

# UC Irvine

## UC Irvine Previously Published Works

### Title

Improving Simulation Efficiency of MCMC for Inverse Modeling of Hydrologic Systems With a Kalman-Inspired Proposal Distribution

### Permalink

<https://escholarship.org/uc/item/5637177v>

### Journal

Water Resources Research, 56(3)

### ISSN

0043-1397

### Authors

Zhang, Jiangjiang  
Vrugt, Jasper A  
Shi, Xiaoqing  
[et al.](#)

### Publication Date

2020-03-01

### DOI

10.1029/2019wr025474

Peer reviewed



# Water Resources Research

## RESEARCH ARTICLE

10.1029/2019WR025474

### Key Points:

- MCMC methods remain rather inefficient in exploring high-dimensional target distributions
- This paper introduces a Kalman-inspired proposal distribution to improve the simulation efficiency of MCMC
- The Kalman-inspired proposal distribution can be conveniently embedded in any adequate MCMC method

### Supporting Information:

- Supporting Information S1

### Correspondence to:

J. A. Vrugt and L. Zeng,  
 jasper@uci.edu;  
 lingzao@zju.edu.cn

### Citation:

Zhang, J., Vrugt, J. A., Shi, X., Lin, G., Wu, L., & Zeng, L. (2020). Improving simulation efficiency of MCMC for inverse modeling of hydrologic systems with a Kalman-inspired proposal distribution. *Water Resources Research*, 56, e2019WR025474. <https://doi.org/10.1029/2019WR025474>

Received 2 MAY 2019

Accepted 19 FEB 2020

Accepted article online 20 FEB 2020

## Improving Simulation Efficiency of MCMC for Inverse Modeling of Hydrologic Systems With a Kalman-Inspired Proposal Distribution

Jiangjiang Zhang<sup>1</sup>, Jasper A. Vrugt<sup>2,3,1</sup>, Xiaoqing Shi<sup>4</sup>, Guang Lin<sup>5</sup>, Laosheng Wu<sup>6</sup>, and Lingzao Zeng

<sup>1</sup>Zhejiang Provincial Key Laboratory of Agricultural Resources and Environment, Institute of Soil and Water Resources and Environmental Science, College of Environmental and Resource Sciences, Zhejiang University, Hangzhou, China, <sup>2</sup>Department of Civil and Environmental Engineering, University of California, Irvine, CA, USA, <sup>3</sup>Department of Earth System Science, University of California, Irvine, CA, USA, <sup>4</sup>Key Laboratory of Surficial Geochemistry of Ministry of Education, School of Earth Sciences and Engineering, Nanjing University, Nanjing, China, <sup>5</sup>Department of Mathematics and School of Mechanical Engineering, Purdue University, West Lafayette, IN, USA, <sup>6</sup>Department of Environmental Sciences, University of California, Riverside, CA, USA

**Abstract** Bayesian analysis is widely used in science and engineering for real-time forecasting, decision making, and to help unravel the processes that explain the observed data. These data are some deterministic and/or stochastic transformations of the underlying parameters. A key task is then to summarize the posterior distribution of these parameters. When models become too difficult to analyze analytically, Monte Carlo methods can be used to approximate the target distribution. Of these, Markov chain Monte Carlo (MCMC) methods are particularly powerful. Such methods generate a random walk through the parameter space and, under strict conditions of reversibility and ergodicity, will successively visit solutions with frequency proportional to the underlying target density. This requires a proposal distribution that generates candidate solutions starting from an arbitrary initial state. The speed of the sampled chains converging to the target distribution deteriorates rapidly, however, with increasing parameter dimensionality. In this paper, we introduce a new proposal distribution that enhances significantly the efficiency of MCMC simulation for highly parameterized models. This proposal distribution exploits the cross covariance of model parameters, measurements, and model outputs and generates candidate states much alike the analysis step in the Kalman filter. We embed the Kalman-inspired proposal distribution in the Differential Evolution Adaptive Metropolis algorithm during burn-in and present several numerical experiments with complex, high-dimensional, or multimodal target distributions. Results demonstrate that this new proposal distribution can greatly improve simulation efficiency of MCMC. Specifically, we observe a speedup on the order of 10–30 times for groundwater models with more than 100 parameters.

## 1. Introduction and Scope

Mathematical modeling has become an integral part of the scientific method. Computer simulation is particularly useful for hypothesis testing, decision making, to gain (new) insights and understanding of system functioning, and to predict system behavior into the space (interpolation) and time (forecasting) domain. The complexity of hydrologic systems poses significant modeling challenges, in particular how to characterize adequately water flow and storage in the presence of (often) incomplete and insufficient observations, process knowledge, and system characterization. This necessitates a systematic framework for uncertainty quantification of model simulations. These uncertainties originate from model conceptualization and data collection and include model structural errors, measurement errors of the initial conditions, forcing data and model output data, and uncertainty in the model parameters (Kavetski et al., 2006a; 2006b; Refsgaard et al., 2012; Renard et al., 2011; Vrugt et al., 2005; Wagener & Gupta, 2005; Zheng et al., 2018).

Let us assume that the data-generating process of some arbitrary real-world system can be written as

$$\tilde{\mathbf{d}} = f(\boldsymbol{\theta}) + \boldsymbol{\epsilon}, \quad (1)$$

where  $\tilde{\mathbf{d}} = \{\tilde{d}_1, \dots, \tilde{d}_n\}$  is an  $n$ -vector of measurements,  $f(\cdot)$  signifies a computer model of the system of interest,  $\boldsymbol{\theta} = \{\theta_1, \dots, \theta_k\}$  denotes a  $k$ -vector of model parameters, and  $\boldsymbol{\epsilon} = \{\epsilon_1, \dots, \epsilon_n\}$  represents an  $n$ -vector of measurement errors, respectively. The Bayesian paradigm treats the parameters in equation (1) as random variables with joint probability density function. This multivariate distribution, the so-called posterior parameter distribution,  $p(\boldsymbol{\theta}|\tilde{\mathbf{d}})$ , is the consequence of two antecedents, a prior distribution,  $p(\boldsymbol{\theta})$ , which captures our initial degree of beliefs in the values of the model parameters, and a likelihood function,  $L(\boldsymbol{\theta}|\tilde{\mathbf{d}}) \equiv p(\tilde{\mathbf{d}}|\boldsymbol{\theta})$ , which quantifies by the rules of probability theory the level of confidence in the parameter values,  $\boldsymbol{\theta}$ , in light of the observed data,  $\tilde{\mathbf{d}}$ . Bayes's theorem expresses mathematically the relationship between the prior, conditional, and posterior distribution of the parameters,  $\boldsymbol{\theta}$ , as follows:

$$p(\boldsymbol{\theta}|\tilde{\mathbf{d}}) = \frac{p(\boldsymbol{\theta})p(\tilde{\mathbf{d}}|\boldsymbol{\theta})}{\int p(\boldsymbol{\theta})p(\tilde{\mathbf{d}}|\boldsymbol{\theta})d\boldsymbol{\theta}} \propto p(\boldsymbol{\theta})L(\boldsymbol{\theta}|\tilde{\mathbf{d}}), \quad (2)$$

where the denominator,  $p(\tilde{\mathbf{d}}) = \int p(\boldsymbol{\theta})p(\tilde{\mathbf{d}}|\boldsymbol{\theta})d\boldsymbol{\theta}$ , is the so-called Bayesian evidence, evidence, marginal likelihood, or model likelihood. This normalization constant guarantees that  $p(\boldsymbol{\theta}|\tilde{\mathbf{d}})$  integrates to unity, that is,  $\int_{\Theta} p(\boldsymbol{\theta}|\tilde{\mathbf{d}})d\boldsymbol{\theta} = 1$ , where  $\Theta \subseteq \mathbb{R}^k$  signifies the feasible parameter space and  $p(\boldsymbol{\theta}|\tilde{\mathbf{d}}) \geq 0$  for all  $\boldsymbol{\theta} \in \Theta$ . The evidence is of great importance for hypothesis testing via model selection (Cao et al., 2018; Volpi et al., 2017; Zeng, Ye, et al., 2018) but can be discarded if our interest lies in estimation of the posterior distribution of the parameters. Henceforth, one often removes the denominator from Bayes's theorem and works instead with the unnormalized density, that is, the right-hand side of equation (2). When the measurement errors,  $\boldsymbol{\epsilon}$ , are normally distributed with zero mean and  $n \times n$  covariance matrix,  $\mathbf{R}$ , that is,  $\boldsymbol{\epsilon} \sim \mathcal{N}_n(\mathbf{0}, \mathbf{R})$ , the likelihood function can be expressed as follows:

$$L(\boldsymbol{\theta}|\tilde{\mathbf{d}}) = \frac{1}{(2\pi)^{n/2}|\mathbf{R}|^{1/2}} \exp \left\{ -\frac{1}{2} [\tilde{\mathbf{d}} - f(\boldsymbol{\theta})]^T \mathbf{R}^{-1} [\tilde{\mathbf{d}} - f(\boldsymbol{\theta})] \right\}, \quad (3)$$

where  $|\cdot|$  signifies the determinant operator. Nevertheless, in some cases, one has to consider correlated, heteroscedastic, and non-Gaussian measurement errors. Then other forms of likelihood functions can be adopted, for example, the formal likelihood function proposed by Schoups and Vrugt (2010).

A key task is now to summarize the posterior parameter distribution,  $p(\boldsymbol{\theta}|\tilde{\mathbf{d}})$ . In most practical cases,  $p(\boldsymbol{\theta}|\tilde{\mathbf{d}})$  cannot be derived by analytical means nor by analytical approximation, and Monte Carlo methods can be used to sample the target distribution. Of these, Markov chain Monte Carlo (MCMC) methods have become increasingly popular in the past decades. Such methods generate a (quasi-)random walk through the parameter space and collect the sampled solutions in one or more Markov chains. Under strict conditions of reversibility and ergodicity, the chain(s) will gradually converge to an equilibrium distribution equivalent to the target distribution. This means that if one looks at the archived values of  $\boldsymbol{\theta}$  in the chain(s) sufficiently far from the arbitrary initial state(s), that is, after the so-called burn-in period, then these successively generated states will be distributed according to  $p(\boldsymbol{\theta}|\tilde{\mathbf{d}})$ , the unknown target distribution of  $\boldsymbol{\theta}$ . In practice, it is necessary to monitor convergence of the sampled chain(s) with some diagnostic metrics (Brooks & Gelman, 1998). Then one discards the Markov chain states in the burn-in period to minimize the effect of initial values on the posterior inference and use the remaining states to derive the desired statistics of  $p(\boldsymbol{\theta}|\tilde{\mathbf{d}})$ .

The earliest MCMC method, that is, the random walk Metropolis (RWM) algorithm, was developed by Metropolis et al. (1953) and works as follows. First, a candidate,  $\boldsymbol{\theta}_p$ , is sampled from a symmetric proposal distribution,  $q(\cdot)$ , which is centered around the current state,  $\boldsymbol{\theta}_{(t-1)}$ , of the Markov chain. Next, the candidate is accepted with Metropolis probability

$$P_{\text{acc}}(\boldsymbol{\theta}_{(t-1)} \rightarrow \boldsymbol{\theta}_p) = \min \left[ 1, \frac{p(\boldsymbol{\theta}_p)L(\boldsymbol{\theta}_p|\tilde{\mathbf{d}})}{p(\boldsymbol{\theta}_{(t-1)})L(\boldsymbol{\theta}_{(t-1)}|\tilde{\mathbf{d}})} \right] = \min \left[ 1, \frac{p(\boldsymbol{\theta}_p|\tilde{\mathbf{d}})}{p(\boldsymbol{\theta}_{(t-1)}|\tilde{\mathbf{d}})} \right]. \quad (4)$$

Finally, if the candidate is accepted, the chain moves to  $\boldsymbol{\theta}_p$ ; otherwise, the chain remains at its current state,  $\boldsymbol{\theta}_{(t-1)}$ . Repeated application of these three steps results in a Markov chain with equilibrium distribution,  $p(\boldsymbol{\theta}|\tilde{\mathbf{d}})$ . Hastings (1970) generalized the RWM algorithm to asymmetric proposal distributions when the probability density of the forward jump,  $q(\boldsymbol{\theta}_p|\boldsymbol{\theta}_{(t-1)})$ , does not equal the probability density of the backward

jump,  $q(\theta_{(t-1)}|\theta_p)$ . A simple correction of the acceptance probability will neutralize this imbalance in jump probabilities

$$p_{\text{acc}}(\theta_{(t-1)} \rightarrow \theta_p) = \min \left[ 1, \frac{p(\theta_p|\tilde{\mathbf{d}})}{p(\theta_{(t-1)}|\tilde{\mathbf{d}})} \cdot \frac{q(\theta_{(t-1)}|\theta_p)}{q(\theta_p|\theta_{(t-1)})} \right], \quad (5)$$

where  $q(\theta_{(t-1)}|\theta_p)$  and  $q(\theta_p|\theta_{(t-1)})$  signify the conditional probabilities of trial moves from  $\theta_p$  to  $\theta_{(t-1)}$  and from  $\theta_{(t-1)}$  to  $\theta_p$ , respectively. Equation (5) is also known as the Metropolis-Hastings (MH) algorithm and has enabled the practical application of Bayesian inference to a very large class of models and data. Indeed, the MH algorithm has revolutionized the field of computational statistics and because of this reason was elected as one of the top 10 most important algorithms of the twentieth century (Beichl & Sullivan, 2000).

The efficiency of the RWM and MH algorithms depends in large part on the scale and orientation of the proposal distribution,  $q(\cdot)$ , used to create trial moves (transitions) in the Markov chain. When the proposal distribution is too dispersed, a disproportionately large number of candidate states will be rejected, and the chain will only converge slowly to the target distribution. On the other hand, when the proposal distribution is too narrow (underdispersed), most candidate states will be accepted, but the chain will also travel slowly to the target distribution as the update is very small, and MCMC will experience a very long burn-in period.

Much research has been devoted in the past decades to improving the efficiency of MCMC methods (Brooks et al., 2011; Calderhead, 2014; Gilks & Roberts, 1996). This includes the use of kernel adaptation (Gilks et al., 1994; Haario et al., 1999, 2001; Kuczera & Parent, 1998; Vrugt et al., 2003), Hamiltonian dynamics (Duane et al., 1987; Hoffman & Gelman, 2014; Neal, 2011), multitype proposals (Laloy & Vrugt, 2012; Liu et al., 2000), Langevin dynamics (Girolami & Calderhead, 2011; Roberts & Rosenthal, 1998; Roberts & Stramer, 2002), parallel marginalization (Weare, 2007), delayed rejection (Haario et al., 2006), bacterial kernels (Yang & Rodríguez, 2013), kernel coupling and multiple chain simulation (Craiu et al., 2009; ter Braak, 2006; ter Braak & Vrugt, 2008; Vrugt, ter Braak, et al., 2008; Vrugt et al., 2009), early rejection (Laloy et al., 2013; Solonen et al., 2012), parallelization (Calderhead, 2014; Neiswanger et al., 2013), and prefetching (Brockwell, 2006; Strid, 2010). These algorithms should, at least in theory, converge asymptotically to the target distribution. Yet their performance can differ tremendously over a finite number of model evaluations.

Notwithstanding these methodological advances, MCMC methods still remain rather inefficient in exploring high-dimensional parameter spaces. This severely impairs their practical application to complex system models. The goal of this paper is twofold. First, we introduce a new proposal distribution that enhances significantly the efficiency of MCMC simulation for highly parameterized system models. This proposal distribution was suggested in Vrugt et al. (2013) and generates candidate states much alike the analysis step in the Kalman filter (Kalman, 1960). However, Vrugt et al. (2013) only foresaw the potential of this new idea and it has never been actually implemented and systematically examined. Second, to turn this raw idea into a reliable product, we incorporate the Kalman-inspired proposal distribution with a selection probability,  $p_K \in (0, 1)$ , in various MCMC algorithms, especially the Differential Evolution Adaptive Metropolis (DREAM) algorithm (Vrugt, 2016). The DREAM algorithm has shown to work well in a large array of inverse problems, involving high-dimensional and/or multimodal target distributions with complex multivariate parameter dependencies (Bikowski et al., 2012; Laloy & Vrugt, 2012; Laloy et al., 2013; Muleta et al., 2012; Ramin et al., 2014; Shi et al., 2012, 2014; Wöhling & Vrugt, 2011; Zhang et al., 2018). As the Kalman-inspired proposal can introduce asymmetry to the sampled chains, it is suggested to restrict its use only during burn-in (recommended) or use the MH algorithm or a randomized move direction to neutralize the asymmetry. The resulting DREAM<sub>(KZS)</sub> algorithm incorporates the Kalman-inspired proposal distribution and is a sibling of the DREAM family (Vrugt, 2016).

