Principal Component Geostatistical Approach for Large-Dimensional Inverse Problems

P. K. Kitanidis and J. Lee[1]

P. K. Kitanidis, Department of Civil and Environmental Engineering, Stanford University, Stanford, CA, 94305, USA. (peterk@stanford.edu)

J. Lee, Department of Civil and Environmental Engineering, Stanford University, Stanford, 94305, USA. (jonghyun@stanford.edu)

[1]Department of Civil and Environmental Engineering, Stanford University, Stanford, 94305, USA.

Abstract. The quasilinear geostatistical approach is for weakly nonlinear underdetermined inverse problems, such as Hydraulic Tomography and Electrical Resistivity Tomography. It provides best estimates as well as measures for uncertainty quantification. However, for its textbook implementation, the approach involves iterations, to reach an optimum, and requires the determination of the Jacobian matrix, *i.e.*, the derivative of the observation function with respect to the unknown. Although there are elegant methods for the determination of the Jacobian, the cost is high when the number of unknowns, m, and the number of observations, n, is high. It is also wasteful to compute the Jacobian for points away from the optimum. Irrespective of the issue of computing derivatives, the computational cost of implementing the method is generally of the order of m^2n, though there are methods to reduce the computational cost. In this work, we present an implementation that utilizes a matrix-free in terms of the Jacobian matrix Gauss-Newton method and improves the scalability of the geostatistical inverse problem. For each iteration, it is required to perform K runs of the forward problem, where K is not just much smaller than m but can be smaller that n. The computational and storage cost of implementation of the inverse procedure scales roughly linearly with m instead of m^2 as in the textbook approach. For problems of very large m, this implementation constitutes a dramatic reduction in computational cost compared to the textbook approach. Results illustrate the validity of the approach and provide insight in the conditions under which this method perform best.

1. Introduction

Hydrology and geophysics are among the many fields where one encounters large-dimensional parameter estimation or inversion problems. The relation between the unknown parameters and the observations is typically governed by partial differential equations. An important feature of such problems is that the information in the data does not suffice to constrain the answer to a unique solution that is not overly sensitive to the data. Examples include hydraulic tomography [*Butler et al.*, 1999; *Yeh and Liu*, 2000; *Cardiff et al.*, 2009, 2013] and electrical resistivity tomography [*Slater et al.*, 2000; *Linde et al.*, 2006; *Pollock and Cirpka*, 2012], which lead to the formulation of such problems when a fine enough grid is used to discretize the unknown function. To summarize the main ideas consider the linear problem

$$\boldsymbol{y} = \boldsymbol{H}\boldsymbol{s} \tag{1}$$

where the n by m observation matrix $\boldsymbol{H}$ has rank less than m in a numerical sense. Given $\boldsymbol{y}$, there are multiple solutions for $\mathbf{s}$ or the solution may vary a lot in response to small changes in the data. The common approach to solving such problems is to introduce additional requirements, such as finding the flattest or smoothest solution consistent with the data, in a precisely specified sense.

This kind of problem can be addressed within a Bayesian framework [*Gavalas et al.*, 1976; *Neuman*, 1980; *Kitanidis and Vomvoris*, 1983; *Carrera and Neuman*, 1986; *Woodbury and Rubin*, 2000; *Rubin et al.*, 2010], where the multiplicity of solutions is explicitly recognized and each possible solution is assigned a probability based on how well it agrees

with the data (the likelihood function criterion) and how well it agrees with other information (the prior probability criterion). A solution is ascribed a high probability if it meets both criteria (logical conjunction) to a significant degree. There are many different Bayesian approaches but the emphasis here is on the geostatistical approach (GA). Within the context of Bayesian inference approaches, GA is an objective [*Berger*, 2006, describes the basic ideas] and empirical Bayes [e.g., *Carlin and Louis*, 2000; *Kitanidis*, 2010] method. Objective is a technical term that refers to the need to rely as much on data as the nature of the problem allows while empirical means that both the prior and the likelihood functions are adjustable models informed by data.

Let us briefly review how the method works by considering just the best estimate. The result is

$$\hat{\boldsymbol{s}} = \boldsymbol{\Lambda} \boldsymbol{y} \tag{2}$$

where $\boldsymbol{\Lambda}$ acts as a pseudo-inverse of $\boldsymbol{H}$. This $\boldsymbol{\Lambda}$ depends not just on $\boldsymbol{H}$ but also on an m by p drift matrix $\boldsymbol{X}$, an m by m covariance $\boldsymbol{Q}$ of $\boldsymbol{s}$ and an n by n covariance $\boldsymbol{R}$ of measurement error, as will be explained later.

1. The pseudo-inverse satisfies $\boldsymbol{\Lambda}\mathbf{H}\mathbf{X} = \mathbf{X}$, which means that $\boldsymbol{\Lambda}$ behaves like the true inverse of $\mathbf{H}$ w.r.t. $\mathbf{X}$. For example, in some applications (e.g., s is hydraulic head) the datum is modeler dependent. By choosing $\mathbf{X} = ones(m, 1)$ we make sure that if the datum is changed, the solution is not affected. In some other problems, we work with $s = \log K$ (where K is hydraulic or electrical conductivity). By using $\mathbf{X} = ones(m, 1)$ we ascertain that the solution is scaled solely on the basis of data. The introduction of

the drift means that p characteristics of the solution depend only on the data and not on the prior.

2. The selection of $\mathbf{Q}$ describes the variability and continuity in the solution, in the part not covered by the drift. A solution may be chosen to be smooth on the basis of what is known about the unknown function. Or, based on the understanding that small scale features cannot be estimated with certainty from data, one may choose to seek solutions that are sufficiently smooth.

3. The selection of $\boldsymbol{R}$ informs about the type and magnitude of observation error or "noise" and allows to properly weight measurements. If $\boldsymbol{R}$ is doubled, for example, the data will receive less weight, the best solution will reproduce the data less closely, and thus the effect of data noise on the result will be reduced.

4. In combination, $\boldsymbol{Q}$ and $\boldsymbol{R}$ have a critical role in error quantification. For example, if both $\boldsymbol{Q}$ and $\boldsymbol{R}$ are multiplied by 4, the credible interval (Bayesian confidence interval) doubles.

The geostatistical approach for nonlinear problems [*Kitanidis*, 1995] uses a Gauss-Newton iterative method to obtain the "best estimate" and linearized uncertainty quantification about the best estimate. However, the textbook implementation of this approach is suited for problems where m and n are of modest size. By modest size we mean that linear systems and matrix multiplications are computable within seconds with computational facilities available. At this point in time, this would mean $10^3 - 10^4$. However, we are increasingly faced with problems where the number of unknowns is large, often larger than 10^6, whereas the number of observations is modest (say 10^3). For such problems, each iteration requires about n runs of the forward problem (to construct the Jacobian

matrix via the adjoint-state method) plus operations of the order $m^2 n$. Furthermore, covariance matrices, of size m^2, are very expensive to compute and store. Overall, the cost of implementation increases roughly with m^2, which means that the method becomes impractical to use for a sufficiently large m. To address this problem, methods using FFT or hierarchical matrices and fast multiple methods have been applied [*Nowak et al.*, 2003; *Saibaba et al.*, 2012; *Ambikasaran et al.*, 2013] that speed up the cost of multiplication of $\boldsymbol{Q}$ with a vector from $O(m^2)$ to $O(m \log m)$, where the big O denotes order of. Such methods reduce the overall cost quite significantly. Additionally, various other works examine aspects of reducing the computational cost [e.g., *Li and Cirpka*, 2006; *Liu and Kitanidis*, 2011; *Saibaba and Kitanidis*, 2012; *Yadav and Michalak*, 2013].

This work presents an implementation, termed Principal Component Geostatistical Approach (PCGA), that is mainly suited for cases with the following characteristics:

1. In terms of problem size, m is huge whereas n is of modest size.

2. The prior covariance corresponds to a reasonably smooth unknown function. More technically the spectrum of the covariance matrix (*i.e.*, the eigenvalues) drops rapidly.

3. The information content of the observations is limited, due to observation error and the ill-conditioning of the inverse problem. We will explain what this means in more precise terms later.

Many problems encountered in practice, such as hydraulic tomography, fit this description. We will later discuss how one can relax 1. and 2. but requirement 3. is needed for the method to perform most efficiently and to be competitive with other approaches. The approach uses the idea expressed in *Bui-Thanh et al.* as follows: "... we exploit the fact that the data are typically informative about low-dimensional manifolds of parameter

space to construct low rank approximations of the covariance matrix of the posterior pdf via a matrix-free randomized method." The method presented in this work also uses low-rank matrix approximations and also avoids the computation and storage of the complete Jacobian matrix but is different from the method in the aforementioned reference.

The main contribution of this paper that it presents a methodology for solving static inverse problems, using the Geostatistical Approach, with the following characteristics: (a) It is fully numerical and scales in computational cost and storage requirements roughly linearly with m, the size of the vector of the unknowns. (b) It is a "matrix-free" approach in the sense that required derivatives are computed by performing matrix-vector multiplications, without computing the complete Jacobian matrix. Computing derivatives is the bulk of computations in inverse methods. The method we present here allows significant savings, particularly at early iterations when computing a full Jacobian is usually a waste of effort. (c) The method requires making calls to a forward model and does not require an adjoint-state model. (d) The method has spectral convergence with the number of components used, which, for many though not necessarily all problems, can be much faster than the $1/\sqrt{K}$ convergence of ensemble or Monte-Carlo based methods.

The next section reviews GA and discuss some important features regarding the standard or "textbook" computational implementation of GA. Then we present the PCGA algorithm, which is suited for problems for which the number of unknowns, m, is very large, the number of observations, n, is moderate, and the evaluation of derivatives is too expensive to be performed a very large number of times. Then, we present results that illustrate the applicability of the method.

