

4. Discussion

579 Some remarks about the presented study are in order. Due to time constraints, the
580 different MCMC and MAD trials were only performed once. Repeated sampling runs
581 with different random seeds would provide a more accurate benchmark of our inversion
582 methodology. Nevertheless, the results presented herein inspire confidence in the effec-
583 tiveness and efficiency of our proposed Bayesian inversion method.

584 The computational requirement is an important issue. Considering a serial calculation
585 framework, we find that the serial implementation of our approach outperforms both SGS
586 and MAD for the case study considered herein. The DREAM_(ZS) sampler is however
587 designed specifically so that it is embarrassingly parallel and thus can take maximum
588 advantage of multi-processor resources. We did so herein using an 8-core workstation,
589 assigning each of the eight interacting Markov chains to a different processor. This resulted
590 in a six times speed up of the calculations. The advantages of the proposed multi-core
591 approach are hence evident. Note also that if more parallel cores are available, then the
592 search efficiency can be further increased by using the multi-try variant of DREAM_(ZS)
593 [*Laloy and Vrugt, 2012*].

594 Another crucial point is the inevitable trade-off between model truncation (dimension-
595 ality reduction) and the accuracy of the posterior field realizations. The MCMC search is
596 performed within a truncated model space that constitutes only a subset of the original
597 model space. By construction, not all possible posterior models can therefore be repre-
598 sented by the truncated posterior pdf. Model truncation might also bias the posterior
599 distribution by shifting the probability mass away from the original MAP values. In ad-
600 dition, for a given truncation level the peakedness of the likelihood function will influence

601 the quality of approximation of the original posterior distribution by the truncated pos-
602 terior distribution. A dimensionality reduction with a factor 40 was shown to work well
603 for the considered case study. The resulting posterior distribution was found to be in
604 good agreement with the distribution stemming from a dimensionality reduction with a
605 factor 10. As model space truncation or peakedness of the likelihood further increases, the
606 distribution and associated parameter uncertainties will nevertheless be increasingly cor-
607 rupted. The combined effects of dimensionality reduction and measurement data quality
608 on the accuracy of the estimated target distribution deserve further analysis.

609 The *Gelman and Rubin* [1992] potential scale reduction factor was computed using the
610 last 90% of the generated samples in each Markov chain evolved by $\text{DREAM}_{(ZS)}$. This
611 value differs from the default of 50% used in $\text{DREAM}_{(ZS)}$, but is warranted in each of our
612 case studies because the joint chains converge to stable RMSE values within less than 10%
613 of the assigned computational budget. Indeed, the sampled RMSE (and thus likelihood)
614 values appropriately converge within 30,000 - 40,000 FEs. The use of 90% of the chains
615 is equivalent to a burn-in of 50,000 (trial with 250 DR variables) and 100,000 (trial with
616 1000 DR variables) samples, well beyond when the posterior distribution has been located.

617 This study considers a two-dimensional flow and transport modeling domain. Extension
618 of the proposed approach to 3D domains is straightforward and will be investigated in
619 future work. Extension to pluri-Gaussian simulation [e.g., *Lantuéjoul et al.*, 2002] for
620 inference of categorical conductivity fields also seems promising.

5. Conclusions

621 This paper presents a novel Bayesian inversion scheme for the simultaneous estima-
622 tion of high-dimensional multi-Gaussian conductivity fields and associated geostatistical

623 properties from indirect hydrological data. Our method merges Gaussian process gener-
624 ation via circulant embedding [*Dietrich and Newsam, 1997*] to decouple the variogram
625 from grid cell specific values, with dimensionality reduction by interpolation to facili-
626 tate Markov chain Monte Carlo (MCMC) simulation with the DREAM_(ZS) algorithm [*ter*
627 *Braak and Vrugt, 2008; Vrugt et al., 2009; Laloy and Vrugt, 2012*]. We use the Matérn
628 variogram model to infer the conductivity values simultaneously with field smoothness (or
629 Matérn shape parameter) and other geostatistical parameters (mean, sill, integral scales
630 and anisotropy factor(s)). The proposed dimensionality reduction approach systematically
631 honors the prescribed variogram and is shown to outperform the Karhunen-Loève [*Loève,*
632 *1977*] transform. Our inverse method is demonstrated using synthetic, error corrupted,
633 data from a flow and transport model involving a fairly heterogeneous 10,000-dimensional
634 multi-Gaussian conductivity field. Despite a reduction of the parameter space by a factor
635 of 40, the measurement data were fitted to the prescribed noise level while the derived
636 posterior parameter distributions always included the true geostatistical parameter values.
637 A comparison between the posterior distributions derived for a 257- (40-times reduction)
638 and 1007-dimensional (10-times reduction) parameter space, respectively, indicated that
639 the bias introduced by the dimensionality reduction in the posterior estimates is rather
640 small. This inspires confidence in the effectiveness of the approach. The posterior field
641 uncertainty encompassed a large range of different geostatistical models, which calls into
642 question the common practice in hydrogeology of fixing the variogram model before inver-
643 sion. For the considered case study, the serial version of our method appears to be more
644 computationally efficient than both the SGS algorithm of *Hansen et al.* [2012, 2013a, b]
645 and the MAD method of *Rubin et al.* [2010]. The advantages of the proposed approach

are even more apparent when executed on a distributed computing network. A 6-times
reduction in CPU-time was observed with DREAM_(ZS) using parallel evaluation of the
eight different Markov chains. Future work will investigate the application of the pro-
posed approach to 3D modeling domains, and to pluri-Gaussian simulations for inference
of categorical field structures.

Appendix A

Conditioning an unconditional simulation of a random field can be easily performed via
kriging. Kriging-based geostatistical methods are extensively described in the literature
[e.g., *Chilès and Delfiner, 1999; Ren et al., 2005; Huang et al., 2011*].

If the values of $Y(\mathbf{x})$ are known at locations $\mathbf{x}_i, i = 1 \cdots n_y$, the value of $Y(\mathbf{x})$ at any
arbitrary location \mathbf{x} , $Y_{\text{kr}}(\mathbf{x})$, can be predicted unbiasedly as

$$Y_{\text{kr}}(\mathbf{x}) = \sum_{i=1}^{n_y} \lambda_i Y(\mathbf{x}_i) \quad (\text{A1})$$

in which the kriging weights λ_i depend solely on the prescribed variogram. Now if we
have an unconditional realization, $Y_{\text{uc}}(\mathbf{x})$, the corresponding random field conditioned to
the n_y observed values, Y_{c} , is given by

$$Y_{\text{c}}(\mathbf{x}) = Y_{\text{uc}}(\mathbf{x}) + [Y_{\text{kr}}(\mathbf{x}) - Y_{\text{kr-u}}(\mathbf{x})], \quad (\text{A2})$$

where $Y_{\text{kr-u}}$ is obtained by kriging (equation (A1)) using the unconditional simulated
values at the n_y data locations, and the same kriging weights are used for determining
 $Y_{\text{kr}}(\mathbf{x})$ and $Y_{\text{kr-u}}(\mathbf{x})$. Equation (A2) implies that, at each conditioning data location,
the unconditional simulated value is taken out and replaced by the conditioning datum.
In the vicinity of a conditioning data location, the kriging operator smooths the change
between the conditioning data and the unconditional simulated values outside the range

667 of kriged values. The conditioning is therefore exact at data locations whereas beyond
 668 the correction range, the conditional simulated values will be the unconditional simulated
 669 values.

670 Using matrix notation, the λ_i in equation (A1) are obtained for a random field with
 671 dimension $m \times n$ as

$$672 \quad \boldsymbol{\lambda} = \mathbf{C}_{\text{df}}^T \mathbf{C}_{\text{dd}}^{-1}, \quad (\text{A3})$$

673 where \mathbf{C}_{df} is the $n_y \times (m \times n)$ matrix of covariances between data and target field values,
 674 \mathbf{C}_{dd} is the data-to-data covariance matrix, and the superscripts T and -1 denote transpose
 675 and inverse matrix operations, respectively.

676 **Acknowledgments.**

677 We would like to thank the Associate Editor Olaf Cirpka and three anonymous review-
 678 ers for their useful comments and suggestions which significantly helped to improve the
 679 manuscript. We are grateful to Thomas Mejer Hansen and coworkers for sharing online
 680 their mGstat and SIPPI toolboxes. We also like to thank Rouven Künze for providing us
 681 with the MaFloT simulator. A MATLAB code of the approach proposed in this study
 682 is available from the first author (elaloy@sckcen.be). The general-purpose DREAM_(ZS)
 683 algorithm is available from the fourth author (jasper@uci.edu).