The remainder of this paper is organized as follows. In section 2, we introduce elements and theory of the Kalman-inspired proposal distribution and provide a detailed recipe and discussion of the DREAM<sub>(KZS)</sub> algorithm. Section 3 presents the results of the DREAM<sub>(KZS)</sub> algorithm for several numerical studies with complex, high-dimensional, or multimodal target distributions. In this section we are especially concerned with benchmark analysis of the DREAM<sub>(KZS)</sub> algorithm. Moreover, to demonstrate that the Kalman-inspired proposal distribution can be conveniently embedded in any adequate MCMC method, we further test its performance in two plainer MCMC algorithms (Haario et al., 2001, 2006). Finally, we conclude this paper and provide some further discussion in section 4.

## 2. Methods

In this section, we introduce the Kalman-inspired proposal distribution and discuss its implementation in the DREAM algorithm that adopts a mix of parallel direction and snooker proposal distributions, that is, the DREAM<sub>(ZS)</sub> algorithm (Vrugt, 2016).

### 2.1. MCMC Simulation

The core of the MH algorithm can be written in just a few lines (see Algorithm 1). This algorithm simulates a single chain trajectory using the proposal distribution,  $q(\cdot)$ , a random number generator on the (0, 1) interval,  $z \sim \mathcal{U}(0, 1)$ , and the target density,  $p(\theta|\tilde{\mathbf{d}})$  in equation (2).

---

#### Algorithm 1 Metropolis-Hastings

---

```

1: Define the desired length of the Markov chain,  $T$ .
2: Draw the initial chain state,  $\theta_{(1)}$ , from the prior distribution,  $\theta_{(1)} \sim p(\theta)$ .
3: Evaluate the unnormalized target density,  $p(\theta_{(1)}|\tilde{\mathbf{d}}) = p(\theta_{(1)})L(\theta_{(1)}|\tilde{\mathbf{d}})$ , at  $\theta_{(1)}$ .
4: for  $t = 2, \dots, T$  do
5:   Sample a candidate,  $\theta_p$ , from a proposal distribution,  $q(\cdot|\theta_{(t-1)})$ ,  $\theta_p \sim q(\theta|\theta_{(t-1)})$ .
6:   Compute the unnormalized target density,  $p(\theta_p|\tilde{\mathbf{d}}) = p(\theta_p)L(\theta_p|\tilde{\mathbf{d}})$ , at  $\theta_p$ .
7:   Calculate the acceptance probability,  $p_{\text{acc}}(\theta_{(t-1)} \rightarrow \theta_p)$ , in equation (5).
8:   Draw a random sample,  $z \sim \mathcal{U}(0, 1)$ .
9:   if  $p_{\text{acc}}(\theta_{(t-1)} \rightarrow \theta_p) \geq z$  then
10:     Set  $\theta_{(t)} = \theta_p$ ,
11:   else
12:     Set  $\theta_{(t)} = \theta_{(t-1)}$ .
13:   end if
14: end for

```

---

The user is free to select the proposal distribution as the acceptance probability in equation (5) preserves the underlying stationary distribution of the Markov chain. A common choice is the multivariate normal distribution,  $q = \mathcal{N}_k(\cdot, s_k \Sigma)$ , with covariance matrix,  $\Sigma$ , and scaling factor,  $s_k$ . To enhance search efficiency and protect against an inadequate selection of the proposal distribution, we can update the scale and orientation of the proposal distribution every  $m \gg 1$  iterations using all past samples stored in the chain,  $\Sigma = \text{Cov}(\{\theta_{(1)}, \dots, \theta_{(t-1)}\}) + \varphi \mathbf{I}_k$ , where  $\mathbf{I}_k$  denotes the  $k \times k$  identity matrix and  $\varphi = 10^{-6}$  is a small scalar that prevents the collapse of the sample covariance matrix to singularity (jumps become zero). This is the so-called Adaptive Metropolis (AM) algorithm of Haario et al. (2001). As a basic choice, the scaling factor is often chosen to be  $s_k = 2.38^2/k$ , which has proven optimal for Gaussian target and proposal distributions (Gelman et al., 1996; Roberts et al., 1997), and should result in an acceptance rate close to 0.44 for  $k = 1$ , 0.28 for  $k = 5$ , and 0.23 for a large  $k$ . It should be evident that the choice of the proposal distribution determines in large part the sampling efficiency and convergence speed of the Markov chain. What would be desirable is a proposal distribution that can transit the chain to the target distribution quickly.

### 2.2. The Kalman-Inspired Proposal Distribution

This section introduces an alternative proposal distribution designed to accelerate the movement of the chain to the posterior distribution,  $p(\theta|\tilde{\mathbf{d}})$ . In this section, we present the theory and formulas of the so-called Kalman-inspired proposal distribution.

This new proposal distribution is inspired by the state analysis step in the Kalman filter (Kalman, 1960). The Kalman-inspired proposal distribution uses explicitly the distance of the  $n$ -vector of simulated model outputs,  $\mathbf{d} = f(\theta)$ , to the observed data,  $\tilde{\mathbf{d}}$ ,

$$\begin{aligned}
 \theta_p &= \theta_{(t-1)} + \mathbf{C}_{\theta\mathbf{d}}(\mathbf{C}_{\mathbf{d}\mathbf{d}} + \mathbf{R})^{-1} \left[ \tilde{\mathbf{d}} + \epsilon_{(t-1)} - f(\theta_{(t-1)}) \right] \\
 &= \theta_{(t-1)} + \mathbf{K}\mathbf{r}_{(t-1)} + \mathbf{K}\epsilon_{(t-1)} \\
 &= \theta_{(t-1)} + \Delta\theta_{(t-1)},
 \end{aligned} \tag{6}$$

where  $\mathbf{C}_{\theta\mathbf{d}} = \text{Cov}(\theta, \mathbf{d})$  denotes the  $k \times n$  cross-covariance matrix of model parameters and model outputs,  $\mathbf{C}_{\mathbf{d}\mathbf{d}} = \text{Cov}(\mathbf{d}, \mathbf{d})$  signifies the  $n \times n$  autocovariance matrix of model outputs,  $\mathbf{R}$  is the  $n \times n$  covariance matrix of measurement errors in equation (3),  $\epsilon_{(t-1)}$  is an  $n$ -vector of random draw from the distribution of

measurement errors,  $\epsilon_{(t-1)} \sim \mathcal{N}_n(\mathbf{0}, \mathbf{R})$ ,  $\mathbf{r}_{(t-1)} = \tilde{\mathbf{d}} - f(\theta_{(t-1)})$  is the residual vector of parameters  $\theta_{(t-1)}$ , and  $\mathbf{K}$  constitutes the so-called Kalman gain

$$\mathbf{K} = \mathbf{C}_{\theta_d}(\mathbf{C}_{dd} + \mathbf{R})^{-1}. \quad (7)$$

It is noted here that the Kalman-inspired proposal uses covariances that converge, after burn-in, toward posterior covariances. Besides, this proposal is only used to generate a candidate point,  $\theta_p$ , from the current chain state,  $\theta_{(t-1)}$ , not to directly estimate the posterior distribution,  $p(\theta|\tilde{\mathbf{d}})$ . In other words, the purpose and implementation of the Kalman-inspired proposal distribution are quite different from the Kalman filter or its Monte Carlo variants.

The jump vector of the Kalman-inspired proposal,  $\Delta\theta_{(t-1)}$ , in equation (6) is made up of two different components: a deterministic displacement vector,  $\mathbf{K}\mathbf{r}_{(t-1)}$ , which orients the jump toward the “true” parameters,  $\theta^*$ , of the data-generating process, and a random displacement vector,  $\mathbf{K}\epsilon_{(t-1)}$ , with zero mean normally distributed variables, which introduce randomness into the sampled candidate states. This latter term of the jump vector,  $\Delta\theta_{(t-1)}$ , is of crucial importance as it enables the Kalman-inspired proposal to sample, with nonzero probability, all possible states of the target distribution. This ensures ergodicity of the sampled chain. From equation (6), it is clear that the width of the Kalman-inspired proposal distribution, which is determined by the random displacement part,  $\mathbf{K}\epsilon_{(t-1)}$ , should keep relatively stable with time, as the distribution of measurement errors is constant. On the contrary, the deterministic jump distance,  $\mathbf{K}\mathbf{r}_{(t-1)}$ , should decrease over time (until convergence has been achieved), as the difference between the simulated model outputs,  $f(\theta)$ , and measurement data,  $\tilde{\mathbf{d}}$ , is expected to decrease with the evolution of the Markov chain.

As mentioned above, the deterministic component,  $\mathbf{K}\mathbf{r}_{(t-1)}$ , of the jump vector,  $\Delta\theta_{(t-1)}$ , will guide the candidate states to the “true” parameters,  $\theta^*$ , of the data-generating process. This will help shorten burn-in yet introduce asymmetry in the sampled candidate states. We must account for this asymmetry of the Kalman-inspired proposal distribution to preserve the unique stationary distribution of the Markov chain. Equation (5) will help to remedy the dissimilar forward  $q(\theta_p|\theta_{(t-1)})$  and backward  $q(\theta_{(t-1)}|\theta_p)$  jump probabilities of equation (6). Alternatively, we can follow the suggestion made by Vrugt et al. (2013) to randomize the direction of the jump vector,  $\Delta\theta_{(t-1)}$ , as follows:

$$\begin{aligned} \theta_p &= \theta_{(t-1)} \pm (\mathbf{K}\mathbf{r}_{(t-1)} + \mathbf{K}\epsilon_{(t-1)}) \\ &= \theta_{(t-1)} \pm \Delta\theta_{(t-1)}. \end{aligned} \quad (8)$$

This modification enforces symmetry of the Kalman-inspired proposal distribution with equal selection probability of the direction of the jump vector,  $\Delta\theta_{(t-1)}$ .

The two approaches described above are easy to implement in practice and will guarantee detailed balance of the sampled Markov chain. These so-called reversibility patches do have an undesired side effect, that is, they deteriorate considerably the sampling efficiency of the Kalman-inspired proposal distribution, although they are theoretically appealing. This is easily demonstrated with numerical experiments. For example, if we use equation (5) to account for the asymmetry of the Kalman-inspired proposal distribution, then trial moves from  $\theta_{(t-1)}$  to  $\theta_p$  will be punished heavily as  $q(\theta_p|\theta_{(t-1)}) \gg q(\theta_{(t-1)}|\theta_p)$ . The Hastings correction will therefore decrease the acceptance rate of candidate states. Unfortunately, the symmetric Kalman kernel of equation (8) provides no solace. Many of the trial moves of this proposal distribution will go to waste (especially when  $\theta_p = \theta_{(t-1)} - \Delta\theta_{(t-1)}$ ) as they lead the chain away from the target distribution. Thus, both reversibility patches may defeat the purpose of the Kalman-inspired proposal distribution. To take optimal advantage of the Kalman-inspired proposal distribution, we limit its application to only the first  $T_K$  steps of the Markov chain. During this prescribed, relatively short burn-in period, the Kalman-inspired proposal will guide the chain to the target distribution and the resulting samples are discarded. At the end of this period we will switch the chain to a reversible proposal distribution that uses only information from past and/or present chain states to generate trial moves.

The Kalman-inspired proposal distribution can be conveniently embedded in any adequate MCMC method. In the next section, we present the implementation of this proposal distribution in the DREAM<sub>(ZS)</sub> algorithm and an algorithmic outline of the resulting DREAM<sub>(KZS)</sub> algorithm.

### 2.3. The DREAM<sub>(KZS)</sub> Algorithm

The DREAM<sub>(KZS)</sub> algorithm is an adaptive, multiple chain MCMC method, which exploits information from past sampled states of model parameters and outputs to rapidly explore the parameter space in pursuit of the target distribution,  $p(\theta|\tilde{\mathbf{d}})$ . This method is an extension of the DREAM<sub>(ZS)</sub> algorithm, which in turn has its roots within DE-MC<sub>(ZS)</sub> of ter Braak and Vrugt (2008). The next two subsections introduce the DREAM<sub>(KZS)</sub> algorithm and provide an algorithmic recipe of this new MCMC sampler.

#### 2.3.1. DREAM<sub>(KZS)</sub>: In Words

Let  $\Theta_{(1)} = \{\theta_{(1)}^1, \dots, \theta_{(1)}^N\}$  be an  $N \times k$  matrix with the initial states of the  $N$  chains. These states are drawn from the prior distribution. Similarly, let  $\mathbf{Z} = \{\theta^1, \dots, \theta^{m_0}\}$  be an archive with  $m_0$  ( $m_0 \gg N$ ) draws from the prior distribution. If  $\mathcal{A}$  is a subset of  $k^-$  dimensions of the original parameter space,  $\mathbb{R}^{k^-} \subseteq \mathbb{R}^k$ , then a candidate,  $\theta_p^i$ , in the  $i$ th chain,  $i = 1, \dots, N$ , at iteration  $t$ ,  $t = 2, \dots, T$ , is calculated from the samples in the archive,  $\mathbf{Z}$ , using a mix of parallel direction (Price et al., 2006; Storn & Price, 1997; Vrugt, ter Braak, et al., 2008; Vrugt et al., 2009)

$$\Delta\theta_{l \in \mathcal{A}}^i = \zeta_l + (1 + \lambda_l)\gamma_{P(\delta, k^-)} \sum_{j=1}^{\delta} (\mathbf{Z}_l^{a_j} - \mathbf{Z}_l^{b_j}) \quad \text{if } u \in [0, p_P], \quad (9)$$

$$\Delta\theta_{l \notin \mathcal{A}}^i = 0$$

snooker (Laloy & Vrugt, 2012; ter Braak & Vrugt, 2008; Vrugt, 2016)

$$\Delta\theta^i = \zeta + (1 + \lambda)\gamma_S(\mathbf{Z}_\perp^{a_1} - \mathbf{Z}_\perp^{b_1}) \quad \text{if } u \in (p_P, p_P + p_S], \quad (10)$$

and Kalman trial moves

$$\Delta\theta^i = \gamma_K [\mathbf{K}_{(t-1)} \mathbf{r}_{(t-1)}^i + \mathbf{K}_{(t-1)} \epsilon_{(t-1)}^i] \quad \text{if } u \in (p_P + p_S, 1] \text{ and } t \leq T_K, \quad (11)$$

where  $u$  is a random draw from a standard uniform distribution,  $u \sim \mathcal{U}(0, 1)$ ;  $p_P$  and  $\gamma_P = 2.38/\sqrt{2\delta k^-}$ ,  $p_S$  and  $\gamma_S \sim \mathcal{U}(1.2, 2.2)$ , and  $p_K = 1 - p_P - p_S$  and  $\gamma_K = 1$  signify the selection probabilities and jump rates of the parallel direction, snooker, and Kalman-inspired proposal distributions, respectively;  $\delta$  denotes the number of chain pairs of the parallel direction proposal;  $\mathbf{a} = \{a_1, \dots, a_\delta\}$  and  $\mathbf{b} = \{b_1, \dots, b_\delta\}$  are  $\delta$ -vectors with integers randomly drawn from  $\{1, \dots, m\}$  without replacement;  $\lambda$  and  $\zeta$  are sampled independently from  $\mathcal{U}_k(-\nu, \nu)$  and  $\mathcal{N}_k(0, \beta)$ , respectively, with  $\nu = 0.05$  and  $\beta = 10^{-12}$  are small compared to the width of the target distribution;  $\mathbf{Z}_\perp^{a_1}$  and  $\mathbf{Z}_\perp^{b_1}$  are orthogonal projection points of the samples  $\mathbf{Z}^{a_1}$  and  $\mathbf{Z}^{b_1}$  onto the line going through the current state of the  $i$ th chain,  $\theta_{(t-1)}^i$ , and sample  $\mathbf{Z}^c$  of the external archive, where  $c$  is an integer randomly drawn from  $\{1, \dots, m\}$  and  $c \neq a_1 \neq b_1$ ;  $\mathbf{r}_{(t-1)}^i$  and  $\epsilon_{(t-1)}^i$  are residual vector of parameters  $\theta_{(t-1)}^i$  and a random draw from the distribution of measurement errors. Furthermore, to enhance the probability of a direct jump between disconnected modes of the target distribution, we use  $\gamma_P = 1$  in 20% of the parallel direction proposals.

