Delft University of Technology

A stochastic simplex approximate gradient (StoSAG) for optimization under uncertainty

Fonseca, Rahul; Chen, B; Jansen, Jan Dirk; Reynolds, Albert C.

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

A Stochastic Simplex Approximate Gradient (StoSAG) for optimization under uncertainty

Rahul Rahul-Mark Fonseca[1], Bailian Chen[2], Jan Dirk Jansen[1,*,†] and Albert Reynolds[2]

[1]*Delft University of Technology, Stevinweg 1, 2628CN, Delft, The Netherlands*
[2]*University of Tulsa, 800 South Tucker Drive, Tulsa, OK 74104, USA*

SUMMARY

We consider a technique to estimate an approximate gradient using an ensemble of randomly chosen control vectors, known as Ensemble Optimization (EnOpt) in the oil and gas reservoir simulation community. In particular, we address how to obtain accurate approximate gradients when the underlying numerical models contain uncertain parameters because of geological uncertainties. In that case, 'robust optimization' is performed by optimizing the expected value of the objective function over an ensemble of geological models. In earlier publications, based on the pioneering work of Chen *et al.* (2009), it has been suggested that a straightforward one-to-one combination of random control vectors and random geological models is capable of generating sufficiently accurate approximate gradients. However, this form of EnOpt does not always yield satisfactory results. In a recent article, Fonseca *et al.* (2015) formulate a modified EnOpt algorithm, referred to here as a Stochastic Simplex Approximate Gradient (StoSAG; in earlier publications referred to as 'modified robust EnOpt') and show, via computational experiments, that StoSAG generally yields significantly better gradient approximations than the standard EnOpt algorithm. Here, we provide theoretical arguments to show why StoSAG is superior to EnOpt. © 2016 The Authors. *International Journal for Numerical Methods in Engineering* Published by John Wiley & Sons Ltd.

1. INTRODUCTION

1.1. Reservoir simulation and optimization

Reservoir simulation is a branch of fluid mechanics concerned with the numerical simulation of multiphase flow through heterogeneous porous media in the form of subsurface hydrocarbon (oil and gas) reservoirs; for example, [1] or [2]. Oil and gas are produced by drilling wells into these reservoirs, which, after an initial period of dry oil production, normally start to co-produce increasing amounts of water that is always naturally present in any hydrocarbon reservoir. To maintain sufficient reservoir pressure and increase ultimate oil recovery, it is often necessary to also drill dedicated injection wells and inject gas, water, or chemicals (e.g., polymers to increase the water viscosity) into the reservoir. The flow through injection and production wells is controlled by specifying flow rates, pressures, or valve settings (either at surface or downhole) in the wells, which may vary over time. The optimal operation of these wells is, therefore, a time-dependent multivariable optimization problem that typically aims to maximize the economic benefit (e.g., total oil production minus water injection and disposal costs) over the producing life of the reservoir [3]. The most efficient

---

*Correspondence to: Jan Dirk Jansen, Delft University of Technology, Stevinweg 1, 2628CN, Delft, The Netherlands.
†E-mail: j.d.jansen@tudelft.nl

optimization techniques are gradient based where the gradient vector (i.e., the search direction in control space) is computed using an adjoint formulation, a method also known as optimal control [4]. For applications to reservoir production optimization, see for example, [5–10] and further references therein. The beauty of the adjoint method is that it is independent of the number of control variables and generates a gradient of the objective function with respect to the entire control vector (i.e., the controls for all wells over all time steps) in a single forward simulation plus one additional backward or adjoint simulation. Once a gradient has been computed, it can be used in a gradient-based algorithm of choice, ranging from a simple steepest ascent to more complex quasi-Newton algorithms with line search or trust region approaches [11]. Because of geological uncertainties, the spatial distribution of the flow-relevant parameters of subsurface reservoirs is only approximately known, and therefore, the simulations are often performed using an ensemble of reservoir flow models, an approach referred to as robust optimization [12]. In robust optimization, the optimal control is defined as the expected value of the objective function over the ensemble of geological models, each of which needs to be simulated over the entire producing reservoir life to compute its corresponding objective function value. In an adjoint-based implementation of robust optimization, the number of simulations to compute a single gradient vector therefore scales linearly with the number of geological ensemble members.

1.2. Ensemble optimization

Unfortunately, implementation of the adjoint in a reservoir simulation code is time consuming and code intrusive, and therefore, various alternative optimization methods are being used. One of these, a non-intrusive approximate gradient method known as Ensemble Optimization (EnOpt), has gained considerable popularity over the past years in the reservoir simulation community following the pioneering work of Chen *et al.* [13, 14]. In the simplest form of EnOpt, an approximate gradient is computed by simulating the response of a single (deterministic) reservoir model for an ensemble of stochastically generated different well-control vectors, where each control vector contains the settings for the set of all wells over all time steps. Do and Reynolds [15] provided an analysis of the deterministic EnOpt method and demonstrated its close connection to other stochastic gradient estimation techniques such as simultaneous perturbation stochastic approximation (SPSA, [16, 17]) and the simplex gradient method [18]. Stordal *et al.* [19] recently demonstrated that deterministic EnOpt can also be interpreted as a special case of a machine learning algorithm known as Gaussian mutation (Amari *et al.* [20]). Note that in this deterministic implementation of EnOpt, the number of simulations to compute a single gradient vector scales linearly with the number of control vectors in the control ensemble.

The concept of generating an approximate gradient from an ensemble of control vectors for a single reservoir model has its roots in the work of Lorentzen *et al.* [21], but the name EnOpt was coined by Chen *et al.* [13] (see also [14]). However, in most cases, when people refer to EnOpt, they refer to the implementation of Chen *et al.* [13], which pertains to robust optimization. In a naive robust implementation of EnOpt, one expects to need an ensemble of control vectors for each member of the ensemble of geological models. That would imply that the number of simulations to compute a single gradient vector would scale linearly with the product of the number of control vectors and the number of geological models. However, Chen *et al.* [13] propose to pair each random control vector with a single randomly chosen geological model. In that way, the number of simulations is dramatically reduced, and just like for the robust adjoint method, the number of simulations to compute a single gradient scales linearly with the number of geological models. Obviously, the quality of the approximate gradient obtained with EnOpt will have an effect on the gradient-based optimization procedure in which it is applied, and a few studies that have been published suggest that the one-to-one pairing of control vectors and geological realizations, as proposed by Chen *et al.* [13], is therefore not necessarily the best choice to be used in reservoir optimization studies. Raniolo *et al.* [22] observed that EnOpt did not provide a good search direction for an optimization problem related to the injection of polymers to enhance oil recovery. They indicate that improved performance of the optimization procedure can be obtained by a procedure in which five control perturbations are paired with each geological model. Unfortunately, the authors did not compare their results with those obtained with EnOpt; moreover, they only used a selection of representative geological models rather than the entire ensemble. Fonseca *et al.* [23] studied the effect of different ratios of control perturbations to geological models with the aid of hypothesis testing. They tested a variety of robust

EnOpt formulations and concluded that 'In general, irrespective of the gradient formulation, using higher ratios (i.e. 1:10, 1:20, 1:50 etc.) always leads to better quality gradient estimates'.

1.3. Modified EnOpt formulations

The original version of (robust) EnOpt, as proposed by Chen *et al*. [13] in 2009, implicitly produces a smoothed version of the approximate gradient because it can be interpreted as a gradient pre-multiplied with the covariance matrix used in generating the random controls. Chen *et al*. [13] also propose a doubly smoothed version, which is obtained by an additional (explicit) multiplication with the covariance matrix. Thereafter, various alternative formulations have been proposed, with or without smoothing, while the theoretical connections between EnOpt and other stochastic optimization methods also have been investigated. Do [24] and Yan and Reynolds [25] observe that, for a deterministic reservoir description, very similar estimates of the economic objective are obtained regardless of whether G-SPSA (a variety of classic SPSA), EnOpt, or a simplex gradient is used to estimate optimal well controls. For the deterministic case, Do and Reynolds [15] provide a theoretical argument to establish why optimization algorithms that use seemingly different stochastic approximations of the gradient as the search direction in an ascent or descent algorithm should yield similar results provided the same degree of the smoothing of the gradient is used in each algorithm, thereby explaining the observations in [24] and [25].

Do and Reynolds [15] also present a modified version of EnOpt that requires fewer assumptions to formulate a search direction than the one used in standard (deterministic) EnOpt, and they demonstrate that this modified EnOpt gradient and the original version perform virtually identically for the case of a single (i.e., deterministic) reservoir model. However, in a robust optimization framework, Fonseca *et al*. [23, 26, 27] observe that a generalization of the modified EnOpt gradient used in [15] results in an algorithm that usually performs substantially better than the standard EnOpt formulation of Chen *et al*. [13]. In the remainder of this paper, we will refer to this improved algorithm as the Stochastic Simplex Approximate Gradient formulation (StoSAG; in earlier publications referred to as 'modified robust EnOpt').

The objective of this paper is to provide a theoretical foundation for the results observed by Fonseca *et al*. [23, 26, 27] and the aforementioned observations of Raniolo *et al*. [22]. In particular, we show that the standard EnOpt formulation cannot be expected to obtain a reliable search direction in the robust case unless all the reservoir models are very similar.

2. ROBUST OPTIMIZATION

We consider the problem of finding $\mathbf{u}$, which maximizes the expectation over $\mathbf{m}$ of a nonlinear functional of the form $J = J(\mathbf{m}, \mathbf{u}, \mathbf{y}(\mathbf{m}, \mathbf{u}))$, where $\mathbf{m}$ is a random vector with a known probability density function (pdf), $\mathbf{u}$ is a vector of system inputs (controls), and $\mathbf{y} = \mathbf{y}(\mathbf{m}, \mathbf{u})$ is the output vector, which is, in turn, a function of the system response to input vector $\mathbf{u}$ for a given model $\mathbf{m}$ [28]. In many applications of interest, for a given $\mathbf{m}$ and $\mathbf{u}$, $\mathbf{y}(\mathbf{m}, \mathbf{u})$ is the solution of a system of nonlinear partial differential equations for an initial-boundary-value problem. As we assume that this system of equations is deterministic, J is a function of $\mathbf{u}$ and $\mathbf{m}$ only, that is, $J = J(\mathbf{m}, \mathbf{u})$. Often, the solution of the initial-boundary-value problem can only be approximated via the solution of a system of nonlinear discrete equations generated from a finite volume of finite element formulation as is the case for the problems considered in [13, 23, 26, 27] as well as the examples presented here.