References

- 684 Box, G. E. P., and G. C. Tiao (1973), *Bayesian Inference in Statistical Analysis*, Addison-
 685 Wesley.
- 686 Brockwell, A. E. (2006), Parallel Markov chain Monte Carlo simulation by pre-fetching,
 687 *J. Comp. Graph. Stat.*, 15, 246–261, doi:10.1198/106186006X100579.

- 688 Chan G., and A.T.A. Wood (1997), Algorithm AS 312: An algorithm for simulating
689 stationary Gaussian random fields, *J. Roy. Stat. Soc. C-App.*, *46*, 171–181.
- 690 Chilès, J.-P., and P. Delfiner (1999), *Geostatistics: Modeling Spatial Uncertainty*, Wiley.
- 691 Dietrich, C. R., and G. H. Newsam (1997), Fast and exact simulation of stationary Gaus-
692 sian processes through circulant embedding of the covariance matrix, *SIAM J. Sci.*
693 *Comput.*, *18*, 1088–1107.
- 694 Evensen, G. (2003), The ensemble Kalman filter: Theoretical formulation and practical
695 implementation, *Ocean Dyn.*, *53*, 343–367.
- 696 Fu, J., and J. J. Gómez-Hernández (2009), A blocking Markov chain Monte Carlo
697 method for inverse stochastic hydrogeological modeling. *Math. Geosci.*, *41*, 105–128.
698 doi:10.1007/s11004-008-9206-0
- 699 Gelman, A. G., and D. B. Rubin (1992), Inference from iterative simulation using multiple
700 sequences, *Stat. Sci.*, *7*, 457–472.
- 701 Hansen, T. M., K. C. Cordua, and K. Mosegaard (2012), Inverse problems with non-
702 trivial priors: efficient solution through sequential Gibbs sampling, *Comput. Geosci.*,
703 *16*, 593–611, doi:10.1007/s10596-011-9271-1.
- 704 Hansen, T. M., K. C. Cordua, M. C. Looms, and K. Mosegaard (2013a), SIPPI: a Matlab
705 toolbox for sampling the solution to inverse problems with complex prior information:
706 Part 1 - Methodology, *Comput. Geosci.*, *52*, 470–480, doi:10.1016/j.cageo.2012.09.004.
- 707 Hansen, T. M., K. C. Cordua, M. C. Looms, and K. Mosegaard (2013b), SIPPI: a Matlab
708 toolbox for sampling the solution to inverse problems with complex prior information:
709 Part 2 - Application to crosshole GPR tomography, *Comput. Geosci.*, *52*, 481–492,
710 doi:10.1016/j.cageo.2012.10.001.

- 711 Hendricks-Franssen, H.J., A. Alcolea, M. Riva, M. Bakr, N. van der Wiel, F. Stauffer,
712 and A. Guadagnini (2009), A comparison of seven methods for the inverse modelling of
713 groundwater flow: application to the characterisation of well catchments. *Adv. Water*
714 *Resour.*, *32*(6), 851–72, doi:10.1016/j.advwatres.2009.02.011.
- 715 Huang, J. W., G. Bellefleur, and B. Milkereit (2010), CSimMDMV: A parallel program for
716 stochastic characterization of multi-dimensional, multi-variant, and multi-scale distri-
717 bution of heterogeneous reservoir rock properties from well log data, *Comput. Geosci.*,
718 *37*, 1763–1776, doi:10.1016/j.cageo.2010.11.012.
- 719 Jafarpour B., and M. Tarrahi (2011), Assessing the performance of the ensemble Kalman
720 filter for subsurface flow data integration under variogram uncertainty, *Water Resour.*
721 *Res.*, *47*(5), doi:10.1029/2010WR009090.
- 722 Jardani, A., J. P. Dupont, A. Revil, M. Massei, M. Fournier, and B. Laignel (2012),
723 Geostatistical inverse modeling of the transmissivity field of a heterogeneous alluvial
724 aquifer under tidal influence, *J. Hydrol.*, *472*, 287–300.
- 725 Jeffreys, H. (1946), An invariant form for the prior probability in estimation problems,
726 *Proc. R. Soc. London*, *A186*, 453–461.
- 727 Kitanidis, P. (1995), Quasi-linear geostatistical theory for inversing, *Water Resour. Res.*,
728 *31*(10), 2411–9, doi:10.1029/95WR01945.
- 729 Kroese, D. P., and Z. I. Botev (2013), Spatial Process Generation, in *Lectures on Stochastic*
730 *Geometry, Spatial Statistics and Random Fields, Volume II: Analysis, Modeling and*
731 *Simulation of Complex Structures*, edited by V. Schmidt, Springer-Verlag.
- 732 Künze R., and I. Lunati (2012), An adaptive multiscale method for density-driven insta-
733 bilities, *J. Comput. Phys.*, *231*, 5557–5570.

- 734 Laloy, E., and J. A. Vrugt (2012), High-dimensional posterior exploration of hydrologic
735 models using multiple-try DREAM_(ZS) and high performance computing, *Water Resour.*
736 *Res.*, *48*, W01526, doi:10.1029/2011WR010608.
- 737 Laloy, E., N. Linde, and J. A. Vrugt (2012), Mass conservative three-dimensional water
738 tracer distribution from Markov chain Monte Carlo inversion of time-lapse ground-
739 penetrating radar data, *Water Resour. Res.*, *48*, W07510, doi:10.1029/2011WR011238.
- 740 Laloy, E., B. Rogiers, J. A. Vrugt, D. Mallants, and D. Jacques (2013), Efficient posterior
741 exploration of a high-dimensional groundwater model from two-stage Markov Chain
742 Monte Carlo simulation and polynomial chaos expansion, *Water Resour. Res.*, *49*, 2664–
743 2682, doi:10.1002/wrcr.20226.
- 744 Laloy, E., J. A. Huisman, and D. Jacques (2014), High-resolution moisture profiles
745 from full-waveform probabilistic inversion of TDR signals, *J. Hydrol.*, *519*, 2121–2135,
746 doi:10.1016/j.jhydrol.2014.10.005.
- 747 Lantuéjoul, C. (2002), *Geostatistical simulation: models and algorithms*, Springer.
- 748 Le Ravalec, M., B. Noetinger, B., and L. Y. Hu (2000), The FFT moving average (FFT-
749 MA) generator: an efficient numerical method for generating and conditioning Gaussian
750 simulations, *Math. Geol.*, *32*(6), 701–723.
- 751 Li, W. and O. A. Cirpka (2006), Efficient geostatistical inverse methods for structured
752 and unstructured grids, *Water Resour. Res.*, *42*, W06402, doi:10.1029/2005WR004668.
- 753 Linde, N., and J. A. Vrugt (2013), Distributed soil moisture from crosshole ground-
754 penetrating radar travel times using stochastic inversion, *Vadose Zone J.*, *12*(1),
755 doi:10.2136/vzj2012.0101.

- 756 Liu, J. S., F. Liang, and W. H. Wong (2000), The multiple-try method and lo-
757 cal optimization in Metropolis sampling, *J. Am. Stat. Assoc.*, *95*(449), 121–134,
758 doi:10.2307/2669532.
- 759 Lochbühler, T., S. J. Breen. R. L. Detwiler, J. A. Vrugt, and N. Linde (2014), Probabilistic
760 electrical resistivity tomography for a CO2 sequestration analog, *J. Appl. Geophys.*, *107*,
761 80–92, doi:10.1016/j.jappgeo.2014.05.013.
- 762 Lochbühler, T., J. A. Vrugt, M. Sadegh, and N. Linde (2015), Summary statistics from
763 training images as prior information in probabilistic inversion, *Geophys. J. Int.*, *201*,
764 157–171. doi:10.1093/gji/ggv008.
- 765 Loève, M. (1977), *Probability Theory, Fourth edition*, Springer-Verlag, New York.
- 766 Mariethoz, G., P. Renard, and J. Caers (2010), Bayesian inverse problem and op-
767 timization with iterative spatial resampling, *Water Resour. Res.*, *46*, W11530,
768 doi:10.1029/2010WR009274.
- 769 Matérn, B. (1960), Spatial Variation, *Meddelanden fran Statens Skogsforskningsinstitut*,
770 *495* (second ed., 1986, Lecture notes in Statistics, no. 36, Springer).
- 771 Murakami, H., X. Chen, M. S. Hahn, Y. Liu, M. L. Rockhold, V. R. Vermeul, J.
772 M. Zachara, and Y. Rubin (2010), Bayesian approach for three-dimensional aquifer
773 characterization at the Hanford 300 Area, *Hydrol. Earth Syst. Sci.*, *14*, 19892001,
774 doi:10.5194/hess-141989-2010.
- 775 Nowak, W., F. P. J. de Barros, and Y. Rubin (2010), Bayesian geostatistical design:
776 Task-driven optimal site investigation when the geostatistical model is uncertain, *Water*
777 *Resour. Res.*, *46*, W03535, doi:10.1029/2009WR008312.