The candidate state of chain  $i$  at iteration  $t$  then becomes

$$\theta_p^i = \theta_{(t-1)}^i + \Delta\theta^i, \quad (12)$$

and the Metropolis ratio of equation (4) is used to determine whether to accept this candidate or not. If the acceptance probability,  $p_{\text{acc}}(\theta_{(t-1)}^i \rightarrow \theta_p^i)$ , is larger than or equal to a uniform random label,  $u \sim \mathcal{U}(0, 1)$ , then the candidate state is accepted and the  $i$ th chain moves to the new position, that is,  $\theta_{(t)}^i = \theta_p^i$ ; otherwise,  $\theta_{(t)}^i = \theta_{(t-1)}^i$ . The snooker candidates demand a multiplicative, Hastings-type correction,  $\alpha(\theta_{(t-1)}^i \rightarrow \theta_p^i)$ , of the acceptance probability

$$\alpha(\theta_{(t-1)}^i \rightarrow \theta_p^i) = \left( \frac{\|\theta_p^i - \mathbf{Z}^c\|}{\|\theta_{(t-1)}^i - \mathbf{Z}^c\|} \right)^{(k-1)}, \quad (13)$$

where  $\|\cdot\|$  signifies the Euclidean operator. The use of the corrected acceptance probability,  $\alpha(\theta_{(t-1)}^i \rightarrow \theta_p^i)p_{\text{acc}}(\theta_{(t-1)}^i \rightarrow \theta_p^i)$ , will negate bias in the jump direction of the snooker move and guarantee reversibility of the sampled chains (ter Braak & Vrugt, 2008).

The convergence speed of DREAM<sub>(KZS)</sub> is largely determined by the samples in the archive  $\mathbf{Z}$ . In each  $K$  generations ( $K \gg 1$ ), we augment this archive of past states with the current states,  $\Theta_{(t)} = \{\theta_{(t)}^1, \dots, \theta_{(t)}^N\}$ , of the

$N$  chains. This turns the  $\text{DREAM}_{(\text{KZS})}$  algorithm into an adaptive MCMC method with scale and orientation of the proposal that depend on the cumulative search history of the sampled chains. This proposal adaptation violates the Markovian properties, that is, the next state of the chain should depend only on the current one. Nevertheless, adaptation has been employed in various MCMC algorithms (just name a few Haario et al., 2001, 2006; Vrugt et al., 2003), to enhance search efficiency and protect against an inadequate selection of the proposal distribution. These algorithms are in theory not Markovian either. In practice, when the target distribution is bounded and the adaptation is performed under some regularity conditions, the adaptive MCMC algorithms can still retain the desired stationary distribution. In  $\text{DREAM}_{(\text{KZS})}$ , if  $m$  denotes the current number of samples in the archive, then appending  $N$  states changes  $\mathbf{Z}$  by an order of  $K/t$ , which decreases with iteration  $t$ . The three proposal distributions used in  $\text{DREAM}_{(\text{KZS})}$  thus become invariant as the length of the thinned past increases without bound. This so-called diminishing adaptation ensures ergodic chains that converge to the exact target distribution (Roberts & Rosenthal, 2007).

The  $k^-$  members of the subset  $\mathcal{A}$  are drawn at random from the entries  $\{1, \dots, k\}$  with the help of a crossover probability,  $\eta \in (0, 1]$ . Each time a proposal is generated, we draw a  $k$ -vector,  $\mathbf{u} = \{u_1, \dots, u_k\}$ , of standard uniform labels,  $\mathbf{u} \sim \mathcal{U}_k(0, 1)$ . All entries  $j$  of  $\mathbf{u}$  which satisfy  $u_j \leq \eta$  are stored in the subset  $\mathcal{A}$  and used to span the search subspace that will be sampled using equation (9).  $\mathcal{A}$  must at least have one element; otherwise, the jump vector will have zero length. To enhance search efficiency, we use a geometric series of  $n_\eta$  crossover values,  $\boldsymbol{\eta} = \{1/n_\eta, 2/n_\eta, \dots, 1\}$ , and sample the value of  $\eta$  from a discrete multinomial distribution,  $\mathcal{M}(\boldsymbol{\eta}, \mathbf{p}_\eta)$ , on  $\boldsymbol{\eta}$  with selection probabilities  $\mathbf{p}_\eta$ . The values of  $\mathbf{p}_\eta$  are tuned adaptively during burn-in by maximizing the traveled distance of the  $N$  chains. This adaptation scheme is described in detail in Vrugt, Diks, et al. (2008) and Vrugt et al. (2009). The use of a vector of crossover probabilities enables single-site Metropolis ( $\mathcal{A}$  has one element), Metropolis-within-Gibbs ( $\mathcal{A}$  has one or more elements), and regular Metropolis sampling ( $\mathcal{A}$  has  $k$  elements) and enables chains to sample outside the subspace spanned by their current positions. The default setting of  $n_\eta = 3$  has shown to work well in practice.

We are now left with a numerical implementation of the Kalman-inspired proposal distribution. We use this proposal distribution only during the first  $T_K$  generations, after which  $p_K = 0$ , and the selection probabilities of the parallel direction and snooker proposal distributions,  $p_P$  and  $p_S$ , respectively, are renormalized so that their values add up to one. The Kalman gain,  $\mathbf{K}_{(t-1)}$ , in equation (11) is computed from an archive  $\mathbf{Z}_K$  with  $m_K$  samples drawn from the chain history, as follows:

$$\mathbf{K}_{(t-1)} = \text{Cov}(\mathbf{Z}_K, f(\mathbf{Z}_K)) [\text{Cov}(f(\mathbf{Z}_K), f(\mathbf{Z}_K)) + \mathbf{R}]^{-1}, \quad (14)$$

where  $f(\mathbf{Z}_K)$  is an  $m_K \times n$  matrix with model outputs of the  $m_K$  samples in the archive  $\mathbf{Z}_K$ .

It should be noted here that although the Kalman filter is based on Gaussian assumption, an MCMC algorithm (e.g., DREAM in this work) that adopts the Kalman-inspired proposal distribution can be used to explore complex and non-Gaussian (e.g., multimodal) target distributions (see section 3.4). Due to the fact that they are both based on the Kalman formula, one might confuse the Kalman-inspired proposal distribution that only generates a candidate at a time, with ensemble Kalman filter (EnKF) and its variants (Evensen, 2009) that use an ensemble of parameter samples updated with the Kalman formula to directly approximate the posterior. In EnKF and its variants, one should use the prior covariances to calculate the Kalman gain, then update the prior ensemble with the Kalman formula, and finally use the updated ensemble to approximate the mean and covariance of the posterior. These methods are restricted to problems with Gaussian parameter distributions. However, in our work, the Kalman-inspired proposal is not designed to directly approximate the posterior but to suggest where the Markov chain might move at the next step by generating a candidate. At the very beginning of the MCMC simulation, it is very likely that the candidate states are far from the posterior mean. Implemented in an AM sampler, the Kalman-inspired proposal actually uses covariances that undergo adaptation. The covariances can converge, after burn-in, toward posterior covariances. To our best knowledge, proposal distributions used in various popular MCMC algorithms also adopt Gaussian forms, which does not mean that the MCMC algorithms are restricted to problems with normal posterior distributions.

Nevertheless, in complicated, multimodal problems, the Kalman-inspired proposal faces the risk of missing secondary modes. To improve the capability of the Kalman-inspired proposal in multimodal cases, one can modify the original Kalman-inspired proposal by adopting the local Kalman update strategy proposed by Zhang, Lin, et al. (2018), which is specifically designed to solve inverse problems with (possible) multimodal



parameter distributions. The basic idea behind the local Kalman update strategy is simple: Although globally the parameter distribution might be non-Gaussian, or even multimodal, one can still use a Gaussian distribution to describe the local parameter distribution. For example, when overall the parameter distribution is multimodal, if one only looks at the neighborhood of a single mode, the local parameter distribution can be still close to Gaussian. Using an integrated measure of distance proposed by Zhang, Lin, et al. (2018), one can find the local ensemble of each sample in the archive for the Kalman-inspired proposal. Through updating all the local ensembles separately, and generating a candidate randomly from the multiple updated local ensembles, the modified proposal distribution will have a larger chance to sample the secondary modes. As the original Kalman-inspired proposal distribution already works well in multimodal cases (see section 3.4), to prevent the length of the paper from being too long, the local Kalman update strategy is not tested here. We refer the reader to Zhang, Lin, et al. (2018) for a detailed description of the related method. Nevertheless, in the Kalman-inspired proposal, distribution of measurement errors still needs to be Gaussian. To enable a proper use of the Kalman-inspired proposal, one can transform (possible) non-Gaussian measurement errors to Gaussian variables with some transformation method (Chou et al., 1998; Sakia, 1992).

### 2.3.2. DREAM<sub>(KZS)</sub>: Algorithmic Recipe

This section provides an algorithmic recipe of the DREAM<sub>(KZS)</sub> algorithm (see Algorithm 2). This recipe translates the initial population,  $\Theta_{(1)} = \{\theta_{(1)}^1, \dots, \theta_{(1)}^N\}$ , into samples from the target distribution.

The parallel direction, snooker, and Kalman-inspired proposal distributions are used interchangeably during the first  $T_K$  generations to propose candidate states in the  $N$  chains. After this prescribed burn-in period, the selection probability of the Kalman-inspired proposal distribution is set to zero, that is,  $p_K = 0$ , and the selection probabilities of the parallel direction and snooker proposal distributions,  $p_P$  and  $p_S$ , are normalized to sum to unity. Thus, after the relatively short burn-in period, the sampled chains will maintain detailed balance.

Convergence of the sampled chains of the DREAM<sub>(KZS)</sub> algorithm can be monitored with a variety of different metrics, including within-chain and between-chain statistics. Of these, the univariate,  $\hat{R}_j$ ,  $j = 1, \dots, k$ , and multivariate,  $\hat{R}^k$ , scale reduction factors of Gelman and Rubin (1992) and Brooks and Gelman (1998), respectively, are most widely used to assess convergence for multiple chain methods. These two diagnostics compare for each parameter individually, or for the distribution as a whole, the within-chain and between-chain variances or covariance matrices, respectively. Values of  $\hat{R}_j \leq 1.2, \forall j \in \{1, \dots, k\}$  and  $\hat{R}^k \leq 1.2$  demonstrate convergence of the sampled chains to a stationary distribution.

---

#### Algorithm 2 DREAM<sub>(KZS)</sub>

---

- 1: Define  $\boldsymbol{\eta} = \{1/n_\eta, 2/n_\eta, \dots, 1\}$  and  $\mathbf{p}_\eta = \mathbf{1}_{n_\eta}/n_\eta$ . Set  $n_\eta = 3$ ,  $m_0 = 10k$  and  $K = 10$ .
  - 2: Draw an initial archive,  $\mathbf{Z}$ , from the prior distribution,  $\theta^j \sim p(\theta)$ ,  $j = 1, \dots, m_0$ .
  - 3: Draw the initial states of the  $N$  chains from the prior distribution,  $\theta_{(1)}^i \sim p(\theta)$ ,  $i = 1, \dots, N$ .
  - 4: Initialize at zero the traveled distance,  $D_l = 0$ , of each crossover value,  $l = 1, \dots, n_\eta$ .
  - 5: **for**  $t = 2, \dots, T$  **do**
  - 6: Compute  $\Delta_{(t-1)}$ , the  $k$ -vector of standard deviations of  $\Theta_{(t-1)} = \{\theta_{(t-1)}^1, \dots, \theta_{(t-1)}^N\}$ .
  - 7: Draw  $N$  labels,  $\mathbf{u} = \{u^1, \dots, u^N\}$ , from a standard uniform distribution,  $u^i \sim \mathcal{U}(0, 1)$ .
  - 8: Draw an integer,  $J$ , from the multinomial distribution,  $\mathcal{M}(\{1, 2, 3\}, \{p_P, p_S, p_K\})$ .
  - 9: If  $J = 1$ , draw a crossover value,  $\eta_l \sim \mathcal{M}(\boldsymbol{\eta}, \mathbf{p}_\eta)$ , otherwise,  $l = n_\eta$ , thus  $\eta_l = 1$ .
  - 10: **for**  $i = 1, \dots, N$  **do**
  - 11: Create a candidate,  $\theta_p^i$ , using the parallel direction ( $J = 1$ ), snooker ( $J = 2$ ) or Kalman-inspired ( $J = 3$ ) proposal distribution.
  - 12: Calculate the Metropolis ratio,  $p_{\text{acc}}(\theta_{(t-1)}^i \rightarrow \theta_p^i)$ , using equation (4).
  - 13: If  $J = 2$ , compute  $\alpha(\theta_{(t-1)}^i \rightarrow \theta_p^i)$  using equation (13), otherwise,  $\alpha(\theta_{(t-1)}^i \rightarrow \theta_p^i) = 1$ .
  - 14: If  $\alpha(\theta_{(t-1)}^i \rightarrow \theta_p^i)p_{\text{acc}}(\theta_{(t-1)}^i \rightarrow \theta_p^i) \geq u^i$ , set  $\theta_{(t)}^i = \theta_p^i$ , otherwise, set  $\theta_{(t)}^i = \theta_{(t-1)}^i$ .
  - 15: **end for**
  - 16: Update the traveled distance,  $D_l = D_l + \sum_{i=1}^N |(\theta_{(t)}^i - \theta_{(t-1)}^i)/\Delta_{(t-1)}|$  for  $\eta_l$ .
  - 17: If  $\text{mod}(t, K) = 0$ , append the current chain states,  $\Theta_{(t)} = \{\theta_{(t)}^1, \dots, \theta_{(t)}^N\}$ , to  $\mathbf{Z}$  and set  $m = m + N$ .
  - 18: If  $t \leq T_K$ , update  $\mathbf{p}_\eta$  using the cumulative traveled distance,  $\mathbf{D} = \{D_1, \dots, D_{n_\eta}\}$ , of  $\boldsymbol{\eta}$ , otherwise, set  $p_K = 0$ ,  $p_T = p_P + p_S$ ,  $p_P = p_P/p_T$  and  $p_S = p_S/p_T$ .
  - 19: **end for**
-

**Table 1**  
Description of the hmodel Parameters and Their Minimum and Maximum Values That Define the Prior Uncertainty Ranges

Parameter	Symbol	Min.	Max.	Unit	True
Maximum interception	$I_{\max}$	0.5	10	mm	3.84
Soil water storage capacity	$S_{\max}$	10	1,000	mm	776.40
Maximum percolation rate	$Q_{\max}$	0	100	mm/d	26.60
Evaporation parameter	$\alpha_E$	0	100	—	61.61
Runoff parameter	$\alpha_F$	−10	10	—	−4.18
Time constant, fast reservoir	$K_F$	0	10	days	6.01
Time constant, slow reservoir	$K_S$	0	150	days	111.67

Note. The last column lists the “true” parameter values used in our numerical experiment.

### 3. Illustrative Case Studies


In this section, we analyze, test, and evaluate the performance of the DREAM<sub>(KZS)</sub> algorithm by application to several different case studies with complex, high-dimensional, or multimodal target distributions. These studies involve mathematical models of rainfall-runoff transformation, groundwater flow, and contaminant transport. We use simulated data of known parameter values to evaluate the performance of the DREAM<sub>(KZS)</sub> algorithm and benchmark its sampling efficiency against the DREAM<sub>(ZS)</sub> algorithm. The DREAM<sub>(ZS)</sub> algorithm has found widespread application and use, and many published studies corroborate its excellent ability to rapidly sample complex, high-dimensional target distributions (Bikowski et al., 2012; Muleta et al., 2012; Ramin et al., 2014; Shi et al., 2014; Wöhling & Vrugt, 2011; Zhang et al., 2018). In the numerical experiments, we suggest to set  $p_P = 0.6$ ,  $p_S = 0.1$ ,  $p_K = 0.3$ , and  $T_K = 0.3T$  in the DREAM<sub>(KZS)</sub> algorithm, where  $T$  signifies the maximum number of chain generations. Moreover, to demonstrate that the Kalman-inspired proposal distribution can be conveniently embedded in any adequate MCMC method, in section 3.1, we further test the performance of the Kalman-inspired proposal in two plainer MCMC algorithms, that is, AM (Haario et al., 2001) and Delayed Rejection AM (DRAM, Haario et al., 2006).

#### 3.1. Case Study 1: A Rainfall-Runoff Model

Our first case study considers application of the DREAM<sub>(KZS)</sub> algorithm to modeling of the rainfall-discharge relationship of the Guadalupe River basin (Duan et al., 2006). We use the seven-parameter hmodel of Schoups and Vrugt (2010) to simulate daily discharge records of the Guadalupe River at Spring Branch, Texas, using basin average estimates of precipitation and potential evapotranspiration. The hmodel transforms rainfall into runoff at the watershed outlet using four different control volumes and simulation of interception, throughfall, evaporation, runoff generation, and percolation through surface and subsurface routings. The seven parameters of the hmodel and their prior ranges are listed in Table 1. We refer the reader to Schoups and Vrugt (2010) for a detailed description of the hmodel, including model structure, process specification, and parameterization.

