Computers & Geosciences 40 (2012) 19–27

Contents lists available at ScienceDirect

Computers & Geosciences

journal homepage: www.elsevier.com/locate/cageo

An agent-based approach to global uncertainty and sensitivity analysis

Dylan R. Harp *, Velimir V. Vesselinov

Earth and Environment Science Division, Los Alamos National Laboratory, Los Alamos, NM, USA

ARTICLE INFO

Article history:
Received 2 February 2011
Received in revised form
6 May 2011
Accepted 29 June 2011
Available online 28 July 2011

Keywords:
Agent-based
Global uncertainty analysis

ABSTRACT

A novel sampling approach to global uncertainty and sensitivity analyses of modeling results utilizing concepts from agent-based modeling is presented (Agent-Based Analysis of Global Uncertainty and Sensitivity (ABAGUS)). A plausible model parameter space is discretized and sampled by a particle swarm where the particle locations represent unique model parameter sets. Particle locations are optimized based on a model-performance metric using a standard particle swarm optimization (PSO) algorithm. Locations producing a performance metric below a specified threshold are collected. In subsequent visits to the location, a modified value of the performance metric, proportionally increased above the acceptable threshold (i.e., convexities in the response surface become concavities), is provided to the PSO algorithm. As a result, the methodology promotes a global exploration of a plausible parameter space, and discourages, but does not prevent, reinvestigation of previously explored regions. This effectively alters the strategy of the PSO algorithm from optimization to a sampling approach providing global uncertainty and sensitivity analyses. The viability of the approach is demonstrated on 2D Griewank and Rosenbrock functions. This also demonstrates the set-based approach of ABAGUS as opposed to distribution-based approaches. The practical application of the approach is demonstrated on a 3D synthetic contaminant transport case study. The evaluation of global parametric uncertainty using ABAGUS is demonstrated on model parameters defining the source location and transverse/longitudinal dispersivities. The evaluation of predictive uncertainties using ABAGUS is demonstrated for contaminant concentrations at proposed monitoring wells.

1. Introduction

Inverse approaches are routinely used to identify appropriate values of model parameters that provide simulations with the highest degree of consistency with existing observations. These approaches can be considered to provide answers to the question, "What do the observations and model tell us about the parameters?" An often neglected question is, "What do the observations and model have the ability to tell us about the parameters?" An answer to the second question is needed to properly evaluate the significance and uncertainty of the answers to the first question. Approaches that answer the second question explore the effect of changes in parameter values on a performance metric and are considered model-based uncertainty analysis (UA) approaches.

UA is often based on sensitivity analysis techniques. Local sensitivity analyses evaluate the sensitivities surrounding a solution by calculating derivatives of model simulations with respect to model parameters (Vecchia and Cooley, 1987; Cooley, 1993) or adjoint solutions of the governing equation (Neuman, 1980; Sykes et al., 1985; Li and Yeh, 1998). Local sensitivity analysis approaches are computationally efficient, requiring relatively few model calls operating under the assumption that parameter probability distributions are normally distributed. These techniques are commonly utilized in gradient-based optimization strategies for parameter estimation. The information provided by these techniques in a UA is limited to a region surrounding the current parameter values, to models with a continuous parameter space, and by the assumption of normally distributed parameter uncertainty.

Null-space Monte Carlo (NSMC) combines concepts from error variance analysis theory and Monte Carlo (MC) sampling to perform UA on highly parameterized models (Tonkin and Doherty, 2009). The null space is defined from local sensitivities of a calibrated model. For a given set of best model parameter estimates, the null space is a subspace of the parameter space composed of parameter combinations that have negligible impact on the performance metric. An MC sampling is utilized to produce parameter realizations by modifying parameter values within the calibration null space. If, in the process of MC sampling, a parameter realization produces an uncalibrated model, parameters in the calibration solution space are reestimated to recalibrate the model. This produces a local UA capable of

* Corresponding author.
E-mail addresses: dharp@lanl.gov (D.R. Harp), vvv@lanl.gov (V.V. Vesselinov).

0098-3004/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cageo.2011.06.025

reducing the computational burden imposed by a large numbers of parameters.

Most global sensitivity analysis approaches are based on evaluating the relative contribution of individual and combinations of parameters to the variance of a performance metric (Sobol, 2001; van Werkhoven et al., 2008; Wagener et al., 2009). These approaches provide scalar indices of global sensitivity. This information indicates parameters of interest and correlated parameter estimates. Such analyses do not provide specific information about sensitivities at any specific point in the parameter space.

Evaluation of the global uncertainty of a model is typically based on global sampling approaches. Vrugt et al. (2008) introduced a Markov chain Monte Carlo (MCMC) approach entitled DiffeRential Evolution Adaptive Metropolis (DREAM). This approach provides estimates of posterior density functions of parameters in a formal Bayesian framework. An informal Bayesian approach to global UA is the Generalized Likelihood Uncertainty Analysis (GLUE) developed by Beven and Binley (1992). This approach performs an MC analysis using a statistically informal likelihood function to rank model performance. Recently, Harp and Vesselinov (in press) developed a sampling approach to global UA of stochastic models of flow medium heterogeneity, introducing the concept of an acceptance probability of a stochastic parameter set. Sampling approaches have the ability to provide detailed information directly addressing the UA. The drawback of such approaches is that the number of model calls is often too large for many practical applications involving process-based models (Keating et al., 2010).

The approach presented here aims to provide an alternative to existing UA approaches that will be useful for complex problems for which a local UA is known to be incomplete and for which the model runs are too computationally intensive for a rigorous sampling-based inference approach. We refer to this approach as Agent-Based Analysis of Global Uncertainty and Sensitivity (ABAGUS). Concepts from agent-based modeling such as particle swarm optimization (PSO) (Kennedy and Eberhart, 1995; Clerc, 2006) and ant colony optimization (Dorigo and Stützle, 2004) have been utilized extensively in optimization algorithms. However, to our knowledge, their direct application to global UA has not been explored. The ABAGUS computational framework is based on integrating concepts of agent-based social simulation with the Standard PSO 2006 (SPSO2006) algorithm (Particle Swarm Central, 2006), effectively altering the strategy of SPSO2006 from optimization to global UA. SPSO2006 is chosen here, as it implements a parsimonious and efficient version of particle swarm optimization that is well-known and freely available for download.