- 778 Ortiz, J. C, and C. L. Deutsch (2002), Calculation of uncertainty in the variogram, *Math.*
779 *Geol.*, *34*(2), 169–183.
- 780 Over, M., W., U. Wollschläger, C. A. Osorio-Murillo, and Y. Rubin (2015),
781 Bayesian inversion of Mualem-van Genuchten parameters in a multi-layer soil pro-
782 file: a data-driven, assumption-free likelihood function, *Water Resour. Res.*, *51*,
783 doi:10.1002/2014WR015252.
- 784 Pardo-Iguzquiza, E., and M. Chica-Olmo (2008), Geostatistics with the Matern semivar-
785 iogram model: A library of computer programs for inference, kriging and simulation.
786 *Comput. Geosci.*, *34*(9), 1073–1079.
- 787 Ren, W., L. Cunha, and C. V. Deutsch (2005), Preservation of multiple point structure
788 when conditioning by kriging, in *Geostatistics Banff 2004*, Quantitative Geology and
789 Geostatistics Vol. 4, edited by O. Leuangthong and C.V. Deutsch, Springer.
- 790 Robert, C. P., and G. Casella (2004), *Monte Carlo statistical methods*, second edition.
791 Springer.
- 792 Rosas-Carbajal, M., N. Linde, T. Kalscheuer, and J. A. Vrugt (2014), Two-dimensional
793 probabilistic inversion of plane-wave electromagnetic data: Methodology, model con-
794 straints and joint inversion with electrical resistivity data, *Geophys. J. Int.*, *196*, 1508–
795 1524, doi: 10.1093/gji/ggt482.
- 796 Rubin, Y. (2003), *Applied Stochastic Hydrology*, Oxford University Press.
- 797 Rubin, Y., X. Chen, H. Murakami, and M. Hahn (2010), A Bayesian approach for inverse
798 modeling, data assimilation, and conditional simulation of spatial random fields, *Water*
799 *Resour. Res.*, *46*, W10523, doi:10.1029/2009WR008799.

- 800 ter Braak, C. J. F., and J. A. Vrugt (2008), Differential evolution Markov chain with
801 snooker updater and fewer chains, *Stat. Comput.*, *18*, 435–46, doi:10.1007/s11222-008-
802 9104-9.
- 803 Vrugt, J. A., C. J. F. ter Braak, C. G. H. Diks, D. Higdon, B. A. Robinson, J. M. Hyman
804 (2009), Accelerating Markov chain Monte Carlo simulation by differential evolution with
805 self-adaptive randomized subspace sampling., *Int. J. Nonlinear Sci. and Numer. Simul.*,
806 *10*, 273–290.
- 807 Whittle, P. (1954), On stationary processes in the plane, *Biometrika*, *41*, 439–449.
- 808 Zhang, D., and Z. Lu (2004), An efficient, high-order perturbation approach for flow
809 in random porous media via Karhunen-Loève and polynomial expansions, *J. Comput.*
810 *Phys.*, *194*(2), 773–794.
- 811 Zhang, Z., and Y. Rubin (2008), MAD: A new method for inverse modeling of spatial
812 random fields with applications in hydrogeology, *Eos Trans. AGU*, *89*(53), Fall Meet.
813 Suppl., Abstract H44C07.
- 814 Zhou, H., J. Gómez-Hernández, and L. Liangping (2014), Inverse methods in
815 hydrogeology: evolution and recent trends, *Adv. Water Resour.*, *63*, 22–37,
816 doi:10.1016/j.advwatres.2013.10.014.

Table 1. Bounds of the Jeffreys (J), uniform (U) and standard normal (N) prior distributions used in our case study. The last column lists the true values of σ_e and the geostatistical parameters.

Parameter	Units	Prior	Prior range	True value
σ_e	kg m^{-3}	J	0.01 - 0.3	0.039
m	-	U	(-4) - (-2)	-3
ν	-	J	0.5 - 2	1
I_M	m	U	0.2 - 2	0.67
A	degree	U	60 - 120	75
R_I	-	U	0.1 - 0.5	0.25
ν	-	J	0.1 - 5	0.5
\mathbf{r}	-	N	-	-

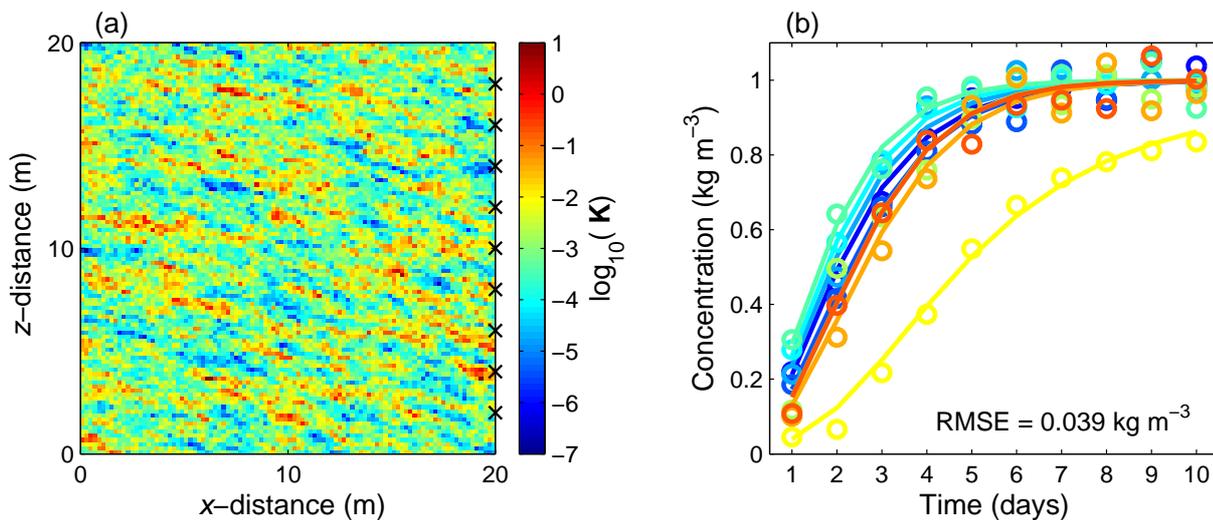


Figure 1. (a) Reference log conductivity field and (b) simulated transport data (b) used in the inversions. Each black cross in the right-hand side borehole denotes a sampling location. Each color in the right plot corresponds to a different sampling location (black cross) in the right-hand side borehole. Lines represent the uncorrupted data and circles signify the noise-contaminated data that are used as measurements for the inversion. The noise level is 0.039 kg m^{-3} .