Throughout, we assume that the uncertainty in $\mathbf{m}$ can be represented by sampling its pdf to obtain an ensemble of N_e realizations (i.e., geological models in our setting), $\mathbf{m}_i$, $i = 1, 2, \cdots N_e$ and thus approximate the expectation of J with respect to $\mathbf{m}$ as the mean of the set $\{J(\mathbf{m}_i, \mathbf{u})\}_{i=1}^{N_e}$, that is,

$$J_E(\mathbf{u}) = \frac{1}{N_e} \sum_{i=1}^{N_e} J(\mathbf{m}_i, \mathbf{u}), \tag{1}$$

where $J_E(\mathbf{u})$ denotes the approximation of the expectation $E_{\mathbf{m}}[J(\mathbf{m}, \mathbf{u})]$. As our objective is theoretical understanding, it suffices to assume that the only constraints on the problem are upper and

lower bounds on the control vector, $\mathbf{u}$, and denote the vector of upper bounds by $\mathbf{u}^{up}$ and the vector of lower bounds by $\mathbf{u}^{low}$. Thus, the robust optimization problem is stated as

$$
\max_{\mathbf{u}} J_E(\mathbf{u}) = \max_{\mathbf{u}} \frac{1}{N_e} \sum_{i=1}^{N_e} J(\mathbf{m}_i, \mathbf{u}),
$$
$$
\text{s.t.} \quad \mathbf{u}^{low} \leq \mathbf{u} \leq \mathbf{u}^{up}.
$$
(2)

3. ENSEMBLE OPTIMIZATION

All the methods that generate a stochastic gradient approximation ([13, 15, 21, 23, 25–27]) use the stochastic gradient $\mathbf{d}_\ell$ or a (singly or doubly) smoothed version of it to generate a search direction. For concreteness, we assume that the steepest ascent algorithm is used for optimization; although, this is not pertinent to the theoretical discussion. For the robust optimization problem of Eq. (2), we write the steepest ascent algorithm as

$$
\mathbf{u}_{\ell+1} = \mathbf{u}_\ell + \alpha_\ell \mathbf{d}_\ell
$$
(3)

for $\ell = 1, 2, \cdots$ until convergence where ℓ is the iteration index. Note that we take the liberty of referring to Eq. (3) as the steepest ascent algorithm even if $\mathbf{d}_\ell$ represents an approximation of the gradient premultiplied by a covariance matrix to obtain a preconditioned steepest ascent direction. In Eq. (3), $\mathbf{u}_1$ is the initial guess, and α_ℓ is the step-size obtained via a line search in the search direction $\mathbf{d}_\ell$ at iteration ℓ. In EnOpt, $\mathbf{u}$ is considered to be a random vector with a Gaussian distribution centered at $\mathbf{u}_\ell$ and covariance matrix $\mathbf{C}_{\mathbf{U}_\ell}$, that is, at iteration ℓ, $\mathbf{u} \sim \mathcal{N}(\mathbf{u}_\ell, \mathbf{C}_{\mathbf{U}_\ell})$. This notation indicates that it is possible to generate a new covariance matrix, for example, by covariance matrix adaptation [19, 29–31]. For our purposes, however, it suffices to consider the covariance matrix as fixed so we assume throughout that $\mathbf{u} \sim \mathcal{N}(\mathbf{u}_\ell, \mathbf{C}_{\mathbf{U}})$ at iteration ℓ.

At iteration ℓ of EnOpt, we generate N_p independent control perturbations (random samples), $\hat{\mathbf{u}}_{\ell,j}$, $j = 1, 2, \cdots N_p$ from the distribution $\mathcal{N}(\mathbf{u}_\ell, \mathbf{C}_{\mathbf{U}})$. The sample mean is denoted by $\overline{\hat{\mathbf{u}}_\ell}$ and given by the following equation:

$$
\overline{\hat{\mathbf{u}}_\ell} = \frac{1}{N_p} \sum_{j=1}^{N_p} \hat{\mathbf{u}}_{\ell,j}.
$$
(4)

At each iteration ℓ, Chen *et al.* [13] define $\overline{J(\mathbf{m}, \mathbf{u}_\ell)}$ by

$$
\overline{J(\mathbf{m}, \mathbf{u}_\ell)} \equiv \frac{1}{N_p} \sum_{j=1}^{N_p} J(\mathbf{m}_j, \hat{\mathbf{u}}_{\ell,j}).
$$
(5)

Note that this basic EnOpt formulation uses a 1:1 ratio of control perturbations $\hat{\mathbf{u}}_{\ell,j}$ to reservoir models $\mathbf{m}_j$, that is, $N_e = N_p$. Although Chen *et al.* did not give a specific interpretation of Eq. (5), the right side of Eq. (5) can be interpreted as an approximation of the mean of $J(\mathbf{m}, \mathbf{u}_\ell)$. With basic EnOpt, the search direction is then computed as

$$
\mathbf{d}_{\ell,ss-eno} = \frac{1}{N_p - 1} \sum_{j=1}^{N_p} (\hat{\mathbf{u}}_{\ell,j} - \overline{\hat{\mathbf{u}}_\ell})(J(\mathbf{m}_j, \hat{\mathbf{u}}_{\ell,j}) - \overline{J(\mathbf{m}, \mathbf{u}_\ell)}),
$$
(6)

where the meaning of the subscipts in $\mathbf{d}_{\ell,ss-eno}$ will be clarified shortly, and where the right side of Eq. (6) is the standard estimator of the cross covariance between $\mathbf{u}_\ell$ and $J(\mathbf{m}, \mathbf{u}_\ell)$. Chen *et al.* [13] indicate that the search direction in Eq. (6) provides a reasonable approximation to the 'singly-smoothed' or preconditioned true gradient $\mathbf{C}_{\mathbf{U}} \nabla_{\mathbf{u}} J_E(\mathbf{u}_\ell)$, where $J_E(\mathbf{u}_\ell) = E_{\mathbf{m}}[J(\mathbf{m}, \mathbf{u}_\ell)]$ is the approximation of the expectation of $J(\mathbf{m}, \mathbf{u}_\ell)$ with respect to $\mathbf{m}$ as defined by Eq. (1). The most commonly used EnOpt implementation is obtained by multiplying the preceding equation for the search direction by $\mathbf{C}_{\mathbf{U}}$ to obtain the 'standard' EnOpt search direction given by

$$\mathbf{d}_{\ell,ds-eno} = \mathbf{C_U}\left(\frac{1}{N_p - 1}\sum_{j=1}^{N_p}(\hat{\mathbf{u}}_{\ell,j} - \overline{\hat{\mathbf{u}}_\ell})(J(\mathbf{m}_j, \hat{\mathbf{u}}_{\ell,j}) - \overline{J(\mathbf{m}, \mathbf{u}_\ell)}))\right), \qquad (7)$$

which Chen *et al.* interpret as an approximation to the 'doubly smoothed' true gradient $\mathbf{C_U^2}\nabla_\mathbf{u}J_E(\mathbf{u}_\ell)$.

As there are several possible formulations for a search direction [23], we digress to define some notation. In the remainder of our paper, any optimization algorithm that uses a search direction obtained by some modification of Eq. (6) will be designated by using 'EnOpt' in its name. Whenever an EnOpt search direction approximates $\mathbf{C_U}$ times the gradient, its name will include 'singly-smoothed', and the associated search direction will include *ss* in its subscript. If the EnOpt search direction approximates $\mathbf{C_U^2}$ times the gradient, its name will include 'doubly-smoothed', and the associated search direction will include *ds* in its subscript. With this notation, the steepest ascent search direction that uses the formula of Eq. (6) to generate the search direction will be referred to as 'singly-smoothed EnOpt', or simply as 'ss-EnOpt'; sometimes, the terminology 'basic EnOpt' will also be used for this algorithm. The *eno* in the subscript refers to EnOpt, and the ℓ still indicates the iteration index. The steepest ascent algorithm that employs $\mathbf{C_U}$ times $\mathbf{d}_{\ell,ss-eno}$ as the search direction, that is, Eq. (7), will be referred to as 'doubly-smoothed EnOpt' or 'ds-EnOpt' (or 'standard EnOpt'), and the associated search direction will be denoted by $\mathbf{d}_{\ell,ds-eno}$. Later, we will introduce StoSAG algorithms. One way to define an StoSAG, search direction will be with a formula that, similar to Eq. (6), has the form of an approximate cross-covariance (cc) and where the search direction is an approximation of $\mathbf{C_U}$ times the gradient. Following the notation we introduced for EnOpt but adding 'cross-covariance' to its name, the method will be named 'singly-smoothed cross-covariance StoSAG' or simply 'ss-cc-StoSAG', and the search direction at iteration ℓ of the associated steepest ascent algorithm will be denoted by $\mathbf{d}_{\ell,ss-cc-sto}$, where *sto* in the subscript indicates StoSAG. Multiplying $\mathbf{d}_{\ell,ss-cc-sto}$ by $\mathbf{C_U}$ gives the search direction of 'doubly-smoothed cross-covariance StoSAG' (ds-cc-StoSAG) with the resulting search direction denoted by $\mathbf{d}_{\ell,ds-cc-sto}$. As shown later, by starting from the direction $\mathbf{d}_{\ell,ss-cc-sto}$ and using a matrix pseudo inverse, we can obtain the simplex gradient that provides an approximation of the true gradient. Using the simplex gradient as the search direction in steepest ascent therefore gives us another way to define a StoSAG search direction, which, at iteration ℓ, will be denoted by $\mathbf{d}_{\ell,sto}$. This search direction may also be premultiplied by $\mathbf{C_U}$ or $\mathbf{C_U^2}$, which leads to the names 'singly-smoothed StoSAG' (ss-StoSAG) and 'doubly-smoothed StoSAG' (ds-StoSAG) with corresponding search directions $\mathbf{d}_{\ell,ss-sto}$ and $\mathbf{d}_{\ell,ds-sto}$, respectively. Table I summarizes the various names and notations.

In an attempt to show that $\mathbf{d}_{\ell,ss-eno}$ (defined in Eq. (6)) indeed represents a reasonable approximation of $\mathbf{C_U}$ times the true gradient $\nabla_\mathbf{u}J_E(\mathbf{u}_\ell)$, Chen *et al.* [13] then make two assumptions, the first of which is that

$$\overline{\hat{\mathbf{u}}_\ell} = \frac{1}{N_p}\sum_{j=1}^{N_p}\hat{\mathbf{u}}_{\ell,j} \approx \mathbf{u}_\ell. \qquad (8)$$

Because the set $\{\hat{\mathbf{u}}_{\ell,j}\}_{j=1}^{N_p}$ are samples from $\mathcal{N}(\mathbf{u}_\ell, \mathbf{C}_{\mathbf{U}_\ell})$, $\overline{\hat{\mathbf{u}}_\ell}$ defined in Eq 4 is an unbiased, consistent estimator of $\mathbf{u}_\ell$, and therefore, Eq. (8) is a good approximation for N_p sufficiently large.