We use the hmodel parameter values listed in Table 1 to simulate an  $n = 1,827$  record of daily discharge values,  $\mathbf{d}$ . This record is corrupted with heteroscedastic measurement errors by drawing from the  $n$ -variate normal distribution,  $\tilde{\mathbf{d}} \sim \mathcal{N}_n(\mathbf{d}, \mathbf{R})$ , with values of  $(0.05\mathbf{d})^2$  on the main diagonal of the covariance matrix,  $\mathbf{R}$ , and zero entries elsewhere. We now use the perturbed record,  $\tilde{\mathbf{d}}$ , to infer the posterior distribution of the hmodel parameters with the DREAM<sub>(KZS)</sub> algorithm using  $N = 4$  chains with  $T = 6,000$  samples in each chain. We assume a uniform prior parameter distribution over the ranges listed in Table 1 and use the Gaussian likelihood function of equation (3) with known measurement error covariance matrix,  $\mathbf{R}$ .

Figure 1 presents a trace plot of the univariate convergence diagnostic for each of the  $k = 7$  hmodel parameters,  $\hat{R}_j$ ,  $j = 1, \dots, k$ , and the multivariate convergence diagnostic,  $\hat{R}^k$ , using the DREAM<sub>(KZS)</sub> (left panel) and DREAM<sub>(ZS)</sub> (right panel) algorithms. The different hmodel parameters are coded with different colors. The horizontal black dashed line in each panel demarcates the threshold of 1.2 below which the chains are assumed to have converged to a stationary distribution. The DREAM<sub>(KZS)</sub> algorithm requires about 1,800 iterations (i.e., 7,200 model evaluations) to satisfy the stipulated convergence threshold of 1.2. The DREAM<sub>(ZS)</sub> algorithm, on the contrary, needs a substantially larger number of about 3,360 iterations (i.e., 13,440 model

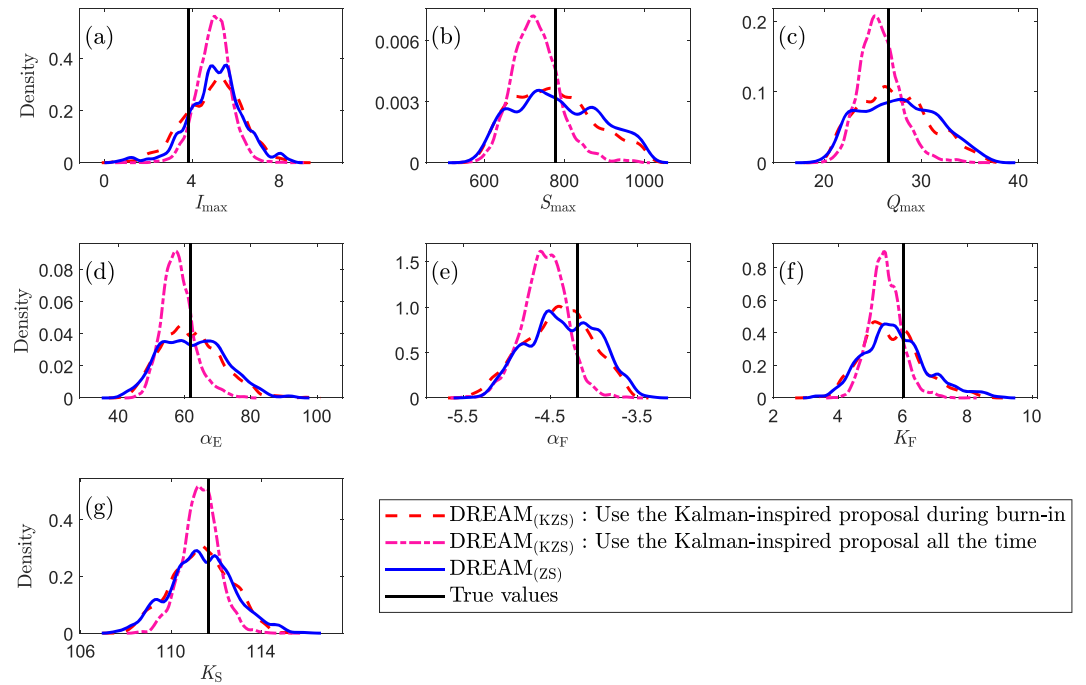


**Figure 1.** Traces of the  $\hat{R}$ -statistics of the seven hmodel parameters using the (a)  $\text{DREAM}_{(KZS)}$  and (b)  $\text{DREAM}_{(ZS)}$  algorithms. The horizontal black dashed lines depict the threshold value of 1.2 used to declare convergence to a stationary distribution.

evaluations) to converge to the stationary distribution. These findings suggest an improvement in simulation efficiency on the order of 1.87 times.


These initial results are encouraging, yet they would not mean much if the  $\text{DREAM}_{(KZS)}$  algorithm did not approximate correctly the target distribution. In Figure 2, we plot marginal posterior distributions of the seven hmodel parameters derived from the  $\text{DREAM}_{(KZS)}$  (red dashed lines) and  $\text{DREAM}_{(ZS)}$  (blue lines) algorithms. The densities are estimated with a normal kernel function using the last 1,000 samples in each of the  $N = 4$  Markov chains. The “true” values of the hmodel parameters are separately indicated in each panel with a vertical black line. The approximated marginal distributions of both algorithms appear remarkably similar and center nicely on the “true” values of the hmodel parameters. This provides evidence for the claim that the  $\text{DREAM}_{(KZS)}$  algorithm successfully approximates the target distribution.

Now one question still remains, that is, what if we do not switch the Kalman-inspired proposal distribution off after burn-in? To answer this question, we further implement the  $\text{DREAM}_{(KZS)}$  algorithm that uses the Kalman-inspired proposal all the time. From Figure S1 (in the supporting information) we can find that the  $\hat{R}$ -diagnostic plots for the two  $\text{DREAM}_{(KZS)}$  approaches can reach the threshold value of 1.2 with a similar number of iterations. This is not surprising as the two approaches use the same settings during burn-in. After burn-in, if we continue to use the Kalman-inspired proposal, we can obtain slightly smaller  $\hat{R}$  values than the approach that switches off the Kalman-inspired proposal after burn-in. However, as shown in Figure 2, using the Kalman-inspired proposal all the time in MCMC will make the final estimation of posterior (magenta dash-dotted lines) slightly narrower than the reference results obtained by  $\text{DREAM}_{(ZS)}$  (blue lines), while the recommended approach that uses the Kalman-inspired proposal only during burn-in obtains reliable results (red dashed lines). To investigate the performance of the Kalman-inspired proposal, we further analyze the acceptance rates of the two  $\text{DREAM}_{(KZS)}$  approaches. In Figure S2, we depict the evolution of the numbers of accepted candidate states in the two  $\text{DREAM}_{(KZS)}$  approaches. During the prescribed burn-in period (i.e., the first 30% of the MCMC simulation), the average acceptance rate (using the parallel direction, snooker and Kalman-inspired proposals with suggested selection probabilities) is about 0.17; if we only account for the candidate states generated by the Kalman-inspired proposal, the acceptance rate can be as high as about 0.40, which indicates that the Kalman-inspired proposal can generate “good” candidates that are less likely to be rejected. After burn-in, if we switch off the Kalman-inspired proposal, the average acceptance rate (using the other two proposals) during the last 70% of the MCMC simulation will be about 0.10; if the Kalman-inspired proposal is still left on, the average acceptance rate (using the three proposals) can be about 0.33. The above results indicate that the Kalman-inspired proposal can shorten burn-in, but it introduces asymmetry to the sampled chains. To maintain detailed balance of the Markov chains, we can adopt the three strategies proposed in section 2.2, among which the simplest way is to restrict the Kalman-inspired proposal distribution to the burn-in period, which is our default setting in the following tests.

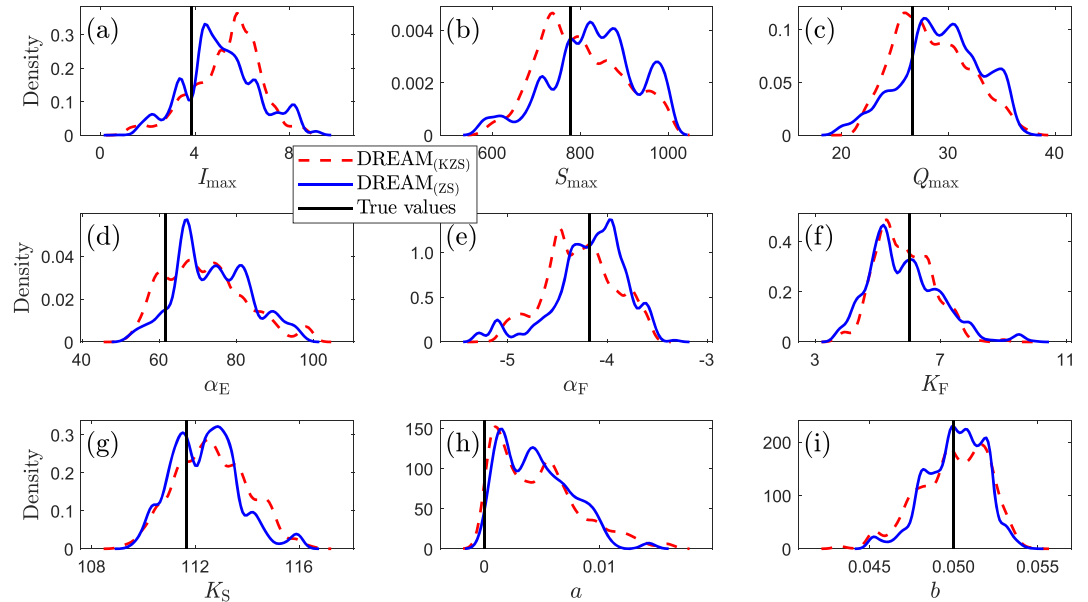


**Figure 2.** (a–g) Marginal posterior distributions of the seven hmodel parameters obtained by  $DREAM_{(KZS)}$  that uses the Kalman-inspired proposal during burn-in (default setting, red dashed lines),  $DREAM_{(KZS)}$  that uses the Kalman-inspired proposal all the time (magenta dash-dotted lines), and  $DREAM_{(ZS)}$  (reference results, blue lines), respectively. The vertical black lines display the “true” values of the hmodel parameters.

To generalize our findings, we repeat our numerical experiment using six sets of other hmodel parameter vectors drawn randomly from the uniform prior distribution. To negate sampling variability, we repeat the simulation of the  $DREAM_{(KZS)}$  and  $DREAM_{(ZS)}$  algorithms five times for each parameter vector. The results of our analysis are presented in Figure 3, which presents traces of the  $\hat{R}^k$ -diagnostic for the five repetitions (within a graph) of each parameter vector (between graphs) using the  $DREAM_{(KZS)}$  and  $DREAM_{(ZS)}$  algorithms. Thus, each subplot corresponds to a different hmodel parameter vector and contains a separate trace of the  $\hat{R}^k$ -convergence diagnostic for each of its five repetitions with the same set of measurement data. Color coding differentiates between the  $DREAM_{(KZS)}$  (red dashed lines) and  $DREAM_{(ZS)}$  (blue lines)



**Figure 3.** Evolution of the multivariate  $\hat{R}^k$ -convergence diagnostic derived from the  $DREAM_{(KZS)}$  (red dashed lines) and  $DREAM_{(ZS)}$  (blue lines) algorithms for six discharge data sets (between plots) each with five repetitions (within plot). The horizontal black dashed line in each plot displays the threshold of 1.2 used to declare convergence to a stationary distribution.

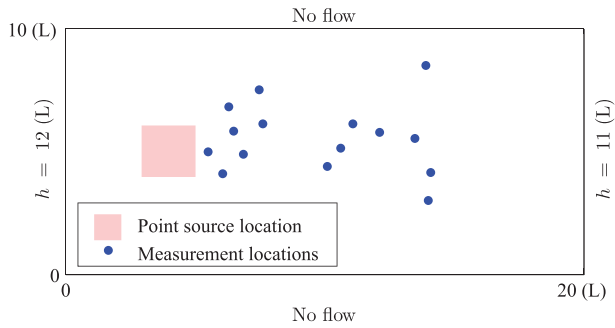


**Figure 4.** (a–i) Marginal posterior distributions of the seven hmodel parameters and two measurement error model coefficients derived from the  $DREAM_{(KZS)}$  (red dashed lines) and  $DREAM_{(ZS)}$  (blue lines) algorithms. The “true” values of the  $k = 9$  parameters are separately indicated in each panel with a vertical black line.

algorithms. As we plot the multivariate  $\hat{R}^k$ -statistic, a single line suffices for each combination of algorithm, parameter vector and repetition simulation. The results in Figure 3 generalize our earlier findings. Indeed, the  $DREAM_{(KZS)}$  algorithm consistently requires fewer model evaluations to converge to the posterior distribution. Note, that the speedup in sampling efficiency of the  $DREAM_{(KZS)}$  algorithm is hardly impressive. This is not surprising due to the rather low dimensionality of the target distribution.

Up to now, we have conveniently assumed the covariance matrix of measurement errors to be known, that is,  $\epsilon \sim \mathcal{N}_n(\mathbf{0}, \mathbf{R})$ , where  $\sqrt{\mathbf{R}}$  was diagonal with entries equal to  $1/20$  of the “true” model responses,  $\mathbf{d}$ . In practice, the distribution of measurement errors may not be known *a priori*. In those cases, we can define a measurement error model; say, the standard deviation of the measurement errors is expressed as,  $\sigma = a + b\tilde{\mathbf{d}}$ . The coefficients  $a$  and  $b$  can be treated as nuisance variables whose values can be inferred simultaneously with the model parameters (see e.g., Schoups & Vrugt, 2010; Vrugt, 2016). As a proof of concept, Figure 4 presents marginal posterior distributions of the hmodel parameters and coefficients  $a$  and  $b$  derived from the  $DREAM_{(KZS)}$  (red dashed lines) and  $DREAM_{(ZS)}$  (blue lines) algorithms. The estimated marginal distributions of both algorithms are in close agreement with each other and cover the assumed values (vertical black lines) of the hmodel parameters and measurement error model coefficients ( $a = 0$  and  $b = 0.05$ ).

It is noted here that the Kalman-inspired proposal distribution is not tied to a particular MCMC algorithm, for example, DREAM, but is straightforwardly extendable to any adequate MCMC method. To support this claim, we also introduce the Kalman-inspired proposal distribution to two widely used, single chain MCMC methods, that is, the AM and DRAM algorithms (Haario et al., 2001, 2006). When implementing the AM algorithm, we first start the Markov chain from the prior mean values and set the chain length as  $T = 24,000$ . Introducing the Kalman-inspired proposal distribution to the AM algorithm produces a modified method, which is termed the  $AM_{(K)}$  algorithm here. In the  $AM_{(K)}$  algorithm, the Kalman-inspired proposal distribution is used during a prescribed burn-in period (the first 30% of the MCMC simulation) with a selection probability of  $p_K = 0.3$ . With the same settings, we then run the DRAM and corresponding  $DRAM_{(K)}$  algorithms to infer the posterior distribution of the seven hmodel parameters. As shown in Figure S3, conditioned on the same set of measurement data (generated from one set of random parameter vector and corrupted with a zero-mean normally distributed error), the AM (black lines),  $AM_{(K)}$  (black dashed lines), DRAM (blue lines),  $DRAM_{(K)}$  (blue dashed lines), and  $DREAM_{(KZS)}$  (red dashed lines) algorithms can obtain very consistent estimates of the posterior distribution. From the traces of the seven hmodel parameters obtained by the AM (red dots in Figure S4),  $AM_{(K)}$  (blue dots in Figure S4), DRAM (red dots in Figure S5), and  $DRAM_{(K)}$  (blue dots in Figure S5) algorithms, it is evident that the  $AM_{(K)}$  and  $DRAM_{(K)}$



**Figure 5.** Schematic overview of the simulated spatial domain with imposed boundary conditions. The area with light red signifies the potential location of contaminant release from a point source, and the blue dots represent the wells at which measurements of the steady-state hydraulic head and transient contaminant concentration are collected.

algorithms require fewer iterations to reach the stationary regime than their original counterparts, and the Kalman-inspired proposal can improve the performance of AM more than DRAM. After systematic evaluations, we find that the initial state, or starting point, of the Markov chain can have a significant impact on the performance of the AM algorithm. As shown in Figure S6, starting from a random state drawn from the prior distribution, the AM algorithm (red dots) performs much worse than the AM algorithm starting from the prior mean state (red dots in Figure S4). However, starting from the same random state, the AM<sub>(K)</sub> algorithm (blue dots in Figure S6) exhibits a more desirable performance than the AM algorithm. Starting from the same random initial state, the DRAM algorithm (red dots in Figure S7) performs better than the AM algorithm, and adopting the Kalman-inspired proposal in DRAM (i.e., DRAM<sub>(K)</sub>, blue dots in Figure S7) can further improve the performance. It is noted here that the use of delayed rejection in DRAM (or DRAM<sub>(K)</sub>) will make the actual number of model evaluations much larger than the

prescribed chain length of  $T = 24,000$ . To improve the performance of MCMC, it is a common practice to first run an optimizer to obtain a “good” initial state. Here we adopt a simplex search method developed by Lagarias et al. (1998) (the build-in function “fminsearch” in MATLAB) to obtain a head start. The optimizer explores the parameter space from the prior mean values and calls 1,406 model evaluations in total. With the optimized initial state, we run the AM algorithm and plot the traces (red dots) of the seven hmodel parameters in Figure S8. In this figure, we again present the results of the AM<sub>(K)</sub> algorithm starting from the prior mean state (blue dots). Without extra model evaluations, the AM<sub>(K)</sub> algorithm still works slightly better than the AM algorithm with an initial optimizer. From the above simulation results we can conclude that the starting point of the Markov chain can have a considerable effect on the convergence speed of a plain MCMC algorithm like AM. Thus, finding a good starting point, for example, using an optimization method, can improve the simulation efficiency of the MCMC algorithm. Nevertheless, without resorting to looking for a good starting point, one can also employ an effective proposal distribution, for example, the Kalman-inspired proposal distribution formulated in this work, to generate “good” candidates that can accelerate the movement of the chain to the target region. To enhance search efficiency, adaptation of covariance(s) based on the chain history is always performed.