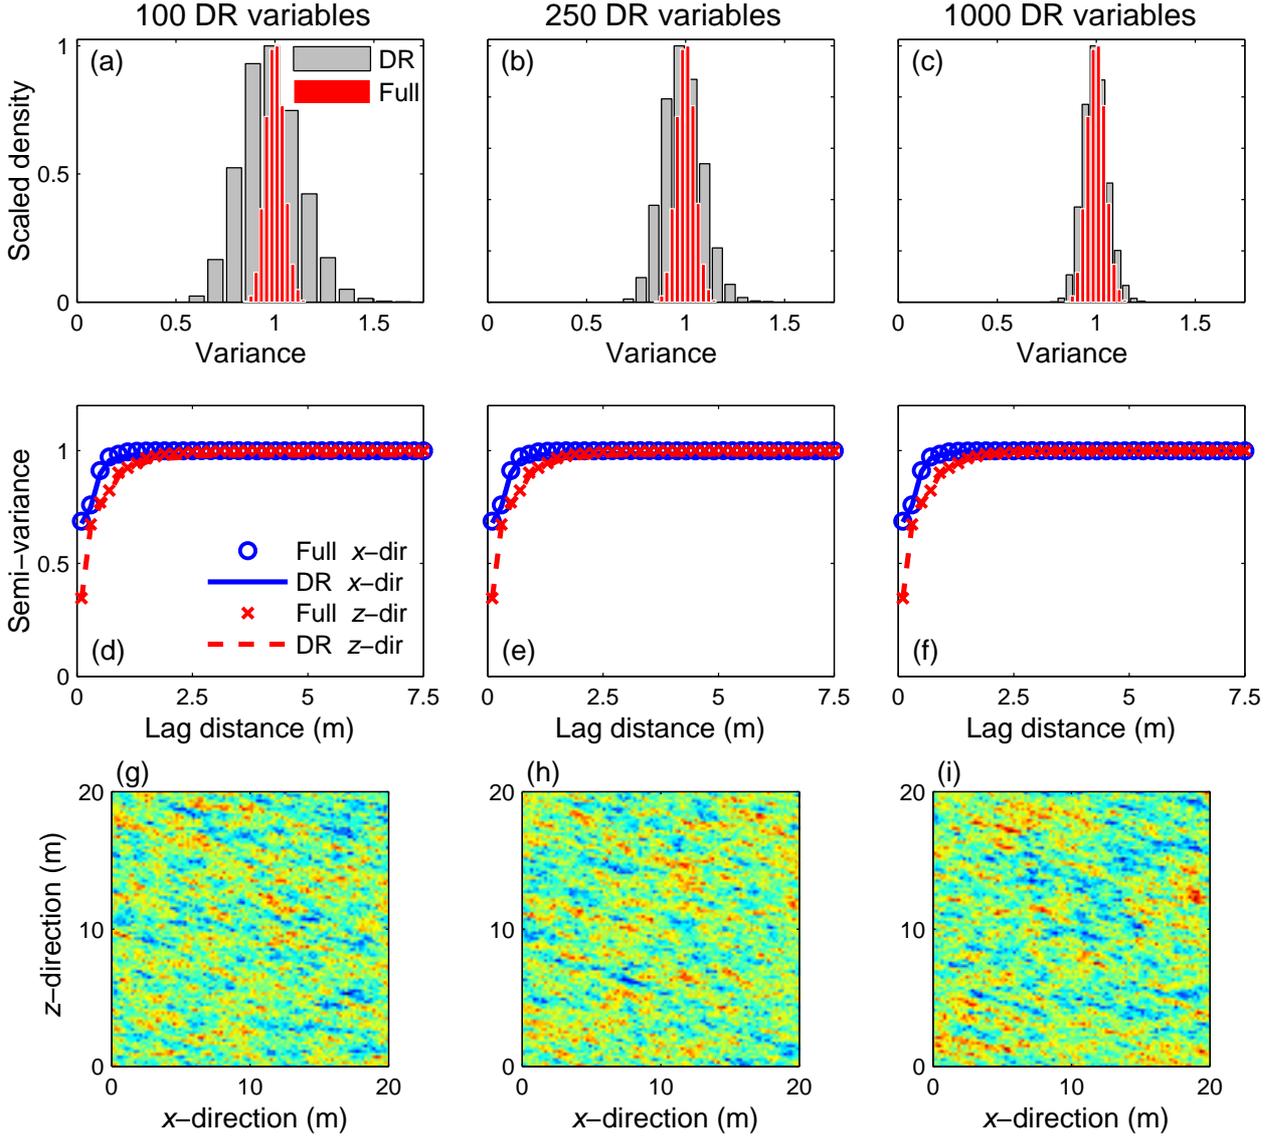


Figure 2. Grid point variance distribution (subplots a, b, c) and experimental variograms in both the horizontal (x -dir) and vertical (z -dir) directions (subplots d, e, f) calculated from 1000 realizations of an autocorrelated multi-Gaussian field with zero-mean and the reference variogram used for the inversions (see last of column Table 1 and Figure 1a), using (a,d) 100-dimensional, (b-e) 250-dimensional and (c-f) 1000-dimensional $\mathbf{r} = [\mathbf{r}_1, \mathbf{r}_2]$ vectors, and randomly chosen field realizations (subplots g-i) derived with (g) 100-dimensional, (h) 250-dimensional and (i) 1000-dimensional $\mathbf{r} = [\mathbf{r}_1, \mathbf{r}_2]$ vectors. The gray and red bins in subplots a, b, and c denote the variance distributions obtained from our dimensionality reduction approach and from directly generating autocorrelated multi-Gaussian fields with the prescribed variogram, respectively. The red histograms in subplots a, b and c thus represent the natural statistical fluctuations. The (plain and dashed) lines and (circle and cross) symbols in subplots d-f signify average experimental variogram values in the x - (blue color) or z -direction (red color) for each of the considered lags, derived from our dimensionality reduction approach (DR) and from directly generating autocorrelated multi-Gaussian fields with prescribed variogram (Full), respectively.

