ITERATIVE IMPORTANCE SAMPLING ALGORITHMS FOR PARAMETER ESTIMATION*

MATTHIAS MORZFELD[†], MARCUS S. DAY[‡], RAY W. GROUT[§], GEORGE SHU HENG
PAU[‡], STEFAN A. FINSTERLE[‡], AND JOHN B. BELL[‡]

Abstract. In parameter estimation problems one computes a posterior distribution over uncertain parameters defined jointly by a prior distribution, a model, and noisy data. Markov chain Monte Carlo (MCMC) is often used for the numerical solution of such problems. An alternative to MCMC is importance sampling, which can exhibit near perfect scaling with the number of cores on high performance computing systems because samples are drawn independently. However, finding a suitable proposal distribution is a challenging task. Several sampling algorithms have been proposed over the past years that take an iterative approach to constructing a proposal distribution. We investigate the applicability of such algorithms by applying them to two realistic and challenging test problems, one in subsurface flow, and one in combustion modeling. More specifically, we implement importance sampling algorithms that iterate over the mean and covariance matrix of Gaussian or multivariate t -proposal distributions. Our implementation leverages massively parallel computers, and we present strategies to initialize the iterations using “coarse” MCMC runs or Gaussian mixture models.

Key words. importance sampling, parameter estimation, Bayesian inverse problem

AMS subject classifications. 62F15, 65C05, 65Y05

DOI. 10.1137/16M1088417

1. Introduction. Predicting the behavior of complex physical systems is a key requirement in science and engineering. One approach to predicting the behavior of such systems is through high fidelity simulation. However, for many systems, model uncertainties limit predictive capability. The idea in parameter estimation is to use experimental data to reduce model uncertainties and, therefore, improve predictions of overall system behavior as follows. One represents uncertainties in the model by assuming that model parameters are random variables with given “prior” distributions, and one represents mismatch between model and data by a “likelihood.” By Bayes’ rule, the prior and likelihood define a posterior distribution that describes how well the parameters describing the system can be determined from the available data—which aspects are tightly bounded and which are less precise. The posterior distribution is thus regarded as the solution of a parameter estimation problem [16, 86, 87].

The central theme of this paper is to present a Monte Carlo sampling method that is efficient for parameter estimation problems for complex systems and well suited to

*Submitted to the journal’s Computational Methods in Science and Engineering section August 8, 2016; accepted for publication (in revised form) November 13, 2017; published electronically March 1, 2018.

<http://www.siam.org/journals/sisc/40-2/M108841.html>

Funding: This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics Program under contract DE-AC02-05CH11231, by the National Science Foundation under grant DMS-1619630, and by the Alfred P. Sloan Foundation through a Sloan Research Fellowship awarded to MM. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under contract DE-AC02-05CH11231.

[†]University of Arizona, Tucson, AZ 85721 (mmo@math.arizona.edu).

[‡]Lawrence Berkeley National Laboratory, Berkeley, CA 94720 (msday@lbl.gov, gpau@lbl.gov, safinsterle@lbl.gov, jbbell@lbl.gov).

[§]National Renewable Energy Laboratory, Golden, CO 80401 (Ray.Grout@nrel.gov).

extreme-scale (parallel) computer architectures. We first formulate the Bayesian parameter estimation problem and review the literature about Monte Carlo approaches to its numerical solution (section 2). We then motivate our use of iterative importance sampling algorithms (ISA) and present its implementation on high performance computing systems, as well as effective strategies to initialize the iterations (section 3). ISA has been discussed in the literature extensively over the past decades (see, e.g., [15, 66, 67]), but it has not yet been applied to “realistic” test problems with complex systems. Our main goal in this work is to demonstrate that ISA is indeed applicable in such situations by applying our ISA implementation to two test cases: estimation of permeabilities in a subsurface flow problem (section 4) and estimation of kinetic parameters for combustion simulation (section 5).

2. Background and problem formulation. We consider a set of parameters θ of a numerical model $\mathcal{M}(\theta)$. Here θ is an n_θ -dimensional vector, and the model $\mathcal{M} : R^{n_\theta} \rightarrow R^{n_z}$ maps the parameter vector to an n_z -dimensional vector of outputs z . In the examples discussed below, θ represents collection of reaction rate parameters and third-body coefficients that define a combustion model, or θ defines a permeability field of a subsurface flow problem (see sections 4 and 5 below). Uncertainty in the model is represented by uncertainty in the parameters, i.e., θ is a random vector with prior probability distribution p_θ . The outputs of the model can be measured, and additional uncertainty in the approximation of the measurements are represented by a random variable v with given distribution p_v . We assume that this uncertainty is additive,

$$(1) \quad z = \mathcal{M}(\theta) + v,$$

but this assumption is not central to our approach. The above equation defines the likelihood $p(z|\theta) = p_v(z - \mathcal{M}(\theta))$. By Bayes’ rule, prior and likelihood define the posterior distribution

$$p(\theta|z) \propto p_\theta(\theta)p(z|\theta),$$

which describes the probability of the parameters θ given the data z . The posterior distribution allows for a complete description of the uncertainty in the parameters given the data. In particular, one can compute least-squares optimal estimates by computing the posterior mean as well as “error bars” based on posterior covariances.

If the prior and the errors v are Gaussian, and if, in addition, the model $\mathcal{M}$ is linear, then the posterior is also Gaussian. In this case, it is sufficient to compute its mean and covariance. This can be done, for example, by minimizing the negative logarithm of the posterior distribution

$$F(\theta) = -\log p_\theta(\theta) - \log p(z|\theta) + C.$$

The minimizer of F is the posterior mean, and the Hessian, $H_{ij} = \partial^2 F / \partial \theta_i \partial \theta_j$, $i, j = 1, \dots, n_\theta$ is the inverse of the posterior covariance matrix [87]. Linear algorithms have been developed that are highly efficient for large-scale problems of this type [13, 32].

When the model is nonlinear, the posterior distribution is no longer Gaussian, even if the prior distribution for the parameters and the model errors, v , remain Gaussian. In this case, minimizing the negative log-posterior $F(\theta)$ results in the most likely state given the data, which is called “posterior mode.” However, there may be

more than one posterior mode (local maxima of the posterior distribution), in which case it is difficult to argue that a (local) posterior mode is a meaningful solution of the parameter estimation problem.

2.1. Markov chain Monte Carlo for inverse problems. Markov chain Monte Carlo (MCMC) methods are routinely used for solving nonlinear parameter estimation problems. The basic idea is to choose a specific proposal distribution to propose a “move,” i.e., a new set of parameters. This move is accepted or rejected based on a suitable accept/reject criterion [47]. The chain then consists of a number of correlated samples and converges in the sense that averages over the samples converge to expected values with respect to the posterior distribution as the ensemble size goes to infinity. Thus, MCMC provides a more complete description of the uncertainty than computing modes of the posterior distribution. In practice, the beginning of an MCMC chain is often regarded as a “burn-in” period, and discarded in the subsequent analysis. In addition, if the correlation “time” of the chain is too large, then prohibitively long chains are needed to effectively sample the posterior. The various MCMC algorithms in the literature differ in their proposal distributions and approaches to reduce burn-in and auto-correlation time.

There are several “classical” MCMC algorithms, e.g., Metropolis–Hastings [47], and many of them are known to be slow in high-dimensional and highly nonlinear problems; see, e.g., [40, 44]. Recent MCMC methods that address these problems include affine invariant ensemble samplers [33, 40], adaptive MCMC [42, 43], and differential evolution MCMC [9, 10]. Another class of MCMC algorithms relies on using local geometry of the posterior to construct proposals. Algorithms of this type include the Metropolis adjusted Langevin algorithm [77], Hamiltonian MCMC [31, 64], and Riemann manifold MCMC [35]. The stochastic Newton MCMC algorithm [56, 63] also falls within this class. In stochastic Newton, one first finds the posterior mode by minimizing the negative log-posterior, $\min_{\theta} F(\theta)$, and then starts an MCMC chain using low-rank approximations of Hessians, evaluated along the chain. Adjoints are used to speed up gradient and Hessian computations. Another related method is “randomize-then-optimize” [4], where a stochastic version of the negative log-posterior optimization problem is solved (repeatedly) to generate samples.

In many situations, the model $\mathcal{M}$ is a discretization of a partial differential equation and the unknown parameter is a discretization of a field. There is a large literature addressing the fact that MCMC methods for such problems should have a well-defined limit as the mesh is refined; see, e.g., [7, 23, 79, 86]. Moreover, indicators of computational requirements (burn-in or auto-correlation times) of such “function space” MCMC algorithms should not increase as the dimension increases when the mesh is refined. In fact these properties should be invariant under mesh-refinement, once grid convergence is reached. For that reason, function space MCMC methods are also often called “mesh-independent” or “dimension-independent.”