2. Overview of Geostatistical Approach

For the sake of completeness and for ease in referencing results, we will review the quasilinear geostatistical approach (GA) [*Kitanidis*, 1995]. The observation equation, which relates the vector of the unknowns $\boldsymbol{s}$ to the vector of the data $\boldsymbol{y}$ is

$$\boldsymbol{y}=\boldsymbol{h}\left(\boldsymbol{s}\right)+\boldsymbol{v} \tag{3}$$

where $\boldsymbol{h}$ is the mapping from $R^{m\times 1}$ to $R^{n\times 1}$, $\boldsymbol{v}$ is Gaussian with mean $\mathbf{0}$ and covariance $\boldsymbol{R}$ (often proportional to the identity matrix). The mapping must be sufficiently smooth so that it can be approximated with a linear relation in the neighborhood where the posterior probability of $\boldsymbol{s}$ is concentrated. The prior probability of $\boldsymbol{s}$ is Gaussian with mean $\boldsymbol{X\beta}$, where $\boldsymbol{X}$ is an m by p known matrix and $\boldsymbol{\beta}$ a p by 1 vector of unknowns (*i.e.*, to be determined by data), and generalized covariance matrix $\boldsymbol{Q}$.

The posterior pdf of $\boldsymbol{s}$ and $\boldsymbol{\beta}$ are obtained through Bayes theorem and its cologarithm (minus the logarithm), $-\ln p''\left(\mathbf{s},\beta\right)$, is

$$\frac{1}{2}\left(\boldsymbol{y}-\boldsymbol{h}\left(\boldsymbol{s}\right)\right)^{T}\boldsymbol{R}^{-1}\left(\boldsymbol{y}-\boldsymbol{h}\left(\boldsymbol{s}\right)\right)+\frac{1}{2}\left(\boldsymbol{s}-\boldsymbol{X\beta}\right)^{T}\boldsymbol{Q}^{-1}\left(\boldsymbol{s}-\boldsymbol{X\beta}\right) \tag{4}$$

The maximum a posteriori or most likely values are obtained by minimizing this expression with respect to $\boldsymbol{s}$ and $\boldsymbol{\beta}$ vectors. This is a nonlinear optimization problem that is commonly solved through iterative methods. For $n << m$, which is usually the case, a convenient form of the Gauss-Newton method is in the form of the so-called cokriging equations, to be described next.

Based on the initial guess (or most recent "good solution") $\boldsymbol{s}_0$, we update to a new solution using a Newton-type iterative approach. First, define the n by m matrix $\boldsymbol{H}$ as the Jacobian matrix of $\boldsymbol{h}$ at $\mathbf{s}_0$:

$$\boldsymbol{H} = \left.\frac{\partial \boldsymbol{h}}{\partial \boldsymbol{s}}\right|_{\boldsymbol{s}=\boldsymbol{s}_0} \tag{5}$$

Then, assuming that the actual $\hat{\boldsymbol{s}}$ is close to $\boldsymbol{s}_0$, approximate

$$\boldsymbol{h}\left(\hat{\boldsymbol{s}}\right) = \boldsymbol{h}\left(\boldsymbol{s}_0\right) + \boldsymbol{H}\left(\hat{\boldsymbol{s}} - \boldsymbol{s}_0\right) \tag{6}$$

and, after some matrix manipulations, one obtains the updated solution to the minimization of (4)

$$\hat{\boldsymbol{s}} = \boldsymbol{X}\hat{\boldsymbol{\beta}} + \boldsymbol{Q}\boldsymbol{H}^T\boldsymbol{\xi} \tag{7}$$

where $\hat{\boldsymbol{\beta}}$ and the $\boldsymbol{n}$ by $\mathbf{1}$ vector $\boldsymbol{\xi}$ are found from the solution of a system of :

$$\begin{bmatrix} \boldsymbol{H}\boldsymbol{Q}\boldsymbol{H}^T + \boldsymbol{R} & \boldsymbol{H}\boldsymbol{X} \\ (\boldsymbol{H}\boldsymbol{X})^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\xi} \\ \hat{\boldsymbol{\beta}} \end{bmatrix} = \begin{bmatrix} \boldsymbol{y} - \boldsymbol{h}\left(\boldsymbol{s}_0\right) + \boldsymbol{H}\boldsymbol{s}_0 \\ \mathbf{0} \end{bmatrix} \tag{8}$$

Note also that the objective function that is minimized can be written, using (7), as

$$\begin{aligned} J &= \frac{1}{2}\left(\boldsymbol{y} - \boldsymbol{h}\left(\boldsymbol{s}\right)\right)^T \boldsymbol{R}^{-1}\left(\boldsymbol{y} - \boldsymbol{h}\left(\boldsymbol{s}\right)\right) + \frac{1}{2}\left(\boldsymbol{s} - \boldsymbol{X}\boldsymbol{\beta}\right)^T Q^{-1}\left(\boldsymbol{s} - \boldsymbol{X}\boldsymbol{\beta}\right) \\ &= \frac{1}{2}\left(\boldsymbol{y} - \boldsymbol{h}\left(\boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{Q}\boldsymbol{H}^T\boldsymbol{\xi}\right)\right)^T \boldsymbol{R}^{-1}\left(\boldsymbol{y} - \boldsymbol{h}\left(\boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{Q}\boldsymbol{H}^T\boldsymbol{\xi}\right)\right) \\ &+ \frac{1}{2}\boldsymbol{\xi}^T\boldsymbol{H}\boldsymbol{Q}\boldsymbol{H}^T\boldsymbol{\xi} \end{aligned} \tag{9}$$

This expression can be used to gauge the progress of the minimization and to make sure that the new solution is not worse than the previous.

Once the optimum is achieved, we can proceed to uncertainty quantification. The linearized approach treats the problem as approximately linear in the proximity of the best estimate. This usually assumes that the posterior error is somehow small and the nonlinearity not too strong so that a linear approximation is good in the neighborhood where $\boldsymbol{s}$ is most likely to be. Under these conditions the posterior is approximately Gaussian.

One approach to uncertainty quantification involves generation of conditional realizations, *i.e.*, samples from the posterior distribution. They are computed as follows:

$$\boldsymbol{s}_i = \boldsymbol{\zeta}_i + \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{Q}\boldsymbol{H}^T\boldsymbol{\xi} \tag{10}$$

where $\boldsymbol{\zeta}_i$ are unconditional realizations, *i.e.*, Gaussian with mean $\mathbf{0}$ and covariance $\boldsymbol{Q}$, and $\hat{\boldsymbol{\beta}}$ and $\boldsymbol{\xi}$ are found from:

$$\begin{bmatrix} \boldsymbol{H}\boldsymbol{Q}\boldsymbol{H}^T + \boldsymbol{R} & \boldsymbol{H}\boldsymbol{X} \\ (\boldsymbol{H}\boldsymbol{X})^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\xi} \\ \hat{\boldsymbol{\beta}} \end{bmatrix} = \begin{bmatrix} \boldsymbol{y} + \boldsymbol{v}_i - \boldsymbol{h}(\boldsymbol{s}_0) + \boldsymbol{H}(\boldsymbol{s}_0 - \zeta_i) \\ \mathbf{0} \end{bmatrix} \tag{11}$$

where $\boldsymbol{v}_i$ is generated Gaussian mean $\mathbf{0}$ and covariance matrix $\boldsymbol{R}$.

An alternative approach is to solve

$$\begin{bmatrix} \boldsymbol{H}\boldsymbol{Q}\boldsymbol{H}^T + \boldsymbol{R} & \boldsymbol{H}\boldsymbol{X} \\ (\boldsymbol{H}\boldsymbol{X})^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\Lambda}^T \\ \mathbf{M} \end{bmatrix} = \begin{bmatrix} \boldsymbol{H}\boldsymbol{Q} \\ \mathbf{X}^T \end{bmatrix} \tag{12}$$

where $\boldsymbol{\Lambda}$ is an $m \times n$ matrix of coefficients and $\mathbf{M}$ is a $p \times n$ matrix of multipliers. Then, the best estimate is:

$$\hat{\boldsymbol{s}} = \boldsymbol{\Lambda}\left(\boldsymbol{y} - \boldsymbol{h}(\boldsymbol{s_0}) + \boldsymbol{H}\boldsymbol{s}_0\right) \tag{13}$$

while the posterior covariance matrix, from the linearized approach, is

$$\boldsymbol{V} = \boldsymbol{Q} - \boldsymbol{F} \tag{14}$$

where

$$\boldsymbol{F} = \boldsymbol{X}\boldsymbol{M} + \boldsymbol{Q}\boldsymbol{H}^T\boldsymbol{\Lambda}^T \tag{15}$$

This result shows that the posterior covariance matrix is the prior covariance minus a low-rank correction. Mathematically, the rank of the correction $\boldsymbol{F}$ cannot be higher than $n + p$, given the dimensions of the factors of the two terms that comprise the correction.

In the textbook approach, the Jacobian matrix must be re-calculated at each Gauss-Newton iteration using the adjoint-state method, which is the method of choice when $n << m$, requiring n solutions of the forward problem. Usually, this is the most computationally expensive part of the method. Furthermore, when m is large, the cost of dealing with covariance matrices is large and can even be prohibitive. The computation of the $\boldsymbol{Q}\boldsymbol{H}^T$ is of the order of m^2n, which is very burdensome when m is large. There are methods to reduce the computational cost, associated with the multiplication of the covariance matrix with vectors [*Nowak et al.*, 2003; *Saibaba et al.*, 2012]. Lastly, the computation and storage of the posterior covariance matrix can be prohibitive.

The method described next is well suited for cases of huge m and moderate n and becomes even better when the effective rank of the correction in (14) is $<< n$. (Or, to put it differently, it can be approximated by a matrix of rank $<< n$.) The method is termed Principal Component Geostatistical Approach (PCGA) because it is an implementation of the Geostatistical Approach that utilizes principal components associated with the covariance matrix and the drift matrix. The method has the potential to reduce markedly

the number of runs of the forward problem and the dominant term in the cost of matrix manipulations drops from $O(m^2)$ to order $O(m)$, which is a dramatic improvement for large m.

3. Important Preliminaries

This section describes some ideas and tools that will be used later.