Table I. Names and notations of different fomulations.

Name	Alternate name	Short name	Search direction
singly-smoothed EnOpt	basic EnOpt	ss-EnOpt	$\mathbf{d}_{\ell,ss-eno}$
doubly-smoothed EnOpt	standard EnOpt	ds-EnOpt	$\mathbf{d}_{\ell,ds-eno}$
unsmoothed StoSAG		StoSAG	$\mathbf{d}_{\ell,sto}$
singly-smoothed StoSAG		ss-StoSAG	$\mathbf{d}_{\ell,ss-sto}$
doubly-smoothed StoSAG		ds-StoSAG	$\mathbf{d}_{\ell,ds-sto}$
singly-smoothed cross-covariance StoSAG		ss-cc-StoSAG	$\mathbf{d}_{\ell,ss-cc-sto}$
doubly-smoothed cross-covariance StoSAG		ds-cc-StoSAG	$\mathbf{d}_{\ell,ds-cc-sto}$

However, this approximation may actually be inaccurate if the upper and/or lower bounds on **u** force truncation of the samples. For example, suppose that the controls are simply the total liquid rates at producers and at iteration ℓ some of these controls are close to the lower bound of zero and some are close to the upper bound. Then a sample $\hat{\mathbf{u}}_{\ell,j}$ may have some negative entries, which have to be truncated to zero in order to run the simulator to evaluate the objective function value. Depending on the implementation and physics, it may also be necessary to truncate entries of $\hat{\mathbf{u}}_{\ell,j}$ that exceed the upper bound. In these circumstances, Eq. (8) becomes an invalid approximation because the truncated versions of the vectors $\hat{\mathbf{u}}_{\ell,j}$ are no longer samples from the distribution $\mathcal{N}(\mathbf{u}_\ell, \mathbf{C_U})$.

The second assumption of Chen *et al.* [13] is that for any **m**,

$$\overline{J(\mathbf{m}, \mathbf{u}_\ell)} = \frac{1}{N_p} \sum_{j=1}^{N_p} J(\mathbf{m}_j, \hat{\mathbf{u}}_{\ell,j}) \approx J(\mathbf{m}, \overline{\hat{\mathbf{u}}_\ell}) \approx J(\mathbf{m}, \mathbf{u}_\ell). \tag{9}$$

The assumption of Eq. (9) is tenuous as it suggests that for any realization $\mathbf{m}_j$ of **m**, $J(\mathbf{m}_j, \mathbf{u}^\ell) = \overline{J(\mathbf{m}, \mathbf{u}_\ell)}$ from which it follows that

$$J(\mathbf{m}_1, \mathbf{u}_\ell) \approx J(\mathbf{m}_2, \mathbf{u}_\ell) \approx \cdots J(\mathbf{m}_{N_e}, \mathbf{u}_\ell), \tag{10}$$

which is clearly an invalid approximation unless the variance in the prior model for the random vector **m** is sufficiently small, so that when the vector of controls $\mathbf{u}_\ell$ is applied to each of these models, the same value of the objective function is obtained. It is important to note that if we assume that Eqs. (9) and (10) hold, then, we must also have $\overline{J(\mathbf{m}, \mathbf{u}_\ell)} = J(\overline{\mathbf{m}}, \mathbf{u}_\ell)$, where $\overline{\mathbf{m}}$ can be taken as the mean of the $\mathbf{m}_j$'s or the mean of the underlying probability density function for $\overline{\mathbf{m}}$. However, we show in Section 6 that using this last result in basic EnOpt does not produce a correct approximation of the gradient of $\mathbf{C_U} \times J_E(\mathbf{u})$ defined in Eq. (1).

Stordal *et al.* [19] recently provided a theoretical analysis of the single and multiple reservoir formulations of standard EnOpt in which they demonstrate that both can be interpreted as special versions of a machine learning algorithm called 'Gaussian mutation optimization' [20]. Stordal *et al.* [19] also show that, although the multiple-reservoir version of EnOpt theoretically converges to an unbiased estimate of the true gradient, the use of a 1:1 ratio of control perturbations to geological models results in an approximate gradient with a much larger variance than the use of independent sets of perturbations for each geological realization.

4. THEORETICAL ARGUMENTS

Instead of relying on the potentially unreliable approximations of Eqs. (8) and (9), it would be far better to use a search direction that does not rely on these approximations. Specifically following Fonseca *et al.* [23, 26, 27], we define a basic foundational search direction by

$$\mathbf{d}_{\ell,ss-cc-sto} = \frac{1}{N_e} \sum_{k=1}^{N_e} \left(\frac{1}{N_p} \sum_{j=1}^{N_p} (\hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell)(J(\mathbf{m}_k, \hat{\mathbf{u}}_{\ell,j}) - J(\mathbf{m}_k, \mathbf{u}_\ell)) \right), \tag{11}$$

where N_p is the number of random control perturbations, and N_e is the number of random geological models, while we assume the two random variables to be mutually uncorrelated. As we will see more clearly later, the preceding equation means that at each iteration of steepest ascent, we use N_p random samples to generate an approximation to $\mathbf{C_U} \nabla_{\mathbf{u}} J(\mathbf{m}_k, \mathbf{u}_\ell)$. We could use different sets of N_p samples for each $\mathbf{m}_k$, or, for computational efficiency, we can use the same N_p samples to generate the approximation of $\mathbf{C_U} \nabla_{\mathbf{u}} J(\mathbf{m}_k, \mathbf{u}_\ell)$ for all k. If we use a different set of control perturbations for each $\mathbf{m}_k$, $k = 1, 2, \cdots N_e$, then, we have to change notation slightly; that is, we let $\hat{\mathbf{u}}_{\ell,j,k}$, $k = 1, 2, \cdots N_e$ denote the set samples used to generate an approximate gradient of $J(\mathbf{m}_k, \mathbf{u}_\ell)$. In this case, Eq. (11) should be replaced by

$$\mathbf{d}_{\ell,ss-cc-sto} = \frac{1}{N_e} \sum_{k=1}^{N_e} \left(\frac{1}{N_p} \sum_{j=1}^{N_p} (\hat{\mathbf{u}}_{\ell,j,k} - \mathbf{u}_\ell)(J(\mathbf{m}_k, \hat{\mathbf{u}}_{\ell,j,k}) - J(\mathbf{m}_k, \mathbf{u}_\ell)) \right). \tag{12}$$

Using the naming conventions introduced earlier, we refer to the search direction of Eq. (12) (or Eq. (11)) as the singly-smoothed cross-covariance StoSAG (ss-cc-StoSAG) search direction or gradient. Just like the ss-EnOpt search direction defined in Eq. (6), the ss-cc-StoSAG search direction can be interpreted as an approximation to the singly-smoothed true gradient $\mathbf{C_U}\nabla_{\mathbf{u}} J_E(\mathbf{u}_\ell)$. Assuming all second derivatives of J are continuous and bounded, the error in the first-order Taylor series

$$J(\mathbf{m}_k, \hat{\mathbf{u}}_{\ell,j,k}) = J(\mathbf{m}_k, \mathbf{u}_\ell) + (\nabla_{\mathbf{u}} J(\mathbf{m}_k, \mathbf{u}_\ell))^T (\hat{\mathbf{u}}_{\ell,j,k} - \mathbf{u}_\ell) \tag{13}$$

is given by

$$e_{j,k} = \frac{1}{2}(\hat{\mathbf{u}}_{\ell,j,k} - \mathbf{u}_\ell)^T \mathbf{H}_{k,\ell}(\hat{\mathbf{u}}_{\ell,j,k} - \mathbf{u}_\ell) = \mathcal{O} \| \hat{\mathbf{u}}_{\ell,j,k} - \mathbf{u}_\ell \|_2^2, \tag{14}$$

where $\mathbf{H}_{k,\ell}$ denotes the Hessian matrix given by

$$\mathbf{H}_{k,\ell} = \nabla_{\mathbf{u}}[(\nabla_{\mathbf{u}} J(\mathbf{m}_k, \mathbf{u}_\ell))^T]. \tag{15}$$

Using Eq. (13) in Eq. (12) gives the approximation

$$\mathbf{d}_{\ell,ss-cc-sto} = \frac{1}{N_e} \sum_{k=1}^{N_e} \left(\frac{1}{N_p} \sum_{j=1}^{N_p} (\hat{\mathbf{u}}_{\ell,j,k} - \mathbf{u}_\ell)(\hat{\mathbf{u}}_{\ell,j,k} - \mathbf{u}_\ell)^T \nabla_{\mathbf{u}} J(\mathbf{m}_k, \mathbf{u}_\ell) \right). \tag{16}$$

We let $E_{\mathbf{u}}$ denote the expectation with respect to $\mathbf{u}$. Then, taking the expectation of Eq. (16) assuming that all $\hat{\mathbf{u}}_{\ell,j,k}$ are samples from $\mathcal{N}(\mathbf{u}_\ell, \mathbf{C_U})$ gives

$$E_{\mathbf{u}}[\mathbf{d}_{\ell,ss-cc-sto}] = \frac{1}{N_e} \sum_{k=1}^{N_e} \left[\frac{1}{N_p} \sum_{j=1}^{N_p} \left(E_{\mathbf{u}}[(\hat{\mathbf{u}}_{\ell,j,k} - \mathbf{u}_\ell)(\hat{\mathbf{u}}_{\ell,j,k} - \mathbf{u}_\ell)^T] \right) \nabla_{\mathbf{u}} J(\mathbf{m}_k, \mathbf{u}_\ell) \right]$$
$$= \mathbf{C_U} \left(\frac{1}{N_e} \sum_{k=1}^{N_e} \nabla_{\mathbf{u}} J(\mathbf{m}_k, \mathbf{u}_\ell) \right) = \mathbf{C_U}\nabla J_E(\mathbf{u}_\ell). \tag{17}$$

For later reference, we note that the error in Eqs. (16) and (17) is $\mathcal{O}(\max_{j,k} \| \mathbf{u}_{\ell,j,k} - \mathbf{u}_\ell \|^2)$, where all norms considered in the paper refer to the ℓ_2 norm. Recall that the motivation of the Chen *et al*. EnOpt formulation is to first obtain a search direction that approximates $\mathbf{C_U}\nabla J_E(\mathbf{u}_\ell)$ and then to provide additional smoothing by multiplying this search direction by $\mathbf{C_U}$; see the discussion centered on Eqs. (9–16) in [13]. It is imperative to note that we have obtained a search direction, which approximates $\mathbf{C_U}\nabla J_E(\mathbf{u}_\ell)$ without invoking the potentially invalid assumptions of Chen *et al*. represented by Eqs. (8) and (9).