### 3.2. Case Study 2: Groundwater Contaminant Source Identification

The second case study considers two-dimensional simulation of steady-state groundwater flow and contaminant transport. We consider a hypothetical rectangular flow domain (see Figure 5) with  $x \in [0, 20]$  (L) and  $y \in [0, 10]$  (L) in units of length in the horizontal and vertical directions, respectively. A constant head of  $h = 12$  (L) and  $h = 11$  (L) at the left and right sides of the domain, respectively, and no-flow condition at the top and bottom of the field impose a steady-state water flux from left to right across the domain. The hydraulic conductivity field,  $\mathbf{K}(x, y)$  ( $LT^{-1}$ ), of the domain is assumed to be spatially heterogeneous and isotropic with covariance function,  $C_{\mathcal{K}}(x_1, y_1; x_2, y_2)$ , of two arbitrary points,  $(x_1, y_1)$  and  $(x_2, y_2)$ , for the log-transformed field,  $\mathcal{K} = \log(\mathbf{K}(x, y))$ , equal to (Zhang & Lu, 2004)

$$C_{\mathcal{K}}(x_1, y_1; x_2, y_2) = \sigma_{\mathcal{K}}^2 \exp\left(-\frac{|x_1 - x_2|}{\lambda_x} - \frac{|y_1 - y_2|}{\lambda_y}\right), \quad (15)$$

where  $\sigma_{\mathcal{K}}^2$  is the variance of the log conductivity field and  $\lambda_x$  (L) and  $\lambda_y$  (L) signify the correlation lengths in the  $x$  and  $y$  directions, respectively.

We sample a reference hydraulic conductivity field (see Figure 8a) with mean logarithmic conductivity,  $\mu_{\mathcal{K}} = 2$ , and values of  $\sigma_{\mathcal{K}}^2 = 1$ ,  $\lambda_x = 10$  (L), and  $\lambda_y = 5$  (L) in equation (15). Next, we simulate the steady-state hydraulic head,  $h$  (L), and corresponding pore water velocity,  $v$  ( $LT^{-1}$ ), in our rectangular domain by solving the following two equations numerically with MODFLOW (Harbaugh et al., 2000)

$$\frac{\partial}{\partial x_i} \left( K_i \frac{\partial h}{\partial x_i} \right) = 0, \quad (16)$$

and

$$v_i = -\frac{K_i}{\phi} \frac{\partial h}{\partial x_i}, \quad (17)$$

where  $\phi = 0.25$  (-) is the aquifer porosity and the subscript  $i$  signifies the respective coordinate axis ( $i = 1, 2$ ).

**Table 2**  
Description of the Eight Contaminant Source Parameters of Our Contaminant Transport Model and the Ranges of Their Multivariate Uniform Prior Distribution

Parameter	Symbol	Min.	Max.	Unit	True
$x$ coordinate of spill	$x_s$	3	5	L	3.52
$y$ coordinate of spill	$y_s$	4	6	L	4.44
Release strength of the first segment	$s_1$	0	8	MT <sup>-1</sup>	5.69
Release strength of the second segment	$s_2$	0	8	MT <sup>-1</sup>	7.88
Release strength of the third segment	$s_3$	0	8	MT <sup>-1</sup>	6.31
Release strength of the fourth segment	$s_4$	0	8	MT <sup>-1</sup>	1.49
Release strength of the fifth segment	$s_5$	0	8	MT <sup>-1</sup>	6.87
Release strength of the sixth segment	$s_6$	0	8	MT <sup>-1</sup>	5.55

Note. The last column lists the “true” values of the parameters used in our numerical experiment.

To simulate contaminant transport, we pollute the water in our flow domain with an unknown point source. The exact location of this point source is assumed unknown, but its spatial coordinates,  $x_s$  and  $y_s$ , must be found within the light red square depicted in Figure 5. As the release strength of the contaminant may be time dependent, we use a simple step function to simulate transient mass-loading rates. This step function is composed of six equidistant time intervals with a constant mass-loading rate,  $s_i$  (MT<sup>-1</sup>),  $i = 1, \dots, 6$ , in each segment. The space-time concentration,  $C(x, y, t)$  (ML<sup>-3</sup>), of the contaminant in our rectangular domain is simulated with MT3DMS (Zheng & Wang, 1999), using numerical solution of the advection-dispersion equation

$$\frac{\partial(\phi C)}{\partial t} = \frac{\partial}{\partial x_i} \left( \phi D_{ij} \frac{\partial C}{\partial x_j} \right) - \frac{\partial}{\partial x_i} (\phi v_i C) + q_a C_s, \quad (18)$$

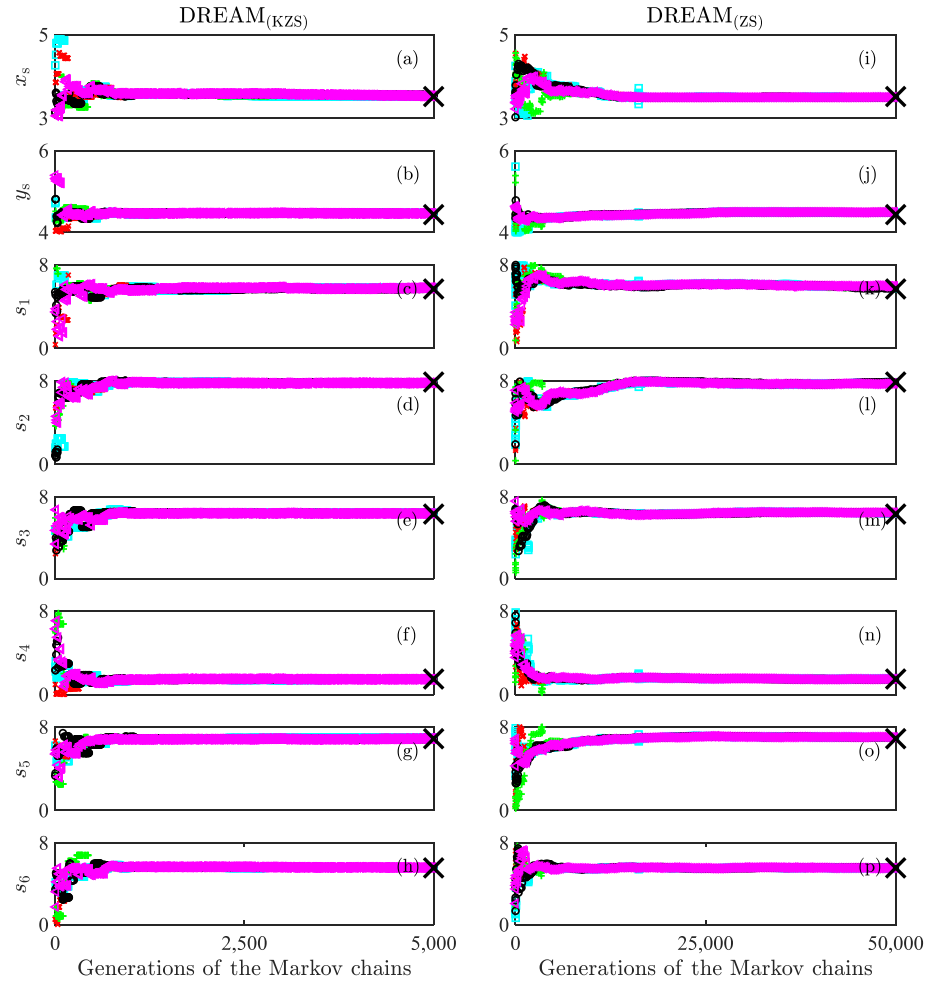
where  $q_a$  (T<sup>-1</sup>) denotes the volumetric flow rate per unit volume of the aquifer,  $C_s$  (ML<sup>-3</sup>) is the concentration of the contaminant source, and  $D_{ij}$  (L<sup>2</sup>T<sup>-1</sup>) signifies the hydrodynamic dispersion tensor. This tensor is made up of the following four components

$$\begin{aligned} D_{xx} &= \frac{1}{\|\mathbf{v}\|} (\alpha_L v_x^2 + \alpha_T v_y^2), \\ D_{yy} &= \frac{1}{\|\mathbf{v}\|} (\alpha_L v_y^2 + \alpha_T v_x^2), \\ D_{xy} &= D_{yx} = \frac{1}{\|\mathbf{v}\|} (\alpha_L - \alpha_T) v_x v_y, \end{aligned} \quad (19)$$

where  $D_{xx}$ ,  $D_{yy}$ ,  $D_{xy}$ , and  $D_{yx}$  (L<sup>2</sup>T<sup>-1</sup>) are the two principal components of the dispersion tensor and their two cross terms, respectively,  $\alpha_L$  and  $\alpha_T$  (L) signify the longitudinal and transverse dispersivity, respectively,  $v_x$  and  $v_y$  are the water flow velocities in the  $x$  and  $y$  directions, respectively, and  $\|\mathbf{v}\| = \sqrt{v_x^2 + v_y^2}$  is the  $\ell^2$  norm, or magnitude, of the velocity vector,  $\mathbf{v}$ .

To test, evaluate, and benchmark the DREAM<sub>(KZS)</sub> algorithm, we create a reference data set as follows. We sample randomly from a multivariate uniform prior distribution of the two coordinates,  $\{x_s, y_s\}$ , of the contaminant source location and six release strengths,  $\{s_1, \dots, s_6\}$ , of the mass-loading rate step function using the ranges listed in Table 2. These sampled values are listed in the last column of Table 2 and used to construct an artificial data set by collecting, at 15 different wells within the flow domain, the simulated steady-state hydraulic heads and transient contaminant concentrations. The location of each of these wells is separately indicated in Figure 5 with a blue dot. We now simulate a reference run of steady-state heads and transient contaminant concentrations at each well using  $\alpha_L = 0.3$  (L) and  $\alpha_T = 0.03$  (L). The simulated heads and contaminant concentrations are subsequently corrupted with a zero mean normally distributed error with standard deviations of 0.005 (L) and 0.005 (ML<sup>-3</sup>), respectively. The final data set now consists of the steady-state hydraulic heads and transient contaminant concentrations at  $t = \{4, 5, \dots, 12\}$  (T) at the 15 wells. This equates to  $n_h = 15$  measurements of the hydraulic head and  $n_C = 135$  observations of the contaminant concentration.

We now use the DREAM<sub>(KZS)</sub> algorithm to reconstruct the hydraulic conductivity field, point source location, and transient mass-loading rates from the measurement data. This requires a parametric definition



**Figure 6.** Trace plots of the sampled values of the point source coordinates,  $\{x_s, y_s\}$ , and the six coefficients,  $\{s_1, \dots, s_6\}$ , of the mass-loading rate step function for the  $DREAM_{(KZS)}$  (a–h) and  $DREAM_{(ZS)}$  (i–p) algorithms. Color and symbol coding are used to differentiate among the sampled chains. The “true” parameter values are separately indicated in each panel with the black crosses.

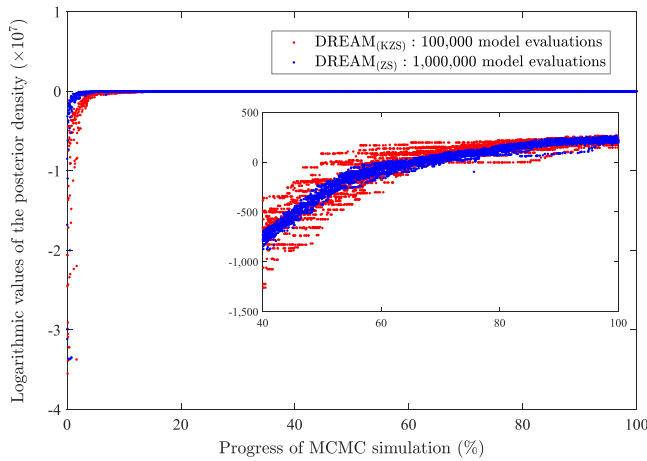
of the hydraulic conductivity field of the rectangular flow domain. A simple Cartesian parameterization of the hydraulic conductivity field would require an excessively large number of parameters to characterize adequately the imposed spatial variability. Instead, we take advantage of the Karhunen-Loève (KL) expansion and approximate the log-hydraulic conductivity field,  $\mathcal{K}$ , using a finite, yet relatively small, number of orthogonal basis functions (Zhang & Lu, 2004)

$$\tilde{\mathcal{K}}(\mathbf{x}) = \mu_{\mathcal{K}}(\mathbf{x}) + \sum_{i=1}^p \sqrt{\tau_i} s_i(\mathbf{x}) \beta_i, \quad (20)$$

where  $\mu_{\mathcal{K}}(\mathbf{x})$  denotes the mean log conductivity,  $\mathbf{x}$  represents the Cartesian coordinates of the flow domain,  $s_i(\mathbf{x})$ ,  $\tau_i$ ,  $i = 1, \dots, p$ , signify the eigenfunctions and eigenvalues of the kernel defined in equation (15), and  $\{\beta_1, \dots, \beta_p\}$  are standard normal random variables, the so-called KL expansion terms. The mean square error of the reconstructed log conductivity field,  $\tilde{\mathcal{K}}(\mathbf{x})$ , will go to zero in the limit of  $p \rightarrow \infty$ . Yet the use of a very large number of KL terms defeats the purpose of this expansion and is not encouraged for statistical inference. Instead, 100 expansion terms preserve about 95% of the variance of the “true” log conductivity field,  $\mathcal{K}(x, y)$ , that is,  $\sum_{i=1}^p \tau_i / \sum_{i=1}^{\infty} \tau_i \approx 0.95$ . Henceforth, we characterize the conductivity field of our flow domain with  $p = 100$  KL terms,  $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_{100}\}$ .

Our numerical experiment now involves inference of  $k = 108$  parameters, namely,  $p = 100$  KL terms,  $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_{100}\}$ , two coordinates of the source location,  $\{x_s, y_s\}$ , and six coefficients,  $\{s_1, \dots, s_6\}$ , of the





**Figure 7.** Evolution of the log-transformed values of unnormalized posterior density, that is,  $\mathcal{P}(\theta|\tilde{\mathbf{d}}) = \log [p(\theta)L(\theta|\tilde{\mathbf{d}})]$ , at  $\theta$ , sampled by both the  $\text{DREAM}_{(\text{KZS})}$  (red dots) and  $\text{DREAM}_{(\text{ZS})}$  (blue dots) algorithms.

