



Fast robust optimization using bias correction applied to the mean model

Lingya Wang^{1,2} · Dean S. Oliver¹

Received: 12 March 2020 / Accepted: 28 October 2020 / Published online: 26 November 2020
© The Author(s) 2020

Abstract

Ensemble methods are remarkably powerful for quantifying geological uncertainty. However, the use of the ensemble of reservoir models for robust optimization (RO) can be computationally demanding. The straightforward computation of the expected net present value (NPV) requires many expensive simulations. To reduce the computational burden without sacrificing accuracy, we present a fast and effective approach that requires only simulation of the mean reservoir model with a bias correction factor. Information from distinct controls and model realizations can be used to estimate bias for different controls. The effectiveness of various bias-correction methods and a linear or quadratic approximation is illustrated by two applications: flow optimization in a one-dimensional model and the drilling-order problem in a synthetic field model. The results show that the approximation of the expected NPV from the mean model is significantly improved by estimating the bias correction factor, and that RO with mean model bias correction is superior to both RO performed using a Taylor series representation of uncertainty and deterministic optimization from a single realization. Use of the bias-corrected mean model to account for model uncertainty allows the application of fairly general optimization methods. In this paper, we apply a nonparametric online learning methodology (learned heuristic search) for efficiently computing an optimal or near-optimal robust drilling sequence on the REEK Field example. This methodology can be used either to optimize a complete drilling sequence or to optimize only the first few wells at a reduced cost by limiting the search depths.

Keywords Robust optimization · Bias correction · Geological uncertainty · Ensemble-based methods · Drilling sequence optimization · Sequential-decision making

1 Introduction

In the development of a hydrocarbon field, optimization is an important process that can substantially improve the profitability through reduction in production and development costs or by increasing or accelerating the production of hydrocarbons. The increase in profitability of a field for a given change in controls is predicted using a reservoir simulation model, in which the geological uncertainties such

as porosity, permeability, and fault transmissibility have critical effects on production forecasts. Ensemble methods are remarkably powerful for quantifying the geological uncertainty. Van Essen et al. [39] demonstrated that robust optimization (RO) performed using the expectation of NPV by averaging over an ensemble of reservoir models outperforms the nominal optimization (NO) from a single realization. However, the straightforward computation of the expected value requires many expensive simulations for each control variable setting.

Many approaches have been proposed to reduce the cost of simulation-based RO. These methods can be mainly classified into two categories: accelerating the reservoir model simulations and reducing the number of simulations required to obtain the maximum. A simplified model such as one obtained from reduced-order modeling and upscaling model parameters [4, 5, 13, 20, 38] is a computationally inexpensive substitute for full reservoir simulation. However, such a proxy model is generally less accurate due to the challenges

✉ Lingya Wang
liwa@norce-research.no

Dean S. Oliver
deol@norce-research.no

¹ NORCE Norwegian Research Centre, P.O. Box 22, 5838 Bergen, Norway

² Department of Mathematics, University of Bergen, P.O. Box 7800, 5020 Bergen, Norway

posed by the high nonlinearity in the reservoir model; consequently, RO performed in a simplified model is likely to yield a suboptimal solution [11]. Techniques for reducing the number of simulations needed for optimization of the expected value of an objective function can be further classified into two main categories: improving the optimization algorithms for certain problems and reducing the cost of evaluating the expected value.

Ensemble-based optimization [6, 14] is an efficient gradient-based approach for fairly general RO problems with continuous variables. The cost of evaluating the gradient of the expected value of the objective function is reduced by using the same ensemble to represent the reservoir model uncertainty and to compute the gradient of the expectation of the objective function with respect to control variable settings. One downside to a gradient-based approach is that it is easily possible to get stuck in a local minimum (or maximum). For optimization problems with discrete control variables, Wang and Oliver [40] proposed a nonparametric online learning methodology with heuristic controls to efficiently solve sequential decision-making problems. This approach can reduce the number of iterations required during the optimization process. Nevertheless, the amount of computation required for RO performed using sample average approximation (SAA) over a large ensemble could still be expensive since the cost increases linearly with the number of realizations required to represent the uncertainty in model properties.

There are several possible ways to compute the expected value at a lower cost than SAA. One is to reduce the representation of uncertainty through model selection, in which a subset of the ensemble members are used for optimization [22, 32, 37]. A small subset of model realizations may not span the uncertainties accurately, leading to a sub-optimal solution. To address this issue, Barros et al. [2] proposed an automated scenario reduction approach for selecting a subset that contains an optimal number of realizations that are able to capture the range of the uncertainties of the full ensemble. A much faster approach is to evaluate the objective function in the mean reservoir model. Chen et al. [7] obtained relatively good results by using the ensemble mean model updated from the ensemble Kalman filter (EnKF) data assimilation for production optimization. Their optimized design increased the expected NPV compared to the reference case, but was not as good as the optimized results obtained using SAA. One major drawback to direct use of the mean model is the large bias in predicted reservoir behavior that may result from the use of a model with reduced heterogeneity, in which case the mean model may offer a poor estimate of the expected NPV. To improve this approximation, one possible approach is to modify the representation of uncertainty using a Taylor series expansion of the objective function

[1, 3, 8–10]. The first term of Taylor expansion can be inexpensively obtained using the mean of reservoir model realizations. To accurately estimate the expected value, the Taylor series expansion generally includes higher-order terms (e.g., a quadratic or cubic term). The computation of higher-order derivatives is impractical for most real reservoir applications, however.

Instead of using higher-order terms to improve the estimation accuracy, we develop a fast and efficient approach to correcting the bias in the NPV obtained from the mean reservoir model by estimating a multiplicative bias correction factor based on the information from distinct controls and individual model realizations. To accurately estimate the expected NPV, we apply distance-based localization to estimate the bias for specific controls, considering the similarity between samples and control variables in terms of the bias correction factor. To avoid poor estimates caused by sampling error due to small sample sizes, a regularization term based on the average value and the variance of the bias correction factor is used to reduce the sensitivity of the estimates to the taper length, thereby allowing more accurate estimates to be generated for a wider range of taper lengths. The initial sampling of control and model realizations is used to create an ensemble of partial correction factors. During the RO phase, when it is necessary to estimate the expected value of NPV for a control, we only require simulation of the control applied to the mean model. The bias correction factor is estimated from partial bias correction factors on similar controls. Hence, after creation of the initial ensemble, RO performed using bias-correction methods requires only one additional simulation of the mean reservoir model at each iteration, which is much less than the effort required in SAA. This bias-correction methodology can be applied to fairly general problems of optimizing the expected value of an objective function. But an appropriate distance metric to measure the similarity of controls in terms of the bias correction factor is required, which is specific to the problem at hand.