The strategy of ABAGUS is to efficiently explore a discretized parameter space by storing information about locations producing simulations consistent with observations. ABAGUS does not require statistical convergence, and the computational expense of the approach can be reduced for initial explorations by coarsening the discretization. The algorithm alters the response surface at the previously sampled locations by increasing the associated performance metric (e.g., objective function, fitness function). As a result, if points within a local area of attraction were already visited by the algorithm, the region appears as a region of concavity (repulsion), as opposed to a region of convexity (attraction), discouraging future exploration. Similarities can easily be drawn between ABAGUS and the Sugarscape agent-based social simulator (Epstein and Axtell, 1996), designed to model the survival of a population on a regenerative resource; however, in ABAGUS, the resource is not regenerative, encouraging global exploration of the parameter space.

The ABAGUS approach differs from many existing sampling-based UA approaches, as it is a set-based approach where all locations below a certain level of consistency with observations are collected without performance-based preference. Therefore, outlying solutions that are marginally acceptable are represented with equal weight with solutions within clustered locations. In a statistical inference scheme, these marginally acceptable outlying solutions can be underrepresented in the results, as the frequency of sampling these isolated locations can be low. These outlying locations can be particularly revealing in the case of long- and heavy-tailed probabilistic distributions (such as nonnormal stable distributions), where the collective probability of a large number of extremely low-probability events is not negligible and cannot be characterized by the second moment of the Gaussian distribution—in other words, in cases where the probability of an extreme event is not negligible, but where the magnitude of the extreme event is uncertain. The set-based approach of ABAGUS provides results in a form that can be utilized by set-based analyses, such as info-gap theory (Ben-Haim, 2006), or within a GLUE framework using a "limits of acceptability" approach (Liu et al., 2009).

The ABAGUS approach is warranted in cases where normal (Gaussian) probabilistic distributions are deemed inappropriate to describe the statistical distribution of a property (e.g., fractal properties) as the statistical moments are ill-defined (stable probabilistic distributions with $\alpha < 2$ have divergent second moments and with $\alpha < 1$ divergent first moments; a Gaussian distribution is a special stable distribution with $\alpha = 2$ (Zolotarev, 1986)). Such situations occur more widely than is often acknowledged, particularly in modeling complicated systems in environmental investigations (Nolte et al., 1989; Neuman, 1990; Dimri, 2000). Therefore, ABAGUS provides an alternative UA approach in cases where a rigorous formal statistical inference scheme is inappropriate due to ill-defined statistical moments. The application of an ABAGUS-type approach in cases where statistical inference is deemed appropriate is ill-advised and would provide an inferior level of detail.

Since the ABAGUS algorithm is based on SPSO2006, a brief discussion of this algorithm is presented in Section 2. The ABAGUS algorithm is discussed in Section 3. Section 4 demonstrates the performance of ABAGUS on 2D Griewank and Rosenbrock functions. Section 5 presents a synthetic five-parameter contaminant transport problem that is utilized to demonstrate the use of ABAGUS on a practical application.

2. Standard PSO 2006 algorithm

SPSO2006 modifies a population of solutions called particles defined by their position and velocity in a D-dimensional parameter space. The position and velocity of the ith particle can be represented as $\mathbf{P}_i = [p_{i,1}, p_{i,2}, \ldots, p_{i,D}]$ and $\mathbf{V}_i = [v_{i,1}, v_{i,2} \ldots, v_{i,D}]$, respectively. An empirical formula of $S = 10 + \sqrt{D}$ for determining the swarm size S has been suggested (Particle Swarm Central, 2006). Particles retain a record of the best location they have visited so far, denoted as $\mathbf{B}_i = [b_{i,1}, b_{i,2}, \ldots, b_{i,D}]$. Particles are also informed of the best location that K other randomly chosen particles have visited, denoted as $\mathbf{G}_i = [g_{i,1}, g_{i,2}, \ldots, g_{i,D}]$. A standard value for K is 3 (Particle Swarm Central, 2006). These networks of informers are reinitialized after iterations with no improvement in the global best location of the swarm. The velocity of the ith particle in the jth dimension is updated from swarm iteration step k to $k+1$ as

$$v_{i,j}(k+1) = w v_{i,j}(k) + c_1 r_1 (b_{i,j} - p_{i,j}(k)) + c_2 r_2 (g_{i,j} - p_{i,j}(k)), \quad k = \{1, \ldots, D\}, \tag{1}$$

where w is a constant referred to as the inertia weight, c_1 and c_2 are constants referred to as acceleration coefficients, and r_1 and r_2 are independent uniform random numbers in [0,1]. The swarm

iteration steps are also referred to as time steps because they represent the progress of swarm development in the parameter space. The parameter w controls the level of influence of a particle's previous displacement on its current displacement; c_1 and c_2 scale the random influence of the particle's memory and its network of informers, respectively. Values of $w=0.72$ and $c_1=c_2=1.2$ have been demonstrated to perform well in many problems (Clerc, 2006). A limitation V_{max} on the magnitude of the velocity is commonly employed. The particle position at each iteration is updated as

$$p_{i,j}(k+1)=p_{i,j}(k)+v_{i,j}(k+1), \quad k=\{1,\ldots,D\}. \tag{2}$$

Additional details on SPSO2006 are available in Clerc (2006) and Cooren et al. (2009). The source code is available for download at Particle Swarm Central (2006).

3. ABAGUS algorithm

Concepts from agent-based modeling have found significant utility in global optimization. The following discusses the first, to our knowledge, utilization of agent-based modeling to perform global UA. A flow diagram of the ABAGUS algorithm is provided in Fig. 1 and discussed below.

As ABAGUS has been developed by modifying SPSO2006, its search algorithm is nearly identical to SPSO2006's; except that parameter space discretization is enforced on the particle movements. This is accomplished by moving proposed particle locations (defined by Eqs. (1) and (2)) to the nearest node of the discretization. The parameter space discretization is based on user-provided parameter-specific resolution (each parameter can be assigned a distinct resolution). The resolution of the analysis can therefore be controlled by the user, depending on computational constraints and/or the desired level of detail. ABAGUS runs can also be nested, using the samples from previous coarser runs as starting points for finer runs. The discretization of the parameter space does not hinder the UA, as the strategy is to identify regions of the parameter space producing simulations indistinguishably consistent with observations, and is not an optimization strategy intended to identify a single optimal solution.