It is also important to note that one of the advantages of ss-EnOpt (Eq. (6)) used by Chen *et al*. [13] is its computational efficiency, which arises from the fact that Eq. (6) uses only one perturbation of the control vector $\mathbf{u}_\ell$ per model realization so that each application of Eq. (6) requires exactly $N_e = N_p$ evaluations of $J(\mathbf{m}, \mathbf{u})$ and hence requires only N_e simulation runs per application. On the other hand, the more theoretically sound search direction of Eq. (12) requires $N_p N_e$ evaluations of $J(\mathbf{m}, \mathbf{u})$ ($N_p N_e$ simulation runs) per application. However, one can select $N_p = 1$ when applying Eq. (12), and in this case, ss-cc-StoSAG has the same computational efficiency as ss-EnOpt. Note that Eq. (17) still holds for $N_p = 1$. Moreover, Fonseca *et al*. [26, 27] show that Eq. (12) with $N_p = 1$ leads to a methodology that results in a higher objective function value than is achievable with ss-EnOpt; in [23], they show that far more accurate estimates of $\mathbf{C_U}\nabla J_E(\mathbf{u}_\ell)$ can be obtained

by using $N_p >> 1$. In the example considered here, we use $N_p = 1$ to obtain StoSAG algorithms that require the same computational effort as ss-EnOpt (or ds-EnOpt) to generate a search direction at each iteration of the steepest ascent algorithm.

5. COMMENTS ON STOCHASTIC SIMPLEX APPROXIMATE GRADIENT

Another interesting aspect of the work of Fonseca *et al.* [23, 26, 27] is that they observe that higher objection function values can be obtained if in place of using Eq. (12) to form a steepest ascent search direction, one instead generates an approximation for $\nabla J_E(\mathbf{u}_\ell)$ to use as a search direction. One can also multiply this approximate gradient by $\mathbf{C_U}$ or $\mathbf{C_U^2}$ to obtain a search direction. While we are not certain that this observation is general, we believe their suggestion may be important when the bounds on $\mathbf{u}$ force the truncation of some components of a sample $\hat{\mathbf{u}}_{\ell,j} \sim \mathcal{N}(\mathbf{u}_\ell, \mathbf{C_U})$. Their suggestion may also be important, whenever it is unclear what correlation length to choose in a covariance function in order to generate the covariance matrix $\mathbf{C_U}$ or when the optimal solution of the well control problem is such that the values of two temporally consecutive controls at a well are radically different, as occurs, for example, when the optimal solution is bang-bang [7, 32, 33]. When the vectors $\hat{\mathbf{u}}_{\ell,j,k}$ used in Eq. (12) are truncated,

$$E_{\mathbf{u}}\left[\frac{1}{N_p}\sum_{j=1}^{N_p}(\hat{\mathbf{u}}_{\ell,j,k} - \mathbf{u}_\ell)(\hat{\mathbf{u}}_{\ell,j,k} - \mathbf{u}_\ell)^T\right] \neq \mathbf{C_U}, \tag{18}$$

and Eq. (17) is no longer valid. However, one can still obtain an approximation of $\mathbf{C_U}\nabla J_E(\mathbf{u}_\ell)$ by first using a simplex gradient [15, 18, 23, 26, 27, 34] to approximate $\nabla J_E(\mathbf{u}_\ell)$ and then multiplying this gradient by $\mathbf{C_U}$. For information on the simplex gradient formulation in general see, for example, [18]. Specifically, we define the $N_u \times N_p$ matrix $\Delta\hat{\mathbf{U}}_{\ell,k}$ and the $N_p \times 1$ vector $\Delta\mathbf{j}_{\ell,k}$, respectively, by

$$\Delta\hat{\mathbf{U}}_{\ell,k} = \left[\, \hat{\mathbf{u}}_{\ell,1,k} - \mathbf{u}_\ell \ \ \hat{\mathbf{u}}_{\ell,2,k} - \mathbf{u}_\ell \cdots \hat{\mathbf{u}}_{\ell,N_p,k} - \mathbf{u}_\ell \,\right], \tag{19}$$

and

$$\Delta\mathbf{j}_{\ell,k} = \left[\, J(\mathbf{m}_k, \hat{\mathbf{u}}_{\ell,1,k}) - J(\mathbf{m}_k, \mathbf{u}_\ell) \ \ J(\mathbf{m}_k, \hat{\mathbf{u}}_{\ell,2,k}) - J(\mathbf{m}_k, \mathbf{u}_\ell) \ \cdots \ J(\mathbf{m}_k, \hat{\mathbf{u}}_{\ell,N_p,k}) - J(\mathbf{m}_k, \mathbf{u}_\ell) \,\right]^T, \tag{20}$$

for $k = 1, 2, \cdots N_e$. Then using the first order Taylor series of Eq. (13) and the definitions of Eqs. (19) and (20), it easily follows that for each k, $k = 1, 2, \cdots N_e$,

$$\sum_{j=1}^{N_p}(\hat{\mathbf{u}}_{\ell,j,k} - \mathbf{u}_\ell)(J(\mathbf{m}_k, \hat{\mathbf{u}}_{\ell,j,k}) - J(\mathbf{m}_k, \mathbf{u}_\ell)) = \Delta\hat{\mathbf{U}}_{\ell,k}(\Delta\mathbf{j}_{\ell,k})$$
$$\approx \left(\Delta\hat{\mathbf{U}}_{\ell,k}(\Delta\hat{\mathbf{U}}_{\ell,k})^T\right)\nabla_{\mathbf{u}}J(\mathbf{m}_k, \mathbf{u}_\ell). \tag{21}$$

Thus, as in [15, 23, 26, 27, 34], an approximate stochastic gradient, denoted by $\nabla_{\mathbf{u},\text{sto}}J(\mathbf{m}_k, \mathbf{u}_\ell)$, is given by

$$\nabla_{\mathbf{u},\text{sto}}J(\mathbf{m}_k, \mathbf{u}_\ell) = \left(\Delta\hat{\mathbf{U}}_{\ell,k}(\Delta\hat{\mathbf{U}}_{\ell,k})^T\right)^+ \Delta\hat{\mathbf{U}}_{\ell,k}(\Delta\mathbf{j}_{\ell,k}) \approx \nabla_{\mathbf{u}}J(\mathbf{m}_k, \mathbf{u}_\ell), \tag{22}$$

for $k = 1, 2, \cdots N_e$ where the superscript $+$ on a matrix denotes the Moore–Penrose pseudo-inverse. By the properties of the Moore–Penrose pseudo-inverse,

$$((\Delta\hat{\mathbf{U}}_{\ell,k})^+)^T = ((\Delta\hat{\mathbf{U}}_{\ell,k})^T)^+ = \left(\Delta\hat{\mathbf{U}}_{\ell,k}(\Delta\hat{\mathbf{U}}_{\ell,k})^T\right)^+ \Delta\hat{\mathbf{U}}_{\ell,k}, \tag{23}$$

and thus, Eq. (22) is equivalent to

$$\nabla_{\mathbf{u},\text{sto}}J(\mathbf{m}_k, \mathbf{u}_\ell) = ((\Delta\hat{\mathbf{U}}_{\ell,k})^+)^T(\Delta\mathbf{j}_{\ell,k}) \approx \nabla_{\mathbf{u}}J(\mathbf{m}_k, \mathbf{u}_\ell), \tag{24}$$

where the pseudo inverse of $(\Delta\hat{\mathbf{U}}_{\ell,k})$ can be obtained by singular value decomposition [35]. Except for the fact that this gradient is stochastic because the $\hat{\mathbf{u}}$'s are random vectors; the term $\nabla_{\mathbf{u},\text{sto}}J(\mathbf{m}_k,\mathbf{u}_\ell)$ is simply the simplex gradient of J at $(\mathbf{m}_k,\mathbf{u}_\ell)$ [18]. It is important to note that the approximations of Eqs. (21) and (24) apply regardless of how the perturbations around $\mathbf{u}_\ell$ are generated. In particular, there is no need to assume that they are generated from a Gaussian distribution. In fact, they could be generated by a Sobol sequence as recently suggested in [36]. Then to obtain smoothing, one could multiply $\nabla_{\mathbf{u},\text{sto}}J(\mathbf{m}_k,\mathbf{u}_\ell)$ defined in Eq. (22) by a true covariance matrix $\mathbf{C_U}$, or by $\mathbf{C_U^2}$ to obtain double smoothing. With the simplex gradient of Eq. (22), the search direction for maximizing the expected objective function value as approximated by Eq. (1) is given by

$$\mathbf{d}_{\ell,sto} = \nabla_{\mathbf{u},sto}J_E(\mathbf{u}_\ell) = \frac{1}{N_e}\sum_{k=1}^{N_e}\nabla_{\mathbf{u},sto}J(\mathbf{m}_k,\mathbf{u}_\ell). \tag{25}$$

Throughout, we refer to the case where just $\mathbf{d}_{\ell,sto}$ is used as the search direction as unsmoothed StoSAG, or, sometimes, simply StoSAG. The algorithms that use $\mathbf{C_U}\mathbf{d}_{\ell,sto}$ and $\mathbf{C_U^2}\mathbf{d}_{\ell,sto}$, respectively, as the search direction in the steepest ascent algorithm are referred to as singly-smoothed StoSAG (ss-StoSAG) and doubly-smoothed StoSAG (ds-StoSAG), respectively. Note that, like ss-cc-StoSAG, ss-StoSAG uses a search direction that is an approximation of $\mathbf{C_U}$ times the true gradient, but the two search directions are distinctly different. Additional smoothing of the gradient as in ds-EnOpt, ds-cc-StoSAG, or ds-StoSAG can provide controls that vary more smoothly with time, which are simpler to impose in a real operational setting. However, there are two potential negative consequences to smoothing. First, one does not know how much temporal smoothing to impose a priori and too much temporal smoothing can yield suboptimal results if the true optimal controls for a well vary significantly from one control step to the next. Such a large change occurs when the optimal solution is bang-bang, [7, 32, 33, 37], that is, the optimal rate control at a well changes from its upper bound on one control step to zero on the next control step; see the first example of [37] for an illustration. The second problem occurs when a perturbed control generated to calculate the stochastic gradient is outside one of its bounds, for example, is less than zero for a rate or an (ICV) setting. In this case, the perturbed control must be truncated back to its bounds. If perturbations are truncated back to bounds, the expectations of the ss-EnOpt, ds-EnOpt ss-cc-StoSAG, and ds-cc-StoSAG search directions no longer approximate $\mathbf{C_U}$ or $\mathbf{C_U^2}$ times the true gradient. In addition to properly accounting for truncation of samples from $\mathcal{N}(\mathbf{u}_\ell,\mathbf{C_U})$, unsmoothed StoSAG also avoids imposing too much temporal smoothness on the controls at a well for a case where the true optimal controls vary significantly from one control step to the next.