3.1. Matrix-Free Approach

Instead of computing the complete Jacobian matrix $\boldsymbol{H}$ and then computing the matrix-vector product $\boldsymbol{Ha}$, one can directly compute this matrix-vector product using a so-called matrix-free method. Consider the nonlinear vector function

$$\boldsymbol{y} = \boldsymbol{h}(\boldsymbol{s}) \tag{16}$$

where $\boldsymbol{s}$ is m by 1 and $\boldsymbol{y}$ is n by 1, with its Jacobian matrix at $\boldsymbol{s} = \boldsymbol{s}_0$ denoted as $\boldsymbol{H}$. Consider that we want to compute

$$\boldsymbol{Ha} \tag{17}$$

where $\boldsymbol{a}$ is a vector with same dimensions as $\boldsymbol{s}$. Rather than follow the expensive process of first evaluating $\boldsymbol{H}$ and then performing the matrix vector multiplication, we can take advantage of the following finite–difference approximation:

$$\boldsymbol{Ha} = \frac{\|\boldsymbol{a}\|}{\delta \|\boldsymbol{s}_0\|} \left[\boldsymbol{h}\left(\boldsymbol{s}_0 + \frac{\|\boldsymbol{s}_0\|}{\|\boldsymbol{a}\|} \delta \boldsymbol{a} \right) - \boldsymbol{h}(\boldsymbol{s}_0) \right] + O(\delta) \tag{18}$$

The ratio of norms (or vector lengths) $\|\boldsymbol{s}_0\| / \|\boldsymbol{a}\|$ is a normalization factor and δ is a small dimensionless number. This approach with δ equal to about the square root of the floating-point relative accuracy (for example, $\delta \approx 10^{-8}$ for "double precision", $\delta \approx 10^{-4}$ for "single precision") is adequate, but one may want to experiment with choosing the δ that

is right for a specific application. The essential point is that each Jacobian times vector computation involves one more run of the forward problem.

Note that this approach makes calls to a solver of the forward problem. If we needed to compute $\boldsymbol{H}^T\boldsymbol{b}$, where $\boldsymbol{b}$ is a vector with the same dimensions as $\boldsymbol{y}$, we could follow the same procedure using a solver that solves the adjoint of the forward problem. The cost of a forward run of the original problem and its adjoint are the same so that there is no real benefit here in using the adjoint, unless the computations are arranged so that the number of forward runs is reduced. In this work, computations have been arranged so that there is no need to use an adjoint state solver, taking into account that most practitioners have access to just forward-problem solvers that they can use in a "black box" fashion.

Note also that instead of using the finite-difference approximation one could solve a linearized version of the forward problem. This approach would be more elegant and would bypass the issue of choosing the right δ. However, again taking into account that it is more practical to use the available forward model as a black box and not have to write another PDE solver, which might be a nontrivial exercise, the finite-difference approximation is recommended as the method of choice. After all, development of forward problem solvers, which are typically PDE solvers, is quite advanced and one can find highly optimized code, often approaching near linear scalability with problem size. The inverse methodology presented herein has been structured to take advantage of such code.

3.2. Normalization of Prior Model

The drift matrix $\boldsymbol{X}$ is m by p, where p is a small number. If $p = 1$, the normalized version of $\boldsymbol{X}$ is $\boldsymbol{U} = \boldsymbol{X}/\sqrt{m}$. For $p > 1$, one may use a singular value decomposition:

$$\boldsymbol{X} = \boldsymbol{U}\boldsymbol{S}_X\boldsymbol{V}_X^T \tag{19}$$

where $\boldsymbol{U}$ is m by p, $\boldsymbol{S}_X$ is diagonal p by p, and $\boldsymbol{V}_X$ is p by p. The $\boldsymbol{S}_X$ and $\boldsymbol{V}_X$ can be discarded. Matrix $\boldsymbol{U}$ is isomorphic to $\boldsymbol{X}$, *i.e.*, mathematically the same result in the geostatistical inversion is obtained whether one defines the drift in terms of $\boldsymbol{X}$ or $\boldsymbol{U}$. However, since $\boldsymbol{U}$ is orthonormal, meaning that $\boldsymbol{U}^T\boldsymbol{U}$ is identity matrix, it is preferable to use in computations. It is recommended to replace $\boldsymbol{X}$ with $\boldsymbol{U}$ when defining the model at the very beginning.

The detrending matrix $\boldsymbol{P}$ is expressed

$$\mathbf{P} = \boldsymbol{I} - \boldsymbol{U}\boldsymbol{U}^T \tag{20}$$

This matrix is symmetric, has $m - p$ eigenvalues equal to 1 and p eigenvalues equal to 0, is a projection to the space orthogonal to the drift matrix, $\boldsymbol{P}\boldsymbol{X} = \mathbf{0}$, and satisfies the idempotence property, meaning that it can be multiplied many times without changing the result beyond the initial multiplication, $\boldsymbol{P}\boldsymbol{P}\boldsymbol{a} = \boldsymbol{P}\boldsymbol{a}$. Furthermore the matrix vector product $\boldsymbol{P}\boldsymbol{a}$ can be computed with $O(m)$ operations using the relation

$$\boldsymbol{P}\boldsymbol{a} = \boldsymbol{a} - \boldsymbol{U}\left(\boldsymbol{U}^T\boldsymbol{a}\right) \tag{21}$$

The generalized covariance matrix $\boldsymbol{Q}$ also has an issue of isomorphism. Whether one uses $\boldsymbol{Q}$ or $\boldsymbol{Q} + \boldsymbol{X}\boldsymbol{b}\boldsymbol{b}^T\boldsymbol{X}$, where $\boldsymbol{b}$ is an arbitrary p by 1 vector, the result of the geostatistical inversion method is mathematically identical. (The same is true of ordinary Kriging, an interpolation method.) This is a powerful, though perhaps not widely appreciated, feature of generalized covariance functions and matrices [*Matheron*, 1973; *Kitanidis*, 1983, 1993]. The essence of isomorphism is that two covariance functions

or matrices may look different but act mathematically in exactly the same way, in the context of defining the prior with a specific drift matrix; and what matters is to identify the criterion for isomorphism and the invariants. The criterion is the following. Two covariance matrices $\boldsymbol{Q}_1$ and $\boldsymbol{Q}_2$ are isomorphic if

$$\boldsymbol{P}\left(\boldsymbol{Q}_1 - \boldsymbol{Q}_2\right)\boldsymbol{P} = \mathbf{0} \tag{22}$$

where $\boldsymbol{P}$ is defined in (20). The definitions above lead to the following: (a) $\boldsymbol{Q}$ and $\boldsymbol{PQP}$ are isomorphic; (b) if $\boldsymbol{Q}_1$ and $\boldsymbol{Q}_2$ are isomorphic, then $\boldsymbol{PQ}_1\mathbf{P}$ is equal to $\boldsymbol{PQ}_2\boldsymbol{P}$ and isomorphic to the original $\boldsymbol{Q}_1$ and $\boldsymbol{Q}_2$; and (c) thus, $\boldsymbol{PQP}$ is the invariant.

One straightforward solution to removing the ambiguity is to replace $\boldsymbol{Q}$ with $\boldsymbol{PQP}$. Note that whereas $\boldsymbol{Q}$ may not be positive definite and may not look like a reasonable covariance matrix of a random vector, $\boldsymbol{PQP}$ is nonnegative definite and has a solid probabilistic interpretation as the covariance of $\boldsymbol{Px}$, where $\boldsymbol{x}$ is a random vector. $\boldsymbol{PQP}$ is "what matters" in the problems we are interested in. But one may avoid performing the multiplication , because it involves $O\left(m^2\right)$ operations, when applying the geostatistical approach. If, for example we multiply with a vector that is already detrended, by taking account of idempotence we have

$$\left(\boldsymbol{Pa}\right)^T\boldsymbol{PQP}\left(\boldsymbol{Pa}\right) = \left(\boldsymbol{Pa}\right)^T\boldsymbol{Q}\left(\boldsymbol{Pa}\right) \tag{23}$$

4. Low-Rank Approximation

A matrix of low rank can be factorized into the product of low-dimensional matrices. For example, a matrix $\boldsymbol{Y}$ of size k by l and rank K can be factorized $\boldsymbol{Y} = \boldsymbol{AB}^T$, where $\boldsymbol{A}$ is k by K and $\boldsymbol{B}$ is l by K. If $K << k, l$, the benefits are potentially tremendous in terms of storage and matrix-vector operations. Instead of storing $k \times l$ entries, one stores

$(k+l)\,K$. The matrix vector product $\boldsymbol{Y}\boldsymbol{y}$ instead of requiring $k \times l$ it takes only $(k+l)\,K$ multiplications, since $\boldsymbol{Y}\boldsymbol{y} = \boldsymbol{A}\left(\boldsymbol{B}^T\boldsymbol{y}\right)$. This crucial idea is the backbone of approaches in dealing with cases of large m.

In, for example, hierarchical matrix methods [*Saibaba et al.*, 2012; *Ambikasaran et al.*, 2013, e.g., for applications in estimation, see], the structure of the dense prior covariance matrix is exploited to factorize it. *Bui-Thanh et al.* factorize a matrix appearing in the Hessian of the linearized problem to compute the low-order correction in the covariance. The method to be used here is closest works that employ orthogonal or Karhunen-Loeve decomposition of the covariance [*Li and Cirpka*, 2006; *Marzouk and Najm*, 2009]

In the approach developed herein, we start with the factorization of the prior covariance $\boldsymbol{Q}$ with low-size matrices. There are many possible methods but we need a method with computational cost roughly linear in m and also one that does not require many repetitions of the procedure. The method we describe next is one of several possible computational approaches.

We start by generating K unconditional realizations; *i.e.*, $\boldsymbol{\zeta}_k$, where $k = 1 : K$, is m by 1 generated with mean $\mathbf{0}$ and generalized covariance function $\boldsymbol{Q}$. For practical situations of generating realizations of random fields in 1, 2, or 3 dimensions using regular grids, one powerful approach is based on FFT or hierarchical matrices. The cost of generating a realization with such methods is $O\left(m \log m\right)$.