ABAGUS collects parameter sets (locations within the discretized space) with a performance metric Φ below a defined threshold ε, and inverts the value of the performance metric as

$$\Phi_{inv}=2\varepsilon-\Phi, \quad \Phi<\varepsilon, \tag{3}$$

where Φ_{inv} is the value of the inverted performance metric. Φ_{inv} is provided to particles on subsequent visits to the location without recomputation of the model run. The value of ε can be defined

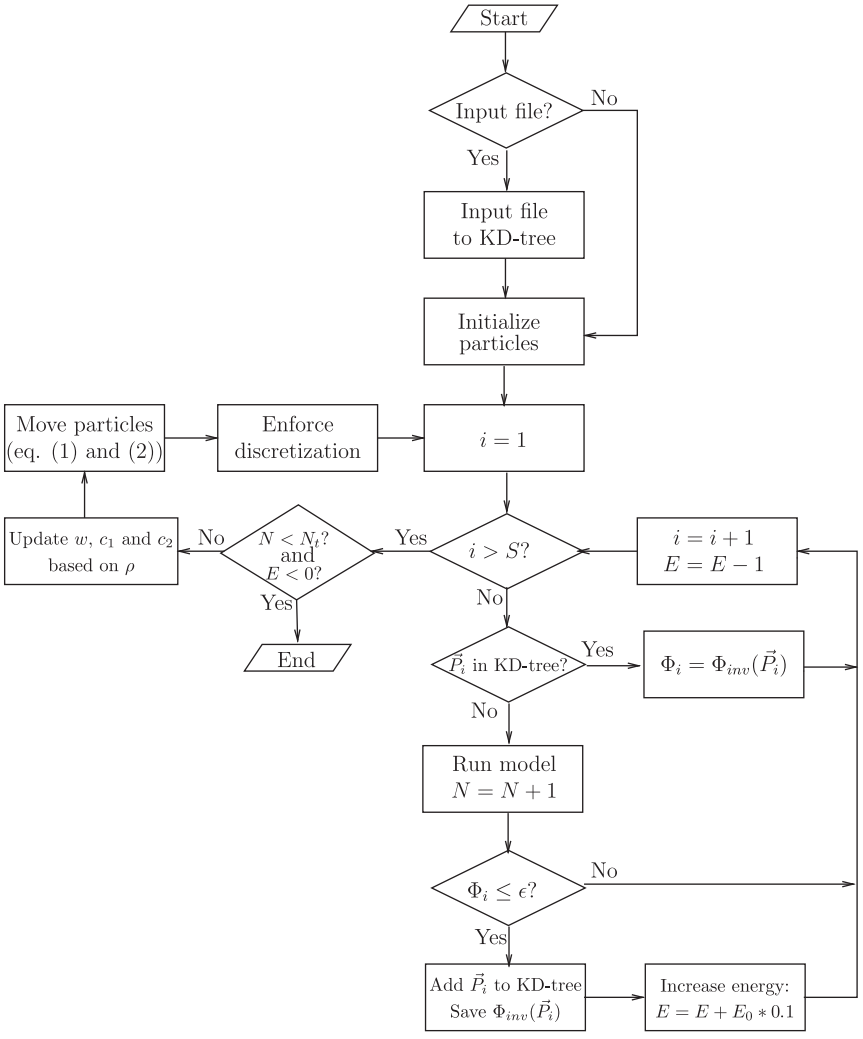

Fig. 1. ABAGUS flow diagram. N is a counter of the current number of function evaluations (model runs), N_t is the total number of allowable function evaluations, S is the number of particles, E is the swarm energy, E_0 is the initial swarm energy, w is the inertia weight, c_1 and c_2 are acceleration coefficients, ρ is the exploration rate metric, $\mathbf{P}_i$ is the current location of the ith particle, ε is the performance metric threshold, Φ_i is the current performance metric for the ith particle, and $\Phi_{inv}(\mathbf{P}_i)$ is the inverted performance metric associated with location $\mathbf{P}_i$.

based on theoretical (e.g., confidence levels under certain assumptions (Vecchia and Cooley, 1987; Cooley, 1993)) or problem-specific considerations (e.g., "limits of acceptability" (Liu et al., 2009)). The potentially large number of locations that must be collected are managed by a KD-tree, allowing the collected locations to be efficiently searched in a binary fashion in a K-dimensional space, where K can be any positive integer (Tsiombikas, 2009). The value of Φ_{inv} associated with the acceptable location is stored to provide to particles on future visits. In the case of ABAGUS, K equals the dimension of the parameter space (D). A nearest-neighbor search of the KD-tree is utilized to identify if a location has been collected previously (Tsiombikas, 2009). If the location has been collected, Φ_{inv} is provided; if not, a forward model run is executed to compute Φ for the location. The details of this process are illustrated in Fig. 1. As a result, revisiting collected locations has a relatively insignificant cost to the algorithm, particularly in cases involving long model execution times.

Eq. (3) effectively adds the discrepancy between ε and Φ to ε and assigns this value as the value of the performance metric associated with the location. The larger the discrepancy, the less attractive the position appears to future visits. As a result, convexities in the response surface become concavities.

As the ABAGUS algorithm progressively identifies and collects acceptable locations in the parameter space, the coefficients w, c_1, and c_2 are dynamically modified to maintain an appropriate balance between exploration and intensification. An exploration rate metric ρ quantifies the level of exploration at each iteration of the ABAGUS run as

$$\rho = N_e/N_r, \tag{4}$$

where N_e is the number of new locations visited this iteration and N_r is the number of revisits to previously collected positions this iteration (therefore, $N_e+N_r=S$ at each iteration). One iteration of ABAGUS involves updating and evaluating the population of solutions (particles). The following rules are used to maintain a reasonable value for ρ:

if $\rho < \rho_0 : w = w(1+a)$,

$c_1 = c_1(1+a)$,

$c_2 = c_2(1+a)$;

if $\rho > \rho_0 : w = w(1-d)$,

$c_1 = c_1(1-d)$,

$c_2 = c_2(1-d)$,

where ρ_0 is set by the user to a value deemed to be reasonable and a and d are constants greater than zero. In the cases investigated here, values of $\rho_0 = 1$ and $a = d = 10^{-5}$ were found to be effective. More complex strategies for controlling ρ by modifying w, c_1, and c_2 are easily conceptualized, and will take time and effort to evaluate on varied response surfaces.