The performance of various bias-correction methods and a linear or quadratic approximation is investigated in two applications. The objective in the first example is to locate an injection well such that flow rate for a fixed pressure in a one-dimensional model is minimized; this problem is small enough that we can evaluate both linear and quadratic Taylor approximations of the objective function. The second example is to maximize the expected NPV in a synthetic field model by optimizing the drilling sequence of wells at fixed locations. Here we use well-position based distance to measure the similarity of drilling sequences for the bias correction factor. We formulate the sequential drilling optimization problem as one of finding a path with the maximum reward in a decision tree, and apply learned heuristic search [40] with mean model bias

correction (MMBC) to compute the RO drilling sequence under geological uncertainty and to optimize only the first few wells at a reduced cost by limiting the search depths.

This paper is organized as follows. Section 2 describes the bias-correction methods for estimating the expected NPV and the learned heuristic search method for optimizing either a complete drilling sequence or only the first few wells. Sections 3 and 4 describe two numerical case studies (i.e., flow optimization in a one-dimensional model and drilling-order problem in a synthetic model). The conclusions are presented in Section 5.

2 Methodology

2.1 Estimation of expected value

Correction factor In this work, we use maximization of the NPV as the objective for an optimal reservoir management problem. Because reservoir characterization is always incomplete, the optimization problem is based on a reservoir model with uncertainty in parameter values. An appropriate approach to account for the uncertainty is to use the expected NPV over an ensemble of reservoir models that have been sampled from the probability distribution for model parameters. The expected value of the objective can be written as:

$$E[f(x, m)] \approx \bar{f}(x) = \frac{1}{N_e} \sum_{j=1}^{N_e} f(x, m_j), \tag{1}$$

where x is a vector of control inputs; $m \in \mathcal{R}^m$ is an m -dimensional vector of uncertain model parameters; j is the index of individual realizations; and N_e indicates the number of reservoir models.

This ensemble-based average value can provide a good approximation of the expected NPV if the ensemble of model realizations is sufficiently large [23, 25]. RO performed using such straightforward estimation of the expected NPV, however, requires many expensive simulations when the ensemble size N_e or the number of iterations needed for the optimization process is large. Instead of computing the expectation of the objective function by using SAA, one could consider optimizing the expectation for a linear or quadratic approximation of the objective function [1, 3, 9, 10]. To second order, the Taylor expansion of the objective function is:

$$f^{\text{quad}}(x, m) = f(x, \bar{m}) + f_m(x, \bar{m})(m - \bar{m}) + \frac{1}{2}(m - \bar{m})^T f_{mm}(x, \bar{m})(m - \bar{m}), \tag{2}$$

where we have neglected higher-order terms in the expansion. f_m and f_{mm} are first and second derivatives of f

with respect to uncertain model parameter m , respectively.

If m is distributed as multivariate Gaussian with mean \bar{m} and covariance C , then the expected value of the quadratic approximation of the objective can be shown [31] to be:

$$E[f^{\text{quad}}(x, m)] = f(x, \bar{m}) + \frac{1}{2} \text{tr}(C^{1/2} f_{mm}(x, \bar{m}) C^{1/2}). \tag{3}$$

A possible advantage of this approach to approximating the expected value of the objective function is that optimization of the expectation does not require evaluation of controls applied to a large number of Monte Carlo samples [1]. Computing $f(x, \bar{m})$ in Eq. 3 will be easy as it only requires the mean of the realizations. Computing $f_{mm}(x, \bar{m})$ is more difficult as it requires the second derivative of the objective function with respect to model parameters. Although it might be possible in some cases to approximate the second derivative, computing second derivatives will be impractical for most reservoir applications. The linear approximation seems more likely to be feasible in practice:

$$f^{\text{lin}}(x, m) = f(x, \bar{m}) + \langle f_m(x, \bar{m}), m - \bar{m} \rangle. \tag{4}$$

Assuming that m is distributed as Gaussian, the expectation over m of the linear approximation to the objective is simply a function of \bar{m} and x :

$$E[f^{\text{lin}}(x, m)] = f(x, \bar{m}). \tag{5}$$

The key point of this is that the mean of the ensemble can be used for the optimization, instead of performing the optimization on the ensemble of realizations.

Such an approach has a major advantage over SAA in that it significantly reduces the number of simulations required for evaluating the expected value. However, since the NPV is generally not a linear function of the uncertain model parameters m , NPV from the mean model is generally not the same as the expected value over uncertainty space, $f(x, E[m]) \neq E[f(x, m)]$. If we want to use the mean model for optimization problems without sacrificing the accuracy of the estimated expected NPV, we need a method for improving the approximation of expected NPV from the mean model.

In this paper, we compute a multiplicative correction factor between the ensemble average of values $\bar{f}(x_i)$ and the value obtained from the mean model $f(x_i, \bar{m})$ at a fixed control x_i :

$$\bar{f}(x_i) = \alpha(x_i) f(x_i, \bar{m}), \tag{6}$$

where $\alpha(x_i)$ is the correction factor for a fixed control x_i . Instead of directly computing the ensemble average

of values of the objective function, we develop an approximation of expected value $\bar{f}(x)$ by estimating the correction factor $\alpha(x_i)$ of control x_i . If it were feasible to compute the value of the NPV at control x_i for all samples of model parameters, then $\alpha(x_i)$ could be computed in the following way:

$$\alpha(x_i) = \frac{1}{N_e} \sum_{j=1}^{N_e} \frac{f(x_i, m_j)}{f(x_i, \bar{m})}, \quad (7)$$

where $f(x_i, m_j)$ is the economic value at control x_i of an individual realization m_j .