The simplest approximation to the covariance $\boldsymbol{Q}$ is through

$$\boldsymbol{Q} \approx \frac{1}{K}\sum_{k=1}^{K} \boldsymbol{\zeta}_k \boldsymbol{\zeta}_k^T \tag{24}$$

This is factorization through the product of two rank K matrices

$$\boldsymbol{Q} \approx \boldsymbol{Z}\boldsymbol{Z}^T, \text{ where } \boldsymbol{Z}_{ij} = \frac{\zeta_{ij}}{\sqrt{K}}, \; i = 1, ..., m, \; j = 1, ..., K \tag{25}$$

This is the same factorization used in ensemble Kalman filtering [*Evensen*, 1994, 2003, 2006], which in some cases has been reported to work well even for K as small as 50 or a 100, but in other cases requires large ensembles or additional enhancements [*Chatterjee et al.*, 2012; *Li et al.*, 2014] Generally, the approximation error in (25) is $O\left(\frac{1}{\sqrt{K}}\right)$. This means that to halve the error from the factorization one needs to quadruple the number or realizations.

An alternative idea is to factorize $\boldsymbol{Q}$ in such a way that the factorization is "optimum" for the specific number K, the rank of the approximation. From matrix theory [*Golub and Van Loan*, 1989], it is known that the smallest achievable error when using a K-rank approximation of a matrix $\boldsymbol{Q}$, where the spectral norm (aka 2-norm and denoted $\| \quad \|$) for matrices is used to quantify the error, is

$$\min_{Q_K} \|\boldsymbol{Q} - \boldsymbol{Q}_K\| = \rho_{K+1} \tag{26}$$

where $\boldsymbol{Q}$ is the full matrix, $\boldsymbol{Q}_K$ is its rank K approximation, and ρ_{K+1} is the $(K+1)$-th largest singular value of $\boldsymbol{Q}$. If the spectrum of $\boldsymbol{Q}$ (*i.e.*, the set of singular values ρ_k, $k = 1 : m$) drops rapidly, it is possible to obtain a low-rank approximation that has a small error. Most importantly, the optimum $\boldsymbol{Q}_K$ is the one obtained from a singular value decomposition of $\boldsymbol{Q}$ where only the first K singular values are kept and the others are set to zero.

To illustrate how the choice of model can make a difference, consider the following example. Consider a one-dimensional domain of length $L = 1$ and points $m = 100$ on a uniform grid and three models:

Model I: $p = 1$, $X_i = 1$, $Q_{ij}\,(x_i - x_j) = \begin{cases} 1, & \text{if } i = j \\ 0, & \text{otherwise} \end{cases}$. This is the classic "nugget-effect" model of geostatistics. No continuity is assumed and using this model in inverse modeling can be interpreted as selecting as best estimate a small variance solution.

Model II: $p = 1$, $X_i = 1$, $Q_{ij} = -\,|x_i - x_j|$. This is the linear generalized covariance function of geostatistics, corresponding to the commonly used linear semivariogram. The unknown is assumed to belong to an ensemble of continuous but not differentiable functions. The use of the model can be interpreted as selecting a best estimate that is a flat solution.

Model III: $p = 2$, $X_{i1} = 1, X_{i2} = x_i$, $Q_{ij} = |x_i - x_j|^3$. This is the cubic generalized covariance function of geostatistics. The unknown is assumed to belong to an ensemble of continuous and differentiable functions and its use can be interpreted as selecting a best estimate that is a smooth solution.

Note that none of these three models has a scale parameter and thus the choice of domain length has no effect. The choice of m has a small effect but does not materially affect the argument we will make. To account for the fact that the drift accounts for some variability and the $\boldsymbol{Q}$ is the variability around the drift, we examine the spectrum of $\boldsymbol{PQP}$, where $\boldsymbol{P}$ is a symmetric projection matrix of rank $m-p$, as previously defined. Whenever we need to compute a matrix vector product, $\boldsymbol{Pa}$, we notice that this is equivalent to detrending $\boldsymbol{a}$; *i.e.*, fitting to $\boldsymbol{a}$ a trend $\boldsymbol{X\beta}$, through least squares techniques [*Golub*, 1965] and then subtracting it from $\boldsymbol{a}$ to arrive at the desired result.

Because we are primarily interested in relative error in approximating $\boldsymbol{Q}$, we will plot the normalized matrix spectrum, $\frac{\rho_k}{\rho_1}$, on figure 1.

The results clearly illustrate that for some covariance matrices, excellent accuracy in approximating $\boldsymbol{Q}$ can be obtained using approximations of low rank. Pretty much the same behavior is for much larger m. In fact, for model III, even if m is 10^6, a matrix with rank as low as 20 suffices to achieve small relative error.

It is important to point out that these results only suggest the maximum number K that is needed. The results will be better than what this analysis suggests, for a number of reasons. Indeed, when K is large and approaches n, the results shown in figure 1 become less relevant. After all, matrices of interest like $\boldsymbol{HQ}$ and $\boldsymbol{HQH}^T$ are of rank not larger than n, and sometimes considerably smaller than n because the low rank of $\boldsymbol{H}$. Thus, the errors in $\boldsymbol{Q}$ become less important. We will later explain how the action of the forward operator and the presence of observation error make it unnecessary to use a large K.

Once one decides what K value to use, one must proceed to compute the low-rank approximation. Among the many methods, each with advantages and disadvantages, we will focus on randomized methods [*Halko et al.*, 2011]. Here, we propose a low cost randomized method to construct an accurate symmetric approximation. The random driver is the m by K matrix $\boldsymbol{Z}$, equation (25), with the additional stipulation that each column (unconditional realization) has been detrended. We can perform a singular value decomposition (or a QR decomposition, it makes no difference in this application) of $\boldsymbol{Z}$,

$$\boldsymbol{Z} = \boldsymbol{U}_z \boldsymbol{S}_z \boldsymbol{V}_z^T \tag{27}$$

Here, $\boldsymbol{U}_z$ is m by K and satisfies $\boldsymbol{U}_z^T\boldsymbol{U}_z = \boldsymbol{I}_K$, $\boldsymbol{S}_z$ is K by K diagonal with nonnegative elements, and $\boldsymbol{V}_z$ is K by K unitary, $\boldsymbol{V}_z^T\boldsymbol{V}_z = \boldsymbol{V}_z\,\boldsymbol{V}_z^T = \boldsymbol{I}_K$. In what follows, only $\boldsymbol{U}_z$ will be needed. Next compute

$$\boldsymbol{C} = \boldsymbol{U}_z^T\,(\boldsymbol{PQP})\,\boldsymbol{U}_z \tag{28}$$

Since $\boldsymbol{Z}$ has been detrended, so has $\boldsymbol{U}_z$ and thus $\boldsymbol{PU}_z = \boldsymbol{U}_z$, which means that we can simplify the notation

$$\boldsymbol{C} = \boldsymbol{U}_z^T\boldsymbol{QU}_z \tag{29}$$

Finally, the low-order approximation of the generalized covariance matrix, $\boldsymbol{PQP}$, is

$$\boldsymbol{PQP} \approx \boldsymbol{U}_z\boldsymbol{C}\boldsymbol{U}_z^T \tag{30}$$

which is an $(m \times k)\,(k \times k)\,(k \times m)$ factorization.

This is a randomized algorithm so one would like to evaluate the effect of randomness. For this purpose the process was repeated twenty times (*i.e.*, each time with a different seed number) for model III and $K = 30$ and the results were plotted on figure 2. By actual we mean the values computed with standard software without approximations about rank. The approximate ones are computed very efficiently from the small matrix $\boldsymbol{C}$. One can see from figure 2 that the error and the randomness are limited to the smaller computed singular values. This, together with other results, suggests that if one wants accuracy at the K-th singular value, one should use an approximation with a rank somewhat higher than that. Randomness, although small in the mean square sense, introduces unsightly small scale variability in the estimate and has a more pronounced material effect in computing variances.

This difficulty is fortunately easy to overcome. In this example, compute for 30 and then keep only the first 15. All that is needed is to drop the last columns of $\boldsymbol{A}$ and the last columns and rows of $\boldsymbol{C}$.

In applications, one can evaluate retroactively how well one has done, without having to compute the singular values of large $\boldsymbol{Q}$, by plotting all the singular values of the small matrix $\boldsymbol{C}$, that should mirror the first K singular values of $\boldsymbol{Q}$. What one wants to see is that they are dropping fast and the smallest is sufficiently small.

Summarizing the algorithm:

1. Generate K_s realizations, cost $O(K_s m \log m)$, create $\boldsymbol{Z}$ matrix. Detrend, if not already detrended, $O(K_s m)$ operations.

2. Do a singular value decomposition of the $\boldsymbol{Z}$ matrix of rank K_s, produce m by K_s matrix $\boldsymbol{U}_z$, with $O(K_s^2 m)$ operations.

3. Compute $\boldsymbol{C}$, with cost only $O(K_s m \log m)$ by using fast matrix vector multiplication like *Fong and Darve* [2009] for 2-D.

4. Discard the "extra" columns and rows to reduce the size of $\boldsymbol{A}$ to K columns and of $\boldsymbol{C}$ to K by K, like $K = K_s - 15$.

5. Principal Component Geostatistical Inversion

The static nonlinear inverse problem, Equation (3), can be solved through a variant to the quasilinear approach, Equations (4)-(15), through the following steps.