ABAGUS is allowed to run to a maximum number of model evaluations (N_t) or until the swarm runs out of energy (E). The initial energy of the swarm (E_0) is specified by the user, where larger values of initial energy will allow more exploration of the parameter space. Each particle move decrements the swarm energy by one. Each identification of an acceptable location increments the swarm's energy. The swarm energy is incremented by 10% of the initial energy for the cases investigated here ($E = E + E_0 \times 0.1$). For an initial investigation of the parameter space, an initial energy of 10,000 is reasonable for the test cases presented here. These steps are illustrated in Fig. 1.

4. Test functions

The performance of ABAGUS is demonstrated on 2D Griewank and Rosenbrock test functions, defined as

$$z = \frac{x^2+y^2}{4000} - \cos\left(\frac{x}{\sqrt{2}}\right)\cos\left(\frac{y}{\sqrt{3}}\right) + 1 \tag{5}$$

and

$$z = (1-x)^2 + 100(y-x^2)^2, \tag{6}$$

respectively. The Griewank and Rosenbrock functions are benchmark problems presenting challenging response surfaces for optimization strategies. The Griewank function contains numerous local minima with a single global minimum of zero at (0,0). The Rosenbrock function contains a large smooth valley with a banana-shaped area of attraction surrounding an ill-defined global minimum of zero at (1,1).

Parameter bounds for x and y are both $[-100,100]$ and the parameter space is discretized to a 0.1 resolution for both functions, resulting in 4×10^6 possible locations. The value of ε is set to 0.1 for the Griewank run and 20 for the Rosenbrock run. The initial swarm energy is set to 10,000 and the number of function evaluations is limited to 2×10^6. Initial values for w, c_1, and c_2 are set according to the constant values commonly utilized by SPSO2006 (Particle Swarm Central, 2006) ($w=0.72$; $c = c_1 = c_2 = 1.2$, and c_1 and c_2 will be referred to collectively as c hereafter). To evaluate the performance of the ABAGUS algorithm, one particle is initialized to the global minimum ((0,0) for the Griewank function, (1,1) for the Rosenbrock function). This eliminates the initial search from random locations prior to the identification of an area of attraction, which, for the ABAGUS algorithm, is identical to SPSO2006 (Particle Swarm Central, 2006). The utilization of ABAGUS in this manner (i.e., beginning the ABAGUS run from a known optimal location obtained by a prior optimization) evaluates the capability of the algorithm to perform UA; the identification of the global minimum can be a separate task.

Fig. 2(a) and (b) presents maps of the response surfaces for the parameter space considered in the ABAGUS runs for the Griewank and Rosenbrock functions, respectively. Fig. 2(c) and (d) presents 3D plots of the structure of the response surfaces near the global minimum for each case. The results of the ABAGUS runs are presented in Fig. 2(e) and (f) as maps of the response surfaces at identified locations. It is apparent that for both test functions, ABAGUS is able to identify the local and global areas of attraction containing acceptable solutions. The set-based nature of the approach, and its lack of a need for distributional assumptions, is evident in these results, as opposed to many UA approaches (e.g., Bayesian approaches). This fundamental difference in approach between ABAGUS and distribution-based approaches makes direct comparison difficult, and is not attempted here. However, it should be apparent that approaches that require assumptions of probabilistic distributions of parameter uncertainty will have difficulties with these types of response surfaces, particularly for the Griewank function.

The Griewank run collected 1552 locations with $\Phi < \varepsilon = 0.1$ from approximately 2.00×10^6 function evaluations with approximately 2.08×10^6 revisits to collected locations. The Rosenbrock run collected 324 locations with $\Phi < \varepsilon = 20$ from 109,060 function evaluations with approximately 2.25×10^5 revisits. The Griewank run took 15 s with approximately 1.3×10^5 function evaluations per second and the Rosenbrock run took 1 s with approximately 1.1×10^5 function evaluations per second on a 2.8 GHz processor.