For each individual model realization m_j and control x_i , we define a partial correction factor $\beta(x_i, m_j, \bar{m})$:

$$\beta_{ij} = \beta(x_i, m_j, \bar{m}) = \frac{f(x_i, m_j)}{f(x_i, \bar{m})}. \quad (8)$$

The correction factor $\alpha(x_i)$ at control x_i of all ensemble realizations (7) can be written in terms of the partial correction factors:

$$\alpha(x_i) = \frac{1}{N_e} \sum_{j=1}^{N_e} \beta_{ij}. \quad (9)$$

We also define the mean value of the correction factor $\alpha(x_i)$:

$$\begin{aligned} \bar{\alpha} &\approx \frac{1}{N_x} \sum_{i=1}^{N_x} \alpha(x_i) = \frac{1}{N_x N_e} \sum_{i=1}^{N_x} \sum_{j=1}^{N_e} \frac{f(x_i, m_j)}{f(x_i, \bar{m})} \\ &= \frac{1}{N_x N_e} \sum_{i=1}^{N_x} \sum_{j=1}^{N_e} \beta_{ij}, \end{aligned} \quad (10)$$

where N_x is the number of relevant controls. Straight-forward application of Eq. 10 requires $N_x \times (N_e + 1)$ evaluations of $f(x, m)$ to compute the average correction factor $\bar{\alpha}$ from N_x different controls on an ensemble of N_e model realizations.

A Monte Carlo estimate of $\bar{\alpha}$ can be obtained at a much lower cost by sampling control x_j uniformly from the space of all possible controls, and sampling reservoir realizations m_k from the space of conditional realizations. Then $\bar{\alpha}$ can be estimated using the following formula:

$$\bar{\alpha} \approx \frac{1}{N_x} \sum_{j=1}^{N_x} \frac{f(x_j, m_j)}{f(x_j, \bar{m})} = \frac{1}{N_x} \sum_{j=1}^{N_x} \beta(x_j, m_j, \bar{m}), \quad (11)$$

where $\beta(x_j, m_j, \bar{m})$ is the observed value at control x_j , which requires two simulations, i.e., apply control x_j to a random individual realization m_j and apply control x_j to the mean reservoir model \bar{m} . Therefore, it would require

$2 \times N_x$ simulations to obtain a set of observations β from N_x distinct controls.

For N_e reservoir model realizations, we can sample N_e distinct controls to obtain observations of β so that each realization will provide one observed value of β . In that case, the number of observed values of β is the same as the ensemble size (i.e., $N_x = N_e$). In some cases, more observations might be needed to obtain a good approximate of $\bar{\alpha}$ (e.g., obtain two observations of β from each realization (i.e., $N_x = 2N_e$) or use a larger ensemble size). If the value of β is similar for most realizations, we can use information from a subset of the ensemble members and a smaller number of distinct controls (i.e., $N_x < N_e$) to estimate the bias.

Although an estimate of $\bar{\alpha}$ can be efficiently obtained from the observed values of β , the accuracy of an estimate of the expected NPV obtained using $\bar{\alpha}$ for correction is limited by the variability in α . Use of the bias correction factor obtained by averaging samples of β (11) will result in the same correction factor being used for all controls, even though the correct values of α for some of the controls may be far from $\bar{\alpha}$. In such a case, the accuracy level of the estimates from different controls is limited by the actual $\bar{\alpha}$ value. Nevertheless, if the variability in $\alpha(x_i)$ is small as control x_i is varied, then it is possible that $\bar{\alpha}$ can provide a useful approximation to $\alpha(x_i)$ for estimation of the value of $\bar{f}(x_i)$ from the value of $f(x_i, \bar{m})$.

Distance-based localization In general, however, we expect that an estimate of the correction factor will be better if it is primarily based on information from similar control variables. Thus, we expect that a weighted estimate will be better than an unweighted estimate. In our work, we use weighted linear estimation. Suppose that N_e distinct controls are applied to N_e individual realizations and the mean model for generating a set of observations β . The weighted linear estimate $\hat{\alpha}(x_i)$ at control x_i is defined as:

$$\hat{\alpha}(x_i) = \frac{\sum_{j=1}^{N_e} \omega(x_i, x_j) \beta(x_j, m_j, \bar{m})}{\sum_{j=1}^{N_e} \omega(x_i, x_j)}, \quad (12)$$

where $\beta(x_j, m_j, \bar{m})$ is the observed correction factor at a random control x_j applied to an individual realization m_j and the mean model \bar{m} , and $\omega(x_i, x_j)$ is the weight for $\beta(x_j, m_j, \bar{m})$. The weights, $\omega(x_i, x_j)$, should depend on a measure of similarity, or distance measure, between controls x_i and x_j .

With an appropriate measure of distance between control sequences, weights are assigned such that $\beta(x_j, m_j, \bar{m})$ at shorter distances will have higher weights while partial correction factors for controls that are more dissimilar will have smaller weights. Lacking information about the correlation of β with distance between control sequences,

we use the Gaspari-Cohn taper function [16] to compute distance-dependent weights:

$$\rho(\delta, L) = \begin{cases} -\frac{1}{4} \left(\frac{\delta}{L}\right)^5 + \frac{1}{2} \left(\frac{\delta}{L}\right)^4 + \frac{5}{8} \left(\frac{\delta}{L}\right)^3 - \frac{5}{3} \left(\frac{\delta}{L}\right)^2 + 1 & \text{for } 0 \leq \delta \leq L \\ \frac{1}{12} \left(\frac{\delta}{L}\right)^5 - \frac{1}{2} \left(\frac{\delta}{L}\right)^4 + \frac{5}{8} \left(\frac{\delta}{L}\right)^3 + \frac{5}{3} \left(\frac{\delta}{L}\right)^2 - 5 \left(\frac{\delta}{L}\right) + 4 - \frac{2}{3} \left(\frac{\delta}{L}\right)^{-1} & \text{for } L < \delta \leq 2L \\ 0 & \text{for } \delta > 2L \end{cases} \quad (13)$$

where $\rho(\delta, L)$ is the distance-based weight and varies from 0 and 1; δ is the distance between controls; L is the taper length determining the distance at which the weighting drops to approximately 0.2, and $2L$ is the critical distance, beyond which the weighting is zero.