Computation of $\boldsymbol{HX}$

For column $\boldsymbol{X}_i$, an $O(\delta)$ finite-difference approximation is

$$\boldsymbol{H}\boldsymbol{X}_i = \frac{\|\boldsymbol{X}_i\|}{\delta \|\boldsymbol{s}_0\|} \left[h\left(\boldsymbol{s}_0 + \delta \frac{\|\boldsymbol{s}_0\|}{\|\boldsymbol{X}_i\|} \boldsymbol{X}_i \right) - h(\boldsymbol{s}_0) \right] \tag{31}$$

Computation of $-\boldsymbol{h}(\boldsymbol{s_0}) + \boldsymbol{H}\boldsymbol{s_0}$

This can be computed from $O(\delta)$ approximation

$$\begin{aligned} & -\boldsymbol{h}(\boldsymbol{s}_0) + \boldsymbol{H}\boldsymbol{s}_0 \\ & = \frac{1}{\delta}\left[h\left((1+\delta)\,\boldsymbol{s}_0\right) - (1+\delta)\,h(\boldsymbol{s}_0)\right] \end{aligned} \tag{32}$$

Computation of $\boldsymbol{HQH^T}$ and $\boldsymbol{QH^T}$

Consider the factorization of the covariance $\boldsymbol{Q}$ is through

$$\boldsymbol{Q} = \boldsymbol{A}\boldsymbol{C}\boldsymbol{A}^T \tag{33}$$

which is the $(m \times k)(k \times k)(k \times m)$ factorization of equation (30) and we defined $\boldsymbol{A} = \boldsymbol{U}_Z$. The idea, of course, is that $\boldsymbol{Q}$ is stored through an $m \times K$ matrix with orthonormal columns and a $K \times K$ symmetric and positive definite matrix, where K is much smaller that m. These matrices are used to perform operations involving $\boldsymbol{Q}$ without ever computing $\boldsymbol{Q}$.

We will employ this approximation to compute $\boldsymbol{HQ}$ and $\boldsymbol{HQH}^T$. Denote the columns of $\boldsymbol{A}$ as $\boldsymbol{a}_k$, $k = 1 : K$,

$$\begin{aligned} \boldsymbol{HQ} &= (\boldsymbol{HA})\,\boldsymbol{C}\boldsymbol{A}^T \\ &= \left(\sum_{k=1}^{K} \boldsymbol{H}\boldsymbol{a}_k\right)\boldsymbol{C}\boldsymbol{A}^T \end{aligned} \tag{34}$$

Let us define

$$\boldsymbol{b}_k = \boldsymbol{H}\boldsymbol{a}_k \tag{35}$$

that can be computed as described previously. Form the n by K matrix $\boldsymbol{B}$ with columns the computed $\boldsymbol{b}_k$. Then,

$$\boldsymbol{HQ} = \boldsymbol{B}\boldsymbol{C}\boldsymbol{A}^T \tag{36}$$

Similarly,

$$\boldsymbol{HQH}^T = (\boldsymbol{HA})\,\boldsymbol{C}\,(\boldsymbol{HA})^T$$

$$= \boldsymbol{BCB}^T \tag{37}$$

One can proceed to solve system (12), obtain the estimate (13) and continue with iterations. There are a number of possible refinements that reduce the computational cost, such as one by taking advantage that some of the matrices we are working with have rank K. For example, the matrix vector product $\left(\boldsymbol{HQH}^T\right)\boldsymbol{a}$ can be computed with $O(Kn)$ multiplications instead of $O(n^2)$. One can also solve a smaller system by effectively compressing the data. Such refinements are important when n is large.

An important additional step, consistent with the objective of improving the scaling of the computational problem with respect to m is the following. Form matrix

$$\boldsymbol{A}_p = [\boldsymbol{U}_z, \boldsymbol{U}]$$

which has $K + p$ orthonormal (i.e., $\boldsymbol{A}_p^T\boldsymbol{A}_p$ is identity matrix). Post-multiply $\boldsymbol{\Lambda}^T$ and $\mathbf{M}$ with $\boldsymbol{A_p}$ so that the system to solve is:

$$\begin{bmatrix} \boldsymbol{HQH}^T\!+\!\boldsymbol{R} & \boldsymbol{HX} \\ (\boldsymbol{HX})^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} (\boldsymbol{\Lambda}^T\boldsymbol{A}_p) \\ (\boldsymbol{M}\boldsymbol{A}_p) \end{bmatrix} = \begin{bmatrix} \boldsymbol{HQA}_p \\ \boldsymbol{X}^T\boldsymbol{A}_p \end{bmatrix} \tag{38}$$

By solving, we obtain the smaller matrices $\left(\boldsymbol{\Lambda}^T\boldsymbol{A}_p\right)$ and $(\boldsymbol{M}\boldsymbol{A}_p)$ .Then approximate $\boldsymbol{\Lambda}^T$ and $\boldsymbol{M}$ through K-rank approximations:

$$\boldsymbol{\Lambda}^T = \left(\boldsymbol{\Lambda}^T\boldsymbol{A}\right)\boldsymbol{A}_p^T \tag{39}$$

$$\boldsymbol{M} = (\boldsymbol{M}\boldsymbol{A})\,\boldsymbol{A}_p^T \tag{40}$$

Thus, to compute the new $\boldsymbol{s}$

$$\hat{\boldsymbol{s}} = \boldsymbol{A}_p \left(\boldsymbol{\Lambda}^T \boldsymbol{A}\right)^T \left(\boldsymbol{y} - \boldsymbol{h}\left(\tilde{\boldsymbol{s}}\right) + \tilde{\boldsymbol{H}}\tilde{\boldsymbol{s}}\right) \tag{41}$$

so that number of multiplications is $O(mK)$ rather than $O(mn)$, which is a slight improvement. The main benefit is in computing the covariance.

At the optimum, the posterior covariance correction satisfies

$$\digamma \boldsymbol{A}_p = \boldsymbol{X}\left(\boldsymbol{M}\boldsymbol{A}_p\right) + \boldsymbol{Q}\boldsymbol{H}^T\left(\boldsymbol{\Lambda}^T\boldsymbol{A}_p\right) \tag{42}$$

Then, the correction to the covariance matrix, approximated through a rank $K+p$ matrix, is

$$\begin{aligned} \digamma &\approx \boldsymbol{X}\left(\boldsymbol{M}\boldsymbol{A}_p\right)\boldsymbol{A}_p^T + \boldsymbol{Q}\boldsymbol{H}^T\left(\boldsymbol{\Lambda}^T\boldsymbol{A}_p\right)\boldsymbol{A}_p^T \\ &= \boldsymbol{X}\left(\boldsymbol{M}\boldsymbol{A}_p\right)\boldsymbol{A}_p^T + \boldsymbol{A}\boldsymbol{C}\boldsymbol{B}^T\left(\boldsymbol{\Lambda}^T\boldsymbol{A}_p\right)\boldsymbol{A}_p^T \end{aligned} \tag{43}$$

The sizes of the matrices involved are shown below

$$(m\times p)(p\times K)(K\times m),\ (m\times K)(K\times K)(K\times n)(n\times K)(K\times m) \tag{44}$$

while $\boldsymbol{Q}$ is approximated through (33). The largest matrix to store is Km. Though the posterior covariance is not computed or stored, one can employ matrix-vector products to compute what is needed out of it. For example, say we want to compute the diagonal element $\digamma(1,1)$ that is needed to find the variance of $\boldsymbol{s}_1$ and compute the Bayesian confidence intervals: One pre-multiplies by $\left[\begin{array}{ccccccc} 1 & 0 & 0 & \ldots & 0 & 0 \end{array}\right]$ and post-multiplies by the transpose of the same. Such operations can be arranged to be done quite efficiently.

Implementation

The approach requires a fast method to generate realizations with generalized covariance matrix $\boldsymbol{Q}$. For practical situations of generating realizations of random fields in 1, 2, or 3 dimensions using regular grids, one approach is based on FFT [*Nowak et al.*, 2003]. For irregular grids, one can use methods from hierarchical matrices [*Saibaba et al.*, 2012; *Ambikasaran et al.*, 2013]. It is also required to have an efficient algorithm for computing unconditional covariance matrix times a vector. The same FFT and hierarchical methods apply here as well. Note that the approach requires to use unconditional realizations and the aforementioned products only at the start, before performing iterations.

If the standard Gauss-Newton method does not converge, one can introduce a between-step search, either a line search or a trust region [for example, see *Zanini and Kitanidis*, 2009]. Assuming that the problem has been formulated correctly, there are really two kinds of non-convergence issues. One is when the radius of convergence is large, compared to the statistical error, but the starting solution is outside of it (the case of bad starting point). In this case, one can start by using intermediate models, such as use a large observation error (this makes the problem more linear-like and increases the radius of convergence – linear problems have infinite radius of convergence) to get a very smooth solution and gradually reduce the variance to the value that it should be. The other case is that the radius of convergence is comparable to the statistical error of the solution (the case of strong nonlinearity), and a plain Gauss-Newton simply does not consistently improve the solution in successive steps even near where it should have converged. In this case, it is essential to use a between-step procedure that guarantees that the objective function

can only improve. In the latter case, it is important to have a way to approximately compute the value of objective function. The objective function to be minimized is

$$J = \frac{1}{2}\left(\mathbf{y}-\mathbf{h}\left(\mathbf{s}\right)\right)^T \mathbf{R}^{-1}\left(\mathbf{y}-\mathbf{h}\left(\mathbf{s}\right)\right) + \frac{1}{2}\left(\mathbf{s}-\mathbf{X}\beta\right)^T \mathbf{Q}^{-1}\left(\mathbf{s}-\mathbf{X}\beta\right)$$
$$\simeq \frac{1}{2}\left(\mathbf{y}-\mathbf{h}\left(\mathbf{s}\right)\right)^T \mathbf{R}^{-1}\left(\mathbf{y}-\mathbf{h}\left(\mathbf{s}\right)\right) + \frac{1}{2}\mathbf{s}^T \boldsymbol{A}\mathbf{C}^{-1}\boldsymbol{A}^T\mathbf{s} \tag{45}$$

This expression is economical to compute, the second terms involving $O\left(Km\right)$ multiplications.

It should be noted that convergence is reached when the value of the J does not decrease much [e.g., *Kitanidis and Lane*, 1985]. For example, in the application shown in the next section, using a change of 0.01 would be enough for practical purposes.

6. Example

In this section, the methodology is tested on a toy problem. Application to a large-dimensional problem are reported elsewhere [*Lee and Kitanidis*, 2014].