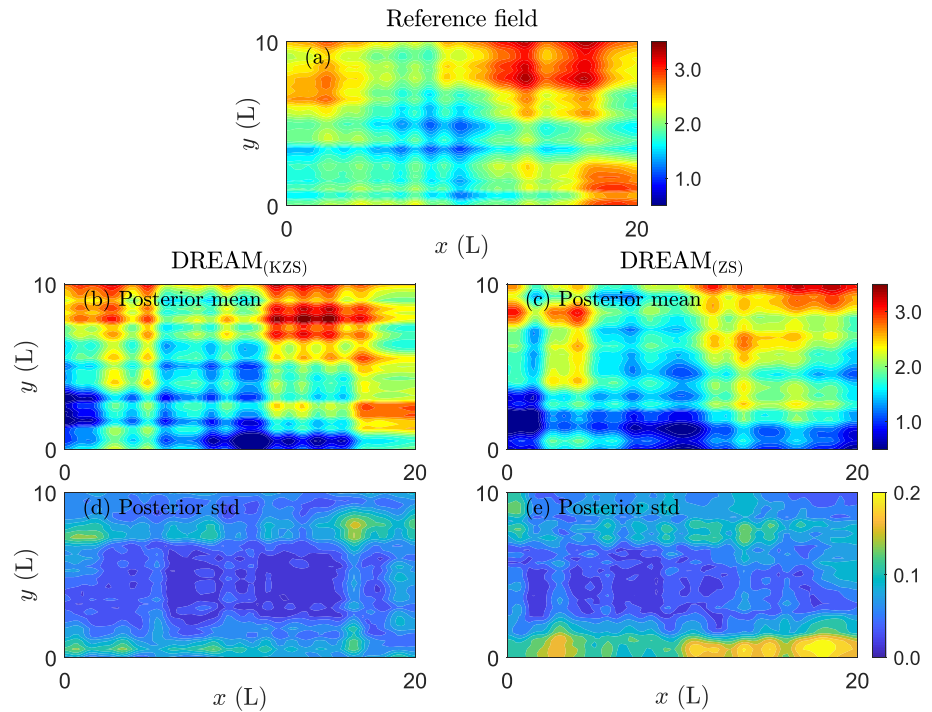
mass-loading rate step function. Thus, the unknown parameters are  $\theta = \{\beta_1, \dots, \beta_{100}, x_s, y_s, s_1, \dots, s_6\}$ . We use a standard Gaussian likelihood function (see equation (3)) and execute the  $\text{DREAM}_{(\text{KZS})}$  and  $\text{DREAM}_{(\text{ZS})}$  algorithms with  $N = 20$  chains using default values of the algorithmic variables.

Figure 6 presents trace plots of the sampled values of the point source coordinates,  $\{x_s, y_s\}$ , and the six coefficients,  $\{s_1, \dots, s_6\}$ , of the mass-loading rate step function for the  $\text{DREAM}_{(\text{KZS})}$  (left column) and  $\text{DREAM}_{(\text{ZS})}$  (right column) algorithms. Here we only draw 5 of the 20 Markov chains so that size of the image file would not be too big. The different chains are coded with different symbols and colors. The “true” value of each parameter (i.e., the last column of Table 2) is separately indicated with a black cross symbol at the right-hand side of each panel. The most important results are as follows. First, the sampled chains of both algorithms converge to the “true” values used to generate the hydraulic head and contaminant concentration observations. Second, the location of the point source and mass-loading rate step function appears well defined with negligible posterior uncertainty compared to the width of the prior distribution. Third, and perhaps the most important, the  $\text{DREAM}_{(\text{KZS})}$  algorithm requires far fewer model evaluations than the  $\text{DREAM}_{(\text{ZS})}$  algorithm to explore the target distribution.

To better understand the search capabilities of the two MCMC algorithms, please consider Figure 7 that presents the evolution of the log-transformed values of unnormalized posterior density, that is,  $\mathcal{P}(\theta|\tilde{\mathbf{d}}) = \log [p(\theta)L(\theta|\tilde{\mathbf{d}})]$ , at  $\theta$ , sampled by the  $\text{DREAM}_{(\text{KZS})}$  (red dots) and  $\text{DREAM}_{(\text{ZS})}$  (blue dots) algorithms, respectively. At early stages of the search with both MCMC algorithms, the sampled chain states exhibit rather small logarithmic values (on the order of  $-10^7$ ) of the posterior density. Gradually, the Markov chains move to the high probability region. During the last 60% of the MCMC simulations,  $\mathcal{P}(\theta|\tilde{\mathbf{d}})$  lies between  $-1,500$  and  $500$ . Finally, both algorithms reach a similar  $\mathcal{P}(\theta|\tilde{\mathbf{d}})$  value of about  $220$ . Note that the total number of model evaluations required by the  $\text{DREAM}_{(\text{KZS})}$  algorithm (5,000 generations  $\times$  20 chains = 100,000 samples) is only one tenth of the  $\text{DREAM}_{(\text{ZS})}$  algorithm (50,000 generations  $\times$  20 chains = 1,000,000 samples). This constitutes a speedup on the order of 10 times. With an average CPU time of 2 s for each evaluation of the integrated model of MODFLOW and MT3DMS, this speedup equates to about 500 hr reduction in CPU cost.

Next, Figure 8 presents log conductivity fields,  $\tilde{\mathcal{K}}(x, y)$ , of the last 10,000 posterior realizations sampled by the  $\text{DREAM}_{(\text{KZS})}$  (left column) and  $\text{DREAM}_{(\text{ZS})}$  (right column) algorithms. Figures 8b and 8c show the posterior mean  $\tilde{\mathcal{K}}(x, y)$  fields obtained by averaging the log conductivity fields of the 10,000 posterior realizations. The bottom row presents maps of the standard deviation of the 10,000 log conductivity fields. Compared with the reference field as shown in Figure 8a, the posterior mean map of the  $\text{DREAM}_{(\text{KZS})}$  algorithm illuminates correctly the areas of high and low conductivity in the flow domain yet underestimates their spatial extent. Altogether, these findings suggest that the hydraulic head and contaminant concentration measurements contain insufficient information to back out exactly the actual conductivity field. A larger and more diverse data set is warranted.

Above we have tested a nonlinear inverse problem with  $k = 108$  unknown model parameters. In this case, even using a state-of-the-art MCMC algorithm, that is,  $\text{DREAM}_{(\text{ZS})}$ , at least one million forward model evaluations are needed, which constitutes a prohibitively high computational cost. However, introducing the Kalman-inspired proposal distribution can bring a speedup of about 10 times in simulation efficiency. The advantage of the Kalman-inspired proposal distribution is obvious in high-dimensional settings. Nevertheless, when a problem has a much larger number of unknown parameters, for example,  $k > 1,000$ , and this problem is highly nonlinear and complex, performance of MCMC methods will deteriorate. In this situation, one can resort to more computationally appealing methods, for example, EnKF and its variants (Chen & Zhang, 2006; Crestani et al., 2013; Evensen, 2009), that assume multi-Gaussian parameter and error distributions. However, these methods will not work properly when the posterior is non-Gaussian or even has multiple modes. To address these issues, several strategies can be adopted, for example, transforming



**Figure 8.** (a) The reference field of log-transformed conductivity,  $\mathcal{K}$ . (b–e) Posterior mean  $\mathcal{K}$  fields (middle row) and associated standard deviations (bottom row) derived from the DREAM<sub>(KZS)</sub> (left column) and DREAM<sub>(ZS)</sub> (right column) algorithms.

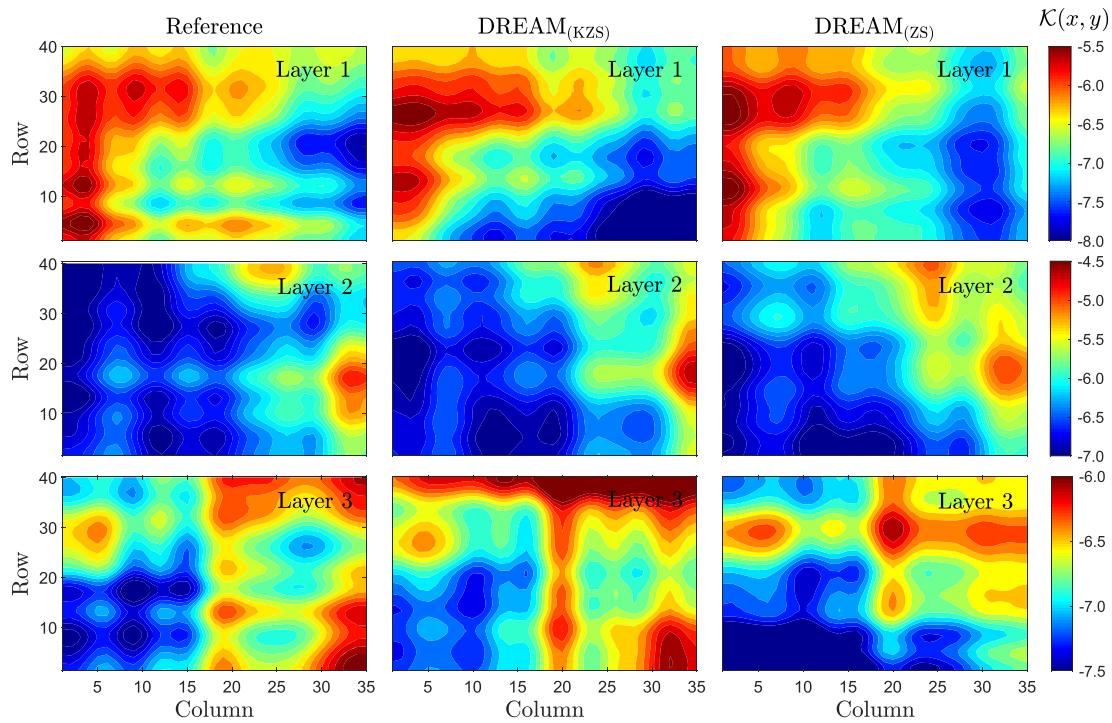
non-Gaussian variables to be Gaussian distributed (Chang et al., 2010; Zhou et al., 2011) or adopting a local Kalman update strategy to handle multimodal posteriors (Zhang, Lin, et al., 2018).

### 3.3. Case Study 3: A 3-D Groundwater Model

The third case study considers application of the DREAM<sub>(KZS)</sub> algorithm to aquifer characterization using the three-dimensional groundwater model of Fienen et al. (2013). This model simulates three horizontal layers of 1.8, 1.4, and 1.8 m thickness that are each discretized into 35 columns and 40 rows with equidistant spacing of 2.0 and 1.5 m in the  $x$  (column) and  $y$  (row) directions, respectively. This equates to a domain of 70 m by 60 m by 5 m and a total of  $N_{\text{node}} = 35 \times 40 \times 3 = 4,200$  nodes. A constant head of 60 m is prescribed at all vertical edges of the flow domain. Furthermore, an artificial well is located at row 18 and column 17, which pumps water from each horizontal layer at a constant rate of 0.01 L/min.


To reduce parametric dimensionality, we use a sparse representation of the conductivity field of each layer. Here, mean and variance of the log conductivity fields are  $\mu_{\mathcal{K}}(\mathbf{x}) = -6.5$  and  $\sigma_{\mathcal{K}}^2 = 0.5$ , and correlation lengths in the  $x$  and  $y$  directions are  $\lambda_x = 37.5$  m and  $\lambda_y = 60$  m, respectively, for all the three layers. About 40 KL terms are deemed sufficient to represent the  $\mathcal{K}$  field of each layer and preserve about 94% of the field variance of each original 1,400 cell conductivity field. Hydraulic head measurements at the three layers from 81 wells located every four rows from Row 3 to Row 35, every three columns from Column 5 to Column 29, are generated from the reference fields depicted in the left column of Figure 9 with additive white noise,  $\epsilon \sim \mathcal{N}_n(0, 0.01^2)$ . We infer the  $k = 120$  unknown KL terms with the DREAM<sub>(KZS)</sub> and DREAM<sub>(ZS)</sub> algorithms using  $N = 20$  chains and  $T = 4,000$  generations. We discard the first 3,500 generations as burn-in and use the 10,000 samples in the last 500 generations of the joint chains to summarize the posterior estimates of the conductivity field of each discretized aquifer layer.

Figure 9 presents the posterior mean log conductivity fields derived from the DREAM<sub>(KZS)</sub> (middle column) and DREAM<sub>(ZS)</sub> (right column) algorithms. The mean log conductivity fields derived from both algorithms capture quite well the main patterns of the reference fields of the three layers (left column). Some discrepancies are visible but appear relatively minor. These results are encouraging but do not convey anything about the efficiency of the two MCMC algorithms. We therefore proceed with analysis of the convergence properties of the DREAM<sub>(KZS)</sub> and DREAM<sub>(ZS)</sub> algorithms.



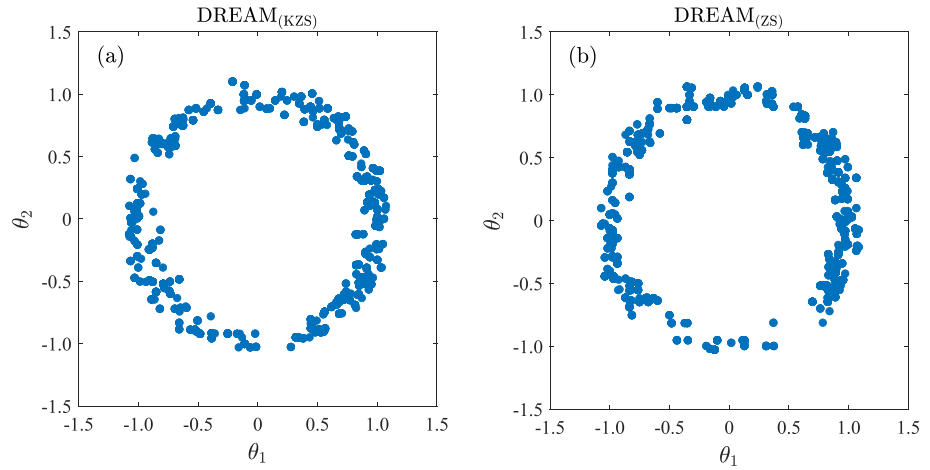
**Figure 9.** Reference field of the log-transformed hydraulic conductivity for each of the three layers of the aquifer (left column) and their posterior mean counterparts inferred from the  $DREAM_{(kzs)}$  (middle column) and  $DREAM_{(zs)}$  (right column) algorithms using the measured steady-state hydraulic heads.

In Figure 10, we present trace plots of the root-mean-square error (RMSE) in the  $N = 20$  Markov chains between the MODFLOW simulated and measured steady-state heads at the 81 wells. We use color coding in red and blue for the  $DREAM_{(kzs)}$  and  $DREAM_{(zs)}$  algorithms, respectively. The results in this figure confirm our earlier conclusions. The  $DREAM_{(kzs)}$  algorithm converges at a much faster pace to the target distribution. Indeed, whereas the  $DREAM_{(zs)}$  algorithm requires about 60,000–80,000 function evaluations to sample values of the RMSE on the order of the Gaussian measurement error of 0.01, the  $DREAM_{(kzs)}$  algorithm needs only about 2,500–3,000 MODFLOW evaluations to minimize the RMSE. This equates to a speedup on the order of 20–30 times. This gives the  $DREAM_{(kzs)}$  algorithm sufficient opportunity to sample thoroughly the target distribution and summarize adequately the posterior moments of the quasi three-dimensional conductivity distribution of the aquifer.



**Figure 10.** Evolution of the sampled RMSE values between the MODFLOW simulated and observed steady-state hydraulic heads. All chains are plotted. Color coding in red and blue differentiates between the chains sampled by the  $DREAM_{(kzs)}$  and  $DREAM_{(zs)}$  algorithms, respectively.

In this case study, although the number of unknown model parameters ( $k = 120$ ) is slightly larger than that in the second case study ( $k = 108$ ), much fewer model evaluations are needed in the MCMC simulations. Using the  $DREAM_{(zs)}$  algorithm, to obtain acceptable results, the total numbers of model evaluations in the second and present case studies are at least 1,000,000 and 60,000, respectively. Using the  $DREAM_{(kzs)}$  algorithm, the corresponding numbers are about 100,000 and 2,500. This comparison indicates that the convergence speed of an MCMC algorithm is determined not only by the number of model parameters but also by the nonlinearity and complexity of the problem. In many situations, the high-dimensionality of unknown parameters (e.g., hydraulic conductivity field) comes from the discretization of the underlying function. Using a standard MCMC algorithm, the convergence speed degrades significantly with mesh refinement. This kind of MCMC algorithm is thus called dimension dependent. To address this kind of dimensionality issue, dimension-independent MCMC methods have been developed,



**Figure 11.** Bivariate scatter plots of posterior parameter samples obtained by (a)  $DREAM_{(KZS)}$  and (b)  $DREAM_{(ZS)}$ .

for example, the preconditioned Crank-Nicolson algorithm proposed by Cotter et al. (2013) that preserves the reference measure with a Crank-Nicolson discretization of the underlying function and the dimension-independent and likelihood-informed algorithm proposed by Cui et al. (2016) that constructs a global likelihood-informed subspace to capture essential features of the posterior distribution. The dimension-independent MCMC methods require that the distributions exhibit certain special structure. Yet sampling generic high-dimensional target distributions with MCMC methods is still a challenging problem.

### 3.4. Case Study 4: Inverse Problems With Multimodal Posteriors

To demonstrate the performance of the  $DREAM_{(KZS)}$  algorithm in solving nonlinear and ill-posed inverse problems, three examples with multimodal posterior distributions are further tested below.