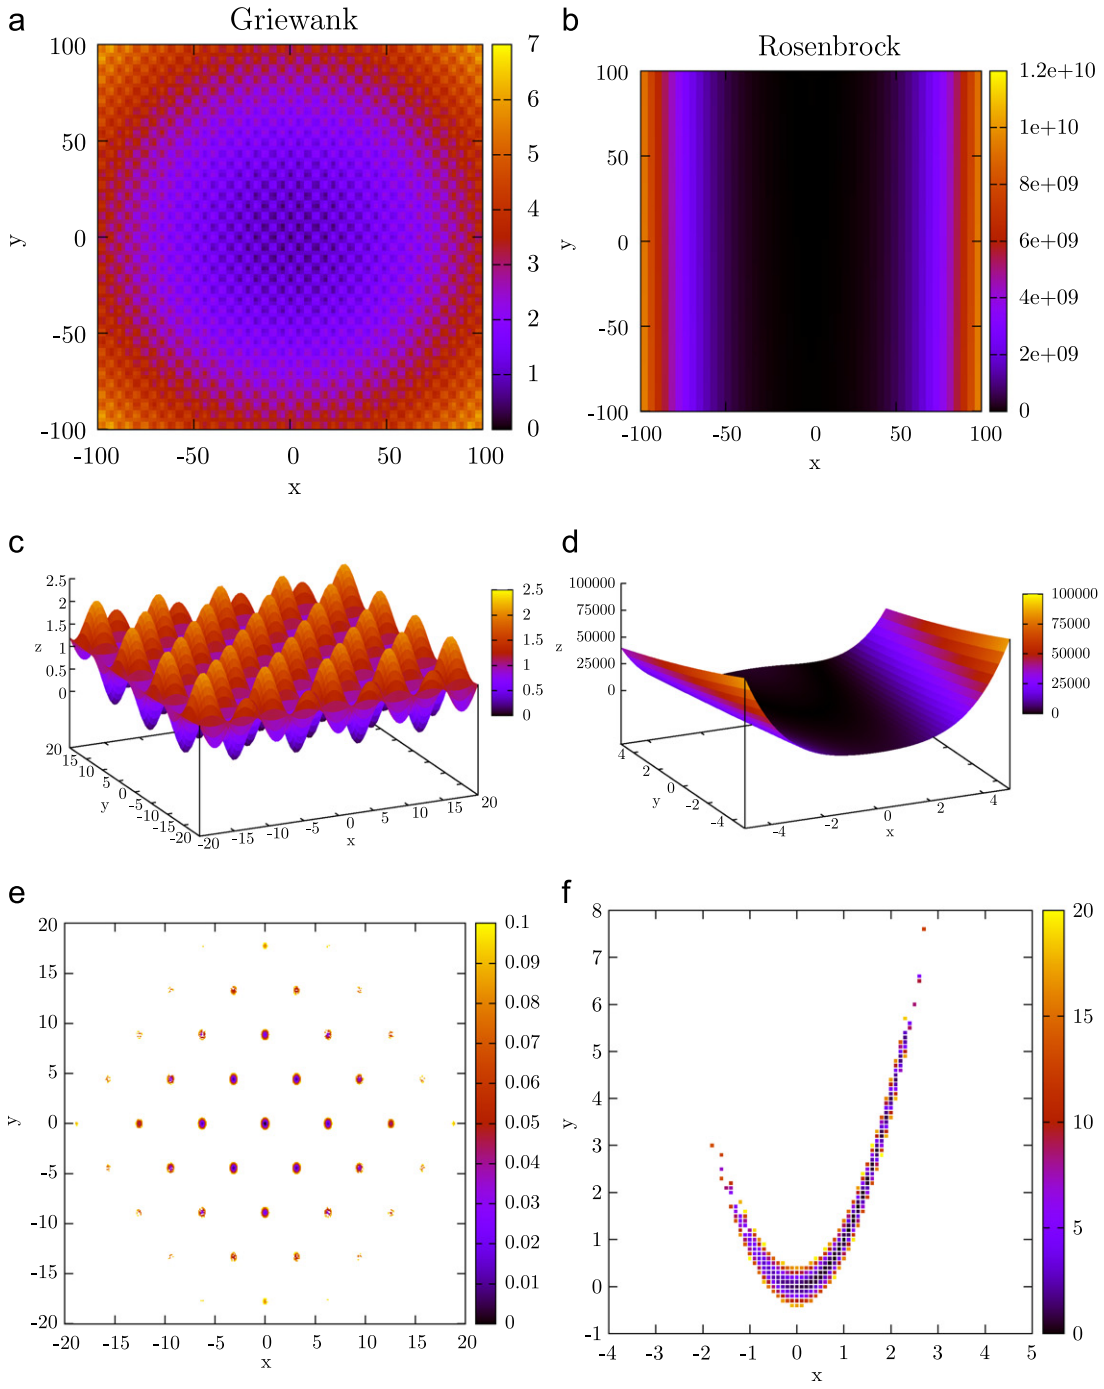

Fig. 2. Griewank and Rosenbrock test function analyses. Maps of the response surface for the full parameter space considered in the search [−100,100] are presented in (a) and (b). (c) and (d) present 3D surfaces of the objective function near the region of the parameter space with values below the cutoff. (e) and (f) present the results of the ABAGUS runs identifying the solutions below predefined cutoffs equal to 0.1 and 20, respectively. A global minimum of 0 exists at (0,0) for the Griewank function and (1,1) for the Rosenbrock function.

5. Contaminant transport case study

The ABAGUS approach is demonstrated on a synthetic contaminant transport problem to explore the model-based uncertainty of distributed contaminant concentrations in an analytical contaminant transport model (Vesselinov and Harp, 2010) considering uncertainty in the plume source location (x_s, y_s) and dispersivities (a_x, a_y, a_z). Flow is in the x-direction. True concentrations are collected from a simulation of the model given true parameter values listed in Table 1. Information regarding the parameters (e.g. value, min, max, and resolution) is presented in Table 2. The collected concentrations have been rounded to values similar in resolution to field-collected measurements. A concentration map of the "truth" at $t=49$ years is presented in Fig. 3.

The performance metric for the contaminant transport case study is a sum of the squared residuals (SSR) expressed as

$$\Phi(\theta) = \sum_{i=1}^{N} (\hat{c}_i(\theta) - c_i)^2, \tag{7}$$

Table 1
Well coordinates, screen top (z_{top}) and bottom (z_{bot}) depths below the water table, and year and value of observed contaminant concentrations.

Well	x (m)	y (m)	z_{top} (m)	z_{bot} (m)	t (y)	c (ppb)
w01	1296	2154	5.57	12.55	49	0.1
w02	1906	1679	36.73	55.14	49	1
w03	212	1150	0	15.04	49	0
w04	1170	1735	13.15	20.41	44	354
					49	392
w05	3062	1274	26.73	33.71	49	0
w06	1906	2494	69.01	83.98	49	0
w07	1879	2484	11.15	18.19	49	0
w08	2563	2320	4.86	11.87	49	0
w09	769	1650	3.66	10.09	49	2140
w10	516	1799	3.32	9.63	49	5
			23.2	26.24	49	2
w11	1644	1568	4.94	7.99	49	48
			32.46	35.48	16	0
w12	1554	1837	3.59	6.64	49	42
			32.51	38.61	12	0
w13	1278	1349	3	6	50	18
			36	42	50	4
d01	496	1579	3	6	–	–
d02	986	1440	3	6	–	–
d03	1236	1945	3	6	–	–
d04	1858	1394	3	6	–	–

Note: Year and concentration are omitted for proposal wells ("d" wells).