6. A CLOSER LOOK AT ENSEMBLE OPTIMIZATION

As noted in Section 3, the generally invalid assumption that Eq. (9) holds implies that

$$\overline{J(\mathbf{m},\mathbf{u}_\ell)} = J(\overline{\mathbf{m}},\mathbf{u}_\ell), \tag{26}$$

where $\overline{\mathbf{m}}$ is the mean of the $\mathbf{m}_j$'s or the mean of the underlying pdf for $\mathbf{m}$. Using Eq. (26) and the assumption that Eq. (8) holds in the ss-EnOpt formula of Eq. (6) gives

$$\mathbf{d}_{\ell,ss-eno-mod} = \frac{1}{N_p}\sum_{j=1}^{N_p}(\hat{\mathbf{u}}_{\ell,j}-\mathbf{u}_\ell)(J(\mathbf{m}_j,\hat{\mathbf{u}}_{\ell,j})-J(\overline{\mathbf{m}},\mathbf{u}_\ell)), \tag{27}$$

where the additional subscript mod indicates 'modified'. It is important to note that we do not actually have to assume that Eqs. (8) and (9) hold. Instead, following [15] for the fixed reservoir model case and [23, 26, 27] for the case with geological uncertainty, we can simply define the search direction of Eq. (27) directly and then try to show that the resulting search direction or its expectation approximates the true gradient or a smoothed version of the true gradient. In Eq. 27, we have replaced the factor $1/(N_p-1)$ in the original basic EnOpt formulation by $1/N_p$ because we no longer approximate the mean of the distribution for $\mathbf{u}$.

We assume that all derivatives of J are continuous and bounded. Then a second order Taylor series terms gives

$$J(\mathbf{m}_j, \hat{\mathbf{u}}_{\ell,j}) - J(\bar{\mathbf{m}}, \mathbf{u}_\ell) = (\nabla_{\mathbf{u}} J(\bar{\mathbf{m}}, \mathbf{u}_\ell))^T (\hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell) + (\nabla_{\mathbf{m}} J(\bar{\mathbf{m}}, \mathbf{u}_\ell))^T (\mathbf{m}_j - \bar{\mathbf{m}})$$
$$+ \frac{1}{2} (\hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell)^T \mathbf{H}_{\mathbf{u},j} (\hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell) + e_{1,j} \quad (28)$$

for $j = 1, 2, \cdots N_p$ with $e_{1,j}$ given by

$$e_{1,j} = (\hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell)^T \left(\nabla_{\mathbf{m}} [(\nabla_{\mathbf{u}} J)^T] \right)^T (\mathbf{m}_j - \bar{\mathbf{m}}) + \frac{1}{2} (\mathbf{m}_j - \bar{\mathbf{m}})^T \mathbf{H}_{\mathbf{m},j} (\mathbf{m}_j - \bar{\mathbf{m}}), \quad (29)$$

where $\mathbf{H}_{\mathbf{u},j}$ is the Hessian with respect to $\mathbf{u}$, $\mathbf{H}_{\mathbf{m},j}$ is the Hessian with respect to $\mathbf{m}$ and these Hessians and all other second degree terms in the previous two equations are evaluated at $(\bar{\mathbf{m}}^T, \mathbf{u}_\ell^T)^T$. For simplicity, we simply neglect terms higher than the second order as the second order terms asymptotically represent the largest error terms neglected in the various stochastic search directions considered. Substituting Eq. (28) into Eq. (27), it follows easily that

$$\mathbf{d}_{\ell,ss-eno-mod} = \left(\frac{1}{N_p} \sum_{j=1}^{N_p} (\hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell)(\hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell)^T \right) \nabla_{\mathbf{u}} J(\bar{\mathbf{m}}, \mathbf{u}_\ell) + \mathbf{e}_2 + \mathbf{e}_3, \quad (30)$$

where $\mathbf{e}_2$ and $\mathbf{e}_3$, respectively, are given by

$$\mathbf{e}_2 = \frac{1}{2N_p} \sum_{j=1}^{N_p} \left[(\hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell) \left((\hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell)^T \mathbf{H}_{\mathbf{u},j} (\hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell) \right) \right] \quad (31)$$

and

$$\mathbf{e}_3 = \frac{1}{N_p} \sum_{j=1}^{N_p} \left[(\hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell) e_{1,j} \right]. \quad (32)$$

The error $\mathbf{e}_2$ satisfies $\| \mathbf{e}_2 \| = \mathcal{O}(\max_j \{ \| \hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell \|^3 \})$ and thus goes to zero as the magnitude of the perturbations goes to zero, that is, as $\hat{\mathbf{u}}_{\ell,j} \to \mathbf{u}_\ell$. However, the error term $\mathbf{e}_3$ contains terms that behave like $\mathcal{O}(\| \hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell \| \| \mathbf{m}_j - \bar{\mathbf{m}} \|^2)$. Because the magnitudes of the $\| \mathbf{m}_j - \bar{\mathbf{m}} \|$ terms are fixed by the uncertainty in the prior model for $\mathbf{m}$ and cannot be changed, $\| \mathbf{e}_3 \| = \mathcal{O}(\max_j \{ \| \hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell \| \})$ and thus goes to zero more slowly than $\| \mathbf{e}_2 \|$ or the StoSAG error , both of which are third order in $\max_j \{ \| \hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell \| \}$; see the comment following Eq. (17) for the magnitude of the term neglected in the StoSAG direction given in Eq. (16).

Following a line of reasoning that Chen originally pursued to justify basic EnOpt in her Ph.D. thesis [38], we note that for each j, $\mathbf{m}_j$, and $\hat{\mathbf{u}}_{\ell,j}$ are independent random variables and thus taking the expectation ($E[.]$) of $\mathbf{e}_3$ with respect to the joint distribution gives $E[\mathbf{e}_3] = \mathbf{0}$, so the expectation of Eq. (30) is given by

$$E[\mathbf{d}_{\ell,ss-eno-mod}] \approx E \left[\left(\frac{1}{N_p} \sum_{j=1}^{N_p} (\hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell)(\hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell)^T \right) \right] \nabla J(\bar{\mathbf{m}}, \mathbf{u}_\ell), \quad (33)$$

or, for the case where $\hat{\mathbf{u}}_{\ell,j} \sim \mathcal{N}(\mathbf{u}_\ell, \mathbf{C_U})$,

$$E[\mathbf{d}_{\ell,ss-eno-mod}] \approx \mathbf{C_U} \nabla_{\mathbf{u}} J(\bar{\mathbf{m}}, \mathbf{u}_\ell), \quad (34)$$

where the bias in Eqs. (33) and (34) is $\mathcal{O}(\max_j \{ \| \hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell \|^3 \})$. Thus, in expectation, the error in the modified EnOpt search direction of Eq. (27) is $\mathcal{O}(\max_j \{ \| \hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell \|^3 \})$, which is of the same order as the StoSAG error. However, because of the $e_{1,j}$ terms that appear in $\mathbf{e}_3$, the autocovariance

of $\mathbf{e}_3$, $E[(\mathbf{e}_3 - E[\mathbf{e}_3])(\mathbf{e}_3 - E[\mathbf{e}_3])^T]$ may have large entries, so the variances of the components of $\mathbf{e}_3$ may be large; note that this last observation is consistent with the results of [19].

As a tangential aside, the fact that errors in the stochastic search directions involve the term $\mathcal{O}(\max_j \{\| \hat{\mathbf{u}}_{\ell,j} - \mathbf{u}_\ell \|\})$ provides a motivation for choosing a covariance $\mathbf{C}_{\mathbf{U}}^\ell$ for $\mathbf{u}$ such that as ℓ increases, the variances in $\mathbf{C}_{\mathbf{U}}^\ell$ decrease, or more precisely, such that $\det \mathbf{C}_{\mathbf{U}}^\ell$ decreases because the confidence intervals associated with samples of $\mathcal{N}(\mathbf{u}_\ell, \mathbf{C}_{\mathbf{U}})$ are proportional to $\det \mathbf{C}_{\mathbf{U}}^\ell$; [39, 40]. In fact, a decreasing perturbation size is necessary to prove convergence of one of the earliest theoretically sound algorithms that use a stochastic gradient, namely, the original SPSA algorithm [16, 17].

7. NUMERICAL EXAMPLES

7.1. Optimal well control problem

The first example that we consider in this paper is to estimate simultaneously the well controls (pressures and/or flow rates) and downhole ICV settings in a three-phase (oil-gas-water) reservoir to maximize the expectation of an economic objective function in a water-alternating-gas (WAG) enhanced oil recovery (EOR) process. (ICVs allow for the indvidual inflow control into a well from different subsurface layers). For such an application, we let the vector $\mathbf{m}$ represent the flow-relevant rock properties of the reservoir simulation gridblocks. The objective function for a specific reservoir model, $\mathbf{m}_i$, is defined by

$$J_i(\mathbf{m}_i, \mathbf{u}) = \sum_{n=1}^{N_t} \left\{ \frac{\Delta t_n}{(1+b)^{\frac{t_n}{365}}} \left[\sum_{j=1}^{N_P} \left(r_o \cdot \overline{q_{o,j}^n} - c_w \cdot \overline{q_{w,j}^n} \right) - \sum_{k=1}^{N_I} \left(c_{wi} \cdot \overline{q_{wi,k}^n} + c_{gi} \cdot \overline{q_{gi,k}^n} \right) \right] \right\}, \tag{35}$$

for $i = 1, 2, \cdots N_e$. Here, $\mathbf{u}$ is an N_u-dimensional column vector, which contains all the well controls (ICV settings and rates or pressures) at all wells over the production lifetime; n denotes the nth time step of the reservoir simulator; N_t is the total number of time steps; the time at the end of the nth time step is denoted by t_n; Δt_n is the nth time step size; b is the annual discount rate; N_P and N_I denote the number of producers and injectors, respectively; r_o is the oil revenue, in \$/STB (where STB indicates 'stock tank barrel', i.e., 5.61 ft^3 at 60 deg. F and 14.7 psi); c_w, c_{wi}, and c_{gi}, respectively, denote the disposal cost of produced water and the cost of water injection, in units of \$/STB, and the cost of gas injection, in units of \$/Mscf (thousand standard cubic feet, i.e., at 60 deg. F and 14.7 psi); $\overline{q_{o,j}^n}$ and $\overline{q_{w,j}^n}$, respectively, denote the average oil production rate and the average water production rate at the jth producer for the nth time step, in units of STB/day; $\overline{q_{wi,k}^n}$ and $\overline{q_{gi,k}^n}$, respectively, denote the average water injection rate and the average gas injection rate at the kth injector for the nth time step, in units of STB/day. For the three-phase example considered in this work, we neglect the revenue from hydrocarbon gas production and the disposal cost of injected gas (CO_2) that is produced. The economic objective function as defined in Eq. (35) is typically referred to as the Net Present Value (NPV) in the reservoir optimization literature.