Computational requirements of MCMC can be reduced by solving a related, but easier, problem. This idea is implemented by reduced-order modeling: rather than using the model $\mathcal{M}$ for MCMC, one replaces the model by a simpler, “reduced-order” version of it. Since the reduced-order model (ROM) is simpler, it is computationally less expensive, and using it therefore speeds up MCMC. In this context, one can also reduce the dimension of the parameter vector θ , the dimension of the state of the underlying numerical model, $\mathcal{M}$, or both; see, e.g., [24, 25, 52, 84]. Reducing the state dimension is usually achieved by finding a simplified approximation to the mathematical model $\mathcal{M}$. This reduced model however may depend on the same

number of parameters as the “full” model. To reduce this number of parameters one can determine a low-dimensional subspace of parameter combinations which are constrained by data, as in the likelihood-informed methods [24, 25]. In problems where the parameters are spatial fields, e.g., permeability fields in subsurface flow problems, one can use principal component analysis/Karhunen–Loeve expansions [61] to define this field in terms of a small number of dominant modes. (See also section 4, where Kriging is used to reduce the parameter dimension.) In either case, MCMC by reduced-order models requires that the ROM does not introduce large errors, or else the posterior distribution defined by the ROM may be very different from the posterior distribution defined by $\mathcal{M}$. Such ideas were made precise in the context of polynomial chaos expansions in [54]. To reduce error due to model reduction one can introduce a second stage in each Metropolis step as in [21], or one can continually refine the ROM during MCMC sampling as in [22]. Nonetheless, ROM approaches are only effective if the model $\mathcal{M}$ indeed has a low-dimensional representation that can be discovered by ROM techniques.

Another approach to reducing computational cost of MCMC involves the use of optimal transport maps; see, e.g., [62, 68, 75, 89]. The idea is to construct a transport map that converts prior samples to posterior samples. More generally, the map should transform a reference distribution to the posterior distribution [20]. Construction of an optimal map requires solving an infinite-dimensional optimization problem, which is expensive, but approximations to an optimal map may be more readily computable and can be used to “speed up” MCMC. Work in this direction is underway [70, 71].

2.2. Importance sampling for inverse problems. An alternative to MCMC is importance sampling. Importance sampling, just as MCMC, requires a proposal distribution q . But rather than using the proposal distribution for defining a Markov chain, one draws samples from it and attaches to each sample a weight

$$w(\theta) \propto \frac{p(\theta|z)}{q(\theta)}.$$

The weighted samples form an empirical estimate of the posterior distribution in the sense that weighted averages over the samples converge to expected values as the ensemble size goes to infinity; see, e.g., [16, 69]. Moreover, the posterior distribution and the proposal distribution need only be known up to a multiplicative constant (as in MCMC), because the weights can be self-normalized such that their sum is one, i.e., $w_i \leftarrow w_i / \sum_{j=1}^{N_e} w_j$. Advantages of importance sampling compared to MCMC are that there is no burn-in time, and the weighted samples are independent. The latter means in particular that the samples and weights can be computed independently, which can lead to enormous computational advantages on massively parallel computer architectures.

However, the efficiency of importance sampling hinges on constructing a suitable proposal distribution q . The reason is that, even though all samples are independent, each sample carries a different weight and, therefore, the samples are not equally important. For an effective importance sampling method all samples should be similarly important, i.e., the variance of the weights should be small. Indeed, a heuristic “effective sample size” can be defined [2, 6, 30, 39, 53, 90] based on the relative variance $\text{var}(w)/E(w)^2$:

$$(2) \quad N_{\text{eff}} = \frac{N}{R}, \quad R = \frac{\text{var}(w)}{E(w)^2} + 1 = \frac{E(w^2)}{E(w)^2}.$$

FIG. 1. Illustration of iterative construction of a proposal distribution.

The effective sample size describes the size of an unweighted ensemble of size N_{eff} that is equivalent to a weighted ensemble of size N . As an illustration, suppose that one weight is close to one, which implies that all other self-normalized weights are close to zero. In this case, $R \approx N$, so that the effective sample size is one. If this happens, the sampling method is said to have “collapsed” [5]. In contrast, if the proposal distribution is equal to the posterior distribution, the variance of the weights is zero, so that R is unity. For effective sampling, one needs an R that is “not too large,” where what large is depends mainly on how much computing power one is willing to spend.

The collapse of sequential Monte Carlo methods for time-dependent problems, called particle filters [30] has been studied extensively in the meteorological literature, and it was shown that, for a certain class of problems, $N_{\text{eff}} \rightarrow 1$ as the “effective dimension” of the problem increases [5, 8, 17, 20, 81, 82, 83]. The effective dimension can depend on the dimension of the data vector z , the dimension of the parameter vector θ , as well as the distributions of errors in the prior and likelihood and may increase when the dimension of the parameter vector θ increases, however, at a slower rate. Moreover, recent work suggests that particle filters may be efficient even if the effective dimension is large, provided that the filtering problem has a certain “sparse” structure [59, 74].

Perhaps due to these numerical difficulties, applications of importance sampling are rare, but importance sampling schemes have been used for some time-dependent “filtering problems” (see, e.g., [2, 3, 18, 19, 60, 89, 90]) and a few (simple) parameter estimation test problems (see, e.g., [7, 61]).

2.3. Iterative importance sampling. The basic idea in ISA is to iteratively improve the proposal distribution. The iteration starts with a given proposal distribution q^0 , which is used to draw N_e^0 weighted samples $\{\theta_j^1, w_j^1\}$, $j = 1, \dots, N_e^0$. These samples are then used to define a new proposal distribution q^1 , and the process is repeated. The rationale for an iterative approach is to discover, via iteration, the global structure of the posterior distribution without becoming stuck in local minima. A successful iteration is illustrated in Figure 1.

ISA are discussed extensively in the literature and are sometimes called “adaptive importance sampling” or “population Monte Carlo.” The earliest account we could find is the applications paper by Buchner [12], where iterative sampling is used to compute failure probabilities in structural dynamics. The proposal distribution at each step is Gaussian and the mean and covariance of the Gaussian are updated at each step of the iteration. The paper also presents initialization strategies for the iteration that work well for the application. A more general version of ISA is described in [66, 67]. Several stopping rules are presented, and consistency of the method is established, in the sense that expected values computed by ISA are asymptotically

equal to expected values computed with respect to the posterior distribution. A method closely related to ISA is “population Monte Carlo” [15, 46], or “nonparametric importance sampling” [91]. In these methods, the ensemble at each step undergoes Kernel smoothing, i.e., the proposal distribution at iteration k is

$$q^k = \sum_{j=1}^{N_e} \alpha_j^k K_j(\theta, \theta_j^k),$$

where the sum of the α_j is one. Here, one can adapt the weights of the mixture α_j^k [29], or adapt the weights and the Kernels [14, 36] in each iteration. A recent application of population Monte Carlo in geosciences is presented in [85]. The scaling of the computational requirements for one step of population Monte Carlo with respect to the dimension of the parameter vector is discussed in a simplified setting in [49]. (See also our conclusions in section 6.) We wish to point out that the adaptive independence samplers, discussed in [34, 45, 48], use ideas similar to ISA in the context of MCMC/rejection sampling. The “cross-entropy method” also adapts proposal distributions, but in a more sophisticated way [78]. ISA also has connections with sequential Monte Carlo (see [58, section 3.5]).

3. Implementation and initialization of ISA. ISA has been discussed in the literature extensively over the past decades (see above). In this paper we focus on strategies for dealing with the pragmatic aspects of applying ISA to “realistic” problems in science and engineering and discussing the implementation of the algorithms, rather than devising fundamentally new algorithmic ideas. We are targeting posterior distributions that are difficult to deal with by the MCMC/importance sampling technology reviewed above. More specifically, we wish to create ISA that are efficient under the following conditions:

1. The priors are “broad,” i.e., using the prior distribution as proposal distribution leads to the collapse of importance sampling.
2. We can compute first derivatives of the negative log-posterior distribution $F(\theta)$, but not second derivatives. This situation is often encountered in complex computational models, e.g., if adjoints are available, but second-order adjoints are not.
3. The negative log-posterior distribution $F(\theta)$ is not convex and may have multiple minima that are not well separated (see also Figure 1).
4. The effective dimension n_θ of the problem is moderate, perhaps tens, rather than hundreds or even thousands, and the model $\mathcal{M}$ can be evaluated with modest computing resources.

As we will demonstrate, ISA can be effective under these conditions, but other, non-iterative, importance sampling methods or MCMC may be slow to converge.

3.1. The basic algorithm. Our basic ISA, summarized in pseudocode in Algorithm 1, is as follows. We first decide on a class of proposal distributions. In principle, one could choose different classes of proposal distributions at each step in the iteration, e.g., a log-normal distribution during the first step, and a uniform distribution at the next step, and so on. However, here we consider only the case where the class of proposal distributions is fixed to either Gaussians or multivariate t -distributions. The reason is that both distributions are defined by a mean and covariance which we can easily compute from the samples. The utility of multivariate t -proposal distributions is discussed in a more general setting in [69].

Algorithm 1 Iterative importance sampling algorithm

Given: an unweighted ensemble $\{\theta_j^0\}$, $j = 1, \dots, N_e$
 Construct a proposal distribution q^0 from the ensemble.
 Set $k = 0$, and define a tolerance tol
while no convergence reached **do**
 Importance sampling with proposal distribution q^k to get $\{\theta_j^{k+1}, w_j^{k+1}\}$, $j = 1, \dots, N_e^k$
 Compute R^{k+1} by (2) for this set of weighted samples
 if $R^{k+1}/R^k < \text{tol}$ **then**
 Convergence is reached
 else
 Construct proposal distribution q^{k+1} from $\{\theta_j^{k+1}, w_j^{k+1}\}$, $j = 1, \dots, N_e^k$ (see text)
 end if
 Set $k \leftarrow k + 1$
end while

Given a set of initial samples, which we obtain using the initialization strategies discussed below, we compute the sample mean and covariance, which define the initial proposal distribution q^0 . We then generate N_e^0 weighted samples $\{\theta_j^1, w_j^1\}$, $j = 1, \dots, N_e^0$, by importance sampling with the proposal distribution q^0 , and compute the (weighted) sample mean and covariance to define a new proposal distribution q^1 , and repeat. During the iteration, we monitor the quality measure R in (2) and stop the iteration when $|R_{k+1} - R_k|$ drops below a specified tolerance. This stopping criterium is similar to the one discussed in [66, 67].