Table 2
Parameter values and resolution for the contaminant transport case study.

	x_s (m)	y_s (m)	a_x (m) ($\log_{10} a_x$)	a_y (m) ($\log_{10} a_y$)	a_z (m) ($\log_{10} a_z$)
Value	810	1657	70 (1.845)	20 (1.301)	0.6 (−0.222)
Min	400	1000	30 (1.477)	5 (0.699)	0.1 (−1.0)
Max	1200	2000	200 (2.301)	30 (1.477)	5 (0.699)
Resolution	0.5	0.5	1.0 (0.005)	0.16 (0.005)	0.029 (0.01)

Note: Log-transformed dispersivities are presented in parenthesis as these are the values provided to ABAGUS for the case study.

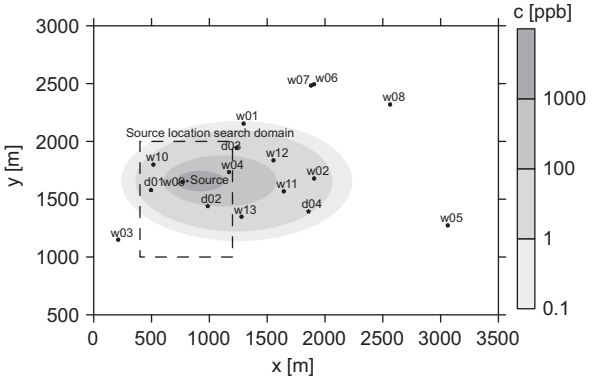

Fig. 3. "True" contaminant concentration map at 49 years. Circles represent monitoring well locations. A dashed-line rectangle indicates the parameter bounds for x_s and y_s. The "true" contaminant source is indicated.

where Φ is the performance metric, θ is a vector containing the parameter values, $\hat{c}(\theta)$ is a vector of simulated concentrations resulting from θ, c is a vector containing the observed concentrations, and N is the number of observations. Due to the rounding of the collected concentrations, a value of $\Phi = 0.14$ is obtained from the true parameter values.

It is assumed that we are interested in collecting parameter sets producing values of Φ below 100 ($\varepsilon = 100$; refer to Eq. (3)). Below the cutoff values, the discrepancies between the model-predicted and observed concentrations are assumed to be due to measurement errors and other factors not captured by the applied model. As a result, all the realizations below the cutoff value are assumed to be equally consistent. The true parameters are provided to define the location of one of the initial particles in the swarm, similarly to providing the optimal location from a previous optimization run. The initial energy is set to 10,000 and the maximum number of model calls is 200,000. As in the test functions, initial values for w and c are set according to values commonly utilized by SPSO2006 (Particle Swarm Central, 2006) ($w = 0.72$; $c = 1.2$).

Fig. 4 presents histograms of the parameter values obtained by the ABAGUS run. This information differs from posterior distributions of a Bayesian analysis in that the histograms are not weighted by the performance metric (i.e., likelihood function). It is also possible to rank the acceptable parameter sets by some model-performance or statistical-interference metric. Each collected discrete parameter set is represented equally within the histogram (i.e., assumed to be equally consistent). The histograms present a frequency analysis only within the context of the samples collected by ABAGUS, which are discrete in nature. The histograms are intended to summarize the results of the ABAGUS run, but should not be considered to be a formal statistical frequency analysis. Within set-based analyses (Ben-Haim, 2006) this representation of parameter uncertainty is appropriate. It is apparent that the histograms include the "true" values for all parameters (Table 2).

Fig. 5 presents a map of the lowest Φ value collected at each source location (x_s,y_s). Multiple Φ values are possible at each source location due to combinations of a_x, a_y, and a_z. While the histograms in Fig. 4 are not centered on the true parameter values, Fig. 5 demonstrates that the lowest Φ values are centered on the true location. This is not apparent in the histograms of Fig. 4, where all collected parameter sets are represented as equally consistent with observations.

Fig. 6 presents histograms of log-transformed predicted concentrations at the proposed well locations (d01, d02, d03, d04) associated with the histograms of collected parameter values in Fig. 4. This constitutes a model-based predictive uncertainty analysis. The histograms indicate varying degrees of predictive uncertainty, with concentrations varying over nine orders of magnitude for d03, and around five orders of magnitude for d02 and d03. The predictive uncertainties are nonparametric, allowing an empirical evaluation unconstrained by any prespecified probabilistic distribution.

The ABAGUS run collected 3590 parameter sets producing $\Phi < \varepsilon = 100$ from 1×10^5 model evaluations. The total number of plausible locations in the discretized parameter space is 4×10^{15}. The ABAGUS run took approximately 23 min on a 2.8 GHz processor, with approximately 117 model evaluations per second. The resolution of x_s and y_s is 10 cm. This level of detail is not likely significant in a practical application, but is used here for demonstration purposes. A coarser level of detail in x_s and y_s would require fewer model calls.

Inspection of the results summarized by Fig. 4 provides information answering the question discussed in the Introduction: "What do the observations and model have the ability to tell us about the parameters?" The summary provided by the histograms in Fig. 4 indicates the frequency of discrete parameter values producing equally consistent simulations of the