The second example considered in this paper pertains to estimating the well controls in a two-phase (oil–water) reservoir to maximize the expectation of the NPV in a 'water flooding' production scenario. Although the physics in this example are somewhat simpler and do not involve hydrocarbon phase-equilibrium computations as required to simulate WAG, the reservoir structure is more complex, and the numerical model contains a much larger number of grid blocks.

7.2. Example 1: Water alternating gas flooding optimization

We consider the optimization of WAG flooding for a 3D synthetic reservoir where the injected gas is carbon dioxide, CO_2. The reservoir simulation model is on a $25 \times 25 \times 3$ grid with grid block dimensions given by $\Delta x = \Delta y = 100$ ft and $\Delta z = 30$ ft. For robust optimization, we generate 15 geological realizations (models) to represent the uncertainty in the reservoir description. Figure 1 shows six realizations of the log-permeability distribution for the first layer. Permeability is the

(a) Realization 2   (b) Realization 4   (c) Realization 6

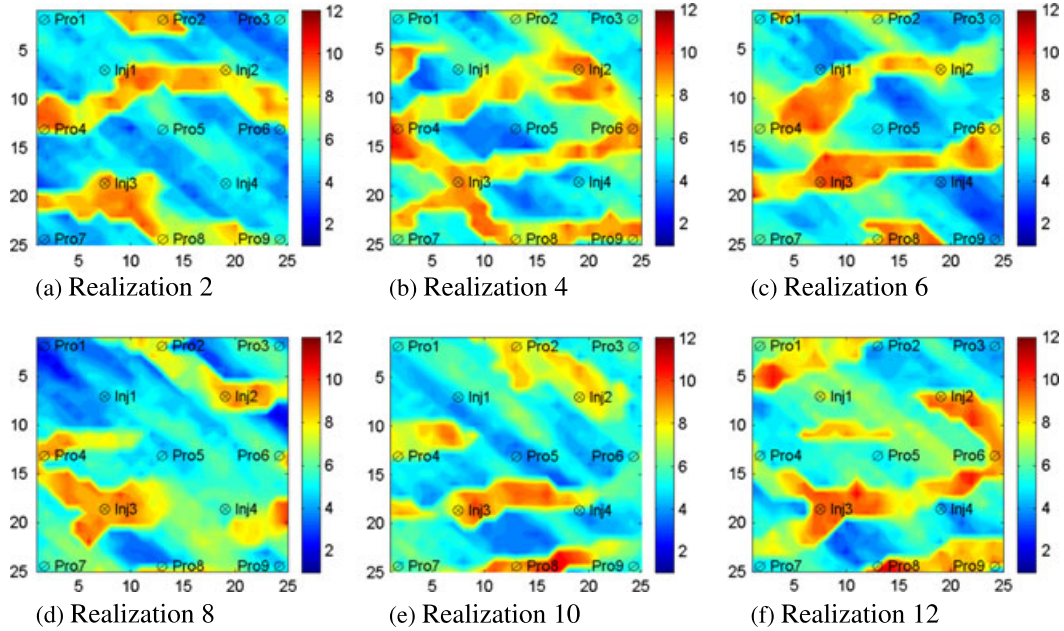(d) Realization 8   (e) Realization 10   (f) Realization 12

Figure 1. Log-permeability distribution of six realizations for the first layer.

inverse resistance to flow of the porous medium expressed in units of mD = 1.06×10^{-14}ft^2. Permeability values are usually highly uncertain and may vary over the reservoir with several orders of magnitude. The set of natural logarithms of all grid block permeability values constitutes the vector **m**, which therefore has $25 \times 25 \times 3 = 1875$ elements. The second and third layers have the same heterogenity features as the first layer but with the permeability field of layer 2 equal to the permeability field of layer 1 multiplied by 0.6; the permeability field of layer 3 is equal to the permeability field of layer 1 multiplied by 0.3. The reservoir contains four injection wells and nine producers. All wells are vertical fully penetrating wells (i.e., being open to flow over the entire reservoir height), and their locations are shown in the realizations of the permeability fields that are displayed in Figure 1. Each of the wells is equipped with an ICV in each of the three layers. The porosity is homogeneous with $\phi = 0.2$. No capillary pressures are included and the reservoir rock is assumed to be incompressible. The initial reservoir pressure is 4500 psi at the top of the reservoir. The producing reservoir lifetime is set equal to 2880 days. A compositional reservoir simulator (i.e., a simulator capable of performing phase equilibrium calculations for complex hydrocarbon mixtures), GEM (Computer Modeling Group, Calgary, Canada) (Version 2009.10), is used for reservoir simulation [41].

Because of the density difference between water, oil, and CO_2, phase segregation will occur resulting in 'over ride' and 'under ride' of the injected CO_2 and water, respectively, resulting in early 'break through' of these undesirable fluids in the production wells. The ICVs provide the capabilities to partly counteract these 'gravity segration' effects, and the physical mechanisms that drive the optimization are therefore both the heterogeneity in the permeability and density difference between the reservoir fluids. To optimize $J_E(u)$, the oil price is set equal to \$ 80.0/STB; the water injection cost is \$5.0/STB; the gas injection cost is \$1.5/Mscf; the disposal cost of produced water is \$5.0/STB; the annual discount rate is 0.1. Here, we neglect the revenue of hydrocarbon gas production and the cost of disposing of injected CO_2 gas that is produced. Within the producing lifetime, we can set each of the well control or ICV settings 16 times in control intervals of $2880/16 = 180$ days. We let $\mathbf{u}^o$ denote the n_0-dimensional column vector ($n_o = (9 + 4) \times 16 = 208$) of all well controls (rates or pressures), and $\mathbf{u}^v$ denote the n_v-dimensional column vector of ICV settings ($n_v = (9 + 4) \times 3 \times 16 = 624$) so that the vector of all control settings is given by

$$\mathbf{u} = \begin{bmatrix} \mathbf{u}^o \\ \mathbf{u}^v \end{bmatrix}. \tag{36}$$

In the EnOpt and StoSAG algorithms, we generate perturbations from a covariance matrix. This covariance matrix is block diagonal and has the form

$$
\mathbf{C_U} = \begin{bmatrix} \mathbf{C_U^o} & \mathbf{O} \\ \mathbf{O} & \mathbf{I} \end{bmatrix}, \tag{37}
$$

where the $\mathbf{O}$'s denote null matrices, and $\mathbf{I}$ is the $n_v \times n_v$ identity matrix. The matrix $\mathbf{C_U^o}$ is a block diagonal matrix, where each kth matrix on the diagonal block is denoted by $\mathbf{C_{U_k^o}}$ and is associated with the subvector $\mathbf{u}_k^o$ of $\mathbf{u}^0$ that contains all the well controls for well k. Here, to impose some temporal smoothness on the the well controls (rates or pressures) for well k, $\mathbf{C_{U_k^o}}$ is generated from a spherical covariance function [42] with variance equal to $(0.1)^2$ and time correlation length equal to 720 days, which is equivalent to four control steps. However, gas injection rates are not correlated with water injection rates, and so effectively, when sampling from $\mathcal{N}(\mathbf{u}_\ell, \mathbf{C_U})$, only two consecutive water rates are correlated, and only two consecutive gas rates are correlated. The appearance of the identity matrix in Eq. (37) indicates that there is no correlation between ICV settings. In this example, five types of search directions (out of the seven listed in Table I) are used to estimate the optimal well controls and ICV settings that maximize $J_E(\mathbf{u})$:

(1) ds-EnOpt, $\mathbf{d}_{\ell,ds-eno}$ defined in Eq. (7);
(2) StoSAG, $\mathbf{d}_{\ell,sto}$, defined in Eq. (25);
(3) ss-StoSAG, $\mathbf{d}_{\ell,ss-sto} = \mathbf{C_U} \mathbf{d}_{\ell,sto}$;
(4) ss-cc-StoSAG, $\mathbf{d}_{\ell,ss-cc-sto}$ defined in Eq. (12), with $N_p = 1$;
(5) ds-cc-StoSAG, $\mathbf{d}_{\ell,ds-cc-sto} = \mathbf{C_U} \mathbf{d}_{\ell,ss-cc-sto}$.

We use a steepest ascent optimization algorithm with simple backtracking as line search. At iteration ℓ of the steepest ascent, we generate $N_e = 15$ samples of the Gaussian random control vector $\mathbf{u}$ with each sample corresponding to one geological model, where $\hat{\mathbf{u}}_{\ell,j} \sim \mathcal{N}(\mathbf{u}_\ell, \mathbf{C_U})$. In steepest ascent, the initial step size is 1.0, and the maximum number of step-size cuts (each with a factor 0.5) is set equal to 5. If five steps sizes do not yield an increase in J_E, we set the new estimate, $\mathbf{u}_{\ell+1}$ of the optimal $\mathbf{u}$ equal to the $\mathbf{u}$ that gave the largest value of J_E obtained during the line search. We allow a maximum of 4000 total reservoir simulation runs. For this relatively small example, the simulations took about 3.5 minutes per simulation run so that the entire optimization procedure could be performed sequentially. For more realistic, large-scale reservoir simulations, which may have run times in the order of several hours, parallel processing will be a practical necessity. Fortunately, the StoSAG algorithms, just like the EnOpt algorithms, are embarrassingly parallel.