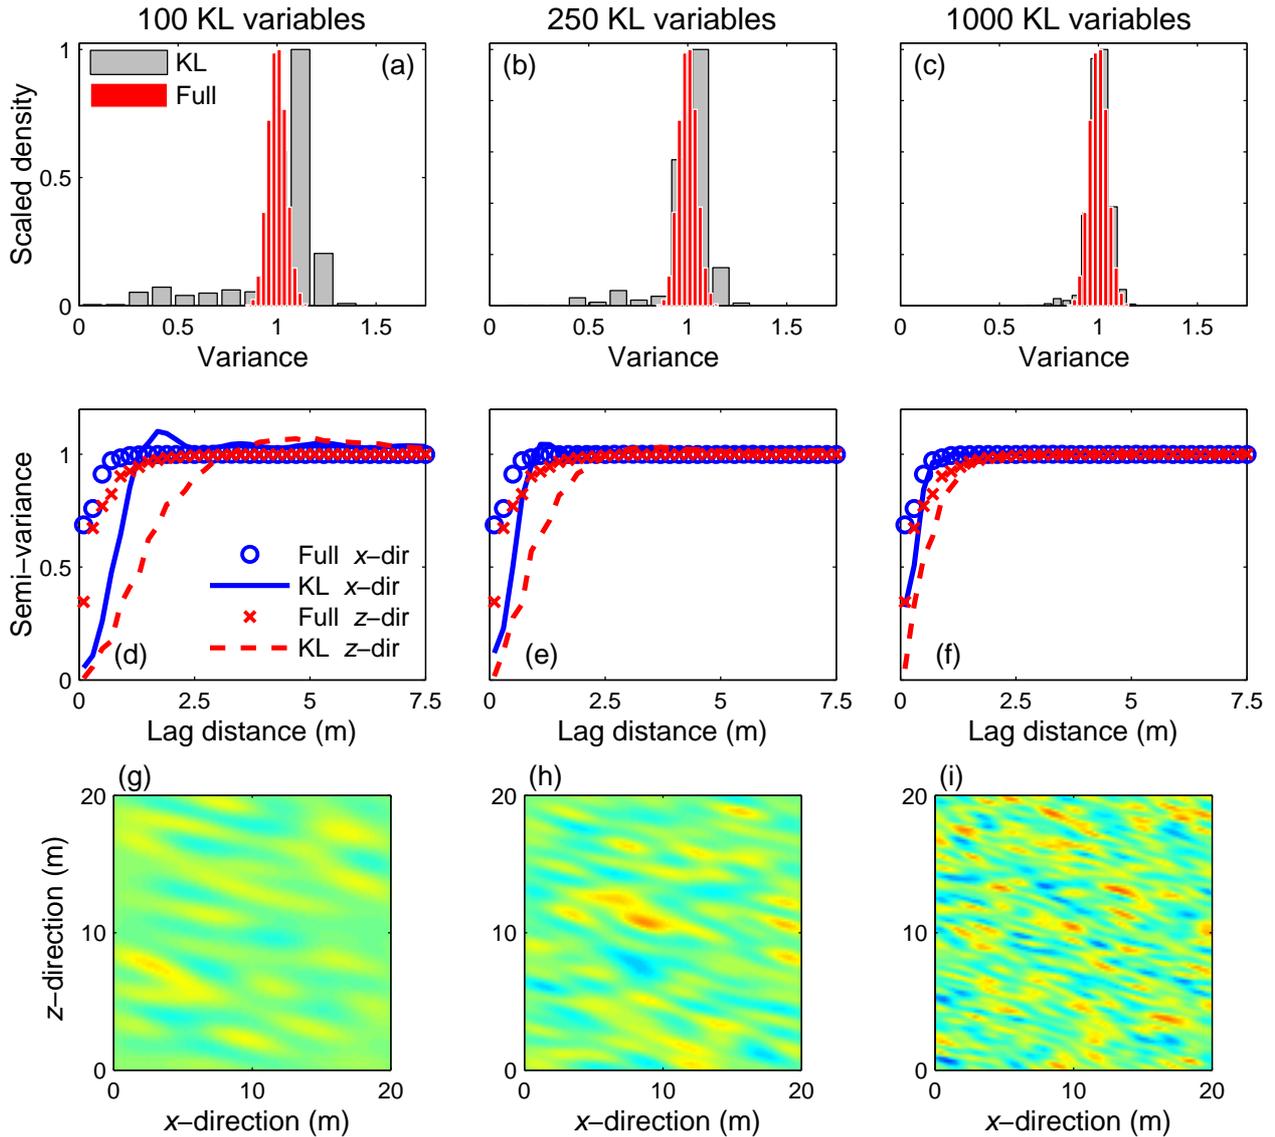


Figure 3. Grid point variance distribution (subplots a, b, c) and experimental variograms in both the horizontal (x -dir) and vertical (z -dir) directions (subplots d, e, f) calculated from 1000 realizations of an autocorrelated multi-Gaussian field with zero-mean and the reference variogram used for the inversions (see last of column Table 1 and Figure 1a), using (a,d) 100, (b-e) 250 and (c-f) 1000 KL variables, and randomly chosen field realizations (subplots g-i) derived with (g) 100, (h) 250 and (i) 1000 KL variables. The gray and red bins in subplots a, b, and c denote the variance distributions obtained from the KL transform and from directly generating autocorrelated multi-Gaussian fields with prescribed variogram, respectively. The red histograms in subplots a, b and c thus represent the natural statistical fluctuations. The (plain and dashed) lines and (circle and cross) symbols in subplots d-f signify average experimental variogram values in the x - (blue color) or z -direction (red color) for each of the considered lags, derived from KL expansion (KL) and from directly generating autocorrelated multi-Gaussian fields with prescribed variogram (Full), respectively.

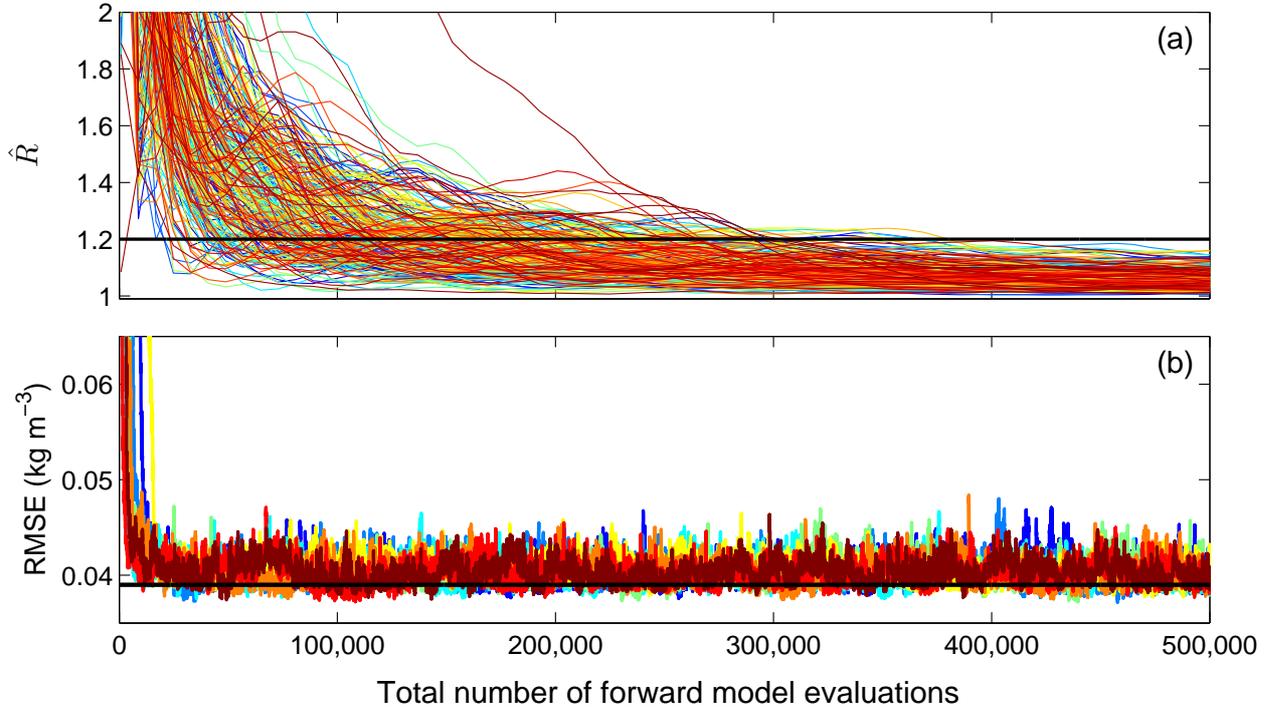


Figure 4. (a) Evolution of the \hat{R} convergence metric [*Gelman and Rubin, 1992*] for the proposed method with 250 DR variables. Each of the 257 sampled parameters is coded with a different color. The chains have officially converged if the plotted lines fall below the horizontal black line. (b) Trace plot of the sampled RMSE values. Each of the eight Markov chains is coded with a different color. The horizontal black line denotes the true RMSE of 0.039 kg m^{-3} .