Consider a problem of steady $1-D$ flow, in a domain from 0 to 1, with variable and unknown conductivity, with constant recharge, and fixed-head conditions. The domain is discretized into 100 blocks, i.e., $m = 100$.

$$\frac{d}{dx}\left(K\frac{d\phi}{dx}\right) = -N \tag{46}$$
$$\phi\left(0\right) = \phi\left(1\right) = 1 \tag{47}$$
$$N = 10^{-5} \tag{48}$$

One may think of these as dimensionless quantities for a certain unit length and unit time. The model for the inversion is cubic generalized covariance function for the prior

and uncorrelated noise

$$Q_{ij} = 200\,|x_i - x_j|^3\,,\; R_{ij} = (0.004)^2\,\delta_{ij} \tag{49}$$

The drift matrix $\boldsymbol{X}$ is m by 2; the first column of $\boldsymbol{X}$ is ones and the second is the location of blocks. The standard error in the 20 observations ($n = 20$) used, 0.004, is about 1% of the difference in head between successive observation points.

For $K = 20$, the results for the textbook solution and the PCGA method are compared against the true field. Between the two methods, the relative difference in norm between the estimates is only 0.2% and in the computed covariance corrections 0.5%. Figure 3 shows the best estimates and the confidence intervals for the two methods.

When the procedure was repeated with fewer terms, like $K = 12$, PCGA gave suboptimal results, but still surprisingly good since there is no way that all the information in 20 measurements can be captured with just 12 components. The results are shown in figure 4. The large-scale features are captured quite well, the method required fewer iterations, and each iteration requires about 40% fewer calls to the forward problem. This is an important advantage of the method. One can start with a small K performing the first few iterations quickly and then to increase K when the largest features have been identified and the solution is close to the optimum.

A strategy to reduce computational cost is to start with a very small number, like $K = 2$, obtain a solution, use it as starting value for $K = 4$, and so on every time doubling the number of components till no improvement can be achieved.

Focusing on the case $K = 20$, we will consider the interaction between the prior covariance and the forward operator, expressed through the Jacobian matrix $\boldsymbol{H}$. We will use the Jacobian matrix that corresponds to the best estimate. There are many ways

to examine interactions between $\boldsymbol{H}$ and $\boldsymbol{Q}$. Here we will compare the eigenvalues of $\boldsymbol{HQH}^T$ with those of $\boldsymbol{R} = \sigma^2\boldsymbol{I}$. Note that $\boldsymbol{HQH}^T$ quantifies variability in the observation attributable to variability in unknowns whereas $\boldsymbol{R}$ variability due to observation error ("noise"). When the first dominates, the measurements are informative about the unknowns or, to put it another way, the measurements can change the estimates more.

In figure 5 we compare the eigenvalues of the three matrices. Eigenvalues of $\boldsymbol{HQH}^T$ that are much smaller than the variance of the measurement error, σ^2, have negligible effect and they might as well be set equal to 0.

This simply is another way to understand why head observation are not informative about small scale fluctuations in the conductivity and why even modest increases in the observation error can affect resolution quite dramatically. In terms of the method proposed here, the figure suggests the reason why, if the measurement error is significant, a small number K of components is sufficient to give the solution. For example, if $\sigma = 0.04$, one would expect that 12 components would be enough even though there are 20 observations. This is verified in figure 6, where one can see that the textbook solution and the approximate with just 12 components are practically the same.

In an application the scaling of $\boldsymbol{Q}$ versus $\boldsymbol{R}$ is a critical part of the process of solving the inverse problem. However, this issue is beyond the scope of this paper that focuses on how to compute the solution.

7. Discussion

We have presented a method with computational cost that has near linear scaling with the size of the vector of the unknowns. Thus, this method is a promising alternative for problems with unknowns in the order of millions. The method has two phases.

In the first phase, the problem is parameterized in terms of orthogonal components. The method proposed herein is purely algebraic and "black box", in the sense that there is no need for analytical expansions. The user needs to provide the geostatistical model in terms of the drift and covariance functions. Assuming that one has access to software tools for fast covariance matrix vector multiplications (like FFT and hierarchical) that are becoming increasingly easier to find, this phase involves computations with near linear scaling. This is done once for every geostatistical model and involves no forward runs and Gauss-Newton iterations.

The second phase is Gauss-Newton iterations. One advantage of the method is that the number of runs of the forward model (typically a PDE solver) is reduced compared to more standard approaches. Only the forward model used as a black box is required and, in this version, there is no need for adjoint-state solvers. A second advantage is that the overall computational cost of matrix operations (excluding the forward runs) scales linearly rather than quadratically with the number of unknowns. The same holds true for storage requirements, including the computation of the posterior covariance. For a large problem, these savings can be by orders of magnitude.

One advantage of the approach is that the computations in the second phase, which is usually the most demanding part computationally, are highly parallelizable. Like in ensemble methods, one can run forward solvers in parallel.

The method is not meant as the solution to all types of inverse problems. It is best for problems where we can focus on seeking relatively smooth solutions and the measurements have limited information content, a situation that is often encountered.

This work focuses on the mathematics of the method. Only a simple and small problem is presented here, to illustrate that the mathematical equations work as they are supposed to and to give some insights into the problem. The scaling follows from a theoretical analysis of the cost of each operation. Future work will present application to large problems with analysis of actual cost and effectiveness.

Acknowledgments.

This material is based upon work supported by the Department of Energy under Award Number DE-FE0009260. Additional support was received from the National Science Foundation through its ReNUWIt Engineering Research Center (www.renuwit.org; NSF EEC-1028968).

References

Ambikasaran, S., Y. Li, E. Darve, and P. K. Kitanidis (2013), Large-scale stochastic linear inversion using hierarchical matrices, *Computational Geosciences*, *published online*, doi: 10.1007/s10596-013-9364-0.

Berger, J. O. (2006), The case for objective Bayesian analysis, *Bayesian Analysis*, *1*(3), 385–402.

Bui-Thanh, T., C. Burstedde, O. Ghattas, J. Martin, G. Stadler, and L. C. Wilcox (), Extreme-scale UQ for Bayesian inverse problems governed by PDEs, in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, pp. 1–11, IEEE Computer Society Press.

Butler, J. J. J., C. D. McElwee, and G. C. Bohling (1999), Pumping tests in networks of multilevel sampling wells; motivation and methodology, *Water Resources Research*,

35(11), 3553–3560.

Cardiff, M., W. Barrash, P. Kitanidis, B. Malama, A. Revil, S. Straface, and E. Rizzo (2009), A potential-based inversion of unconfined steady-state hydraulic tomography, *Ground Water*, *47*(2), 259–270, doi:10.1111/j.1745-6584.2008.00541.x.

Cardiff, M., W. Barrash, and P. K. Kitanidis (2013), Hydraulic conductivity imaging from 3-d transient hydraulic tomography at several pumping/observation densities, *Water Resour. Res.*, *49*, 7311–7326, doi:10.1002/wrcr.20519.

Carlin, B. P., and T. A. Louis (2000), *Bayes and empirical Bayes methods for data analysis*, 2nd edition ed., Chapman and Hall CRC, Boca Raton.

Carrera, J., and S. P. Neuman (1986), Estimation of aquifer parameters under transient and steady state conditions, 1. Maximum Likelihood method incorporating prior information, *Water Resour. Res.*, *22*(2), 199–210.

Chatterjee, A., A. M. Michalak, J. L. Anderson, K. L. Mueller, and V. Yadav (2012), Toward reliable ensemble Kalman filter estimates of CO2 fluxes, *Journal of Geophysical Research: Atmospheres*, *117*(D22), D22,306, doi:10.1029/2012jd018176.

Evensen, G. (1994), Sequential data assimilation with a non-linear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, *J. Geophys Res*, *99*('C5'), 10.143–10.162, doi:10.1029/94JC00572.

Evensen, G. (2003), The ensemble Kalman filter: Theoretical formulation and practical implementation, *Ocean Dynamics*, *53*, 343–367, doi:10.1007/s10236-003-0036-9.

Evensen, G. (2006), *Data Assimilation: The Ensemble Kalman Filter*, Springer.

Fong, W., and E. Darve (2009), The black-box fast multipole method, *Journal of Computational Physics*, *228*(23), 8712–8725, doi:10.1016/j.jcp.2009.08.031.

Gavalas, G. R., P. C. Shah, and J. H. Seinfeld (1976), Reservoir history matching by Bayesian estimation, *Soc. Petrol. Eng. J.*, *16*, 337–350, box 3C.

Golub, G. H. (1965), Numerical methods for solving linear least squares problems, *Numer. Math.*, *17*, 206–216.

Golub, G. H., and C. F. Van Loan (1989), *Matrix Computations*, 462 pp., Johns Hopkins Univ. Press, Baltimore.

Halko, N., P. G. Martinsson, and J. A. Tropp (2011), Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Review*, *53*(2), 217–288.

Kitanidis, P. K. (1983), Statistical estimation of polynomial generalized covariance functions and hydrologic applications, *Water Resources Research*, *19*(4), 909–921, box 5B.

Kitanidis, P. K. (1993), Generalized covariance functions in estimation, *Mathematical Geology*, *25*(5), 525–540.

Kitanidis, P. K. (1995), Quasilinear geostatistical theory for inversing, *Water Resour. Res.*, *31*(10), 2411–2419.

Kitanidis, P. K. (2010), Bayesian and geostatistical approaches to inverse problems, in *Large-scale inverse problems and quantification of uncertainty*, edited by L. Biegler, G. Biros, O. Ghattas, M. Heinkenschloss, D. Keyes, B. Mallick, Y. Marzouk, L. Tenorio, B. van Bloemen Waanders, and K. Willcox, pp. 71–85, Wiley.

Kitanidis, P. K., and R. W. Lane (1985), Maximum Likelihood parameter estimation of hydrologic spatial processes by the Gauss-Newton method, *J. Hydrology*, *79*, 53–71, box 4A.

Kitanidis, P. K., and E. G. Vomvoris (1983), A geostatistical approach to the inverse problem in groundwater modeling (steady state) and one-dimensional simulations, *Water Resour. Res.*, *19*(3), 677–690.

Lee, J., and P. K. Kitanidis (2014), Large-scale hydraulic tomography and joint inversion of head and tracer data using the PCGA method, *Water Resour. Res.*, *submitted*.