Figure 2 shows the expectation of NPV (Eq.(1)) versus the number of reservoir simulation runs using steepest ascent with the five different search directions enumerated in the aforementioned text.


Figure 2. Expectation of Net Present Value versus number of simulation runs for different search direction formulations; yellow: ds-EnOpt (standard); green: StoSAG; dark blue: ss-StoSAG; purple: ss-cc-StoSAG; red: ds-cc-StoSAG.

(a) StoSAG

(b) ds-cc-StoSAG

Figure 3. Estimated well controls for producers at different control steps calculated from StoSAG and ds-cc-StoSAG; units are psi.
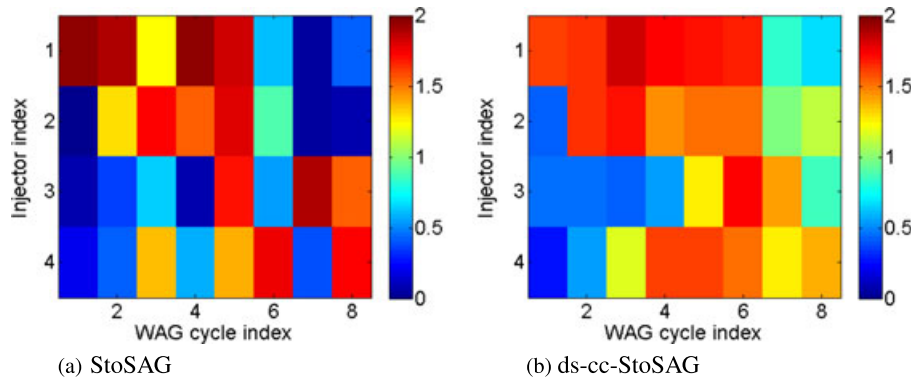

(a) StoSAG

(b) ds-cc-StoSAG

Figure 4. Estimated well controls for gas injectors at different water-alternating-gas cycles calculated from StoSAG and ds-cc-StoSAG; units are $\times 10^7$ scf/day.

It is clear that ds-EnOpt (standard EnOp) generates the lowest expected NPV while the StoSAG search direction proposed by Fonseca *et al.* [23, 26] and ds-cc-StoSAG generate the highest average NPV. Note that although both ds-cc-StoSAG and ds-EnOpt impose the same degree of smoothing, ds-cc-StoSAG results in a much higher NPV than is obtained with ds-EnOpt. It is curious to note that StoSAG (unsmoothed), which uses a stochastic approximation to the gradient as the search direction, and ds-cc-StoSAG, which utilizes an approximation of $\mathbf{C}_U^2$ times the gradient as the search direction, yield very similar final values of NPV. However, as expected, the well controls and ICV settings obtained with ds-cc-StoSAG are much smoother than those obtained with StoSAG (Figures 3–6). In these figures, the vertical axes refer to the well number (recall that there are nine producers and four injectors), while the horizontal axis in Figures 3 and 5 refer to the control interval (16 in total). In Figures 4 and 6, the horizontal axis has a related but slightly different intepretation: although there are 16 control intervals at injection wells, the well controls are gas injection rates at odd numbered control steps and water injection rates at even numbered time steps, so only eight control steps (referred to as 'WAG cycles') are shown in Figures 4 and 6. The color scale in Figure 3 corresponds to the pressures (at the top of the reservoir) in the production wells, while the color scales in Figure 4 refer to the gas rates in the injection wells. In Figures 5 and 6, ICV 1 refers to the ICV setting of layer 1, which is the top layer with the highest permeability distribution; see the discussion of the permeability fields of Figure 1.

In order to quantify the smoothness of the control at an individual well, or, more precisely, the variation in a well's controls over the total prescribed reservoir life, the total variation of the function representing well controls is calculated on a well-by-well basis. As the control function at a well is a piecewise constant function, its total variation is given by the sum of the absolute values of the

(a) StoSAG, ICV 1   (b) ds-cc-StoSAG, ICV 1

Figure 5. Estimated settings of inflow control valve 1 for producers at different control steps calculated from StoSAG and ds-cc-StoSAG on a scale from closed (0) to fully open (1).


(a) StoSAG, ICV 1

(b) ds-cc-StoSAG, ICV 1

Figure 6. Estimated settings of inflow control valve 1 for gas injectors at different water-alternating-gas cycles calculated from StoSAG and ds-cc-StoSAG on a scale from closed (0) to fully open (1).

Table II. Total variation of the estimated well controls for different producers and average total variation across all the producers; StoSAG and ds-cc-StoSAG.

Prod. index	StoSAG ($\times 10^3$ psi)	ds-cc-StoSAG ($\times 10^3$ psi)
1	4.82	3.66
2	5.19	5.35
3	5.46	3.87
4	6.33	3.49
5	4.57	0.78
6	4.31	3.90
7	5.99	2.50
8	4.89	3.62
9	7.27	1.50
Average	5.42	3.19

Table III. Total variation of the estimated well controls for gas injectors and average total variation across all the gas injectors; StoSAG and ds-cc-StoSAG.

Inj. index	StoSAG ($\times 10^7$ scf/day)	ds-cc-StoSAG ($\times 10^7$ scf/day)
1	3.73	1.39
2	3.92	2.30
3	5.64	2.17
4	5.89	1.81
Average	4.79	1.92

Table IV. Total variation of the estimated ICV settings for producers; StoSAG and ds-cc-StoSAG.

Prod. index	StoSAG ICV 1	ds-cc-StoSAG ICV 1	StoSAG ICV 2	ds-cc-StoSAG ICV 2	StoSAG ICV 3	ds-cc-StoSAG ICV 3
1	3.77	1.64	2.86	1.90	3.37	2.11
2	3.43	1.54	3.15	1.62	4.23	1.43
3	2.51	1.78	3.63	1.95	2.24	1.15
4	2.96	1.07	3.10	2.04	2.34	0.46
5	4.03	2.07	3.87	0.50	1.98	1.83
6	3.16	1.40	3.18	2.67	2.57	0.52
7	4.47	1.76	2.79	0.88	2.80	1.22
8	3.83	2.38	4.23	0.89	3.42	1.32
9	2.73	2.23	3.53	1.54	3.15	2.84
Average	3.43	1.76	3.37	1.55	2.90	1.43

differences between two consecutive well controls. The total variation of the function representing each well's ICV settings can be calculated on a well-by-well basis in a similar manner. As illustrated in Tables II–V, the total variations of a well's optimized controls and ICV settings are lower when the optimized values are computed from ds-cc-StoSAG than when the optimized values are estimated from StoSAG. Note the average total variation is also given in Tables II–V. Overall, the average total variation of each ds-cc-StoSAG result is roughly one-half of the corresponding average total variation for the StoSAG result. On average, each total variation computed from ds-EnOpt (standard EnOpt) results that are not shown is within 10% of the corresponding total variation obtained from ds-cc-StoSAG. Results similar to and consistent with those of Figures 4 and 6 and Tables III and V were obtained for the water injection rate controls and ICV settings. Although smoother controls

Table V. Total variation of the estimated ICV settings for gas injectors; StoSAG and ds-cc-StoSAG.

Inj. index	StoSAG ICV 1	ds-cc-StoSAG ICV 1	StoSAG ICV 2	ds-cc-StoSAG ICV 2	StoSAG ICV 3	ds-cc-StoSAG ICV 3
1	1.54	1.00	2.34	1.06	2.40	0.79
2	2.84	1.39	2.69	1.67	1.96	1.28
3	2.50	1.18	1.58	0.99	2.40	1.26
4	1.80	0.97	2.30	1.98	0.94	1.60
Average	2.17	1.14	2.23	1.43	1.92	1.23

Figure 7. One of the egg model realizations displaying the injector (blue) and producer (red) locations [12, 43].


Figure 8. Six realization of the egg model displaying different permeability fields [12, 43].

and ICV settings are simpler to implement operationally, as noted earlier, too much smoothing can lead to suboptimal results.

7.3. Example 2: water flooding optimization

The model, illustrated in Figure 7, was introduced in [12] along with an ensemble of equiprobable geological scenarios, six of which are illustrated in Figure 8. The models represent a geological setting of fossilized meandering channels (high-permeability sand stones) in clayish background (low-permeability shales) in the form of discrete permeability fields modeled with $60 \times 60 \times 7 = 25,200$ grid cells of which 18,553 cells are active (i.e., non-zero). A detailed description including reservoir and fluid properties of a standardized version of this 'egg model' is given in [43]. The uncertainty in the models is described by the different channel orientations and permeability fields of the models. The model is produced using eight peripheral water injection wells and four production wells, which are simply connected to the reservoir in all seven layers without ICVs. No capillary

Figure 9. Expectation of Net Present Value versus number of simulation runs for different search direction formulations; yellow: ds-EnOpt (standard); green: StoSAG; dark blue: ss-StoSAG; purple: ss-cc-StoSAG; red: ds-cc-StoSAG.

pressures are included, and the reservoir rock is assumed to be incompressible. We use a commercial fully implicit black oil simulator ([44]) for the reservoir simulations in combination with EnOpt or StoSAG algorithms to compute the gradient of J with respect to the controls **u**. The gradient is used in a simple steepest ascent scheme, where we determine the step size with the aid of an inexact line search and the Arjimo conditions following a trust region (i.e., step size restriction) approach as described in [11].

In this example, water injection rates are the controls with a maximum allowable injection rate per well fixed at 377 STB/day and a minimum rate of 0 STB/day. The producers are operated at a minimum bottom hole pressure of 5584 psi without rate constraints. The producing life of the reservoir is divided into 40 optimization time intervals of 90 days each, and the control vector **u** has, therefore, $8 \times 40 = 320$ elements. An optimal life-cycle strategy of injection rates for the individual wells is obtained by optimizing the NPV as described in Eq. (35), with $r_o = 20$ \$/STB, $c_w = 3$ \$/STB, and $c_{wi} = 0.8$ \$/STB, while c_{gi} is irrelevant for our example because we only consider the flow of oil and water. The discount rate b is set to 0. The initial strategy (starting point) of the life-cycle optimization is a control vector with maximum injection flow rates at all control times.

Figure 9 illustrates the objective function values obtained from optimization using the same five search directions considered in Example 1. For this example, we observe that StoSAG (green) results in an approximately 25% higher expected objective function value than ds-EnOpt (yellow). Even higher objective function values are obtained for the two single-smoothed varieties of StoSAG: ss-StoSAG (dark blue) and ss-cc-StoSAG (purple). We note that in Figure 9, the values of the objective function are versus the number of reservoir simulation runs, where 25 iterations correspond roughly to 4000 reservoir simulation runs.