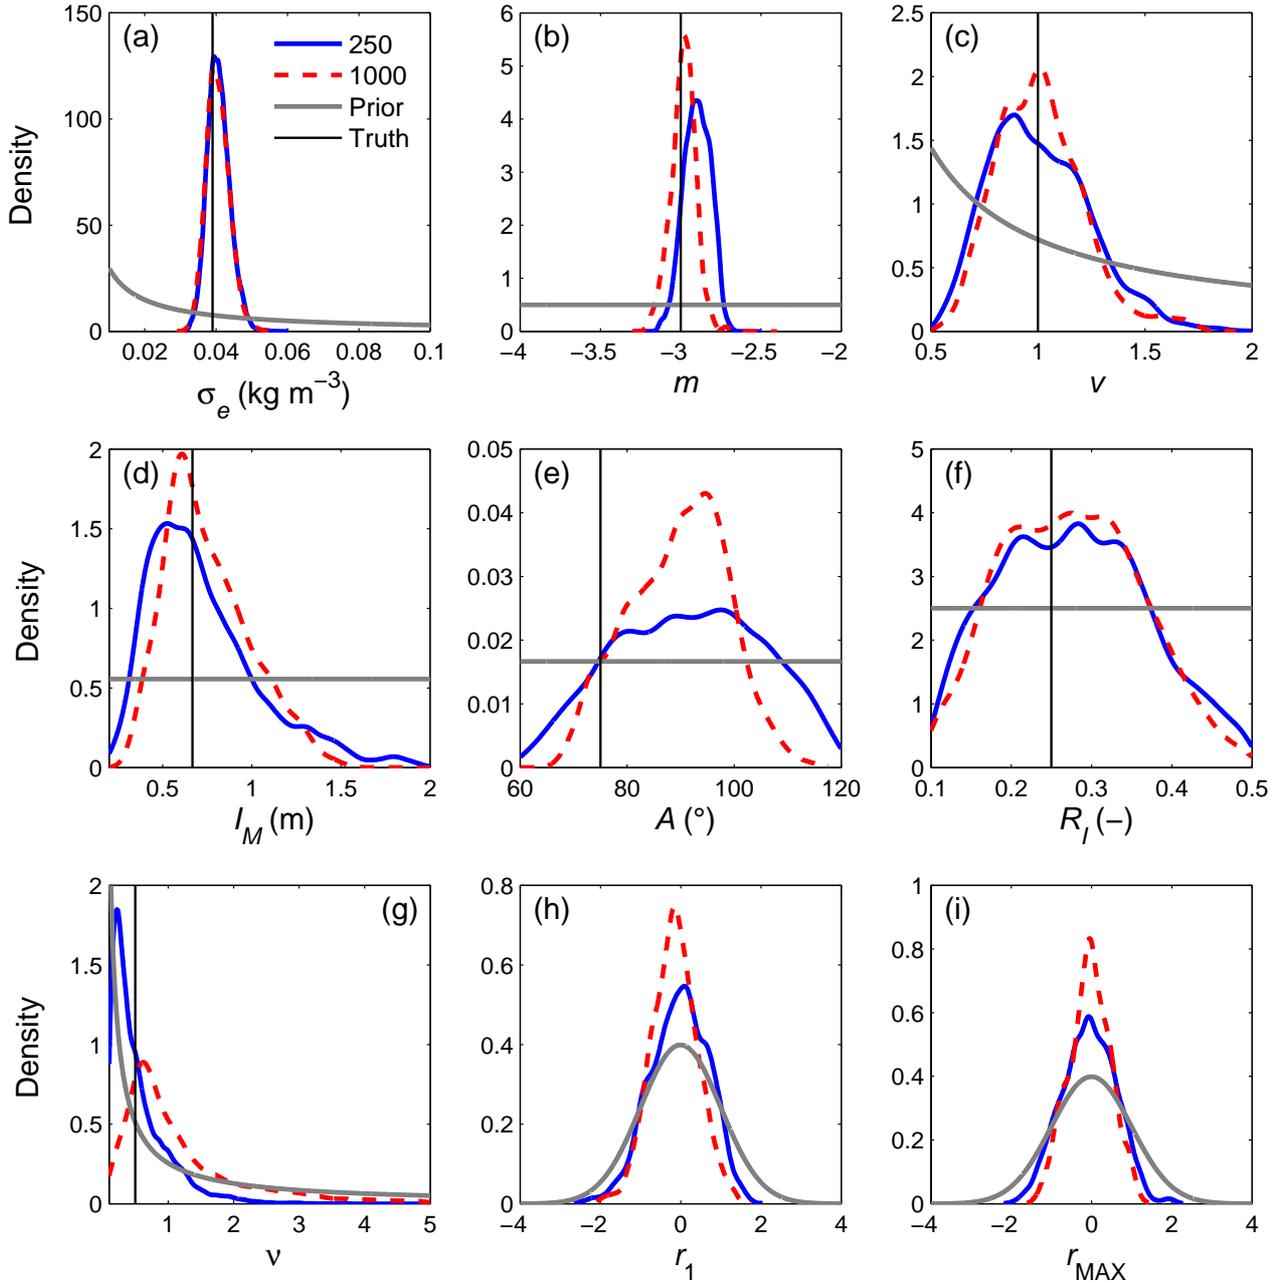


Figure 5. Marginal prior and posterior distributions of (a) the standard deviation of the residual errors, (b-g) the six geostatistical parameters, and (h) the first (r_1), and (i) last (r_{MAX}) elements of the \mathbf{r} vector for the inversion with 250 (blue lines) and 1000 (dashed red lines) DR variables. The distributions are derived from kernel density smoothing using the last 90% of the samples generated by DREAM_(ZS). The parameter values used for creating the reference field are separately indicated with a vertical black line. Since the reference field was generated without dimensionality reduction, there are no true values for r_1 and r_{MAX} . The r_{MAX} parameter is either r_{250} (blue line) or r_{1000} (red dashed line), whereas the respective r_1 and r_{MAX} distributions derived from the inversions with 250 and 1000 DR variables are grouped into the same plots for visual convenience only. There is no reason for these distributions to be similar across inversions.

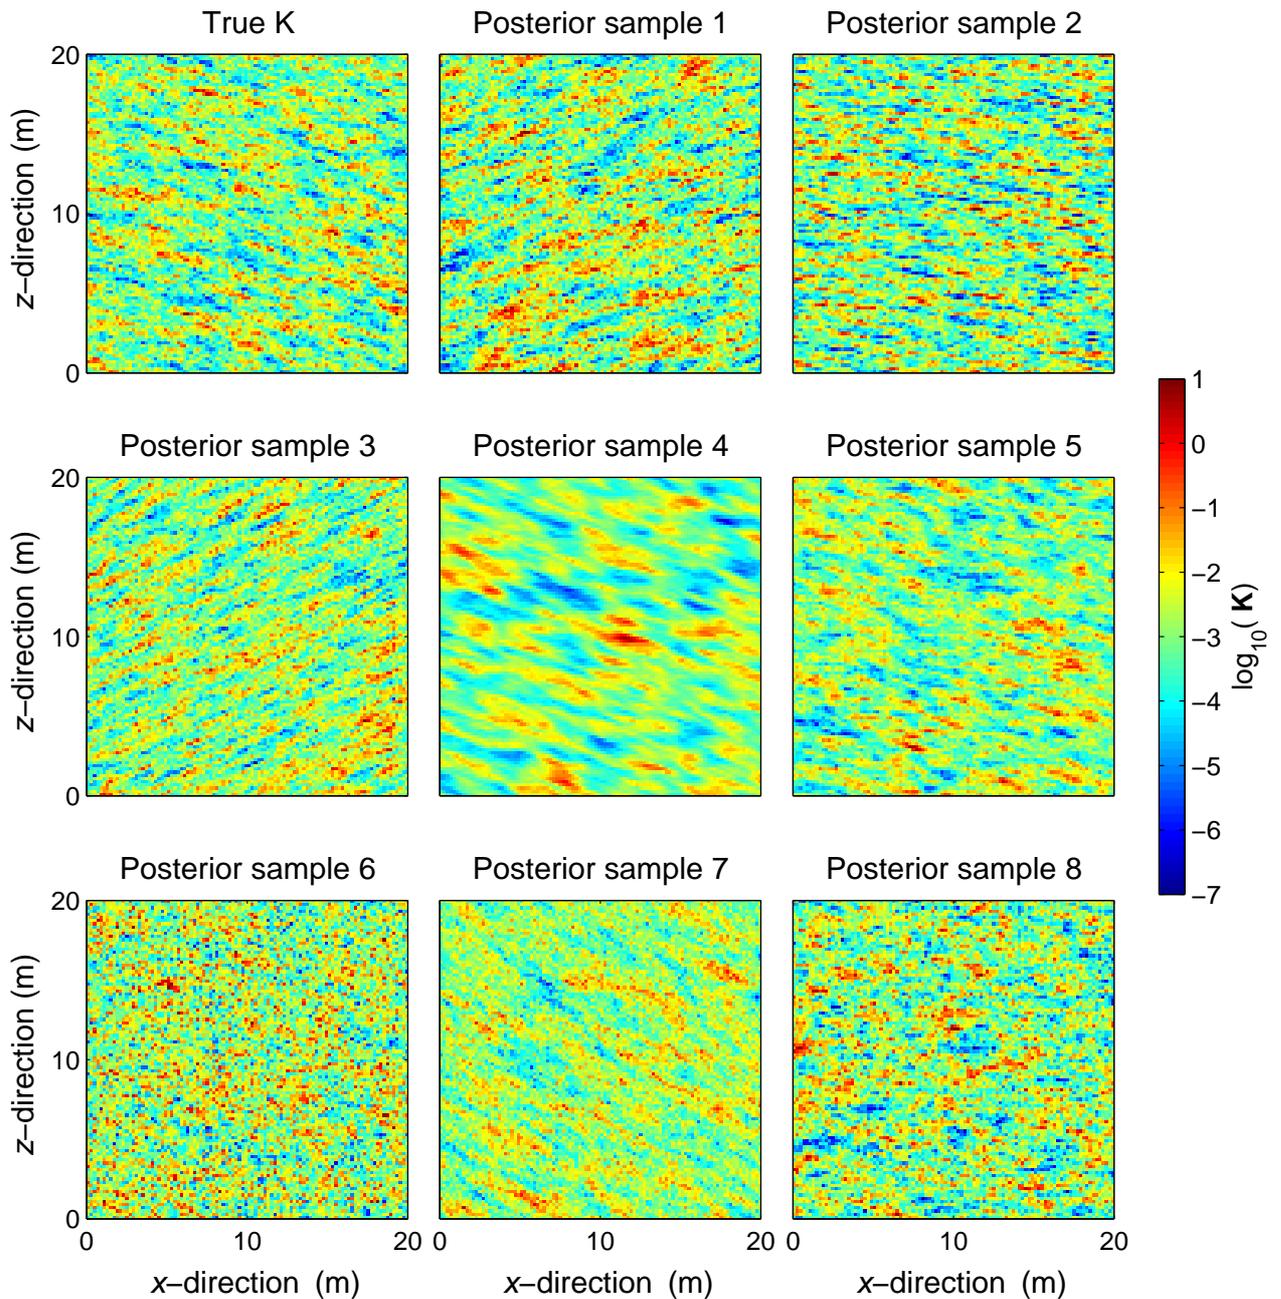


Figure 6. True (reference) model and eight (randomly chosen) realizations of the $\text{DREAM}_{(\text{ZS})}$ derived posterior distribution for 250 DR variables.