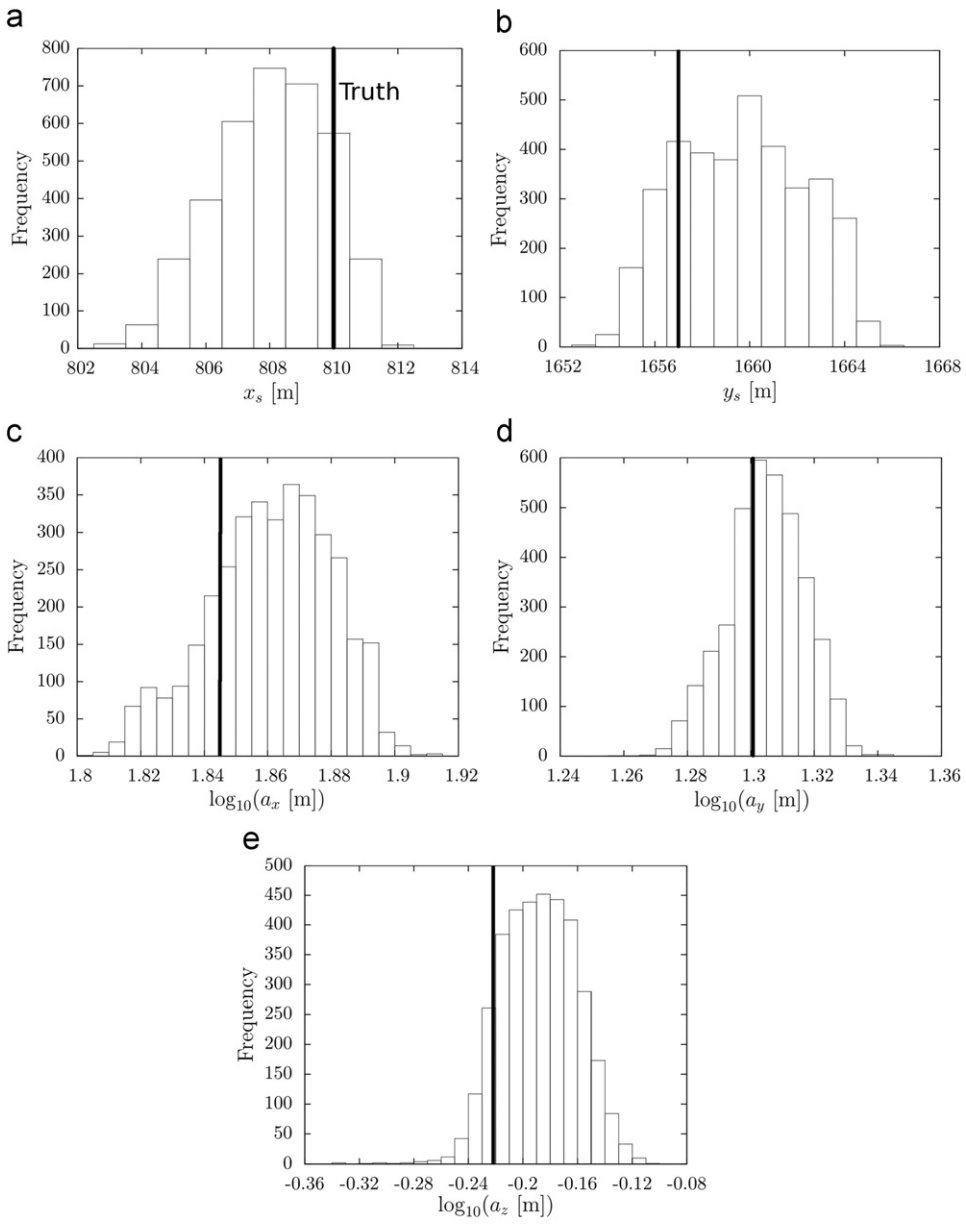

Fig. 4. Histograms of parameter values obtained from ABAGUS evaluation. "True" values are indicated by bold vertical lines.

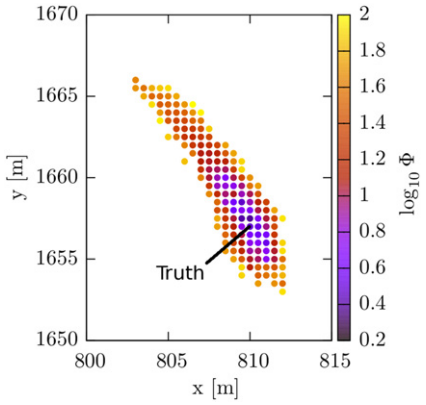

Fig. 5. Map of log-transformed minimum performance metric ($\log_{10}\Phi$) values at identified source locations x_s and y_s. The location of the "true" source is indicated.

observations assuming a value of $\varepsilon = 100$ indicating model parameter uncertainty and sensitivity. The histograms in Fig. 6 provide information about uncertainty and sensitivity related to model predictions.

6. Conclusions

The utilization of concepts from agent-based modeling coupled with the efficiency of KD-tree data storage provide a novel approach to global UA. The efficiency of the approach can be tailored to the computational constraints of a problem by specifying the resolution of the search. ABAGUS does not produce formal posterior distributions of parameter probabilities consistent with Bayes' rule, instead focusing on identifying regions of the parameter space producing simulations acceptably consistent with observations. The performance of ABAGUS is evaluated on two test functions with known

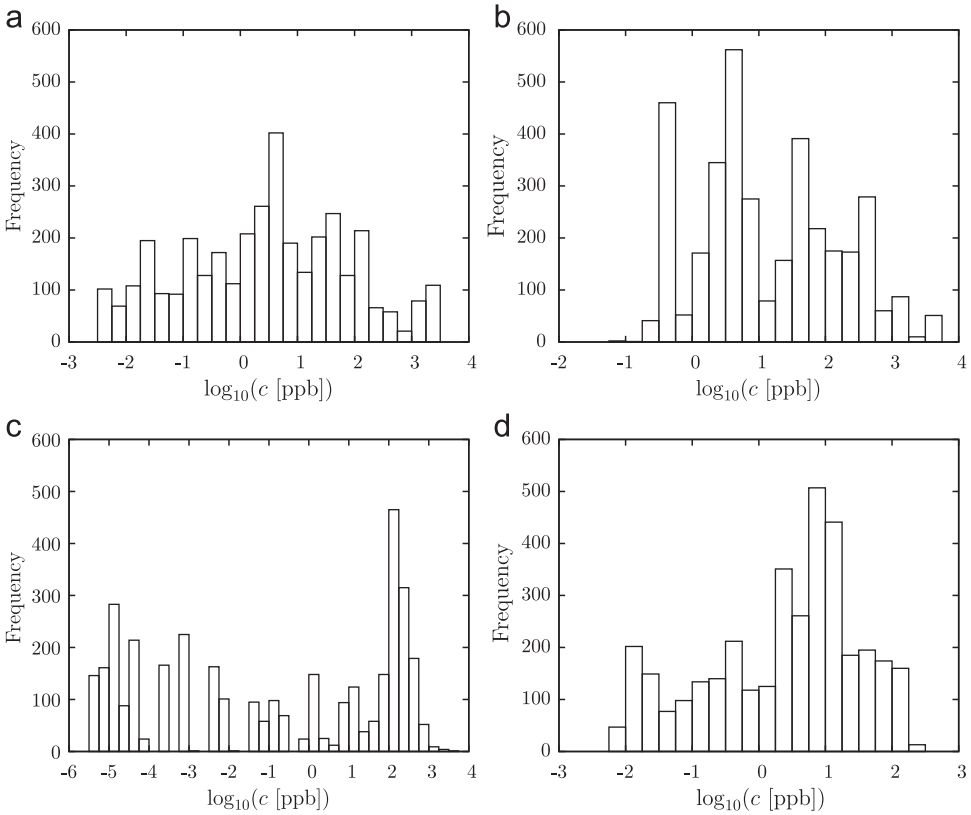

Fig. 6. Histograms of predicted concentrations at proposed monitoring well sites at $t=51$ y. Refer to Fig. 3 for proposal locations.