Li, J. Y., S. Ambikasaran, E. Darve, and P. K. Kitanidis (2014), A Kalman filter powered by H-matrices for quasi-continuous data assimilation problems, *Water Resour. Res.*, *in review*.

Li, W., and O. A. Cirpka (2006), Efficient geostatistical inverse methods for structured and unstructured grids, *Water Resour. Res.*, *42*(W06402), doi:10.1029/2005WR004668.

Linde, N., A. Binley, A. Tryggvason, L. B. Pedersen, and A. Revil (2006), Improved hydrogeophysical characterization using joint inversion of cross-hole electrical resistance and ground-penetrating radar traveltime data, *Water Resources Research*, *42*(12), W12,404, doi:10.1029/2006wr005131.

Liu, X., and P. K. Kitanidis (2011), Large-scale inverse modeling with an application in hydraulic tomography, *Water Resour. Res.*, *47*(2), W02,501, doi:10.1029/2010wr009144.

Marzouk, Y. M., and H. N. Najm (2009), Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems, *Journal of Computational Physics*, *228*(6), 1862–1902, doi:10.1016/j.jcp.2008.11.024.

Matheron, G. (1973), The intrinsic random functions and their applications, *Applied Probability*, *5*, 439–468, box 5B.

Neuman, S. P. (1980), A statistical approach to the inverse problem of aquifer hydrology, 3. Improved method and added perspectives, *Water Resour. Res.*, *16*(2), 331–346.

Nowak, W., S. Tenkleve, and O. Cirpka (2003), Efficient computation of linearized cross-covariance and auto-covariance matrices of interdependent quantities, *Mathematical Geology*, *35*(1), 53–66.

Pollock, D., and O. A. Cirpka (2012), Fully coupled hydrogeophysical inversion of a laboratory salt tracer experiment monitored by electrical resistivity tomography, *Water Resour. Res.*, *48*(1), W01,505, doi:10.1029/2011wr010779.

Rubin, Y., X. Chen, H. Murakami, and M. Hahn (2010), A Bayesian approach for inverse modeling, data assimilation, and conditional simulation of spatial random fields, *Water Resour. Res.*, *46*(10), W10,523, doi:10.1029/2009wr008799.

Saibaba, A. K., and P. K. Kitanidis (2012), Efficient methods for large-scale linear inversion using a geostatistical approach, *Water Resour. Res.*, *48*(5), W05,522.

Saibaba, A. K., S. Ambikasaran, J. Y. Li, P. K. Kitanidis, and E. F. Darve (2012), Application of hierarchical matrices to linear inverse problems in geostatistics, *Oil Gas Sci. Technol.*, *67*(5), 857–875, doi:10.2516/ogst/2012064.

Slater, L., A. M. Binley, W. Daily, and R. Johnson (2000), Cross-hole electrical imaging of a controlled saline tracer injection, *Journal of Applied Geophysics*, *44*(2), 85–102, doi:10.1016/S0926-9851(00)00002-1.

Woodbury, A. D., and Y. Rubin (2000), A full-Bayesian approach to parameter inference from tracer travel time moments and investigation of scale effects at the Cape Cod experimental site, *Water Resour. Res.*, *36*(1), 159–171.

Yadav, V., and A. M. Michalak (2013), Improving computational efficiency in large linear inverse problems: an example from carbon dioxide flux estimation, *Geosci. Model Dev.*, *6*(3), 583–590, doi:10.5194/gmd-6-583-2013.

Yeh, T.-C. J., and X. Liu (2000), Hydraulic tomography: Development of a new aquifer test method, *Water Resour. Res.*, *36*(8), 2095–2105.

Zanini, A., and P. K. Kitanidis (2009), Geostatistical inversing for large-contrast transmissivity fields, *Stochastic Environmental Research and Risk Assessment*, *23*(5), 565–577, doi:10.1007/s00477-008-0241-7.

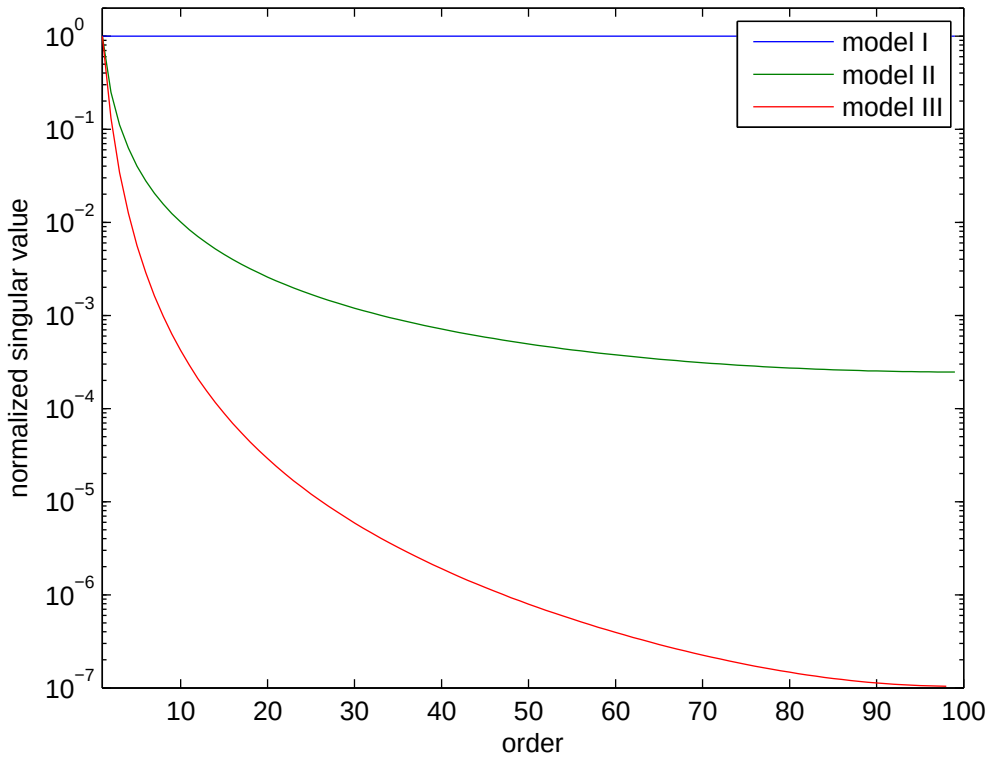

Figure 1. Normalized matrix spectrum (singular values) for three models.

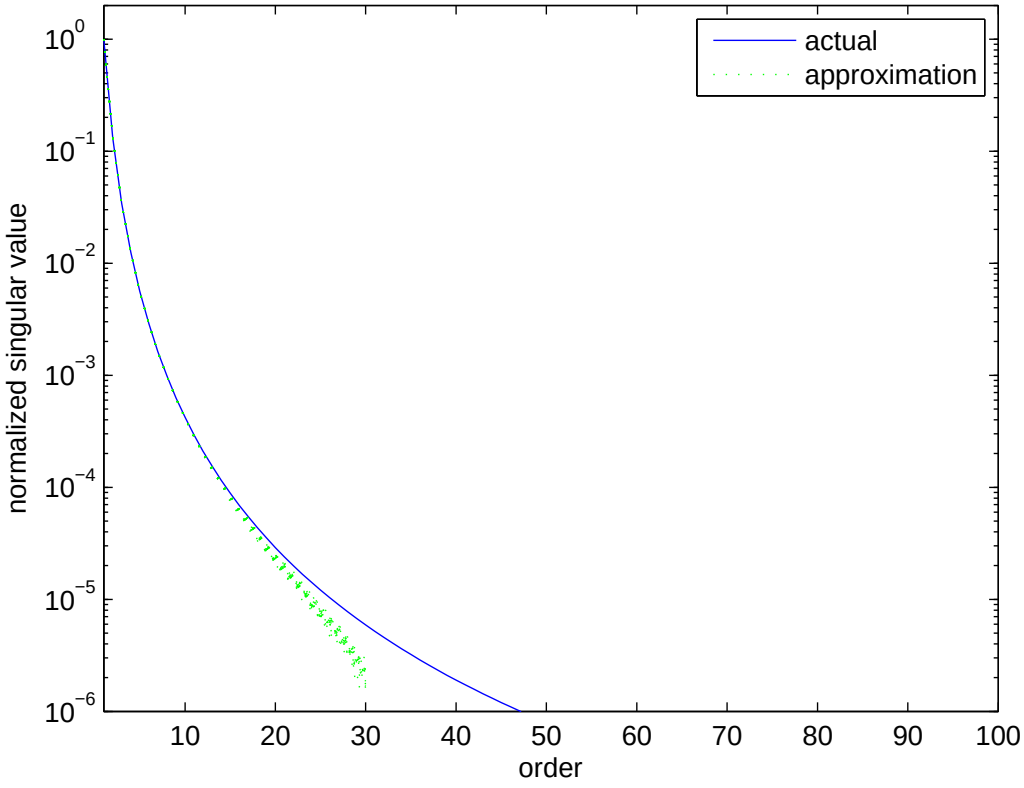

Figure 2. Results of twenty repetitions of the process for model III and $K = 30$.