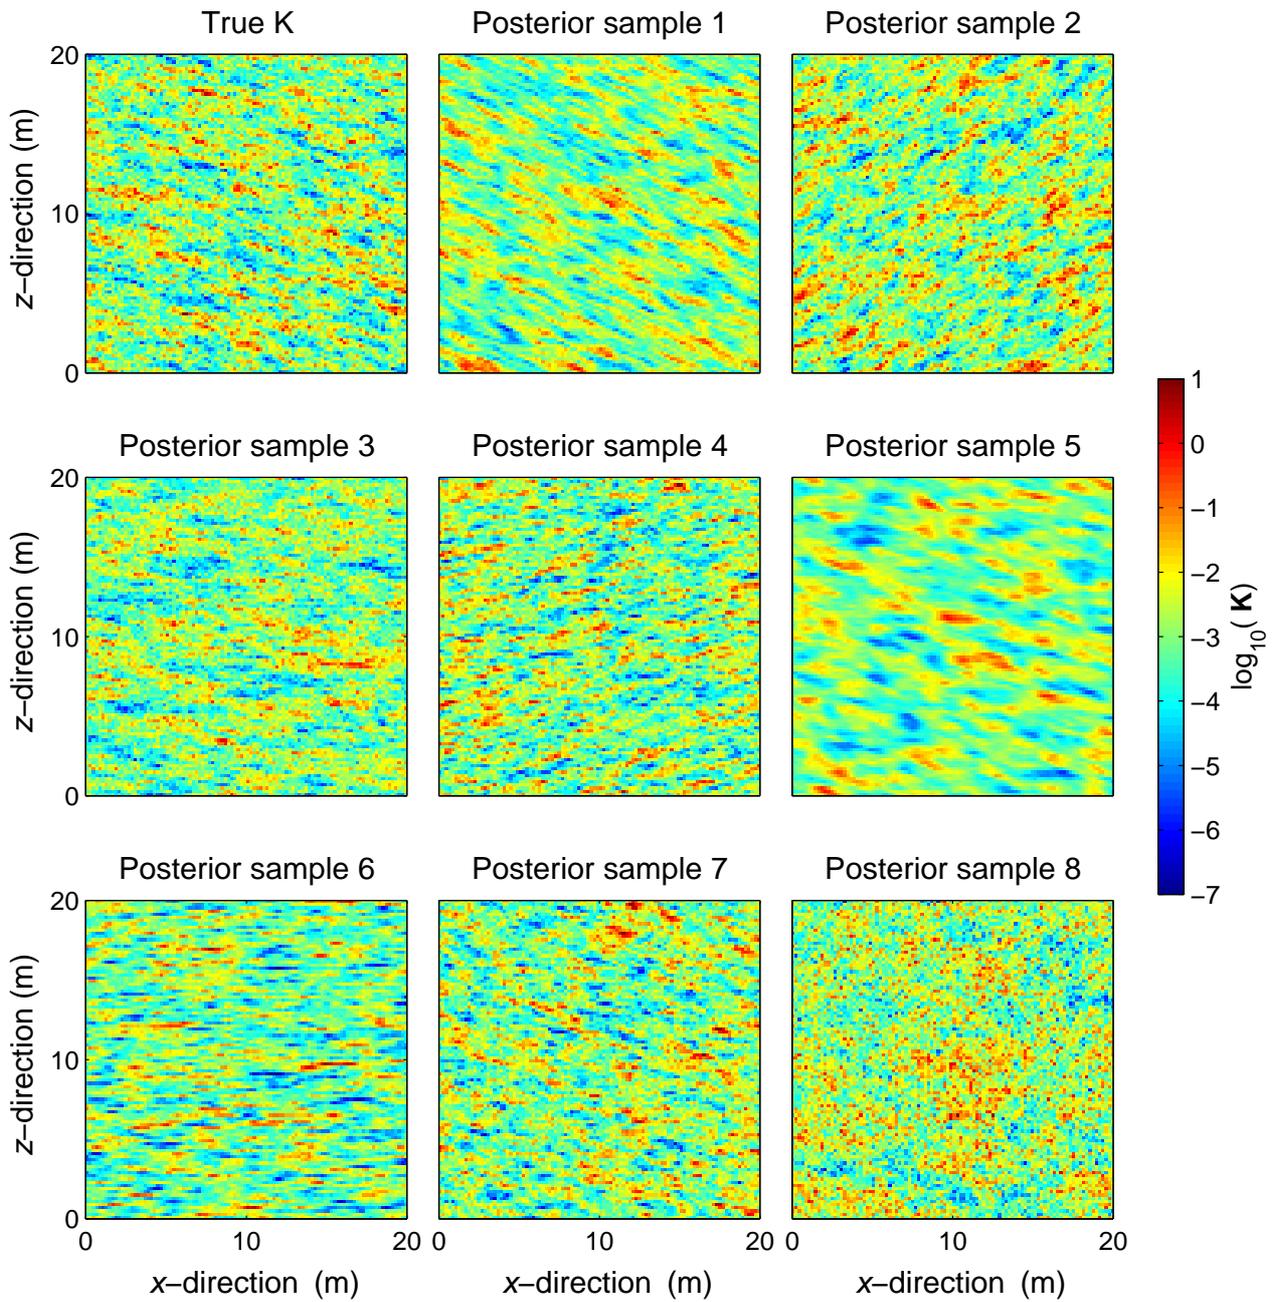


Figure 7. True (reference) model and eight (randomly chosen) realizations of the $\text{DREAM}_{(\text{ZS})}$ derived posterior distribution for 1000 DR variables.