Note that we cannot guarantee that R decreases monotonically during the iteration when N_e^k is finite, in part because we must approximate R in (2) by averaging; thus, we only know R up to sampling error, which can be substantial if R is large and N_e^k is not. Indeed, a large R is always uncertain, since an approximation of R by N_e samples is always less than N_e ; however, the “true” R may be much larger. Since the iteration is stopped when the change in R from one iteration to the next is small, one must consider sampling error when choosing the tolerances.

3.2. Sample sizes and convergence of ISA. One must pick an ensemble size N_e^k at each step of the distribution. We advocate choosing this number as large as possible for two reasons. The first reason is computational. Since the iteration of ISA is serial, we wish to use ISA with a small number of iterations and large N_e at each iteration to leverage parallelism. However, the assumption that a large number of samples can be generated “easily” at each iteration restricts the classes of problems we can solve with ISA. For example, if a supercomputer is needed to evaluate the model $\mathcal{M}$ once, then ISA as described here is not feasible. Perhaps it is fair to state that ISA may enable parameter estimation using massively parallel computers if the forward problem, i.e., evaluation of the model $\mathcal{M}$, can be done using a small fraction (order of 0.1%) of the available resources.

Our second reason for choosing large N_e^k is motivated by the convergence of ISA. Suppose the support of the proposal distribution at iteration k , q^k includes the support of the posterior distribution (which is a routine assumption). Then weighted samples from the proposal distribution q^k generate an “empirical estimate” of the proposal distribution that converges in the sense that weighted averages converge to expected

values with respect to the posterior distribution as $N_e^k \rightarrow \infty$. This follows from the usual importance sampling theory (see, e.g., [16]), even if the proposal distribution depends on past samples (see also [15]). The next iteration is guaranteed to converge in the same sense, and under the same routine assumptions. Thus, ISA converges within one step, provided one generates sufficiently many samples, i.e., if N_e^k is large enough.

The importance of a large N_e^k at each step can be illustrated further under simplifying assumptions. Suppose that the posterior distribution is Gaussian with variance Σ and that the proposal distribution at iteration $k+1$ is also Gaussian. We neglect error in the mean, i.e., we assume that the mean of the proposal distribution is exactly equal to the mean of the posterior distribution, which can be justified because small errors in the mean are less severe than small errors in the covariances [69]. We further assume that the covariance matrix, computed from the weighted samples from the k th iteration, is $\hat{\Sigma} = (1 + \varepsilon)\Sigma$, where ε is a small number. Under these assumptions, the weights at iteration $k+1$ are

$$w^{k+1} \propto \exp\left(-\varepsilon \frac{1}{2} x^t \Sigma^{-1} x\right), \quad x = \sim \mathcal{N}(0, \hat{\Sigma}),$$

which leads to

$$R^{k+1} = \left(\frac{1 + \varepsilon}{\sqrt{1 + 2\varepsilon}}\right)^{n_\theta} \approx 1 + \varepsilon^2 \cdot \left(\frac{n_\theta}{2}\right),$$

where n_θ is the number of parameters we wish to estimate, i.e., the dimension of θ . We can interpret ε as “sampling” error, and assume ε is inversely proportional to the square root of the effective sample size N_{eff} in (2). Thus,

$$R^{k+1} \approx 1 + R^k \frac{C \cdot n_\theta}{2N_e^k},$$

where C is a constant. The above equation indicates that the efficiency of the iteration is directly controlled by how many samples N_e^k one can generate at each step. However, the expansions above are only valid if ε is “small,” which means that sampling error is small and the algorithm is already performing well with small R^k . In fact most theoretical considerations rely on large ensemble sizes N_e^k , but in practice the behavior of the algorithm when the ensemble size is *small* is more interesting. We study these practical issues in the context of examples (see sections 4 and 5 below).

3.3. Initialization. Initialization of ISA has a significant impact on the success of the algorithm. Indeed, the iteration cannot proceed if the initial proposal distribution q^0 yields only one effective sample, so that the variance that defines q^1 is zero, i.e., the algorithm has “collapsed” during the first step. We suggest two strategies for initializing ISA.

The first initialization strategy is to run MCMC. However, if the MCMC has been run long enough to generate sufficiently many samples to characterize the posterior, then the problem has already been solved and refinement of a proposal distribution by iteration is unnecessary. We thus consider “short” MCMC runs, and in the examples and applications below this yielded good results. The initialization strategy combines the robustness of MCMC with the computational advantages of importance sampling. In this context, one can think of ISA as a means to “fill out” an underresolved MCMC computation, while making effective use of massively parallel machines. This can yield dramatic reductions in run times, from weeks or months of serial MCMC on a workstation, to minutes on an HPC system.

The second initialization strategy we propose is to construct a Gaussian mixture model (GMM), for which each mixture component is centered at a mode of the posterior distribution. The various modes can be found in parallel, by initializing a suitable optimization code at various locations in parameter space obtained by, e.g., sampling the prior distribution. The covariances of the mixture are inverses of the approximate Hessians of the negative log posterior distribution $F(\theta)$ at the modes. The weights of the mixture components are

$$(3) \quad \psi_j = \frac{\exp(-\phi_j)}{\sum_{i=1}^n \exp(-\phi_i)},$$

where n is the number of minima we found and ϕ_i , $i = 1, \dots, n$ are the local minima of F . Alternatively, one could also use weights based on normalizing constants of the Gaussians, but we do not pursue this idea further.

One may find that the optimizer, initialized at two different points in parameter space, finds the same minimum. We distinguish minima by considering the distance

$$d_{i,j} = (\mu_i - \mu_j)^T H_i^{-1} (\mu_i - \mu_j),$$

where μ_k are the various minima we find, and where H_k are the corresponding approximate Hessians with elements $(H_k)_{ij} = \partial^2 F / \partial \theta_i \partial \theta_j |_{\mu_k}$, $i, j = 1, \dots, n$. We say that a minimum of F at μ_i is different from the one at μ_j if $d_{i,j}$ is more than an application-specific threshold. Other, more sophisticated assessments, may be needed in other applications to account for asymmetries, since

$$(\mu_i - \mu_j)^t H_i^{-1} (\mu_i - \mu_j) \neq (\mu_i - \mu_j)^t H_j^{-1} (\mu_i - \mu_j).$$

However, in the applications we consider below, our simple strategy is sufficient.

3.4. Illustration by a two-dimensional toy problem. We illustrate ISA by a simple “toy” problem, where we estimate a two-dimensional parameter vector θ . The prior is uniform on the cube $[0, 11] \times [0, 11]$, and the likelihood is such that the posterior is given by

$$p(\theta|z) \propto \exp(-F(\theta)), \quad \text{where } F(\theta) = 10^{-2} \|\theta - [5; 5]\|^4 + 0.2 \sin(5\|\theta\|),$$

and where $\|x\| = \sqrt{x^T x}$ is the 2-norm. The posterior distribution is illustrated in the bottom right panel of Figure 2. Note that there are several modes that carry significant probability mass, but the modes are not well separated.

We apply ISA with Gaussian proposal distributions to obtain samples from this posterior distribution. We initialize the iteration by samples we obtain by the emcee implementation [33] of the affine invariant MCMC ensemble sampler [40]. We use four walkers, initialized by drawing samples from the uniform prior distribution. The first 20 samples are used to initialize ISA. A “triangle plot” of these 20 samples, which shows histograms of all one-dimensional marginals and the two-dimensional PDF, of a given set of samples, is shown in the upper left panel of Figure 2. We perform 5 iterations, each with $N_e = 20,000$ samples. The quality measure R converges quickly to about $R \approx 1.1$, by the sequence $\{R^k\} = \{658, 11.2, 1.13, 1.10\}$. A triangle plot from the Gaussian proposal we obtained after five iterations is shown in the upper right panel of Figure 2. A triangle plot of 20,000 weighted samples of this Gaussian proposal is shown on the lower left panel of the figure, and one can clearly identify multiple modes in the posterior distribution. Moreover, the efficiency of importance sampling

FIG. 2. Illustration of ISA initialized by a short MCMC run. Top row, left: triangle plot of the first 20 MCMC samples. Top row, right: triangle plot of the Gaussian proposal distribution obtained after 5 iterations. Bottom row, left: triangle plot of the posterior distribution obtained by ISA. Bottom row, right: the posterior distribution, bird view.

is almost perfect in this example: since $R \approx 1$, almost all samples are effective samples. Emcee hammer on the other hand is characterized by an integrated auto-correlation time of about 39, which means that only one in 39 samples is effective. We estimated the integrated auto-correlation time from an emcee run with four walkers, each taking 100,000 steps.

The convergence of ISA is further illustrated in Figure 3, where we show histograms of one-dimensional marginals of the proposal distributions q^0, q^1, q^2, q^3 , obtained during the iteration, as well as a histogram of the corresponding marginal of the posterior distribution. We observe that the high probability region of the initial proposal distribution is far from that of the posterior distribution. However, ISA quickly corrects this during iteration and moves the Gaussian proposal distributions closer to the posterior distribution.