The first example is very simple yet highly ill-posed. The underlying function has the following form

$$d = \theta_1^2 + \theta_2^2, \quad (21)$$

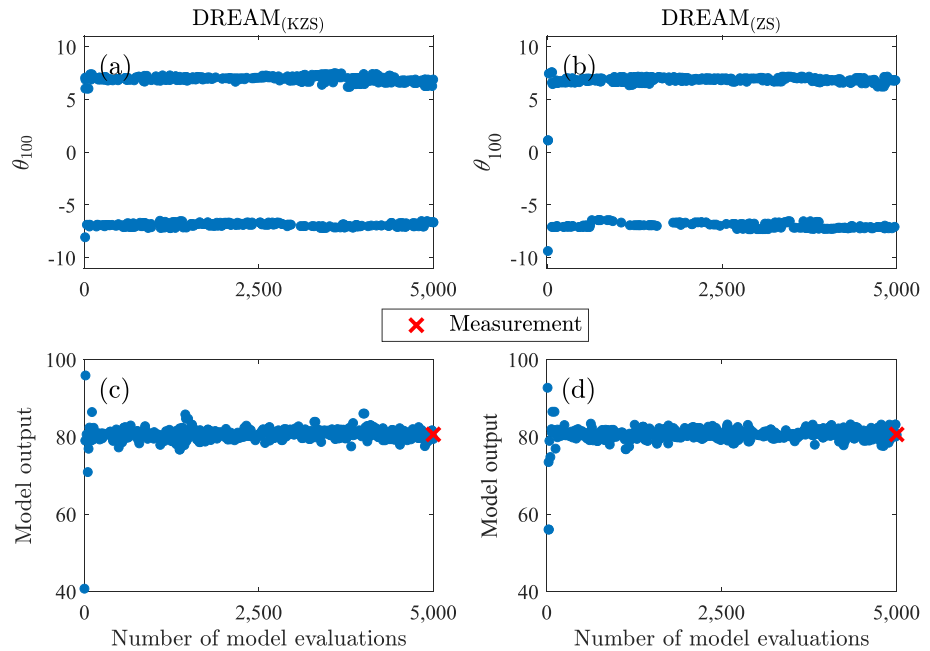
where  $\theta_1 \sim \mathcal{U}(-2, 2)$  and  $\theta_2 \sim \mathcal{U}(-2, 2)$  are all uniformly distributed in the prior. The measurement is  $\tilde{d} = 1$  with measurement error,  $\epsilon \sim \mathcal{N}(0, 0.1^2)$ . Here we are trying to infer the joint posterior distribution of two parameters,  $\theta = \{\theta_1, \theta_2\}$ , from a scalar measurement,  $\tilde{d}$ . It is evident that there exists an infinite number of parameter combinations that can fit the measurement well. Figure 11 shows the bivariate scatter plots of posterior parameter samples obtained by both the  $DREAM_{(KZS)}$  and  $DREAM_{(ZS)}$  algorithms. Here both algorithms evolve  $N = 3$  Markov chains with  $T = 1,000$  samples in each chain. The last 300 samples in each of the three chains are used to draw Figure 11. It is found that both algorithms obtain reasonable results.

Then we extend the above example to a high-dimensional setting, which turns the underlying function to

$$d = \theta_1^2 + \theta_2^2 + \dots + \theta_{100}^2, \quad (22)$$

where  $\{\theta_1, \dots, \theta_{99}\}$  are all uniformly distributed with  $\mathcal{U}(0, 2)$  and  $\theta_{100} \sim \mathcal{U}(-10, 10)$ . There is still only a scalar measurement, whose value is  $\tilde{d} = 80.62$ , and the error is normally distributed,  $\epsilon \sim \mathcal{N}(0, 1^2)$ . Compared to the previous example, inferring the joint distribution of 100-dimensional parameters,  $\theta = \{\theta_1, \dots, \theta_{100}\}$ , from a scalar measurement,  $\tilde{d}$ , becomes more challenging. As  $\theta_{100}$  has a wide prior range, the marginal posterior distribution of  $\theta_{100}$  can be obviously bimodal, as both  $\theta_{100}$  and  $-\theta_{100}$  have the same effect on the function output. As shown in Figure 12, both the  $DREAM_{(KZS)}$  and  $DREAM_{(ZS)}$  algorithms (here  $N = 5$  and  $T = 1,000$ ) can identify the bimodality of  $\theta_{100}$  and fit the measurement quite well.

Finally, we test a groundwater contaminant source identification problem with multiple modes in the posterior. The model settings are the same as the example tested in section 3.2, except that here we consider a homogeneous conductivity field with a known value of  $K = 8 \text{ (LT}^{-1}\text{)}$  and a different parameterization of the contaminant source. The source located at  $(x_s, y_s)$  (L) starts to release from  $t = t_{\text{on}}$  (T) with a constant mass-loading rate of  $S_s \text{ (MT}^{-1}\text{)}$  until  $t = t_{\text{off}}$  (T). Thus, there are five unknown parameters, that is,  $\theta = \{x_s, y_s, S_s, t_{\text{on}}, t_{\text{off}}\}$ , whose prior ranges (uniform distributions) and “true” values are provided in Table 3. Concentration measurements are collected at a single well (the blue dot in Figure S9) at



**Figure 12.** Trace plots of the sampled values of (a, b)  $\theta_{100}$  and (c, d) the model output. The left column is for the DREAM<sub>(KZS)</sub> algorithm, and the right column is for the DREAM<sub>(ZS)</sub> algorithm. The measurement  $\bar{d}$  is represented by a red cross in each panel at the bottom row.

$t = \{6, 8, 10, 12, 14\}(T)$ . The measurement error is normally distributed,  $\epsilon \sim \mathcal{N}(0, 0.01^2)$ . Then we implement the DREAM<sub>(KZS)</sub> and DREAM<sub>(ZS)</sub> algorithms to infer the posterior distribution of the five unknown parameters, respectively. Here, for both algorithms, there are  $N = 4$  parallel chains, and each chain has 5,000 samples. Trace plots of the sampled values of the five contaminant source parameters, that is,  $\theta = \{x_s, y_s, S_s, t_{on}, t_{off}\}$ , obtained by the DREAM<sub>(KZS)</sub> and DREAM<sub>(ZS)</sub> algorithms are depicted in Figures 13 and S10, respectively. Both algorithms can identify the bimodal posterior distribution of  $y_s$ . From the multivariate  $\hat{R}^k$ -diagnostic that monitors the convergence of the Markov chains (Figure S11), it is found that the DREAM<sub>(KZS)</sub> algorithm converges to its stationary regime slightly faster than the DREAM<sub>(ZS)</sub> algorithm.

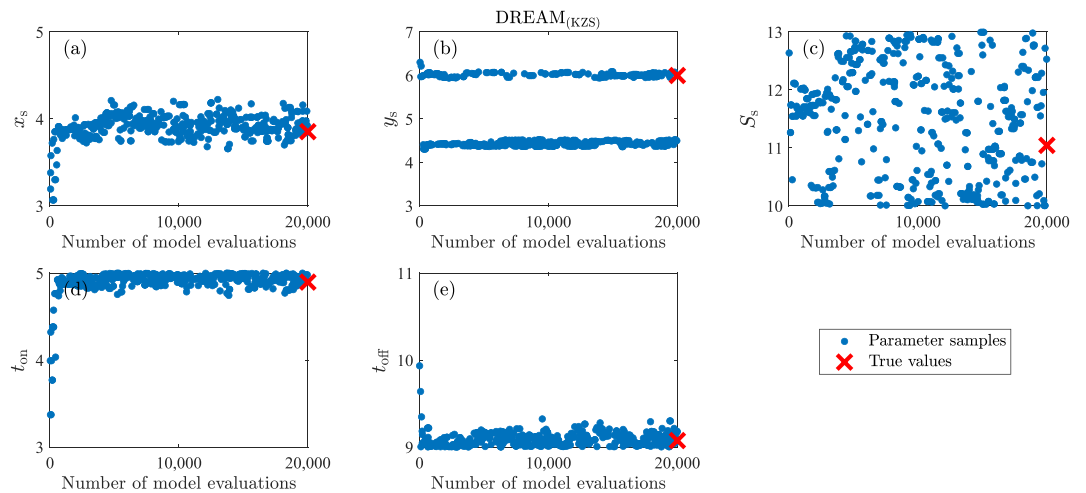
The above three examples demonstrate that the proposed method can still work properly in multimodal cases. When dealing with a nonlinear, non-Gaussian inverse problem, the simulation results of EnKF and its variants will deteriorate significantly, as the updated ensemble of states cannot approximate the complex posterior distribution accurately. However, the Kalman-inspired proposal used in the MCMC simulation can still have a positive effect. In the Kalman-inspired proposal, we only generate a candidate at a time. Although the jump from the current state  $\theta_{(t-1)}$  to the candidate  $\theta_p$  may be not optimal, it can still provide some information about the high posterior density region and thus shorten burn-in. Even in the worst case that most of the Kalman-inspired candidates are rejected (very unlikely though), as we only use the

**Table 3**

*Description of the Five Contaminant Source Parameters in the Multimodal Case and the Ranges of Their Multivariate Uniform Prior Distribution*

Parameter	Symbol	Min.	Max.	Unit	True
x coordinate of spill	$x_s$	3	5	L	3.85
y coordinate of spill	$y_s$	3	7	L	6.00
Constant release strength	$S_s$	10	13	MT <sup>-1</sup>	11.04
Start time of contaminant release	$t_{on}$	3	5	T	4.90
End time of contaminant release	$t_{off}$	9	11	T	9.08

*Note.* The last column lists the “true” values of the parameters used in our numerical experiment.



**Figure 13.** (a–e) Trace plots of the sampled values of the five contaminant source parameters obtained by the  $DREAM_{(KZS)}$  algorithm. The “true” parameter values are separately indicated in each panel with the red crosses.

Kalman-inspired proposal during a relatively short burn-in period (e.g., the first 30 % of the MCMC simulation) with a small selection probability (e.g.,  $p_K = 0.3$ ), only less than 9% of the total computational cost will be wasted because of introducing the Kalman-inspired proposal distribution. The complex posterior can still be explored with other proposal distributions (e.g., the parallel direction and snooker proposal distributions). If we monitor the rejection rate of the Kalman-inspired proposal distribution and find a very high value, we can actually stop using the Kalman-inspired proposal early to avoid further computational waste. Nevertheless, in complicated, multimodal cases, the proposed method still faces the risk of missing secondary modes. To improve the capacity of the Kalman-inspired proposal in multimodal cases, we can adopt the modified Kalman-inspired proposal suggested in the last paragraph of section 2.3.1 The modified method borrows ideas from a new ensemble smoother-based method (Zhang, Lin, et al., 2018) that is specifically designed to solve inverse problems with (possible) multimodal parameter distributions.

#### 4. Discussion and Conclusions

MCMC methods have found widespread application and use to approximate the posterior distribution. Such methods generate a random walk through the parameter space and successively visit solutions with frequency proportional to the density of the underlying target distribution. The speed with which MCMC methods converge to the stationary distribution, however, deteriorates rapidly with increasing target dimensionality.

The power and usefulness of the Kalman analysis step has been demonstrated time and again with application to state estimation in real-time forecasting studies. The use of the analysis state enhances considerably the short-term predictive skill of computer simulation models. The analysis step in data assimilation methods can facilitate parameter estimation as well, which makes possible the application of data assimilation methods such as EnKF and its variants to solving high-dimensional inverse problems. The so-obtained posterior parameter distribution can, at best, only roughly approximate the target distribution.

This paper introduces a Kalman-inspired proposal distribution to improve the efficiency of posterior exploration using MCMC methods. This new proposal distribution exploits the cross covariances of model parameters, measurements, and model outputs and generates candidate states much alike the analysis step in the Kalman filter. The Kalman-inspired proposal distribution was embedded in the DREAM algorithm, and this new sibling of the DREAM family of MCMC methods coined the  $DREAM_{(KZS)}$  algorithm. As the Kalman-inspired proposal distribution is asymmetric, its use is restricted to a relatively short burn-in period, after which a mix of parallel direction and snooker candidate states is used to evolve the chains in the  $DREAM_{(KZS)}$  algorithm. Diminishing adaptation guarantees that the sampled chains converge to the exact target distribution. Numerical experiments with watershed and aquifer models confirm that the Kalman-inspired proposal distribution enhances considerably the efficiency of posterior exploration. Specifically, we observe a speedup on the order of 20–30 times for a three-dimensional groundwater model with

more than 100 unknown parameters. Although we combine the Kalman-inspired idea with the DREAM algorithm, the new proposal distribution is not tied to a specific MCMC algorithm but can be conveniently embedded in any adequate MCMC method.

When the inverse problem is ill-posed and the posterior is multimodal, introducing the Kalman-inspired proposal distribution into an MCMC algorithm can still help. Furthermore, there have been some strategies developed for EnKF and its variants to solve non-Gaussian inverse problems, for example, transforming non-Gaussian variables to be Gaussian distributed (Chang et al., 2010; Zhou et al., 2011) or adopting a local Kalman update strategy to handle multimodal posteriors (Zhang, Lin, et al., 2018). These strategies can also be used in the Kalman-inspired proposal to gain further strength. On a more theoretical note, it may be desirable to enforce symmetry of the Kalman-inspired proposal distribution. This would deteriorate at least somewhat sampling efficiency but make possible its application to the entire chain generations simulated by the MCMC algorithm. These issues will be addressed in our future work.

### Acknowledgments

Computer codes and data used are available online ([https://www.researchgate.net/publication/318645382\\_MATLAB\\_codes\\_of\\_DREAM\\_KZS](https://www.researchgate.net/publication/318645382_MATLAB_codes_of_DREAM_KZS)). This work is supported by the National Key Research and Development Program of China (Grant 2018YFC1800303), the National Natural Science Foundation of China (Grants 41807006 and 41771254), and the China Postdoctoral Science Foundation funded project (Grant 2018M630680). The authors would like to thank the Editor and anonymous reviewers for their constructive comments and suggestions, which significantly improve the quality of this work. The authors would also like to thank Marko Laine from Finnish Meteorological Institute for providing the MATLAB codes of the Adaptive Metropolis and Delayed Rejection Adaptive Metropolis algorithms.