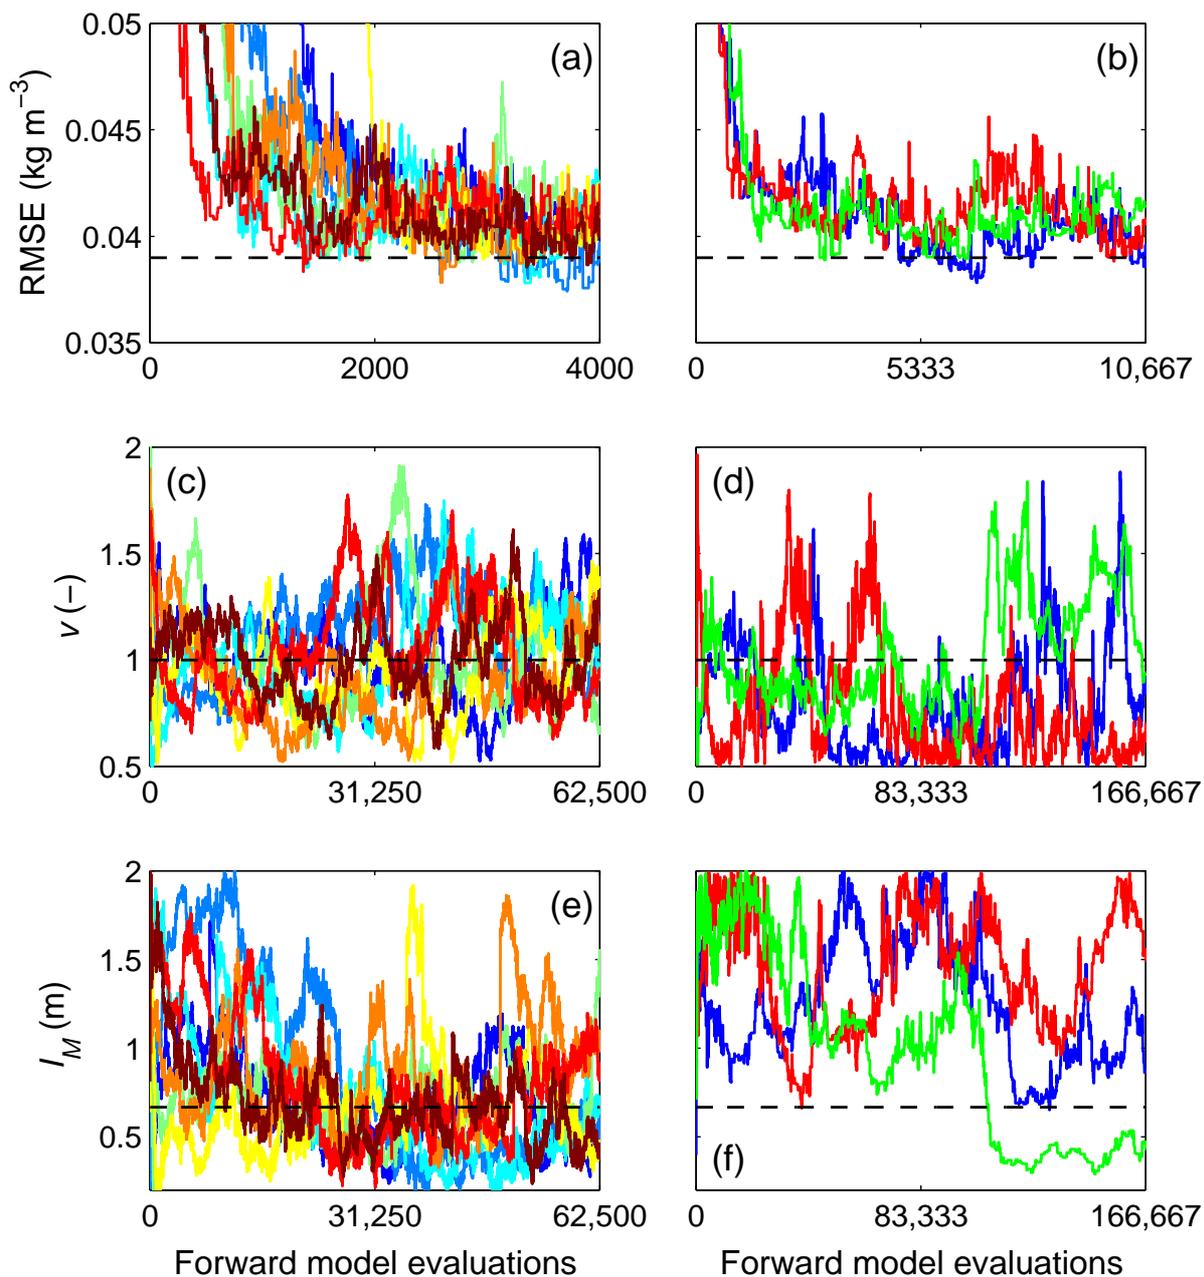


Figure 8. Trace plots of the sampled (a-b) RMSE, (c-d) v and (e-f) I_M values derived from our proposed inversion approach (left column: subplots a, c, and e) and the SGS method (right column: subplots b, d, and f) using 250 DR variables. The dashed horizontal lines depict the true parameter values of the reference log-conductivity field used herein. A visual comparison of the results demonstrates that the DREAM_(ZS) (left-hand side) algorithm of the proposed inversion method exhibits an improved searching behavior compared to SGS (right-hand side).

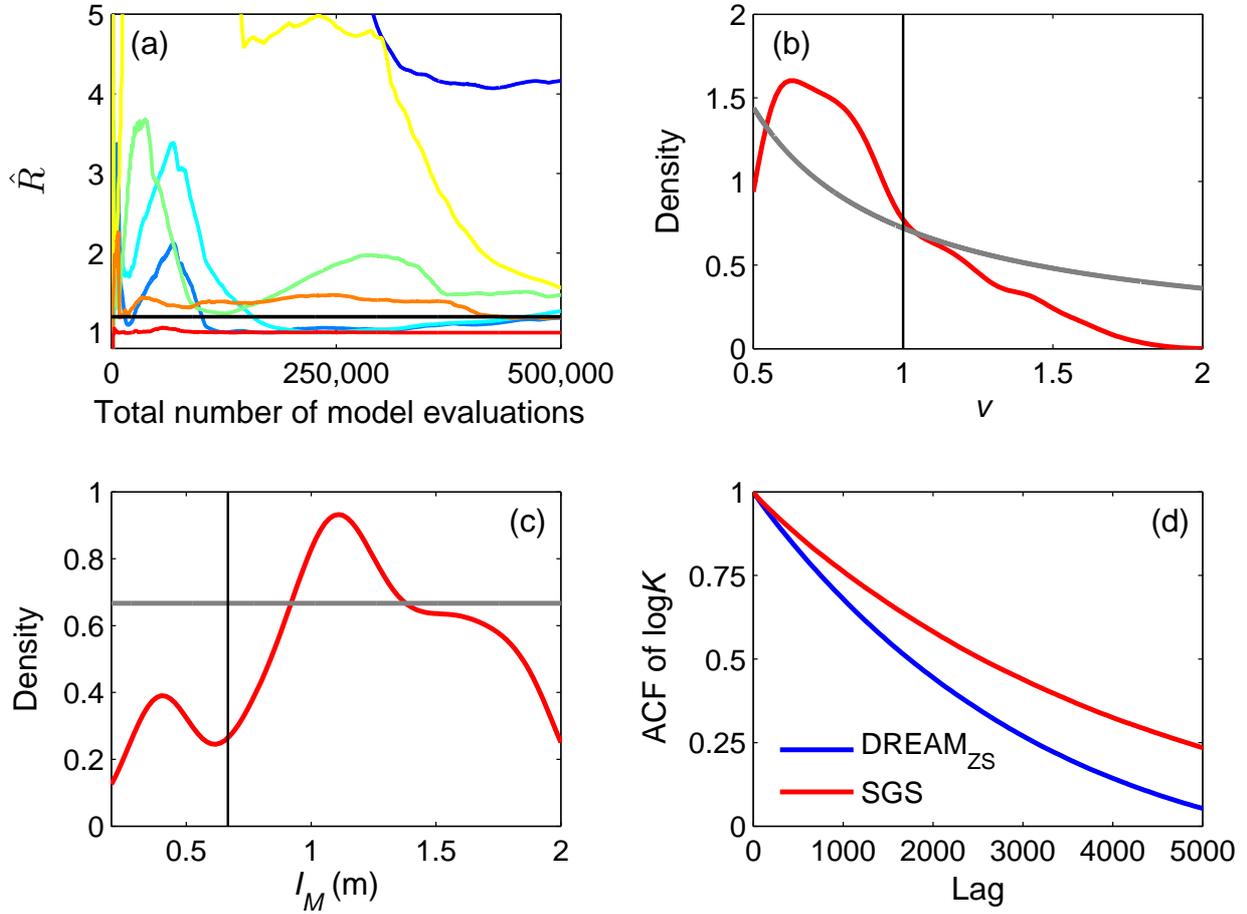


Figure 9. (a) Evolution of the \hat{R} convergence diagnostic [Gelman and Rubin, 1992] for the total computational budget of 500,000 model evaluations of the three SGS trials. We only plot traces (color coded) of the six geostatistical parameters and σ_e . (b) and (c) Marginal posterior distributions (red lines) of v and I_M derived from the last 90% of the samples generated with the three different SGS trials. Kernel density smoothing is applied. The prior distribution is indicated by a gray line, and the true values by the vertical black lines. (d) Mean autocorrelation function (ACF) of the 10,000 log-conductivity grid values derived from DREAM_(ZS) (blue line) or SGS (red line) for lags 0 to 5000. The lag- k autocorrelation is defined as the correlation between draws k lags apart. Listed statistics are computed for the last 50,000 log-conductivity fields sampled in the eight DREAM_(ZS) chains or the three independent SGS chains. The average of the chains is presented.

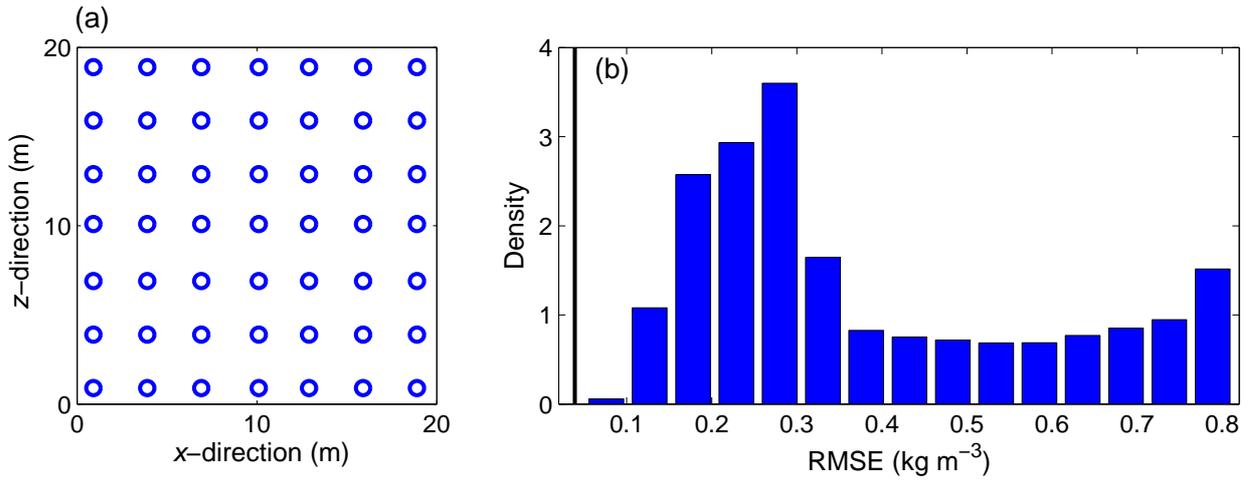


Figure 10. (a) Anchor locations and (b) RMSE distribution of the 600,000 forward runs performed by the MAD trial. The vertical black line in plot (b) denotes the true RMSE value of 0.039 kg m^{-3} . The RMSE values sampled with MAD are much larger than the true value. This simply demonstrates that the method cannot converge properly within the assigned computational budget of 600,000 model evaluations.