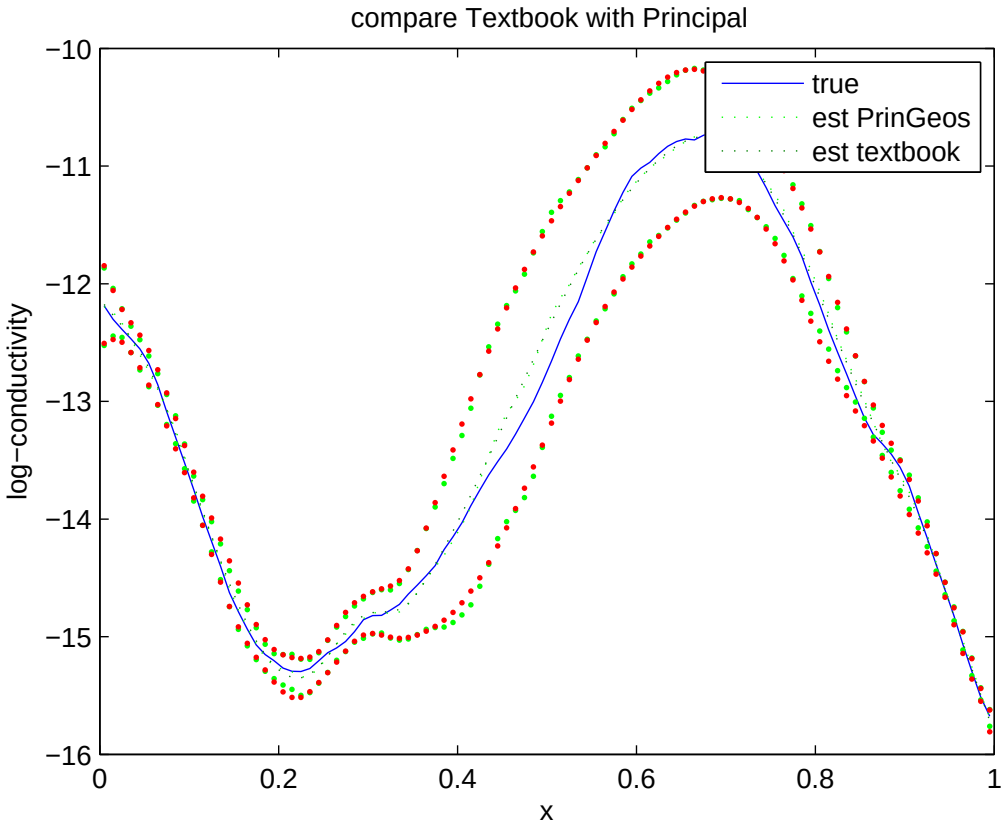

Figure 3. With K=20 components: The true, best estimate, and confidence intervals. Red is for textbook and green for the new method.

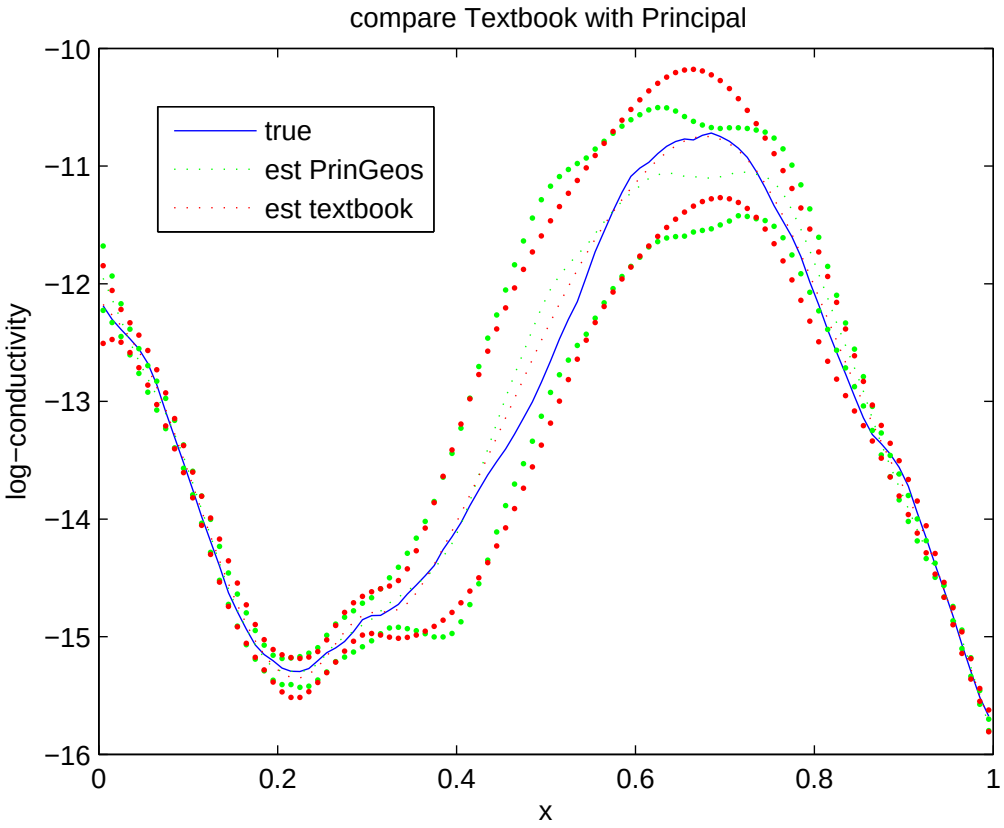

Figure 4. With K=12 components: The true, best estimate, and confidence intervals.

Red is for textbook and green for the new method.

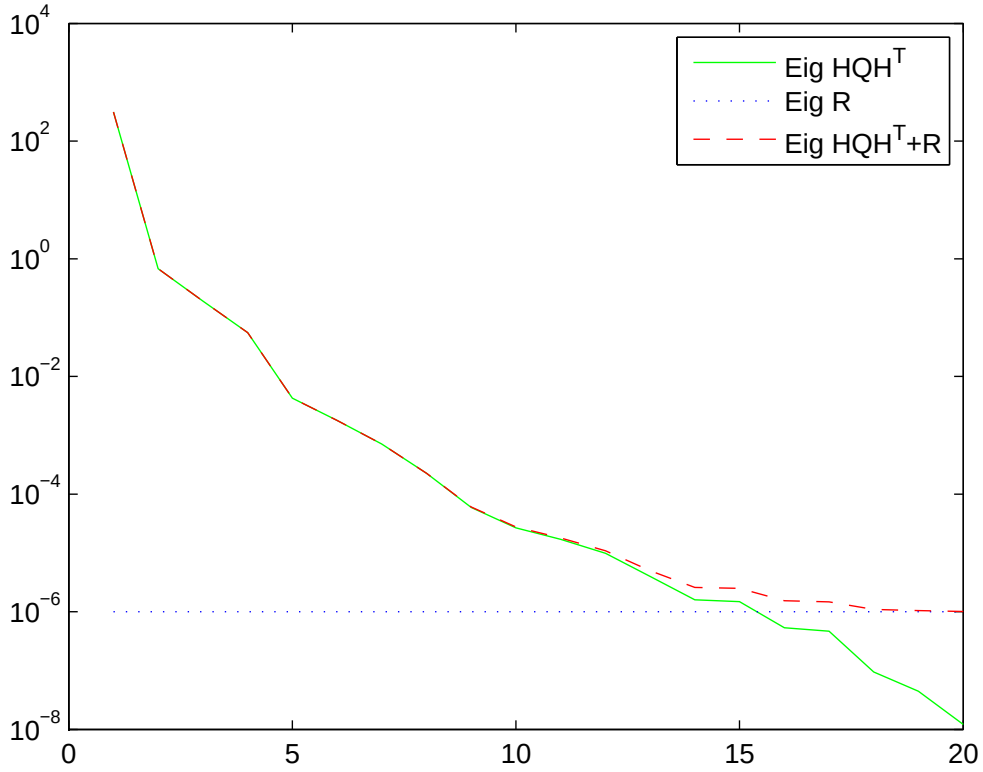

Figure 5. Comparison of eigenvalues of $\boldsymbol{HQH}^T$ and $\boldsymbol{HQH} + \boldsymbol{R}$.

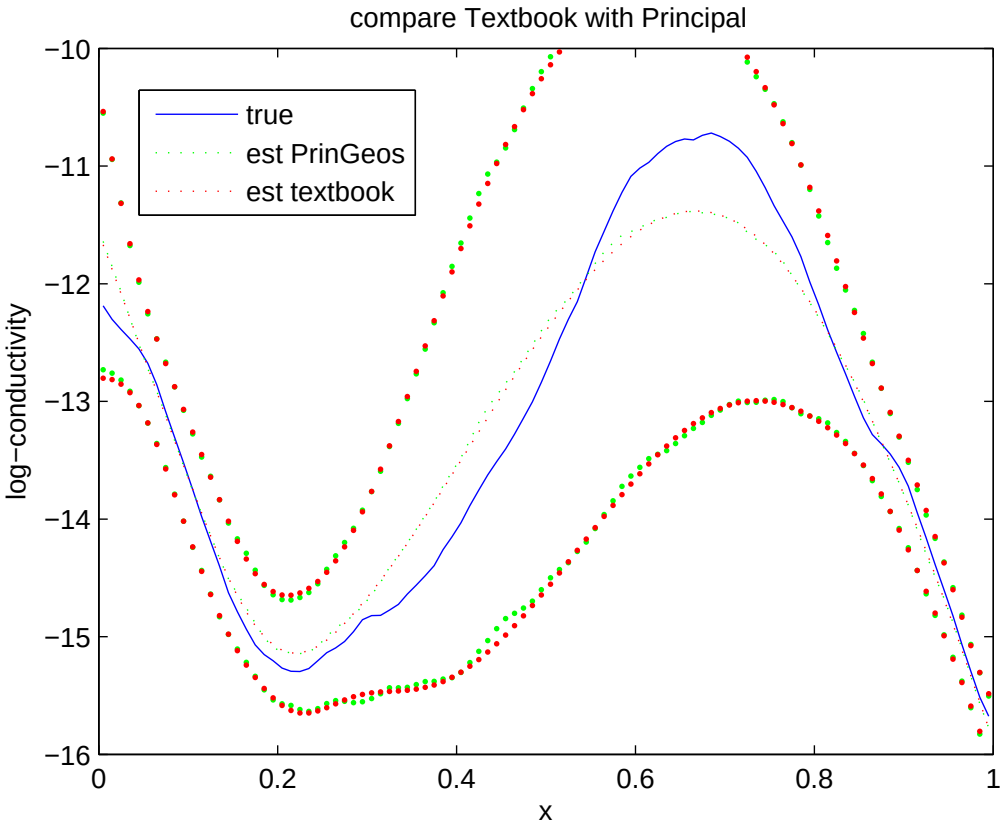

Figure 6. High-noise case with K=12 components: The true, best estimate, and confidence intervals. Red is for textbook and green for the new method.