### References

- Beichl, I., & Sullivan, F. (2000). The Metropolis algorithm. *Computing in Science & Engineering*, 2(1), 65. <https://doi.org/10.1109/5992.814660>
- Bikowski, J., Huisman, J. A., Vrugt, J. A., Vereecken, H., & van der Kruk, J. (2012). Integrated analysis of waveguide dispersed GPR pulses using deterministic and Bayesian inversion methods. *Near Surface Geophysics*, 10(6), 641–652. <https://doi.org/10.3997/1873-0604.2012041>
- Brockwell, A. E. (2006). Parallel Markov chain Monte Carlo simulation by pre-fetching. *Journal of Computational and Graphical Statistics*, 15(1), 246–261. <https://doi.org/10.1198/106186006X100579>
- Brooks, S., & Gelman, A. (1998). General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7(4), 434–455. <https://doi.org/10.1080/10618600.1998.10474787>
- Brooks, S., Gelman, A., Jones, G., & Meng, X.-L. (2011). *Handbook of Markov chain Monte Carlo*. Boca Raton, FL: Chapman & Hall/CRC.
- Calderhead, B. (2014). A general construction for parallelizing Metropolis-Hastings algorithms. *Proceedings of the National Academy of Sciences*, 111(49), 17,408–17,413. <https://doi.org/10.1073/pnas.1408184111>
- Cao, T., Zeng, X., Wu, J., Wang, D., Sun, Y., Zhu, X., et al. (2018). Integrating MT-DREAMzs and nested sampling algorithms to estimate marginal likelihood and comparison with several other methods. *Journal of Hydrology*, 563, 750–765. <https://doi.org/10.1016/j.jhydrol.2018.06.055>
- Chang, H., Zhang, D., & Lu, Z. (2010). History matching of facies distribution with the EnKF and level set parameterization. *Journal of Computational Physics*, 229(20), 8011–8030. <https://doi.org/10.1016/j.jcp.2010.07.005>
- Chen, Y., & Zhang, D. (2006). Data assimilation for transient flow in geologic formations via ensemble Kalman filter. *Advances in Water Resources*, 29(8), 1107–1122. <https://doi.org/10.1016/j.advwatres.2005.09.007>
- Chou, Y.-M., Polansky, A. M., & Mason, R. L. (1998). Transforming non-normal data to normality in statistical process control. *Journal of Quality Technology*, 30(2), 133–141. <https://doi.org/10.1080/00224065.1998.11979832>
- Cotter, S. L., Roberts, G. O., Stuart, A. M., & White, D. (2013). MCMC methods for functions: Modifying old algorithms to make them faster. *Statistical Science*, 28(3), 424–446. <https://doi.org/10.1214/13-STS421>
- Craiu, R. V., Rosenthal, J., & Yang, C. (2009). Learn from thy neighbor: Parallel-chain and regional adaptive MCMC. *Journal of the American Statistical Association*, 104(488), 1454–1466. <https://doi.org/10.1198/jasa.2009.tm08393>
- Crestani, E., Camporese, M., Baù, D., & Salandin, P. (2013). Ensemble Kalman filter versus ensemble smoother for assessing hydraulic conductivity via tracer test data assimilation. *Hydrology and Earth System Sciences*, 17(4), 1517–1531. <https://doi.org/10.5194/hess-17-1517-2013>
- Cui, T., Law, K. J., & Marzouk, Y. M. (2016). Dimension-independent likelihood-informed MCMC. *Journal of Computational Physics*, 304, 109–137. <https://doi.org/10.1016/j.jcp.2015.10.008>
- Duan, Q., Schaake, J., Andreassian, V., Franks, S., Goteti, G., Gupta, H. V., et al. (2006). Model Parameter Estimation Experiment (MOPEX): An overview of science strategy and major results from the second and third workshops. *Journal of Hydrology*, 320(1-2), 3–17. <https://doi.org/10.1016/j.jhydrol.2005.07.031>
- Duane, S., Kennedy, A. D., Pendleton, B. J., & Roweth, D. (1987). Hybrid Monte Carlo. *Physics Letters B*, 195(2), 216–222. [https://doi.org/10.1016/0370-2693\(87\)91197-X](https://doi.org/10.1016/0370-2693(87)91197-X)
- Evensen, G. (2009). *Data assimilation: The ensemble Kalman filter*. Berlin, Germany: Springer.
- Fienen, M. N., D'Oria, M., Doherty, J. E., & Hunt, R. J. (2013). *Approaches in highly parameterized inversion: bgaPEST, a Bayesian geostatistical approach implementation with PEST: Documentation and instructions*. Reston, VA, US: Geological Survey. Retrieved from <https://pubs.usgs.gov/tm/07/c09/pdf/TM7-C9.pdf>
- Gelman, A., Roberts, G., & Gilks, W. (1996). *Bayesian statistics*. UK: Oxford University Press.
- Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4), 457–472. <https://doi.org/10.1214/ss/1177011136>
- Gilks, W. R., & Roberts, G. O. (1996). Strategies for improving MCMC. In W. R. Gilks, S. Richardson, & D. J. Spiegelhalter (Eds.), *Markov chain Monte Carlo in practice*. Boca Raton, FL: Chapman & Hall/CRC.
- Gilks, W. R., Roberts, G. O., & George, E. I. (1994). Adaptive direction sampling. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 43(1), 179–189. <https://doi.org/10.2307/2348942>
- Girolami, M., & Calderhead, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2), 123–214. <https://doi.org/10.1111/j.1467-9868.2010.00765.x>
- Haario, H., Laine, M., Mira, A., & Saksman, E. (2006). DRAM: Efficient adaptive MCMC. *Statistics and Computing*, 16(4), 339–354. <https://doi.org/10.1007/s11222-006-9438-0>
- Haario, H., Saksman, E., & Tamminen, J. (1999). Adaptive proposal distribution for random walk Metropolis algorithm. *Computational Statistics*, 14(3), 375–396. <https://doi.org/10.1007/s001800050022>
- Haario, H., Saksman, E., & Tamminen, J. (2001). An adaptive Metropolis algorithm. *Bernoulli*, 7(2), 223–242.

- Harbaugh, A. W., Banta, E. R., Hill, M. C., & McDonald, M. G. (2000). *MODFLOW-2000, the U. S. Geological Survey modular ground-water model-user guide to modularization concepts and the ground-water flow process*. Reston, VA: U. S. Geological Survey. Retrieved from <https://pubs.usgs.gov/of/2000/0092/report.pdf>
- Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, *57*(1), 97–109. <https://doi.org/10.1093/biomet/57.1.97>
- Hoffman, M. D., & Gelman, A. (2014). The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, *15*(1), 1593–1623.
- Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, *82*(1), 35–45. <https://doi.org/10.1115/1.3662552>
- Kavetski, D., Kuczera, G., & Franks, S. W. (2006a). Bayesian analysis of input uncertainty in hydrological modeling: 1. Theory. *Water Resources Research*, *42*, W03407. <https://doi.org/10.1029/2005WR004368>
- Kavetski, D., Kuczera, G., & Franks, S. W. (2006b). Bayesian analysis of input uncertainty in hydrological modeling: 2. Application. *Water Resources Research*, *42*, W03408. <https://doi.org/10.1029/2005WR004376>
- Kuczera, G., & Parent, E. (1998). Monte Carlo assessment of parameter uncertainty in conceptual catchment models: The Metropolis algorithm. *Journal of Hydrology*, *211*(1–4), 69–85. [https://doi.org/10.1016/S0022-1694\(98\)00198-X](https://doi.org/10.1016/S0022-1694(98)00198-X)
- Lagarias, J. C., Reeds, J. A., Wright, M. H., & Wright, P. E. (1998). Convergence properties of the Nelder-Mead simplex method in low dimensions. *SIAM Journal on Optimization*, *9*(1), 112–147. <https://doi.org/10.1137/S1052623496303470>
- Laloy, E., Rogiers, B., Vrugt, J. A., Mallants, D., & Jacques, D. (2013). Efficient posterior exploration of a high-dimensional groundwater model from two-stage Markov chain Monte Carlo simulation and polynomial chaos expansion. *Water Resources Research*, *49*, 2664–2682. <https://doi.org/10.1002/wrcr.20226>
- Laloy, E., & Vrugt, J. A. (2012). High-dimensional posterior exploration of hydrologic models using multiple-try DREAM<sub>(ZS)</sub> and high-performance computing. *Water Resources Research*, *48*, W01526. <https://doi.org/10.1029/2011WR010608>
- Liu, J. S., Liang, F., & Wong, W. H. (2000). The multiple-try method and local optimization in Metropolis sampling. *Journal of the American Statistical Association*, *95*(449), 121–134. <https://doi.org/10.2307/2669532>
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953). Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, *21*(6), 1087–1092. <https://doi.org/10.1063/1.1699114>
- Muleta, M. K., McMillan, J., Amenu, G. G., & Burian, S. J. (2012). Bayesian approach for uncertainty analysis of an urban storm water model and its application to a heavily urbanized watershed. *Journal of Hydrologic Engineering*, *18*(10), 1360–1371.
- Neal, R. M. (2011). MCMC using Hamiltonian dynamics. In S. Brooks, A. Gelman, G. Jones, & X.-L. Meng (Eds.), *Handbook of Markov chain Monte Carlo*. Boca Raton, FL: Chapman & Hall/CRC.
- Neiswanger, W., Wang, C., & Xing, E. (2013). Asymptotically exact, embarrassingly parallel MCMC. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence* (pp. 623–632) in July, 2014. held in Quebec City, Quebec, Canada. Arlington, VA: AUAI Press.
- Price, K., Storn, R. M., & Lampinen, J. A. (2006). *Differential evolution: A practical approach to global optimization*. Berlin: Springer Science & Business Media.
- Ramin, E., Wágner, D. S., Yde, L., Binning, P. J., Rasmussen, M. R., Mikkelsen, P. S., & Plósz, B. G. (2014). A new settling velocity model to describe secondary sedimentation. *Water Research*, *66*, 447–458. <https://doi.org/10.1016/j.watres.2014.08.034>
- Refsgaard, J. C., Christensen, S., Sonnenborg, T. O., Seifert, D., Højberg, A. L., & Trolldborg, L. (2012). Review of strategies for handling geological uncertainty in groundwater flow and transport modeling. *Advances in Water Resources*, *36*, 36–50. <https://doi.org/10.1016/j.advwatres.2011.04.006>
- Renard, B., Kavetski, D., Leblond, E., Thyer, M., Kuczera, G., & Franks, S. W. (2011). Toward a reliable decomposition of predictive uncertainty in hydrological modeling: Characterizing rainfall errors using conditional simulation. *Water Resources Research*, *47*, W11516. <https://doi.org/10.1029/2011WR010643>
- Roberts, G. O., Gelman, A., & Gilks, W. R. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability*, *7*(1), 110–120.
- Roberts, G. O., & Rosenthal, J. S. (1998). Optimal scaling of discrete approximations to Langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *60*(1), 255–268. <https://doi.org/10.1111/1467-9868.00123>
- Roberts, G. O., & Rosenthal, J. S. (2007). Coupling and ergodicity of adaptive Markov chain Monte Carlo algorithms. *Journal of Applied Probability*, *44*(2), 458–475. <https://doi.org/10.1239/jap/1183667414>
- Roberts, G. O., & Stramer, O. (2002). Langevin diffusions and Metropolis-Hastings algorithms. *Methodology and Computing in Applied Probability*, *4*(4), 337–357. <https://doi.org/10.1023/A:1023562417138>
- Sakia, R. (1992). The Box-Cox transformation technique: A review. *Journal of the Royal Statistical Society: Series D (The Statistician)*, *41*(2), 169–178. <https://doi.org/10.2307/2348250>
- Schoups, G., & Vrugt, J. A. (2010). A formal likelihood function for parameter and predictive inference of hydrologic models with correlated, heteroscedastic, and non-Gaussian errors. *Water Resources Research*, *46*, W10531. <https://doi.org/10.1029/2009WR008933>
- Shi, X., Ye, M., Curtis, G. P., Miller, G. L., Meyer, P. D., Kohler, M., et al. (2014). Assessment of parametric uncertainty for groundwater reactive transport modeling. *Water Resources Research*, *50*, 4416–4439. <https://doi.org/10.1002/2013WR013755>
- Shi, X., Ye, M., Finsterle, S., & Wu, J. (2012). Comparing nonlinear regression and Markov Chain Monte Carlo methods for assessment of prediction uncertainty in vadose zone modeling. *Vadose Zone Journal*, *11*(4), vjz2011.0147.
- Solonen, A., Ollinaho, P., Laine, M., Haario, H., Tamminen, J., & Järvinen, H. (2012). Efficient MCMC for climate model parameter estimation: Parallel adaptive chains and early rejection. *Bayesian Analysis*, *7*(3), 715–736.
- Storn, R., & Price, K. (1997). Differential evolution—A simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, *11*(4), 341–359. <https://doi.org/10.1023/A:100820282>
- Strid, I. (2010). Efficient parallelisation of Metropolis-Hastings algorithms using a prefetching approach. *Computational Statistics & Data Analysis*, *54*(11), 2814–2835. <https://doi.org/10.1016/j.csda.2009.11.019>
- ter Braak, C. J. (2006). A Markov chain Monte Carlo version of the genetic algorithm Differential Evolution: Easy Bayesian computing for real parameter spaces. *Statistics and Computing*, *16*(3), 239–249. <https://doi.org/10.1007/s11222-006-8769-1>
- ter Braak, C. J., & Vrugt, J. A. (2008). Differential evolution Markov chain with snooker updater and fewer chains. *Statistics and Computing*, *18*(4), 435–446. <https://doi.org/10.1007/s11222-008-9104-9>
- Volpi, E., Schoups, G., Firmani, G., & Vrugt, J. A. (2017). Sworn testimony of the model evidence: Gaussian Mixture Importance (GAME) sampling. *Water Resources Research*, *53*, 6133–6158. <https://doi.org/10.1002/2016WR020167>
- Vrugt, J. A. (2016). Markov chain Monte Carlo simulation using the DREAM software package: Theory, concepts, and MATLAB implementation. *Environmental Modelling & Software*, *75*, 273–316. <https://doi.org/10.1016/j.envsoft.2015.08.013>



- Vrugt, J. A., Diks, C. G., & Clark, M. P. (2008). Ensemble Bayesian model averaging using Markov chain Monte Carlo sampling. *Environmental Fluid Mechanics*, 8(5-6), 579–595. <https://doi.org/10.1007/s10652-008-9106-3>
- Vrugt, J. A., Diks, C. G., Gupta, H. V., Bouten, W., & Verstraten, J. M. (2005). Improved treatment of uncertainty in hydrologic modeling: Combining the strengths of global optimization and data assimilation. *Water Resources Research*, 41, W01017. <https://doi.org/10.1029/2004WR003059>
- Vrugt, J. A., Gupta, H. V., Bouten, W., & Sorooshian, S. (2003). A Shuffled Complex Evolution Metropolis algorithm for optimization and uncertainty assessment of hydrologic model parameters. *Water Resources Research*, 39(8), 1201. <https://doi.org/10.1029/2002WR001642>
- Vrugt, J. A., ter Braak, C. J., Clark, M. P., Hyman, J. M., & Robinson, B. A. (2008). Treatment of input uncertainty in hydrologic modeling: Doing hydrology backward with Markov chain Monte Carlo simulation. *Water Resources Research*, 44, W00B09. <https://doi.org/10.1029/2007WR006720>
- Vrugt, J. A., ter Braak, C. J., Diks, C., Robinson, B. A., Hyman, J. M., & Higdon, D. (2009). Accelerating Markov chain Monte Carlo simulation by differential evolution with self-adaptive randomized subspace sampling. *International Journal of Nonlinear Sciences and Numerical Simulation*, 10(3), 273–290. <https://doi.org/10.1515/IJNSNS.2009.10.3.273>
- Vrugt, J. A., ter Braak, C. J., Diks, C. G., & Schoups, G. (2013). Hydrologic data assimilation using particle Markov chain Monte Carlo simulation: Theory, concepts and applications. *Advances in Water Resources*, 51, 457–478. <https://doi.org/10.1016/j.advwatres.2012.04.002>
- Wagener, T., & Gupta, H. V. (2005). Model identification for hydrological forecasting under uncertainty. *Stochastic Environmental Research and Risk Assessment*, 19(6), 378–387. <https://doi.org/10.1007/s00477-005-0006-5>
- Weare, J. (2007). Efficient Monte Carlo sampling by parallel marginalization. *Proceedings of the National Academy of Sciences*, 104(31), 12,657–12,662. <https://doi.org/10.1073/pnas.0705418104>
- Wöhling, T., & Vrugt, J. A. (2011). Multiresponse multilayer vadose zone model calibration using Markov chain Monte Carlo simulation and field water retention data. *Water Resources Research*, 47, W04510. <https://doi.org/10.1029/2010WR009265>
- Yang, Z., & Rodriguez, C. E. (2013). Searching for efficient Markov chain Monte Carlo proposal kernels. *Proceedings of the National Academy of Sciences*, 110(48), 19,307–19,312. <https://doi.org/10.1073/pnas.1311790110>
- Zeng, X., Ye, M., Wu, J., Wang, D., & Zhu, X. (2018). Improved nested sampling and surrogate-enabled comparison with other marginal likelihood estimators. *Water Resources Research*, 54, 797–826. <https://doi.org/10.1002/2017WR020782>
- Zhang, J., Lin, G., Li, W., Wu, L., & Zeng, L. (2018). An iterative local updating ensemble smoother for estimation and uncertainty assessment of hydrologic model parameters with multimodal distributions. *Water Resources Research*, 54, 1716–1733. <https://doi.org/10.1002/2017WR020906>
- Zhang, D., & Lu, Z. (2004). An efficient, high-order perturbation approach for flow in random porous media via Karhunen-Loève and polynomial expansions. *Journal of Computational Physics*, 194(2), 773–794. <https://doi.org/10.1016/j.jcp.2003.09.015>
- Zhang, J., Man, J., Lin, G., Wu, L., & Zeng, L. (2018). Inverse modeling of hydrologic systems with adaptive multifidelity Markov chain Monte Carlo simulations. *Water Resources Research*, 54, 4867–4886. <https://doi.org/10.1029/2018WR022658>
- Zheng, F., Tao, R., Maier, H. R., See, L., Savic, D., Zhang, T., et al. (2018). Crowdsourcing methods for data collection in geophysics: State of the art, issues, and future directions. *Reviews of Geophysics*, 56, 698–740. <https://doi.org/10.1029/2018RG000616>
- Zheng, C., & Wang, P. P. (1999). MT3DMS: A modular three-dimensional multispecies transport model for simulation of advection, dispersion, and chemical reactions of contaminants in groundwater systems; documentation and user's guide. DTIC Document Retrieved from <http://www.geology.wisc.edu/courses/g727/mt3dmanual.pdf>
- Zhou, H., Gomez-Hernandez, J. J., Franssen, H.-J. H., & Li, L. (2011). An approach to handling non-gaussianity of parameters and state variables in ensemble Kalman filtering. *Advances in Water Resources*, 34, 844–864. <https://doi.org/10.1016/j.advwatres.2011.04.014>