We also tested ISA initialized by a GMM, which we obtain by performing 100 minimizations of F , each initialized by a sample from the uniform prior density. We then compare the minima by the distance $d_{i,j}$ as explained above and find five distinct minima of F . At each distinct minimum, we approximate the Hessian by finite differences. (We chose computationally more efficient methods to approximate Hessians in the applications below.) The distinct minimizers and corresponding Hessians define the GMM with five mixture components, where the weights of each component is given by (3). We use this GMM as the proposal distribution and compute 20,000 samples to obtain $R \approx 2.59$. Using the GMM mixture distribution is thus less effective than using a few iterations of ISA as above. Nonetheless, we can use a few samples of the GMM to initialize ISA. Here we use the 50 samples from the GMM to initialize ISA and then iterate using Gaussian proposal distributions with 20,000 samples at each step. The Gaussian proposal distribution we obtain after three iterations is characterized by a quality measure $R = 1.10$. (The sequence is $R = \{3.27, 1.1, 1.1\}$.) The

FIG. 3. Convergence of the proposal distributions of ISA.

experiments indicate the iterations of ISA yield similar proposal distributions after a few iterations for both initialization strategies.

4. Application 1: Subsurface flow.

4.1. Background and problem set-up. Subsurface flow models are extensively used in hydrological and geosciences applications, such as the prediction of contaminant fate during a chemical spill, and the estimation of oil production from an oil recovery process. Robust predictions of quantities of interest require accurate characterization of the model parameters, such as the material properties and process-related parameters. Difficulties in direct measurements of these parameters necessitate the use of inverse modeling tools to infer them based on sparse measurements of observables that can be modeled.

In this section, we use ISA in a synthetic subsurface flow problem to estimate the permeability distribution, and the associated uncertainty, in the vadose zone from sparse measurements of saturations and changes in water volume in a domain that covers a two-dimensional $4\text{m} \times 3\text{m}$ (width $\times$ depth) spatial domain and is discretized by a mesh with uniform grid cells of size $0.05 \times 0.05\text{m}$. Water is released from a pond into a heterogeneous vadose zone that has a water table at depth 3m. The saturation of the heterogeneous soil is initially in capillary-gravity equilibrium. A one-day ponded infiltration period is followed by a one-day water redistribution period. During infiltration, the water level in the pond is maintained at 2cm, and saturations at 36 locations, shown in the left panel of Figure 4, are measured every two hours (25 snapshots) during the two days of infiltration and redistribution. The total volume of water flowing out of the pond is also measured at the same frequency, yielding a total of 925 observations that we combine into a vector z as in (1).

Our goal is to determine the permeability distribution of the vadose zone based on the above measurements. We describe the permeability by the pilot-point method [38, 73] with fourteen pilot points, located in two columns near the pond, as shown in the left panel of Figure 4. The pilot points serve as adjustable conditioning points for determining the permeability modifier (κ_m) in each grid cell through a Kriging process that utilizes a spherical semivariogram model with fixed parameters [27]. The permeability in each grid cell is then given by $\kappa_{\text{ref}} \times 10^{\kappa_m}$, where $\kappa_{\text{ref}} = 0.316 \times 10^{-11} \text{ m}^2$ is a reference permeability value. The inverse problem we consider amounts to estimating the parameter vector $\theta = \{\kappa_{m,i}, 1 \leq i \leq 14\}$, where $\kappa_{m,i}$ are the κ_m values at the pilot point locations, based on two days worth of saturation and flow data.

FIG. 4. Left: computational domain and locations of the conditioning points for Kriging; the color map illustrates the distribution of the permeability modifier when the values at the pilot points are given by θ_{ref} . Right: quality measure R during ISA iteration.

We assume a broad Gaussian prior with zero mean and a diagonal covariance matrix C_{pp} where all 14 diagonal elements are equal to nine. The likelihood is defined by (1), where the model $\mathcal{M}$ is based on Richards' equation [76]. Numerically, we solve this equation using the finite-volume simulator TOUGH2 [72] with relative permeability and capillary pressure functions as in [88]. The errors v in (1) are Gaussian with mean zero and a diagonal covariance C_{zz} . The standard deviations that define the covariance are assumed to be 0.01 for the saturation measurements and 0.005 m^3 for the water volume measurements. We perform experiments with synthetic data, i.e., we assume a reference permeability field defined by

$$\theta_{\text{ref}} = \{1.5, .5, 1., 1.5, -1.5, 1.5, 0.5, -0.5, -0.5, -1.5, -0.001, 0.5, -1.5, -0.5\},$$

and then perform a simulation of the above infiltration. The reference set of parameters defines the permeability modifier field shown in the left panel of Figure 4. The synthetic observation vector z is then given by $\mathcal{M}(\theta_{\text{ref}}) + \epsilon$, where ϵ is an instance of v in (1).

4.2. Results. We initialize ISA by a single Gaussian. We define this Gaussian by solving the nonlinear least squares problem $\min_{\theta} F(\theta)$, where F is the negative logarithm of the posterior (up to an additive constant)

$$(4) \quad F(\theta) = \frac{1}{2}(z - \mathcal{M}(\theta))^T C_{zz}^{-1}(z - \mathcal{M}(\theta)) + \frac{1}{2}\theta^T C_{pp}^{-1}\theta.$$

Here $\mathcal{M}(\theta)$ is the output of the TOUGH2 simulator. The minimizer of F defines the mean, and the approximate Hessian, evaluated at the minimum, defines the inverse of the covariance matrix. We approximate the Hessian by neglecting contributions from second derivatives, i.e., we set the covariance equal to the inverse of $2J^T J$, where J is the Jacobian matrix of the nonlinear least squares problem. Numerically, we rely on Levenberg–Marquardt [50, 55] (the lsqnonlin function in MATLAB) to solve the optimization problem and use finite differences for computing the first derivatives. For this particular problem, 20 iterations are needed to find a minimum. We note that the evaluations of the negative log-posterior (4) and of the Jacobian J in each iteration can be easily parallelized.

We start ISA using the Gaussian proposal distribution and generate 5,120 samples at each iteration. The likelihoods and weights of the samples are evaluated in parallel

FIG. 5. Triangle plot of posterior distribution obtained by ISA after nine iterations, each with 5,120 samples. The plot shows histograms of the parameters $\{\theta_6, \theta_7, \theta_8, \theta_9, \theta_{10}\}$ (diagonal), as well as two-dimensional histograms of all pairs of these five parameters (lower triangle).

using 160 nodes, each with 32 cores, on Cori, a supercomputer at NERSC. With this set-up, the quality measure R converges from 163 to 1.3 in 9 iterations, as illustrated by the right panel of Figure 4.

We compare ISA to solving the same problem by the affine invariant ensemble sampler [40], as implemented in the emcee code [33]. We use a parallel implementation of the algorithm with 126 walkers, which allows us to use 2 nodes on Cori efficiently. We discard the first 2,500 steps as “burn-in” and take another 2,500 steps per walker. The computational cost thus amounts to a total of 630,000 likelihood evaluations, each requiring an evaluation of the TOUGH2 model that takes approximately 3 minutes. The integrated auto-correlation times for θ vary from 84 to 532, indicating that we have 3,750 to 592 “effective” samples. This should be compared to the quality measure of $R \approx 1.3$ after nine iterations of ISA. This means, in particular, that after nine iterations, almost every Gaussian sample contributes effectively to the numerical approximation of the posterior density. To get there, about 46,080 samples are used, and these can be evaluated in parallel in nine batches of 5,120 samples. Our results demonstrate that ISA can solve this parameter estimation problem effectively, while achieving greater level of parallelism than MCMC, which reduces wall-clock time. Including the cost of the initial minimization of (4), solving the problem by ISA takes about 90 minutes, while the emcee code requires about 20 days.

Note that in the context of ISA, a quality measure R close to 1 can indicate that the posterior is nearly Gaussian. The reason is that we use a Gaussian proposal distribution, and $R \approx 1$ implies that the posterior distribution is also nearly Gaussian. This is illustrated by the triangle plot in Figure 5. The plot shows histograms of the five parameters θ_6 – θ_{10} , as well as two-dimensional histograms of pairwise combinations of these five parameters. Histograms of those parameters that are not shown in the figure are qualitatively similar. The figure suggests that the variables are indeed nearly Gaussian. While we can conclude, after the fact, that other methods specifically designed for nearly Gaussian problems (see the literature review) may perform better on the present problem, the distributions of the parameters are not typically known a priori. On the other hand, the Gaussian proposal distribution whose mean is the posterior mode and whose covariance matrix is derived from the Gauss–Newton

solution if the optimization problem exhibits an R -value much larger than what we find after a few iterations of ISA.

5. Application 2: Combustion modeling.

5.1. Background. Predictive combustion modeling entails solving a set of fluid conservation equations that are supplemented with models for thermodynamics, transport, and chemical kinetics. For this study we assume an ideal gas mixture, a gradient-driven transport, and a set of reversible Arrhenius rate expressions for chemical reactions. Here, we focus on chemical combustion on hydrogen combustion. Compared to fuels involving complex hydrocarbons or biofuels, the hydrogen-oxygen case is composed of a relatively small number of chemical species and reactions. For that reason, and perhaps since hydrogen-oxygen represents an important building block for more complex fuels, the basic form of the numerical model for hydrogen-oxygen systems, including the set of specific fundamental chemical reactions, is well-established. The parameters that define the model, however, remain open to considerable debate, as indicated by the number of models under development (e.g., [26, 37, 51, 57, 65, 80]). Note that the various parameter values reported over the years for each line of model development result from a long history of incremental optimization efforts that incorporate emerging experimental data sets. As a result, while different parameter sets may lead to similar predictive capabilities within a common range of their respective validated conditions, the models they define can generate dramatically different predictions outside that range. (See the discussion, e.g., in [41], particularly related to Figure 1, and in [28], related to Figure 4.)