response surfaces, demonstrating the viability of the approach. The use of ABAGUS on a practical application is evaluated on a five-parameter synthetic contaminant transport case study, demonstrating the approach's ability to identify regions of the response surface producing simulations acceptably consistent with observations surrounding the "true" parameter values. ABAGUS provides a discretized global UA approach filling the gap between local UA approaches and rigorous sampling-based global UA approaches. ABAGUS will be an attractive alternative for complex problems where it is recognized that a local UA is inappropriate, but for which rigorous sampling-based global UA is infeasible due to computational constraints and inappropriate due to required assumptions. The ABAGUS algorithm is included in the MADS toolbox (Vesselinov and Harp, 2010).

Acknowledgments

This work was supported by various projects within the Environmental Programs Directorate of the Los Alamos National Laboratory. We thank the anonymous reviewers for providing insights and comments that improved the quality of the paper.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.cageo.2011.06.025.

References

Ben-Haim, Y., 2006. Info-Gap Decision Theory: Decisions Under Severe Uncertainty, second ed. Academic Press, Oxford.

Beven, K.J., Binley, A.M., 1992. The future of distributed models: model calibration and uncertainty prediction. Hydrological Processes 6, 279–298.

Clerc, M., 2006. Particle Swarm Optimization. ISTE, London.

Cooley, R.L., 1993. Exact Scheffé-type confidence intervals for output from groundwater flow models 1. Use of hydrogeologic information. Water Resources Research 29 (1), 17–33.

Cooren, Y., Clerc, M., Siarry, P., 2009. Performance evaluation of TRIBES, an adaptive particle swarm optimization algorithm. Swarm Intelligence 3, 149–178.

Dimri, V., 2000. Fractal dimension analysis of soil for flow studies. In: Dimri, V. (Ed.), Application of Fractals in Earth Sciences. Balkema, Brookfield VT, pp. 189–193.

Dorigo, M., Stützle, T., 2004. Ant Colony Optimization. MIT Press, Cambridge, MA.

Epstein, J.M., Axtell, R., 1996. Growing Artificial Societies: Social Science From the Bottom Up. MIT Press, Cambridge, MA.

Harp, D.R., Vesselinov, V.V. Analysis of hydrogeological structure uncertainty by estimation of hydrogeological acceptance probability of geostatistical models. Advances in Water Resources, in press. doi:10.1016/j.advwatres.2011.06.007.

Keating, E.H., Doherty, J., Vrugt, J.A., Kang, Q., 2010. Optimization and uncertainty assessment of strongly nonlinear groundwater models with high parameter dimensionality. Water Resources Research 46, W10517. doi:10.1029/2009WR008584.

Kennedy, J., Eberhart, R., 1995. Particle swarm optimization. In: Proceedings of the IEEE International Conference on Neural Networks. IEEE Press, Piscataway, pp. 1942–2948.

Li, B., Yeh, T.-C., 1998. Sensitivity and moment analyses of head in variably saturated regimes. Advances in Water Resources 21, 477–485.

Liu, Y., Freer, J., Beven, K., Matgen, P., 2009. Towards a limits of acceptability approach to the calibration of hydrological models: extending observation error. Journal of Hydrology 367, 93–103.

Neuman, S.P., 1980. A statistical approach to the inverse problem of aquifer hydrology, 3. Improved solution method and added perspective. Water Resources Research 16 (2), 331–346.

Neuman, S.P., 1990. Universal scaling of hydraulic conductivities and dispersivities in geologic media. Water Resources Research 26 (8), 1749–1758.

Nolte, D., Pyrak-Nolte, L., Cook, N., 1989. The fractal geometry of flow paths in natural fractures in rock and the approach to percolation. In: Scholz, C.H., Mandelbrot, B.B. (Eds.), Fractals in Geophysics. Birkhäuser, Basel, pp. 111–138.

Particle Swarm Central, 2006. 〈http://www.particleswarm.info/Standard_PSO_2006.c〉.

Sobol, I., 2001. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. Mathematics and Computers in Simulation 55, 271–280.

Sykes, J., Wilson, J., Andrews, R., 1985. Sensitivity analysis for steady state groundwater flow using adjoint operators. Water Resources Research 21 (3), 359–371.

Tonkin, M., Doherty, J., 2009. Calibration-constrained monte carlo analysis of highly parameterized models using subspace techniques. Water Resources Research 45, W00B10. doi:10.1029/2007WR006678.

Tsiombikas, J., 2009. kdtree: A simple C library for working with KD-trees. 〈http://code.google.com/p/kdtree/〉.

van Werkhoven, K., Wagener, T., Reed, P., Tang, Y., 2008. Characterization of watershed model behavior across a hydroclimatic gradient. Water Resources Research 44, W01429. doi:10.1029/2007WR006271.

Vecchia, A.V., Cooley, R.L., 1987. Simultaneous confidence and prediction intervals for nonlinear regression models with application to a groundwater flow model. Water Resources Research 23 (7), 1237–1250.

Vesselinov, V.V., Harp, D.R., 2010. MADS, modeling analysis and decision support toolkit in C. 〈http://www.ees.lanl.gov/staff/monty/codes/mads.html〉 (accessed on 15 November 2010).

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. doi:10.1029/2007WR006720.

Wagener, T., van Werkhoven, K., Reed, P., Tang, Y., 2009. Multiobjective sensitivity analysis to understand the information content in streamflow observations for distributed watershed modeling. Water Resources Research 45.

Zolotarev, V., 1986. One-dimensional Stable Distributions, vol. 65. American Mathematical Society.