For the drilling-order problem, the control variables x have no physical locations but are permutations of sequences of possible actions, in which case an order-based encoding is appropriate. We have chosen to use the permutation encoding [34] of the drilling sequence. In this encoding, each integer value in the vector encodes the relative ordering of the drilling of a specific well. Consider, for example, two possible control sequences x_i and x_j in which four wells are drilled, i.e., $S_{x_i} = [W_1, W_2, W_3, W_4]$ and $S_{x_j} = [W_3, W_1, W_4, W_2]$. The permutation encodings for these two sequences are $\mathbf{P}_{x_i} = [1, 2, 3, 4]^T$ and $\mathbf{P}_{x_j} = [2, 4, 1, 3]^T$, respectively.

Distance between two control sequences is then measured by the distance between the vectors \mathbf{P}_{x_i} and \mathbf{P}_{x_j} . Appropriate distance measures for ordering problems include the “edit” distance [30], which is the minimum number of operations required to transform one sequence to another sequence, and the “swap” or Jaro-Winkler distance [21], which counts the minimum number of swaps of two elements required to transform one sequence to another. Because computation of swap and edit distances are relatively expensive, it is common to use fitness-distance measures as surrogates for the permutation distances [35]. The Hamming distance [17] between two sequences of equal length is the number of positions at which the corresponding actions are different, i.e., the number of wells that have different positions in the drilling sequence. The Manhattan distance (also known as the “position-based distance”) measures the sum of the absolute differences between positions of the elements. In terms of the permutation encoded vectors, the Manhattan distance is:

$$\delta(x_i, x_j)_{L_1} = \|\mathbf{P}_{x_i} - \mathbf{P}_{x_j}\|_1 = \sum_{k=1}^{N_w} |P_{x_i,k} - P_{x_j,k}|, \quad (14)$$

where the k th elements in \mathbf{P}_{x_i} and \mathbf{P}_{x_j} are $P_{x_i,k}$ and $P_{x_j,k}$, respectively, which are the positions of a fixed well W_k . N_w is the number of wells.

In addition to the Hamming and Manhattan distance measures for permutation encodings, we initially considered the use of two standard distance metrics on vector spaces. The Euclidean distance is defined as:

$$\delta(x_i, x_j)_{L_2} = \|\mathbf{P}_{x_i} - \mathbf{P}_{x_j}\|_2 = \sqrt{\sum_{k=1}^{N_w} |P_{x_i,k} - P_{x_j,k}|^2} \quad (15)$$

and the cosine distance [26, 36] is a correlation-based distance measure defined as:

$$\begin{aligned} \delta(x_i, x_j)_{\cos} &= 1 - \cos(\mathbf{P}_{x_i}, \mathbf{P}_{x_j}) \\ &= 1 - \frac{\mathbf{P}_{x_i} \cdot \mathbf{P}_{x_j}}{\|\mathbf{P}_{x_i}\| \|\mathbf{P}_{x_j}\|} \\ &= 1 - \frac{\sum_{k=1}^{N_w} P_{x_i,k} \cdot P_{x_j,k}}{\sqrt{\sum_{k=1}^{N_w} P_{x_i,k}^2} \sqrt{\sum_{k=1}^{N_w} P_{x_j,k}^2}}. \end{aligned} \quad (16)$$

When the lengths of all sequences are identical, as they are when the sequences are all perturbations of a base sequence, then the cosine distance is simply a scaled version of the Euclidean distance:

$$\frac{\|\mathbf{P}_{x_i} - \mathbf{P}_{x_j}\|_2^2}{2\|\mathbf{P}_{x_i}\|_2^2} = 1 - \cos(\mathbf{P}_{x_i}, \mathbf{P}_{x_j}) \quad (17)$$

consequently, there is no need to consider both the Euclidean and cosine distance measures.

The choice of an appropriate distance measure to use as a measure of similarity is problem specific. In the application section, we show that the Manhattan, Euclidean, and cosine distance metrics are all very similar when applied to the well ordering problem. They are all superior to the Hamming distance in explaining similarity of drilling-order sequences in terms of the bias correction factor.

Taper window selection The performance of localization for estimation of correction factor depends not only on the choice of distance measure but also on the taper parameter L which affects the weights and the effective sample size [24] used for computation of correction factor. A good distance measure will effectively identify control variables with similar correction factors so that the number of realizations used for estimation is maximized and sampling error is reduced.

Suppose that N_e is the number of observed values of β , then:

$$n_{\text{eff}} = \frac{\left(\sum_{j=1}^{N_e} \omega_j\right)^2}{\sum_{j=1}^{N_e} \omega_j^2}, \quad (18)$$

is a common approximation of the effective sample size [12, 33]. In this equation, ω_j is the weight on the j th

partial correction factor. If all weights are identical, then the effective sample size is equal to N_e , while if one of the normalized weights is equal to one and all others are zero, the effective sample size is 1. n_{eff} is a measure of the equivalent number of equally weighted samples. The weights are determined by the distance between the estimation location and the j th drilling sequence.

The accuracy of the estimate of the correction factor for control variable sequence, x_i , is influenced both by the effective sample size, n_{eff} , and the bias resulting from the use of partial correction factors based on random control variables with different values of α . Reducing the taper length will decrease the bias by only including values from control variable sequences with very similar values of the correction factor, but will also increase the sampling error by decreasing the effective sample size. Because the optimal taper length is not known a priori, we generally apply regularization to reduce the effect of a non-optimal choice of taper length.

Regularization The major disadvantage of pure distance-based localization is that use of a taper length that is smaller than optimal one may result in an estimate of α that is far from the correct value due to sampling error resulting from the small number of samples within a small distance of the estimation point. Instead of using a long taper length to avoid such a situation, it is generally possible to improve the accuracy of estimated α value by adding a regularization term based on the average value and the variance of correction factor to reduce the sensitivity of the estimate to taper length while still generating accurate estimates.