The numerical model used here consists of a systems of differential equations whose solution requires a database of parameters that characterize the relevant thermodynamic relationships, transport coefficients, and kinetic rate coefficients for the component gases. Our goal is to use ISA to incorporate measured experimental data into characterization of the uncertainty of a selected set of nine parameters that describe important behavior of this particular numerical model.

To a large extent, our study follows a recent effort to optimize 28 parameters of a syngas combustion model [26]. In that work, the model was tuned specifically to predict measured values of 36 experiments, including constant-pressure flow reactors, constant-volume shock-tubes, and steady, unstrained premixed flames. In the reference, a solution-mapping technique was employed to approximate the response of the model using a second-order polynomial. For each experiment, an a priori sensitivity analysis was used to reduce the set of active parameters in the response surface model. An objective function was constructed and minimized based on the square of model prediction errors, normalized by the uncertainty of the respective measurements. Prior knowledge of the parameters was used only to enforce bounds on the variation of each active parameter, and to inform the construction of the underlying polynomial response surface model. Implicitly, this work also incorporated prior knowledge in the choice of what reactions to include in the kinetic model and the choice of which model parameters to hold fixed during the optimization. In a Bayesian framework, the referenced optimization approach can be interpreted loosely as one based on a uniform prior over a cube whose sides are defined by prior-informed parameter bounds. In the present work, discussed below, we estimate a subset of the parameters considered in [26] and quantify the respective uncertainties using ISA, based on the complete numerical combustion model, rather than solution-mapping. However, we make use of many of the same target experiments in order to quantify our likelihood estimate.

TABLE 1

Arrhenius rate parameters for H_2 - O_2 combustion model (kinetics, and accompanying thermodynamics and transport parameter database taken from [26]). Parameters in bold are active for the present study. The forward rate constant, $K_f = AT^\beta \exp(-E_a/RT)$. *The number in parentheses is the exponent of 10, i.e., $2.65(16) = 2.65 \times 10^{16}$. For the active parameters, standard deviation of the prior, σ , and lower and upper bounds of the truncated Gaussian are provided in the last two columns.

ID	Reaction	A*	β	E_a	σ	Bounds (L:U)
1	$H + O_2 = O + OH$	2.644(16)	-0.6707	17041	2(16)	1(16):5(16)
2	$O + H_2 = H + OH$	4.589(4)	2.7	6260	1(4)	1(4):1(5)
3	$OH + H_2 = H + H_2O$	1.734(8)	1.51	3430	2(8)	4(7):4(8)
4	$2 OH = O + H_2O$	3.973(4)	2.4	-2110		
5	$2 H + M = H_2 + M$	1.780(18)	-1	0		
6	$H + OH + M = H_2O + M$	4.400(22)	-2	0		
7	$O + H + M = OH + M$	9.428(18)	-1	0		
8	$2 O + M = O_2 + M$	1.200(17)	-1	0		
9	$H + O_2 (+M) = HO_2 (+M)$					
	high pressure, $K_{f\infty}$	5.116(12)	0.44	0	5.116(12)	0:1.6(13)
	low pressure, K_{f0}	6.328(19)	-1.4	0	6.328(19)	0:2(20)
	TROE: $F_c = 0.5$					
	Third-body: $O_2(0.85), H_2O(11.89), Ar(0.40)$					
	He(0.46), $H_2(0.75)$					
10	$H_2 + O_2 = HO_2 + H$	5.916(5)	2.433	53502	5.916(5)	0:2(6)
11	$2 OH (+M) = H_2O_2 (+M)$					
	high pressure, $K_{f\infty}$	1.110(14)	-.37	0		
	low pressure, K_{f0}	2.010(17)	-.584	-2293		
	TROE: $F_c = 0.2654 \exp(-T/94)$					
	$+0.7346 \exp(-T/1756) + \exp(-5182/T)$					
	Third-body: $H_2(6), H_2O(6), Ar(0.7), He(0.7)$					
12	$HO_2 + H = O + H_2O$	3.970(12)	0	671		
13	$HO_2 + H = 2 OH$	7.485(13)	0	295	7.485(13)	0:2(14)
14	$HO_2 + O = OH + O_2$	4(13)	0	0	4(13)	0:2(14)
15a	$HO_2 + OH = O_2 + H_2O$	2.375(13)	0	-500	2.375(13)	0:1(14)
15b	$HO_2 + OH = O_2 + H_2O$	1.000(16)	0	17330		
16a	$2 HO_2 = O_2 + H_2O_2$	1.300(11)	0	-1630		
16b	$2 HO_2 = O_2 + H_2O_2$	3.658(14)	0	12000		
17	$H_2O_2 + H = HO_2 + H_2$	6.050(6)	20	5200		
18	$H_2O_2 + H = OH + H_2O$	2.410(13)	0	3970		
19	$H_2O_2 + O = OH + HO_2$	9.630(6)	2	3970		
20a	$H_2O_2 + OH = HO_2 + H_2O$	2.000(12)	0	427		
20b	$H_2O_2 + OH = HO_2 + H_2O$	2.670(41)	-7	37600		

5.2. Prior and forward model, likelihood and data. We use the optimized mechanism of [26], shown in detail in Table 1, to define the mean of a truncated Gaussian prior distribution for the parameters shown in red. Since we wish to quantify the extent to which experimental data (see below) constrains the active parameters, we take the standard deviations of each parameter to be of the order of its mean value (as shown in the table). In this way, we generate a broad Gaussian distribution that extends well beyond the upper and lower bounds specified for each parameter (shown in the rightmost column of the table). We truncate the Gaussians at the bounds to obtain our truncated Gaussian prior distribution. We note that these bounds represent are much broader than those considered in [26].

Our goal is to sharpen the prior knowledge by incorporating a limited set of experimental data. We consider the six flow reactors and eight ignition delay experiments labeled (*ign1a*, *ign1b*, *ign2a*, *ign2b*, *ign3a*, *ign4a*, *ign4b*, *ign5a*, *flw1a*, *flw2a*, *flw3a*, *flw4a*, *flw8a*, *flw8b*) in [26, Table 1]. The data include Gaussian error estimates, taken as the standard deviations, σ , reported in [26, Table 1]. We assume that the errors are

uncorrelated, i.e., the covariance matrix that defines the errors v in (1) is diagonal. The data, assembled into a 14×1 vector z , and the Gaussian error model are used along with our numerical model to define a likelihood by (1).

In (1), each element of the vector z corresponds to one of the 14 experiments. We compute z using point-reactors subject to constant pressure (flow-reactor) or constant volume (shock-tube) constraints. VODE [11] is used to integrate either set of equations over time, and the solution is monitored until the desired diagnostic is obtained. Initial data for the simulations is provided in the description of the experiment in the reference, as is the event defining the measured quantity. For example, experiment *flw1a* reports the mean rate of change of H_2 mole fraction during the time interval when it is between 40 and 60 percent of its initial value. The computed and measured data, along with the sampled and prior mean values of the parameters, and the corresponding prior standard deviations and experimental measurement errors are combined for each sample vector to form the resulting likelihood.

Our combustion model is accessorized with several options for efficiency and robustness so that it can be driven by a number of sampling strategies. In particular, the model is allowed to fail gracefully if any component of a sample (a vector of nine parameters) is generated outside the bounds specified in Table 1. In this case, the likelihood is treated formally as zero. Similarly, some combinations of parameter values lead to predicted values that are far from the corresponding measured data. Ignition delays, for example, may be prohibitively long to compute efficiently, or too short to accurately capture when scanning the solution at a “reasonable” sampling frequency. In either case, the diagnostic analysis routines can identify the failure mode and gracefully exit, signaling the driver code that the corresponding likelihood is to be taken as zero. Extra controls are provided/required in the specification of each experiment to help detect such scenarios. These auxiliary parameters are adjusted to have minimal effect on the presented data, but enable robust performance of the model evaluation over millions of sample sets. For each sample, the list of simulated experiments are evaluated in parallel over compute threads, using an OpenMP task manager. When a large set of independent sample vectors can be generated at once, the likelihood function itself can be evaluated in parallel over samples using MPI, where each MPI rank maintains its own version of the model parameter set and evaluates its likelihood using a local subset of the available cores.

5.3. Results. As explained in section 2, the prior distribution and a likelihood jointly define a posterior distribution, and we apply several variants of ISA algorithms to draw samples from this posterior distribution. We first test initialization by a GMM. To construct the GMM we run 10^4 minimizations of the negative logarithm of the posterior. The starting point for the optimizations are drawn from the broad prior. We use the nonlinear least squares code in the C implementation of “minpack” to carry out the required optimizations. One could also use derivative free optimization, but we decided to use minpack and compute derivatives by finite differences. We used 2,400 cores on NSERC’s supercomputer “Edison” to perform the optimization and of our 10^4 optimization attempts, 2,482 were “successful.” The large number of failures is caused our broad prior: many of the starting parameter sets we tried were outside of the range the numerical model can simulate, which made the optimizer return a failure. We identified 12 minima to be distinct by using the criterion outlined in section 3. Here we use a threshold based on a χ^2 distribution so that the probability of the minima being within the threshold exceeds 95%. The resulting Gaussian mixture consists of 12 components, centered at the minima.