Figure 10 displays plots of the well controls versus time for the five different search direction. The diffrnce in smoothness between the various formulations is clearly visible.

A key difference between the EnOpt and StoSAG search directions is the way the objective function anomalies are calculated. In EnOpt, the anomalies $\Delta \mathbf{j}_{\ell,k}, k = 1, 2, \cdots N_p (= N_e)$ are calculated as $J(\mathbf{m}_k, \hat{\mathbf{u}}_{\ell,k}) - \overline{J(\mathbf{m}, \mathbf{u}_\ell)}$, that is, with respect to the mean objective function value of the (1 to 1 paired) ensembles of controls and geological models (Eq. (7) or Eq. (6)). In the (1 to 1 paired) StoSAG formulations, however, the anomalies are calculated as $J(\mathbf{m}_k, \hat{\mathbf{u}}_{\ell,k}) - J(\mathbf{m}_k, \mathbf{u}_\ell)$, that is, with respect to the individual objective function values of the geological realizations (Eq. (11)). The importance of the objective function anomalies is that they act as weighting factors in the gradient estimate. Figure 11 illustrates that the objective function anomalies for StoSAG (red) are significantly different to those of the ds-EnOpt (blue). For some realizations the 'weights' are even opposite in sign, which explains the large difference between the two formulations.

(a) ds-EnOpt

(b) StoSAG

(c) ss-StoSAG

(d) ss-cc-StoSAG

(e) ds-cc-StoSAG

Figure 10. Well controls versus time for ds-EnOpt en four different StoSAG formulations; units are STB/day.

Figure 11. Comparison of weights used in gradient estimate for the different gradient formulations; blue: ds-EnOpt; red: StoSAG.

8. CONCLUSION

We have provided a theoretical explanation and understanding as to the superior performance of StoSAG algorithms compared with EnOpt algorithms. Theoretically, a steepest ascent algorithm that uses a search direction based on StoSAG should outperform ds-EnOpt or ss-EnOpt. We highlight the reasons why the EnOpt algorithms will produce unsatisfactory results when used for optimization under uncertainty, that is, when using geological ensembles with significantly different ensemble members. We have provided two numerical examples involving different physical problems, which illustrate the advantages of using StoSAG compared with EnOpt algorithms and highlight the main benefits accrued.

REFERENCES

1. Aziz K, Settari A. *Petroleum Reservoir Simulation*. Applied Science Publishers Ltd: London, 1979.
2. Chen Z, Huan G, Ma Y. *Computational Methods for Multiphase Flows in Porous Media*. SIAM: Philadelphia, 2006.
3. Jansen JD, Bosgra OH, van den Hof PMJ. Model-based control of multiphase flow in subsurface oil reservoirs. *Journal of Process Control* 2008; **18**(9):846–855.
4. Bryson AE, Ho YC. *Applied Optimal Control*. Taylor and Francis (Hemisphere): Levittown, 1975.
5. Asheim H. Maximization of water sweep efficiency by controlling production and injection rates. *Paper SPE 18365 presented at the SPE European Petroleum Conference*, London, United Kingdom, October 1988; 16–19.
6. Zakirov IS, Aanonsen SI, Zakirov ES, Palatnik BM. Optimization of reservoir performance by automatic allocation of well rates. *Proceedings of the 5th European Conference on the Mathematics of Oil Recovery (ECMOR V) Leoben*, Austria, September 1996; 3–6.
7. Sudaryanto B, Yortsos YC. Optimization of fluid front dynamics in porous media using rate control. *Physics of Fluids* 2000; **12**(7):1656–1670.
8. Brouwer DR, Jansen JD. Dynamic optimization of water flooding with smart wells using optimal control theory. *SPE Journal* 2004; **9**(4):391–402.
9. Sarma P, Chen WH, Durlofsky LJ, Aziz K. Production optimization with adjoint models under nonlinear control-state path inequality constraints. *SPE Reservoir Evaluation and Engineering* 2008; **11**(2):326–339.
10. Jansen JD. Adjoint-based optimization of multiphase flow through porous media a review. *Computers and Fluids* 2011; **46**(1):40–51.

11. Nocedal J, Wright SJ. *Numerical Optimization* (2nd ed.) Springer: New York, 2006.
12. Van Essen GM, Zandvliet MJ, Van den Hof PMJ, Bosgra OH, Jansen JD. Robust waterflooding optimization of multiple geological scenarios. *SPE Journal* 2009; **14**(1):202–210.
13. Chen Y, Oliver DS, Zhang D. Efficient ensemble-based closed-loop production optimization. *SPE Journal* 2009; **14**(4):634–645.
14. Chen Y, Oliver DS. Ensemble-based closed-loop optimization applied to Brugge field. *SPE Reservoir Engineering and Evaluation* 2010; **13**(1):56–71.
15. Do ST, Reynolds AC. Theoretical connections between optimization algorithms based on an approximate gradient. *Computational Geosciences* 2013; **17**(6):959–973.
16. Spall JC. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions onAutomatic Contol* 1992; **37**(3):332–341.
17. Spall JC. *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control*. Wiley, 2003.
18. Conn AR, Scheinberg S, Vicente LN. *Introduction to Derivative-Free Optimization*. SIAM: Philadelphia, 2009.
19. Stordal AS, Szklarz SP, Leeuwenburgh O. A theoretical look at ensemble-based optimization in reservoir management. *Mathematical Geosciences* 2016; **48**(4):399–417.
20. Amari SI. Natural gradients work efficiently in learning. *Neural Computation* 1998; **10**(2):251–276.
21. Lorentzen RJ, Berg A, Naevdal G, Vefring EH. A new approach for dynamic optimization of waterflooding problems. *Paper SPE-99690-MS presented at the Intelligent Energy Conference and Exhibition*, Amsterdam, The Netherlands, April 2006; 11–13.
22. Raniolo S, Dovera L, Cominelli A, Callegaro C, Masserano F. History matching and polymer injection optimization in a mature field using the ensemble Kalman filter. *Proc. 17th European Symposium Improved Oil Recovery, St*, Petersburg, Russia, April 2013; 16–18.
23. Fonseca RM, Kahrobaei S, Van Gastel LJT, Leeuwenburgh O, Jansen J.D. Quantification of the impact of ensemble size on the quality of an ensemble gradient using principles of hypothesis testing. *Paper 173236-MS presented at the 2015 SPE Reservoir Simulation Symposium*, Houston, USA; 22-25 February 2015.
24. Do S. Application of SPSA-type algorithms to production optimization. *Ph.D. thesis*, The University of Tulsa, Tulsa, USA, 2012.
25. Yan X, Reynolds AC. Optimization algorithms based on combining FD approximations and stochastic gradients compared with methods based only on a stochastic gradient. *SPE Journal* 2014; **19**(5):873–890.
26. Fonseca RM, Stordal AS, Leeuwenburgh O, Van den Hof PMJ, Jansen JD. Robust ensemble-based multi-objective optimization. *Proceedings of the 14th European Conference on the Mathematics of Oil Recovery (ECMOR XIV)*, Catania, Italy, September 2014; 8–11.
27. Fonseca RM, Leeuwenburgh O, Della Rossa E, Van den Hof PMJ, Jansen JD. Ensemble-based multiobjective optimization of on-off control devices under geological uncertainty. *SPE Reservoir Evaluation and Engineering* 2015; **18**(4):1094–6470.
28. Jansen JD. *A Systems Description of Flow Through Porous Media*, SpringerBriefs in Earth Sciences. Springer: Heidelberg, 2013.
29. Bouzarkouna Z, Ding DY, Auger A. Well placement optimization with the covariance matrix adaptation evolution strategy and meta-models. *Computational Geosciences* 2012; **16**(1):75–92.
30. Fonseca RM, Leeuwenburgh O, Van den Hof PMJ, Jansen JD. Improving the ensemble optimization method through covariance matrix adaptation (CMA-EnOpt). *SPE Journal* 2015; **20**(1):155–168.
31. Forouzanfar F, Poquioma WE, Reynolds AC. Simultaneous and sequential estimation of optimal placement and controls of wells using a covariance matrix adaptation algorithm. *SPE Journal* 2016; **21**(2):501–521.
32. Zandvliet MJ, Bosgra OH, Van den Hof PMJ, Jansen JD, Kraaijevanger JFBM. Bang-bang control and singular arcs in reservoir flooding. *Journal of Petroleum Science and Engineering* 2007; **58**:186–200.
33. Wang C, Li G, Reynolds AC. Production optimization in closed loop reservoir management. *SPE Journal* 2009; **14**(3):506–523.
34. Leeuwenburgh O, Egberts PJP, Abbink OA. Ensemble methods for reservoir life-cycle optimization and well placement. *Paper SPE-136916-MS presented at the SPE/DGS Saudi Arabia Section Technical Symposium and Exhibition*, Al-Khobar, Saudi Arabia, April 2010; 4–7.
35. Golub G, Van Loan C. *Matrix Computations* (4th ed.) John Hopkins University Press: Baltimore, 1989.
36. Sarma P, Chen W. Improved estimation of the stochastic gradient with quasi-Monte Carlo methods. *Proc. 14th European Conference on the Mathematics of Oil Recovery (ECMOR XIV), Catania*, Italy, September 2014; 8–11.
37. Oliveira DF, Reynolds AC. An adaptive hierarchical algorithm for estimation of optimal well controls. *SPE Journal* 2014; **9**(4):391Ǔ402.
38. Chen Y. Efficient ensemble based reservoir management. *PhD Thesis*, University of Oklahoma, USA, 2008.
39. Beck JV, Arnold KJ. *Parameter Estimation in Engineering and Science*. John Wiley & Sons: Chichester, 1997.
40. Tavakoli R, Reynolds AC. History matching with parameterization based on the SVD of a dimensionless sensitivity matrix. *SPE Journal* 2010; **15**(12):495–508.
41. *GEM*. (Available from: http://www.cmgl.ca/software/gem2014) [Accessed 13 September 2015].
42. Journel A, Huijbregts CJ. *Mining Geostatistics*. Academic Press: New York, 1978.
43. Jansen JD, Fonseca RM, Kahrobaei SS, Siraj MM, Van Essen GM, Van den Hof PMJ. The Egg Model- A geological ensemble for reservoir simulation. *Geoscience Data Journal* 2014; **1**(2):192–195.
44. (Available from: http://www.software.slb.com/products/foundation/Pages/eclipse.aspx) [Accessed 13 September 2015].