A regularized estimate of $\alpha(x_i)$ is obtained by minimizing an objective function with both a local and a global term:

$$S(\alpha_r(x_i)) = \frac{n_{\text{eff}}}{\sigma_\beta^2} \left(\alpha_r(x_i) - \frac{\sum_{j=1}^{N_e} \omega(x_i, x_j) \beta(x_j, m_j, \bar{m})}{\sum_{j=1}^{N_e} \omega(x_i, x_j)} \right)^2 + \frac{1}{\sigma_\alpha^2} (\alpha_r(x_i) - \bar{\alpha})^2. \tag{19}$$

In Eq. 19, σ_α^2 is an estimate of the variance of α over the domain of interest, σ_β^2 is the variance of β , and n_{eff} is the effective sample size for the observations of β (18). The regularized estimate, $\hat{\alpha}_r(x_i)$, is obtained by solving $\nabla_\alpha S = 0$, obtaining:

$$\hat{\alpha}_r(x_i) = \left(1 + \frac{\sigma_\beta^2}{n_{\text{eff}} \sigma_\alpha^2} \right)^{-1} \left(\frac{\sum_{j=1}^{N_e} \omega(x_i, x_j) \beta(x_j, m_j, \bar{m})}{\sum_{j=1}^{N_e} \omega(x_i, x_j)} + \frac{\sigma_\beta^2}{n_{\text{eff}} \sigma_\alpha^2} \bar{\alpha} \right) = \left(1 + \frac{\sigma_\beta^2}{n_{\text{eff}} \sigma_\alpha^2} \right)^{-1} \left(\hat{\alpha}_{\text{loc}}(x_i) + \frac{\sigma_\beta^2}{n_{\text{eff}} \sigma_\alpha^2} \bar{\alpha} \right). \tag{20}$$

Note that the regularized objective function is a weighted average of the localized estimate $\hat{\alpha}_{\text{loc}}(x_i)$ from Eq. 12 and the mean value of α . When the effective sample size is large compared to the ratio $\sigma_\beta^2/\sigma_\alpha^2$, the regularized estimate will be based primarily on the local samples of β .

Estimation with regularized localization has a major advantage over an approach that relies only on localized estimation: by improving the accuracy of the estimated values that are obtained with an inappropriate distance measure or taper length, regularized estimate $\hat{\alpha}_r$ is potentially more accurate than $\bar{\alpha}$ for a wider range of taper lengths. When the variance of α is unknown, it might be difficult to select the optimal value, but as shown in experiments, results are not strongly sensitive to the exact choice.

Optimal weights Here we show how the optimal weights can be estimated if the covariance of partial correction factors is known. In that case, an estimate of α at a fixed control x_0 is calculated based on a linear combination of observations β_{ij} with weights w_i . In vector notation, the estimate is written as:

$$\hat{\alpha}(x_0) = \frac{\sum_{j=1}^N \omega(x_0, x_j) \beta(x_j, m_j, \bar{m})}{\sum_{j=1}^N \omega(x_0, x_j)} = \mathbf{w}^T \mathbf{b}, \tag{21}$$

where elements in vector \mathbf{b} are the observed values of β_{jk} from random controls and realizations. The collection of observations will be denoted by the vector \mathbf{b} , i.e.:

$$\mathbf{b} = [\beta_{11} \ \beta_{22} \ \cdots \ \beta_{NN}]^T.$$

Although the notation is different, Eq. 21 is identical to Eq. 12.

The quantity $\alpha(x_0)$, that is to be estimated, is defined to be the linear combination of β_{0j} at the estimation location:

$$\alpha(x_0) = \frac{1}{N} \mathbf{1}^T \mathbf{b}_0, \tag{22}$$

where the j th element of \mathbf{b}_0 is β_{0j} at a fixed control x_0 of an individual realization m_j that is sampled from the probability distribution for model parameters, i.e.:

$$\mathbf{b}_0 = [\beta_{01} \ \beta_{02} \ \cdots \ \beta_{0N}]^T.$$

Imposing the constraint $\mathbf{w}^T \mathbf{1} = 1$ provides an unbiased estimate for which the expected error is 0. The optimal weights for estimating $\alpha(x_i)$ from a set of random observations β are obtained by minimizing the expected variance of the estimate, constrained to the unbiasedness condition. For estimation of $\alpha(x_0)$, the variance of the expected error is:

$$S_w(\mathbf{w}) = E \left[\left(\frac{1}{N} \mathbf{1}^T \mathbf{b}_0 - \mathbf{w}^T \mathbf{b} \right)^2 \right] \tag{23}$$

To minimize the variance in estimation error, subject to the constraint that $\mathbf{w}^T \mathbf{1} = 1$, we define a Lagrangian function:

$$S(\mathbf{w}, \lambda) = S_w(\mathbf{w}) - 2\lambda (\mathbf{w}^T \mathbf{1} - 1),$$

where λ is the Lagrangian parameter and $S_w(\mathbf{w})$ is the variance of the estimator error (23). The optimal weights are then obtained by solving for \mathbf{w} that minimizes $S(\mathbf{w}, \lambda)$:

$$\nabla_{\mathbf{w}, \lambda} S = 0.$$

Straightforward computation shows that:

$$\nabla_{\mathbf{w}} S = 2 \text{cov}(\mathbf{b}, \mathbf{b}) \mathbf{w} - \frac{2}{N} \text{cov}(\mathbf{b}, \mathbf{b}_0) \mathbf{1} - 2\lambda \mathbf{1} \tag{24}$$

and

$$\nabla_{\lambda} S = \mathbf{1}^T \mathbf{w} - 1. \tag{25}$$

Then the weights \mathbf{w} can be found from the following systems of linear equations for $\nabla_{\mathbf{w}, \lambda} S = 0$, which can be written in matrix form:

$$\begin{bmatrix} \text{cov}(\mathbf{b}, \mathbf{b}) & -\mathbf{1} \\ -\mathbf{1}^T & 0 \end{bmatrix} \begin{bmatrix} \mathbf{w} \\ \lambda \end{bmatrix} = \begin{bmatrix} \frac{1}{N} \text{cov}(\mathbf{b}, \mathbf{b}_0) \mathbf{1} \\ -1 \end{bmatrix}, \tag{26}$$

where $\text{cov}(\mathbf{b}, \mathbf{b})$ denotes the covariance of the variables β in \mathbf{b} , and each element in $\text{cov}(\mathbf{b}, \mathbf{b}_0)$ is the covariance function between the corresponding observed values of β in \mathbf{b} and \mathbf{b}_0 (Appendices 1 and 2).

2.2 Learned heuristic search

Heuristic function Heuristic search [19] is an efficient approach for solving sequential decision-making problems by repeatedly expanding the partial path with the largest estimated value until a complete path for which the true objective-function value is higher than the estimated values of all evaluated partial paths is found. The estimated value of a partial path could be obtained using an evaluation function $f(n_s)$, which estimates the value of objective-function for the optimal complete path constrained to the previous actions. This estimated objective-function value consists of two elements:

$$f(n_s) = g(n_s) + h(n_s), \tag{27}$$

where $g(n_s)$ is the true reward from the initial state to a specific state n_s through a set of selected actions, and $h(n_s)$ is a heuristic function that estimates the maximum reward from current state n_s to a goal state.