FIG. 6. Iterations of ISA. Left: quality measure R as a function of iteration number for ISA initialized by a GMM, using Gaussian or multivariate t -distributions and 10^5 samples at each iteration. Center: quality measure R as a function of iteration number for ISA initialized by 10^4 samples of an MCMC run, using Gaussian or multivariate t -distributions and 10^5 samples at each iteration. Right: generating Gaussian/multivariate t -proposal distributions from 10^4 , 10^5 , and $6 \cdot 10^5$ samples.

The covariances are computed from approximate Hessians at these minima, using the same Gauss–Newton approximation as in section 4. We generate 10^5 samples from this GMM and evaluate the corresponding likelihoods using 3,072 cores on Edison. Evaluating and weighting 10^5 samples takes less than 15 minutes of wall-clock time and results in the quality measure, $R = 293$. From the weighted samples we generate Gaussian or multivariate t -distributions with shape parameter $\nu = 3$ and iterate as described above. At each step of the iteration we “inflate” the sample covariance by a factor of two to mitigate sampling error and the underestimation of covariances, as is common in numerical weather prediction [1]. The inflation parameter could be tuned further, but we did not pursue this because we obtained good results. The left panel of Figure 6 shows the quality measure R as a function of the iteration number when using a multivariate t -distribution. The figure suggests that ISA produces an effective proposal distribution after the first iteration with an associated R -value of less than 42. The precise sequence during iteration is $R = \{293, 32, 42, 35, 38, 29, 27\}$.

Also shown in the left panel of Figure 6 are results we obtain by Gaussian proposal distributions at each step of ISA. We note that the values of R fluctuate more than when using multivariate t -distributions, and the lowest value of R we obtain is also larger than the one we find by using multivariate t -proposal distributions. The precise sequence is $R = \{293, 94, 140, 54, 66, 105, 47\}$. Nonetheless, the values of R we obtain, in combination with the massive parallelism we can leverage at each step of the iteration, make ISA with Gaussian proposal distributions effective for this problem.

In the center panel of Figure 6, we show results we obtain by ISA initialized by multivariate t or Gaussian distributions computed from 10^4 samples of an MCMC run. As before we use the emcee implementation of the affine invariant ensemble sampler [33, 40]. We describe the details of our MCMC run below, as it also serves as a reference solution. In the context of ISA, we note that the initial Gaussian/multivariate t -distribution yields a significantly larger R than the GMM we constructed above. The iteration however can quickly reduce R to values of around $R = 42$ (Gaussian proposal), and $R = 35$ (multivariate t -proposal). As above, we use 10^5 samples for each iteration. Initializing ISA with only a “few” MCMC samples requires at least two iterations to obtain R -values below 50, whereas we could reach low values of R after only one iteration when we initialize ISA by a GMM. Constructing the GMM is

FIG. 7. Triangle plot of posterior distribution obtained by ISA after six iterations, each with 10^5 samples. The plot shows histograms of reaction IDs $\{2, 3, 10, 14, 15a\}$ (diagonal), as well as two-dimensional marginals of all pairs of these five parameters (lower triangle).

computationally more expensive due to the many optimizations we need to run, but constructing the GMM can leverage parallelism effectively. We conclude that, for this problem and computational set-up, the initialization by GMM is the more effective strategy.

We compare the ISA iterations to constructing Gaussian and multivariate t -proposal distributions from longer MCMC runs. These tests will (i) illustrate the degree to which ISA improves sampling efficiency compared to longer MCMC runs to obtain comparable results and (ii) demonstrate the convergence of ISA on this problem. In the right panel of Figure 6 we show the values of R we compute for proposal distributions as a function of the number of MCMC samples we use to compute the mean and covariances of these distributions. Specifically, we use 10^4 , 10^5 , and $6 \cdot 10^5$ to obtain $R = \{689, 60, 42\}$ for Gaussian proposals and $R = \{471, 130, 34\}$ for multivariate t -proposals. We thus obtain similar values of R as when iterating by ISA, provided that sufficient MCMC samples are used to construct the proposal distribution. This suggests that ISA can indeed find Gaussian or multivariate t approximations of the posterior distribution that can be used to generate posterior samples with efficiency of $R \approx 40$.

For our MCMC runs, we use a serial implementation of the emcee code, but parallelize the likelihood evaluations across the various experiments (flow reactors and ignition delays). In this way, we could effectively utilize the eight cores on a workstation. Generating $6 \cdot 10^5$ samples then requires about one month. For comparison, initializing ISA by 10^4 MCMC samples requires about 10 hours of MCMC for generating the initial ensemble, and then 15 mins per iteration using 3,072 cores. Initializing ISA by GMM further reduces the wall-clock time required to generate the initial ensemble. On the other hand the emcee code also can be parallelized to a limited extent, which would reduce run-times and improve overall performance. Nonetheless, it is difficult to achieve the same level of parallelism as ISA by emcee or other MCMC codes. Moreover, the “shape” of the posterior distribution is difficult for the emcee

FIG. 8. *Negative logarithm of posterior distribution along lines in parameter space.*

code, which results in relatively long integrated auto-correlation times of a few hundred samples. Precise estimates of auto-correlation times are difficult to come by, but here we are satisfied with the estimates provided by the emcee code (which differ significantly across the nine variables). An autocorrelation of a few hundred means that the “efficiency” of the emcee code is about one effective sample for every few hundred we try. This should be compared to the quality measure $R \approx 42$ of ISA (after a few iterations), which suggests that one in about 40 samples is effective.

Finally we show a triangle-plot of the posterior distribution we obtain by ISA after 6 iterations initialized by a GMM. The plot shows histograms of five parameters (reaction IDs $\{2, 3, 10, 14, 15a\}$), as well as two-dimensional marginals of pairwise combinations of these five variables. In the plot, the samples are scaled by the prior mean as shown in Table 1. We note that there is significant probability mass away from the prior mean, i.e., away from the value one. Moreover, the posterior mass is concentrated on narrow ridges in parameter space, which suggests that posterior variances are reduced significantly compared to the broad prior. We also note that there are rather strong correlations among some of the variables, e.g., reaction IDs two and three, and such correlations are absent in the prior. Thus, the experimental data we use here do indeed constrain the parameters we wish to estimate, and ISA can identify the parameter regimes that are consistent with the data we consider.

Finally, we wish to point out that the triangle plot suggests that the posterior has only one mode, whereas our optimization found 12 “distinct” modes. This apparent contradiction is caused by the modes not being “well separated.” In fact we believe that the posterior distribution of this problem is a nine-dimensional analogue of the distribution illustrated in Figure 1, i.e., a plateau with several “shallow dents.” In this situation, the various modes may not be easy to spot in a triangle plot because (i) we may not have sufficient resolution due to limited sample size and (ii) the integration implied by viewing marginals only. We can, however, illustrate the complex structure of the posterior distribution by evaluating the distribution along lines in parameter space as in Figure 8. These lines are obtained as follows. We take seven parameter vectors from our reference MCMC run that result in the smallest negative logarithm of the posterior. We then connect any two of these vectors by a line and evaluate the posterior distribution at 100 equally spaced points along these lines. We observe several dents along these lines. These dents are deep enough to

cause importance sampling algorithms that make use of local geometry to fail. For example, we tried implicit sampling (see, e.g., [19, 60]), in which a proposal distribution is defined by the mode and the Hessian of the negative log-posterior at the mode and found that the method samples efficiently in the vicinity of the local minima, but fails to explore the broader parameter space. The various local minima are also deep enough to cause large integrated auto-correlation times in the MCMC runs (see above).

6. Summary and conclusions. We presented iterative importance sampling schemes, which use samples obtained by importance sampling to construct a new proposal distribution that leads to more efficient importance sampling. These ideas have been discussed before (see our literature review), and the convergence of such iterations has also been studied. It is important in practice that the iteration remains robust when only few effective samples are produced at each step. We studied this behavior in two test problems, one in subsurface flow and one in combustion modeling. In our implementation of ISA we made use of massively parallel computers and found that the iterations can robustly leverage parallelism, which led to significant speed ups when compared to MCMC. We also presented two strategies to initialize the iterations. One strategy uses a “short” MCMC run, and the other makes use of optimization and Gaussian mixture models.

A natural question is this: how do the computational requirements of ISA scale with the dimension of the parameter vector? Our numerical results provide little information here because the dimension of the parameter vectors we considered are relatively small. On the other hand, we argued in section 3.2 that the quality measure R of Gaussian sampling problems depends exponentially on the dimension of the parameter vector (see also [20, 69]). This in turn implies that the number of samples required to achieve a given number of “effective samples” also increases exponentially with dimension. This may limit applicability of ISA, as the method would get stuck during the initialization phase if R were inevitably large in high-dimensional problems. On the other hand, the Gaussian example that exhibits exponential dependence on the dimension does not account for, e.g., sparsity of covariance matrices of the target or proposal distributions. If such sparsity were present, and if one could exploit sparsity during sampling, then high-dimensional problems may come within reach of ISA or other importance sampling methods. The situation is perhaps analogous to numerical linear algebra, where computations with large dense matrices are infeasible, while sparse matrices can be handled easily. In fact, sparsity of covariances is exploited during “covariance localization” of ensemble Kalman filters and is the key to being able to effectively solve estimation problems in numerical weather prediction (where the state dimension is several hundred million). If such strategies could be applied to parameter estimation problems, e.g., by enforcing that some data only inform a subset of the parameters one wants to estimate, then high-dimensional problems may become feasible. In the applications presented here, we avoid this issue by selecting a priori only a small set of parameters.

REFERENCES

- [1] J. ANDERSON, *Spatially and temporally varying adaptive covariance inflation for ensemble filters*, Tellus, 61 (2009), pp. 72–83.
- [2] M. ARULAMPALAM, S. MASKELL, N. GORDON, AND T. CLAPP, *A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking*, IEEE Trans. Signal Process., 50 (2002), pp. 174–188.

- [3] E. ATKINS, M. MORZFELD, AND A. CHORIN, *Implicit particle methods and their connection with variational data assimilation*, Mon. Weather Rev., 141 (2013), pp. 1786–1803.
- [4] J. BARDSLEY, A. SOLONEN, H. HAARIO, AND M. LAINE, *Randomize-then-optimize: A method for sampling from posterior distributions in nonlinear inverse problems*, SIAM J. Sci. Comput., 36 (2014), pp. A1895–A1910.
- [5] T. BENGTTSSON, P. BICKEL, AND B. LI, *Curse of dimensionality revisited: The collapse of importance sampling in very large scale systems*, in Probability and Statistics: Essays in Honor of David A. Freedman, IMS Collections 2, 2008, pp. 316–334.
- [6] N. BERGMAN, ED., *Recursive Bayesian Estimation: Navigation and Tracking Applications*, Ph.D dissertation, Linköping University, Linköping, Sweden, 1999.
- [7] A. BESKOS, A. JASRA, E. MUZAFFER, AND A. STUART, *Sequential Monte Carlo methods for Bayesian elliptic inverse problems*, Stat. Comput., 25 (2015), pp. 727–737, <https://doi.org/10.1007/s11222-015-9556-7>.
- [8] P. BICKEL, T. BENGTTSSON, AND J. ANDERSON, *Sharp failure rates for the bootstrap particle filter in high dimensions*, in Pushing the Limits of Contemporary Statistics: Contributions in Honor of Jayanta K. Ghosh, IMS Collections 3, 2008, pp. 318–329.
- [9] C. T. BRAAK, *A Markov chain Monte Carlo version of the genetic algorithm differential evolution: Easy Bayesian computing for real parameter spaces*, Stat. Comput., 16 (2006), pp. 239–249, <https://doi.org/10.1007/s11222-006-8769-1>.
- [10] C. T. BRAAK AND A. VRUGT, *Differential evolution Markov chain with snooker updater and fewer chains*, Stat. Comput., 18 (2008), pp. 435–446, <http://dx.doi.org/10.1007/s11222-008-9104-9>.
- [11] P. BROWN, G. BYRNE, AND A. HINDMARSH, *Vode: A variable-coefficient ode solver*, SIAM J. Sci. and Stat. Comput., 10 (1989), pp. 1038–1051, <https://doi.org/10.1137/0910062>.
- [12] C. BUCHER, *Adaptive sampling—An iterative fast Monte Carlo procedure*, Struct. Saf., 5 (1988), pp. 119–126.
- [13] T. BUI-THANH, O. GHATTAS, J. MARTIN, AND G. STADLER, *A computational framework for infinite-dimensional Bayesian inverse problems. Part I: The linearized case, with application to global seismic inversion*, SIAM J. Sci. Comput., 36 (2013), pp. A2494–A2523.
- [14] O. CAPPÉ, R. DOUC, A. GUILLIN, J.-M. MARIN, AND C. ROBERT, *Adaptive importance sampling in general mixture classes*, Stat. Comput., 18 (2008), pp. 447–459.
- [15] O. CAPPÉ, A. GUILLIN, J.-M. MARIN, AND C. ROBERT, *Population Monte Carlo*, J. Comput. Graph. Statist., 13 (2004), pp. 907–929.
- [16] A. CHORIN AND O. HALD, *Stochastic Tools in Mathematics and Science*, 3rd ed., Springer, New York, 2013.
- [17] A. CHORIN AND M. MORZFELD, *Conditions for successful data assimilation*, J. Geophys. Res. Atmos., 118 (2013), pp. 11,522–11,533.
- [18] A. CHORIN, M. MORZFELD, AND X. TU, *Implicit particle filters for data assimilation*, Commun. Appl. Math. Comput. Sci., 5 (2010), pp. 221–240.
- [19] A. CHORIN AND X. TU, *Implicit sampling for particle filters*, Proc. Natl. Acad. Sci. USA, 106 (2009), pp. 17249–17254.
- [20] A. J. CHORIN, F. LU, R. MILLER, M. MORZFELD, AND X. TU, *Sampling, feasibility, and priors in data assimilation*, Discrete Contin. Dyn. Syst., 36 (2016), pp. 4227–4246.
- [21] J. CHRISTEN AND C. FOX, *Markov chain Monte Carlo using an approximation*, J. Comput. Graph. Statist., 14 (2005), pp. 795–810.
- [22] P. CONRAD, Y. MARZOUK, N. PILLAI, AND A. SMITH, *Accelerating asymptotically exact MCMC for computationally intensive models via local approximations*, J. Amer. Statist. Assoc., 111 (2016), pp. 1591–1607.
- [23] T. CUI, K. J. H. LAW, AND Y. M. MARZOUK, *Dimension-independent likelihood-informed MCMC*, J. Comput. Phys., 304 (2016), pp. 109–137, <https://doi.org/doi:10.1016/j.jcp.2015.10.008>.
- [24] T. CUI, J. MARTIN, Y. M. MARZOUK, A. SOLONEN, AND A. SPANTINI, *Likelihood-informed dimension reduction for nonlinear inverse problems*, Inverse Problems, 29 (2014), 114015.
- [25] T. CUI, Y. M. MARZOUK, AND K. WILLCOX, *Scalable posterior approximations for large-scale Bayesian inverse problems via likelihood-informed parameter and state reduction*, J. Comput. Phys., 315 (2016), pp. 363–387.
- [26] S. G. DAVIS, A. V. JOSHI, H. WANG, AND F. N. EGOLFOPOULOS, *An optimized kinetic model for H₂/CO combustions*, Proc. Combust. Inst., 30 (2005), pp. 1283–1292.
- [27] C. V. DEUTSCH AND A. G. JOURNEL, *GSLIB: Geostatistical Software Library and User's Guide*, Oxford University Press, Oxford, 1992.
- [28] Y. DONG, A. T. HOLLEY, M. G. ANDAC, F. N. EGOLFOPOULOS, S. G. DAVIS, P. MIDDHA, AND H. WANG, *Extinction of premixed H₂/air flames: Chemical kinetics and molecular diffusion effects*, Combust. Flame, 142 (2005), pp. 374–387.

- [29] R. DOUC, A. GUILLIN, J.-M. MARIN, AND C. ROBERT, *Convergence of adaptive mixtures of importance sampling schemes*, Ann. Statist., 35 (2000), pp. 420–448.
- [30] A. DOUCET, N. DE FREITAS, AND N. GORDON, EDS., *Sequential Monte Carlo Methods in Practice*, Springer, New York, 2001.
- [31] S. DUANE, A. KENNEDY, B. PENDLETON, AND D. ROWETH, *Hybrid Monte Carlo*, Phys. Lett. B, 195 (1987), pp. 216–222.
- [32] H. FLATH, L. WILCOX, V. AKÇELIK, J. HILL, B. VAN BLOEMEN WAANDER, AND O. GHATTAS, *Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations*, SIAM J. Sci. Comput., 33 (2011), pp. 407–432.
- [33] D. FOREMAN-MACKEY, A. CONLEY, W. MEIERJURGEN FARR, D. W. HOGG, D. LONG, P. MARSHALL, A. PRICE-WHELAN, J. SANDERS, AND J. ZUNTZ, *emcee: The MCMC Hammer*. Astrophys. Source Code Library, 2013, <http://arxiv.org/abs/1202.3665>.
- [34] P. GIORDANI AND R. KOHN, *Adaptive independent Metropolis–Hastings by fast estimation of mixtures of normals*, J. Comput. Graph. Statist., 19 (2010), pp. 243–259.
- [35] M. GIROLAMI AND B. CALDERHEAD, *Riemann manifold Langevin and Hamiltonian Monte Carlo methods*, J. R. Stat. Soc. Ser. B Stat. Methodol., 73 (2011), pp. 1–37.
- [36] G. GIVENS AND A. RAFTERY, *Local adaptive importance sampling for multivariate densities with strong nonlinear relationships*, J. Amer. Statist. Assoc., 91 (1996), pp. 132–141.
- [37] P. GLARBERG, M. U. ALZUETA, K. DAM-JOHANSEN, AND J. A. MILLER, *Kinetic modelling of hydrocarbon/nitric oxide interactions in a flow reactor*, Combust. Flame, 115 (1998), pp. 1–27.
- [38] J. J. GOMEZ-HERNANEZ, A. SAHUQUILLO, AND J. CAPILLA, *Stochastic simulation of transmissivity fields conditional to both transmissivity and piezometric data—I. Theory*, J. Hydrology, 203 (1997), pp. 167–174.
- [39] J. GOODMAN, K. LIN, AND M. MORZFELD, *Small-noise analysis and symmetrization of implicit Monte Carlo samplers*, Comm. Pure Appl. Math., 69 (2016), pp. 1924–1951.
- [40] J. GOODMAN AND J. WEARE, *Ensemble samplers with affine invariance*, Commun. Appl. Math. Comput. Sci., 5 (2010), pp. 65–80.
- [41] J. F. GRGAR, J. B. BELL, AND M. S. DAY, *The soot effect in naturally propagating, premixed, lean, hydrogen-air flames*, Proc. Combust. Inst., 32 (2008), pp. 1173–1180.
- [42] H. HAARIO, M. LAINE, A. MIRA, AND E. SAKSMAN, *DRAM: Efficient adaptive MCMC*, Statist. Comput., 16 (2006), pp. 339–354.
- [43] H. HAARIO, E. SAKSMAN, AND J. TAMMINEN, *An adaptive metropolis algorithm*, Bernoulli, 7 (2001), pp. 223–242.
- [44] K. HANSON, *Markov Chain Monte Carlo posterior sampling with the Hamiltonian method*, in Proceedings of Medical Imaging, 2001.
- [45] L. HOLDEN, R. HAUGE, AND M. HOLDEN, *Adaptive independent Metropolis–Hastings*, Ann. Appl. Probab., 19 (2009), pp. 395–413.
- [46] Y. IBA, *Population-based Monte Carlo algorithms*, Trans. Jpn. Soc. Artif. Intell., 16 (2000), pp. 279–286.
- [47] M. KALOS AND P. WHITLOCK, *Monte Carlo Methods*, Vol. 1, 1st ed., John Wiley & Sons, New York, 1986.
- [48] J. KEITH, D. KROESE, AND G. SOFRONOV, *Adaptive independence samplers*, Statist. Comput., 18 (2008), pp. 409–420.
- [49] J. LEE, R. MCVINISH, AND K. MENGENSEN, *Population Monte Carlo algorithm in high dimensions*, Methodol. Comput. Appl. Probab., 13 (2011), pp. 369–389.
- [50] K. LEVENBERG, *A method for the solution of certain problems in least squares*, Quart. Appl. Math., 2 (1944), pp. 164–168.
- [51] J. LI, Z. ZHAO, A. KAZAKOV, AND F. DRYER, *An updated comprehensive kinetic model of hydrogen combustion*, Int. J. Chem. Kinet., 36 (2004), pp. 566–575.
- [52] C. LIEBERMAN, K. WILLCOX, AND O. GHATTAS, *Parameter and state model reduction for large-scale statistical inverse problems*, SIAM J. Sci. Comput., 32 (2010), pp. 2523–2542.
- [53] J. LIU AND R. CHEN, *Sequential Monte Carlo methods for dynamical systems*, J. Amer. Statist. Assoc., 93 (1998), pp. 1032–1044.
- [54] F. LU, M. MORZFELD, X. TU, AND A. CHORIN, *Limitations of polynomial chaos expansions in the Bayesian solution of inverse problems*, J. Comput. Phys., 282 (2015), pp. 138–147.
- [55] D. MARQUARDT, *An algorithm for least-squares estimation of nonlinear parameters*, J. Soc. Indust. Appl. Math., 11 (1963), pp. 431–441.
- [56] J. MARTIN, L. C. WILCOX, C. BURSTEDDE, AND O. GHATTAS, *A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion*, SIAM J. Sci. Comput., 34 (2012), pp. A1460–A1487.