For drilling-order problem, the objective of robust optimization is maximization of expected NPV over uncertainty by optimizing the drilling sequence of wells. The expected NPV computed using mean model bias

correction (MMBC) can be mathematically represented as:

$$J = \alpha(x_i) \left\{ \sum_{j=1}^T \frac{(q_{o,j}^{\bar{m}} r_o - q_{w,j}^{\bar{m}} r_w - q_{wi,j}^{\bar{m}} r_{wi}) \Delta t_j}{(1+b)^{t_j/\tau}} - \sum_{n=1}^{N_w} \frac{W_n}{(1+b)^{t_n/\tau}} \right\}, \tag{28}$$

where $q_{o,j}^{\bar{m}}$, $q_{w,j}^{\bar{m}}$, and $q_{wi,j}^{\bar{m}}$ denote the rates of produced oil, produced water, and injected water, respectively, from the mean model in m^3/day ; r_o , r_w , and r_{wi} are the oil price, water production cost, and water injection cost, respectively; T represents the number of time steps; t_j is the cumulative time in days up to time step j ; Δt_j is the time interval in days; b is the discount rate for a certain reference time τ (365 days); N_w is the total number of drilling wells; W_n denotes the cost of drilling the n th well; t_n is the cumulative time in days up to the open time for each well; $\alpha(x_i)$ is the bias correction factor for the corresponding control x_i .

The cost of finding a strong and admissible heuristic for drilling-problem can be prohibitive since the evaluation of heuristic function requires simulations. In this paper, we use a heuristic function in which all remaining wells are drilled simultaneously at the next step [28, 40], as this estimate can be obtained inexpensively and generally provides an overestimate of the NPV. Such a heuristic is guaranteed to find the true optimal drilling order. However, it might lead to an exhaustive search due to large estimated values.

Online learning techniques To efficiently find a solution that is optimal or near optimal, the evaluation function $f(n)$ in Eq. 27 should be close to the true maximum value $f^*(n)$. It is difficult to design a heuristic function that is accurate in all situations, but a crude heuristic function can be improved by online-learning techniques, i.e., estimate the error of the initial approximate value by learning the observations from previous decision steps. From a set of available online-learning mechanisms $\Phi_1, \Phi_2, \Phi_3, \dots, \Phi_n$, the best-improved evaluation function $\hat{f}_{\Phi}(n)$ with multiple online-learning techniques is defines as:

$$\hat{f}_{\Phi}(n) = \max \left(\hat{f}_{\Phi_1}(n), \hat{f}_{\Phi_2}(n), \hat{f}_{\Phi_3}(n), \dots, \hat{f}_{\Phi_n}(n) \right), \tag{29}$$

which might not be the most accurate value, but it is more likely to overestimate the actual maximum value and guide the heuristic search close to the optimal solution.

Wang and Oliver [40] have proposed two possible online-learning techniques (i.e., single-step adjustment and multiple-time-periods learning) for improving the initial approximate values obtained from heuristic sequences (i.e., all the remaining wells are drilled simultaneously and

opened at the next time step) by estimating a set of forecast errors of $f(n)$ and $h(n)$ for the remaining decision steps.

The single-step adjustment is defined as:

$$\hat{f}(n_s) = f(n_s) \prod_{n=n_s}^{n_{\text{goal}}} \hat{\gamma}_{fn} = f(n_s) \prod_{i=1}^{d(n_s)} \gamma_{fn_s} \bar{\mu}_{n_s}^i$$

$$= f(n_s) \gamma_{fn_s}^{d(n_s)} \bar{\mu}_{n_s}^{\frac{(d(n_s)+1)d(n_s)}{2}}, \tag{30}$$

where $d(n_s)$ is the number of remaining actions at n_s , and $\hat{\gamma}_{fn}$ is the forecast error of $f(n)$ in future decision step n , which is estimated by using the ratio γ_{fn_s} associated with $f(n_s)$ and $f(n_{s-1})$ and the mean single-step ratio $\bar{\mu}_{n_s}$ of $\gamma_{fn_1}, \gamma_{fn_2}, \dots, \gamma_{fn_s}$ along the current optimal path.

Multiple-time-periods learning is calculated by correcting the heuristic values of various time periods simultaneously:

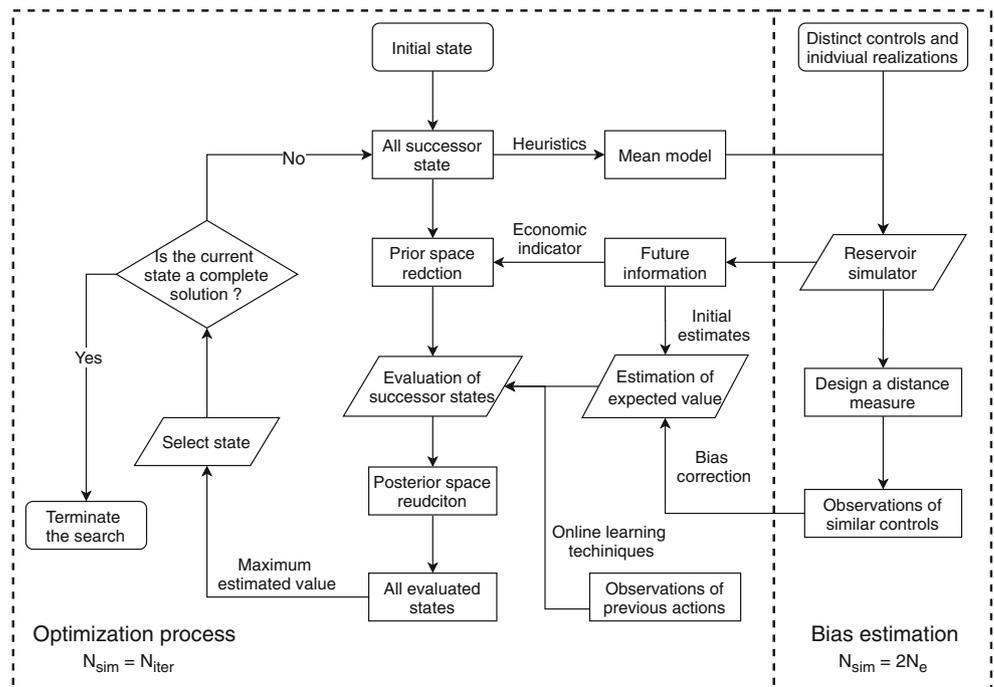
$$\hat{f}(n_s) = h_{\Delta t_1^\phi}(n_s) \prod_{n=n_s}^{n_{\text{goal}}} \hat{\gamma}_{h_{\Delta t_1^\phi}(n)} + \sum_{i=2}^{N_L} \left(h_{\Delta t_i^\phi}(n_s) \prod_{n=n_s}^{n_{\text{goal}}} \hat{\gamma}_{h_{\Delta t_i^\phi}(n)} \right). \tag{31}$$