- [57] J. A. MILLER, *Chemical Mechanism for Methane Combustion*, private communication.
- [58] P. D. MORAL, A. DOUCET, AND A. JASRA, *Sequential Monte Carlo samplers*, J. R. Stat. Soc. Ser. B Stat. Methodol., 68 (2006), pp. 411–436.
- [59] M. MORZFELD, D. HODYSS, AND C. SNYDER, *What the collapse of the ensemble Kalman filter tells us about particle filters*, Tellus, 69 (2017), 1283809.
- [60] M. MORZFELD, X. TU, E. ATKINS, AND A. CHORIN, *A random map implementation of implicit filters*, J. Comput. Phys., 231 (2012), pp. 2049–2066.
- [61] M. MORZFELD, X. TU, J. WILKENING, AND A. CHORIN, *Parameter estimation by implicit sampling*, Commun. Appl. Math. Comput. Sci., 10 (2015), pp. 205–225.
- [62] T. A. MOSELHY AND Y. M. MARZOUK, *Bayesian inference with optimal maps*, J. Comput. Phys., 231 (2012), pp. 7815–7850.
- [63] G. S. N. PETRA, J. MARTIN AND O. GHATTAS, *A computational framework for infinite-dimensional Bayesian inverse problems: Part II. Stochastic Newton MCMC with application to ice sheet flow inverse problems*, SIAM J. Sci. Comput., 36 (2013), pp. A1525–1555.
- [64] R. NEAL, *MCMC using Hamiltonian dynamics*, in Handbook of Markov Chain Monte Carlo, S. Brooks, A. Gelman, G. Jones, and X.-L. Meng, eds., Chapman & Hall, Oxford, 2011.
- [65] M. O’CONNAIRE, H. J. CURRAN, J. M. SIMMIE, W. J. PITZ, AND C. K. WESTBROOK, *A comprehensive modeling study of hydrogen oxidation*, Int. J. Chem. Kinet., 36 (2004), pp. 603–622.
- [66] M.-S. OH AND J. BERGER, *Adaptive Importance Sampling in Monte Carlo Integration*, Technical report 89–19C, Department of Statistics, Purdue University, 1989.
- [67] M.-S. OH AND J. BERGER, *Adaptive Importance Sampling in Monte Carlo Integration*, J. Statist. Comput. Simul., 41 (1992), pp. 143–168.
- [68] D. OLIVER, *Minimization for conditional simulation: Relationship to optimal transport*, J. Comput. Phys., 265 (2012), pp. 1–15.
- [69] A. OWEN, *Monte Carlo Theory, Methods and Examples*, statweb.stanford.edu/~owen/mc (2013).
- [70] M. PARNO AND Y. M. MARZOUK, *Transport map accelerated Markov chain Monte Carlo*, submitted, 2016.
- [71] M. PARNO, T. MOSELHY, AND Y. M. MARZOUK, *A multiscale strategy for Bayesian inference using transport maps*, submitted, 2016.
- [72] K. PRUESS, G. MORIDIS, AND C. OLDENBURG, *TOUGH2 User’s Guide, Version 2.1*, Report LBNL-43134, Lawrence Berkeley National Laboratory, Berkeley, CA, 1999.
- [73] B. RAMARAO, A. LAVENUE, G. D. MARSILY, AND M. MARIETTA, *Pilot point methodology for automated calibration of an ensemble of conditionally simulated transmissivity fields: 1. Theory and computational experiments*, Water Resour. Res., 31 (1995), pp. 475–493.
- [74] P. REBESCHINI AND R. VAN HANDEL, *Can local particle filters beat the curse of dimensionality?*, Ann. Appl. Probab., 25 (2015), pp. 2809–2866.
- [75] N. RECCA, *A New Methodology for Importance Sampling*, Master’s thesis, Courant Institute of Mathematical Sciences, New York University, New York, 2011.
- [76] L. RICHARDS, *Capillary conduction of liquids through porous mediums*, Physics, 1 (1931), pp. 318–333.
- [77] G. ROBERTS AND J. ROSENTHAL, *Optimal scaling of discrete approximations to Langevin diffusions*, J. R. Stat. Soc. Ser. B Stat. Methodol., 60 (1998), pp. 255–268.
- [78] R. RUBINSTEIN AND D. KROESE, *The Cross-Entropy Method*, Springer, New York, 2004.
- [79] S. L. COTTER, G. O. ROBERTS, A. STUART, AND D. WHITE, *MCMC methods for functions: Modifying old algorithms to make them faster*, Statist. Sci., 28 (2013), pp. 424–446.
- [80] G. P. SMITH ET AL., *GRI-Mech 3.0*, <http://combustion.berkeley.edu/gri-mech/>.
- [81] C. SNYDER, *Particle filters, the “optimal” proposal and high-dimensional systems*, in Proceedings of the ECMWF Seminar on Data Assimilation for Atmosphere and Ocean, 2011.
- [82] C. SNYDER, T. BENGTSOON, P. BICKEL, AND J. ANDERSON, *Obstacles to high-dimensional particle filtering*, Month. Weather Rev., 136 (2008), pp. 4629–4640.
- [83] C. SNYDER, T. BENGTSOON, AND M. MORZFELD, *Performance bounds for particle filters using the optimal proposal*, Month. Weather Rev., 143 (2015), pp. 4750–4761.
- [84] A. SPANTINI, A. SOLONEN, T. CUI, J. MARTIN, L. TENORIO, AND Y. M. MARZOUK, *Optimal low-rank approximations of Bayesian linear inverse problems*, SIAM J. Sci. Comput., 37 (2015), pp. A2451–A2487.
- [85] A. STORDAL, *Iterative Bayesian inversion with Gaussian mixtures: Finite sample implementation and large sample asymptotics*, Comput. Geosci., 19 (2015), pp. 1–15.
- [86] A. M. STUART, *Inverse problems: A Bayesian perspective*, Acta Numer., 19 (2010), pp. 451–559.
- [87] A. TARANTOLA, *Inverse Problem Theory and Methods for Model Parameter Estimation*, SIAM, Philadelphia, 2005.

- [88] M. VAN GENUCHTEN, *A closed-form equation for predicting the hydraulic conductivity of unsaturated soils*, Soil Sci. Soc. Am. J., 44 (1980), pp. 892–898.
- [89] P. VAN LEEUWEN, Y. CHEN, AND S. REICH, *Nonlinear Data Assimilation*, Springer, New York, 2015.
- [90] E. VANDEN-EIJNDEN AND J. WEARE, *Data assimilation in the low noise, accurate observation regime with application to the Kuroshio current*, Month. Weather Rev., 2013.
- [91] P. ZHANG, *Nonparametric importance sampling*, J. Amer. Statist. Assoc., 91 (1996), pp. 1245–1253.