Note that estimated value $h_{\Delta t_1^\phi}(n_s)$ of the first learning period Δt_1^ϕ contains both true contribution to the expected NPV at time periods $t_s = \sum_{i=1}^{N_s} \Delta t_i$ from N_s sequentially drilled wells and heuristic values at a certain time period $\Delta t_1^\phi - t_s$ when all wells are open.

To summarize, multi-learned heuristic search with space reduction (MLHS-SR) based on the economic indicator

[15] and improved evaluation function $\hat{f}_\phi(n)$ (29) can be used to find a solution to an optimization problem faster without losing quality [40]. Figure 1 shows the flowchart of using MLHS-SR with the bias-corrected mean model for the drilling-order optimization problem under geological uncertainty. To estimate the bias correction factor α for different drilling sequences, we first sample N_x distinct controls and apply them to individual model realizations to obtain the initial observations of β . This step requires $2 \times N_x$ simulations. Both bias-correction methods (i.e., improve the estimates of expected NPV) based on the observations from distinct controls and model realizations, and online learning techniques (i.e., improve the estimates of maximum expected NPV) based on the observations from previous drilling steps do not require any simulations. Hence, we only need to perform one additional simulation in the mean model at each iteration for evaluating the expected NPV of one specific control. The use of bias correction applied to the mean model allows the application of fairly general optimization methods. During the optimization process, it requires only simulations in the mean model for obtaining initial estimates of expected NPV. For an ensemble with N_e realizations, we assume that the information from N_e distinct controls is used to estimate the bias correction factor, and the optimization process requires N_{iter} iterations (i.e., N_{iter} different controls have to be evaluated to obtain the optimal solution). Then, the total number of simulations required in RO is $N_{\text{tot}} = 2N_e + N_{\text{iter}}$, in which $2N_e$ simulations are performed to obtain N_e observed values of β and N_{iter} simulations are

Fig. 1 Flowchart of using MLHS-SR with mean model bias correction



performed in the mean reservoir model to obtain the initial approximations of expected NP for N_{iter} different controls. If N_{iter} different controls were to be evaluated using SAA, the cost in RO would be $N_{tot} = N_e \times N_{iter}$, which is much more expensive and will increase linearly with the ensemble size N_e and the number of iterations N_{iter} .

Depth-limited search The reservoir model will almost certainly be updated based on information obtained from drilling the first few wells. After model updating, the optimal order of the remaining wells will differ from the initial estimate of optimal order. It is therefore more important to correctly identify the first few wells to be drilled rather than provide an entire drilling sequence. One advantage of using learned heuristic search for the drilling-order problem is that a partial solution of the first wells can be obtained at a reduced cost by cutting off the search at a specified depth.

The most straightforward approach to partial sequence optimization is through the depth-limited search (DLS) [27], in which the learned heuristic search is terminated at a certain depth. To find the solution of the first N_s wells, we prefer to terminate the search at the first-visited best partial path with $N_s + 1$ selected wells, because online learning techniques with more observations along a longer path can further improve the approximations and potentially generate a better solution. Note that the last well along the first selected partial path of $N_s + 1$ wells might not be the optimal well because the search might change direction after evaluating its extended paths.

A faster approach to finding the partial solution is to use iterative depth-limited search (IDLS). It works by iteratively optimizing the next well based on the first selected partial path with two more wells until a solution of $N_s + 1$ selected wells is found. Because only the partial paths extended from previous decisions are considered, this approach can avoid evaluating unnecessary paths along other directions caused by underestimated values. Learned heuristic search with accurately estimated values generally will not change direction frequently, so that the optimized sequence of the first few wells with a limited depth is likely to be near the final optimized drilling sequence.

3 Case study 1: Flow optimization in one-dimensional model

The purpose of this simple example is to thoroughly investigate methods of efficient optimization on a flow problem for which the dependence of the objective function on the model parameters is highly nonlinear. The problem is chosen to be small enough that we can evaluate both linear

and quadratic Taylor series approximations of the objective function, and can evaluate the correct optimal solution.

The objective in this example is to locate an injection well, operating at fixed pressure in a one-dimensional flow domain with uncertain permeability, such that the total flow rate out of the reservoir through fixed pressure boundaries is minimized. The reservoir is discretized with 150 grid cells (i.e., 150 possible injector locations), each of length Δx , and cross-sectional area A . The permeability k_i is distributed as log-gaussian. Pressures at both ends of the grid are fixed at 0. An injector is located at the interface between cells $i = i_w - 1$ and $i = i_w$, where the pressure is fixed at $P = P_w$. Instead of permeability, we use the log-permeability θ_i as a parameter in the problem, $k_i = \exp(\theta_i)$.

In this notation, the total flow rate is the sum of the flow to the left and flow to the right:

$$q(i_w, \theta) = \frac{AP_w}{\mu \Delta x} \left(\sum_{i=0}^{i_w-1} \exp(-\theta_i) \right)^{-1} + \frac{AP_w}{\mu \Delta x} \left(\sum_{i=i_w}^{N_x} \exp(-\theta_i) \right)^{-1} \tag{32}$$

where all variables are in consistent units. Because the permeability is uncertain, we minimize the *expected value* of flow rate q , by adjusting the location of the injection well, i_w . To avoid a problem in which the optimal injector location is at the center of the reservoir because of symmetry, we modify the prior probability for θ by assuming observations of θ at four locations ($x = 10, 30, 50, 70$). The mean of the posteriori distribution for θ is shown as the orange curve in Fig. 2a, along with the true log-permeability (blue). Ten realizations of the log-permeability field from the posteriori distribution are shown in Fig. 2b.

The most straightforward approach to approximating the expected value of the objective function is through the use of SAA or averaging of Monte Carlo samples from the posterior distribution for θ . We use SAA with 400 samples for each control location as a benchmark for other methods. Hence, the SAA method uses 150×400 function evaluations to generate the expected value of $q(i_w, \theta)$ for optimization.

Taylor series expansions of the objective function provide much less expensive approximations of the expected value of $q(i_w, \theta)$. For the linear approximation, the expected value of $q(i_w, \theta)$ is approximated using the mean of the log-permeability field (5). The quadratic approximation is considerably more expensive as it requires computation of the second derivative of the objective function with respect to well location. For this 1D steady flow problem, that is still a manageable computation. Figure 3 compares the linear and quadratic Taylor series approximations of the expected value of flow rate to the sample average approximation.

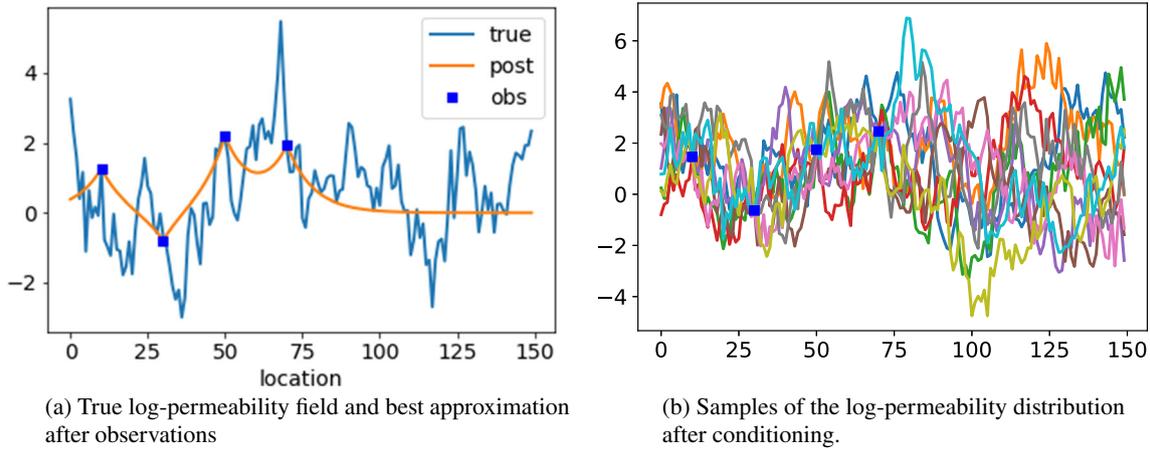


Fig. 2 a, b Light blue curve shows the true θ that generated observations. Blue squares show observations of θ with noise added. Correlation range of prior distribution for log-permeability is 26

Although there is clearly a large bias in the values from both the linear and quadratic approximations, the shapes are quite similar to the SAA, and the location of the minimum for each curve is also approximately the same. For this particular problem (minimizing total flow rate), a uniform bias does not affect the optimization result.

To correct the nonuniform bias in α , we use a regularized localization approach (20) in which higher weights are given to samples that are closer to the control variable for which the bias correction is being estimated. In this optimization problem, the actual distance between well locations is an appropriate measure of similarity. The variability of values of β is quite large (Fig. 4a) because the rate is strongly affected by the occurrence of low permeability values between the injector and the boundary. The variability in α is much smaller than the variability in

β (Fig. 4b). Unfortunately, while the variability in β can be estimated from samples, the variability in α will generally be unknown. Because it is computed by averaging over samples of β , however, we should expect it to generally be smaller than the variability in β , so that the ratio $\sigma_\alpha^2/\sigma_\beta^2$ which appears in Eq. 20 will generally be substantially smaller than 1.

In Figure 5, we compare the optimal well locations obtained using three different methods. In the SAA approach, we use N_e samples of the permeability field for each well location between 35 and 115 to compute an approximation of the expected value of flow rate. In the linear approximation, we simply compute the flow rate using the mean of the log-permeability field, and for the bias-correction approach, we used regularized localization with $\gamma = n_{\text{eff}}\sigma_\alpha^2/\sigma_\beta^2$ and a taper length of 20 to obtain a bias-correction to the linear approximation of expected value of flow rate. Each experiment was repeated 100 times to reduce the effect of sampling error. For both ensemble sizes, the bias-correction approach give better results than the linear approximation, although the difference is more apparent at larger ensemble size because the spread in the results is reduced in that case.

The most important criterion for judging success of the methodology is the ability to actually minimize the flow rate. Figure 6 compares the distribution of total flow rates obtained from the optimal well locations applied to the truth case for each of the methods. The test was repeated 100 times so there are 100 different truth cases and 100 “optimal” locations for each method. Note that the bias-correction method is now clearly superior to the linear approximation, even in this example in which the optimization depends only on relative differences in the value of the objective function. Optimization of the well

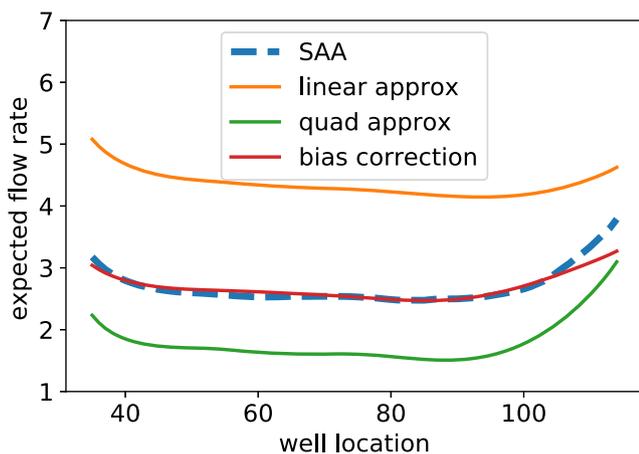


Fig. 3 The expected value of the total injection rate, $E[q(x, \theta)]$, computed using the sample average approximation, and two levels of Taylor series approximation

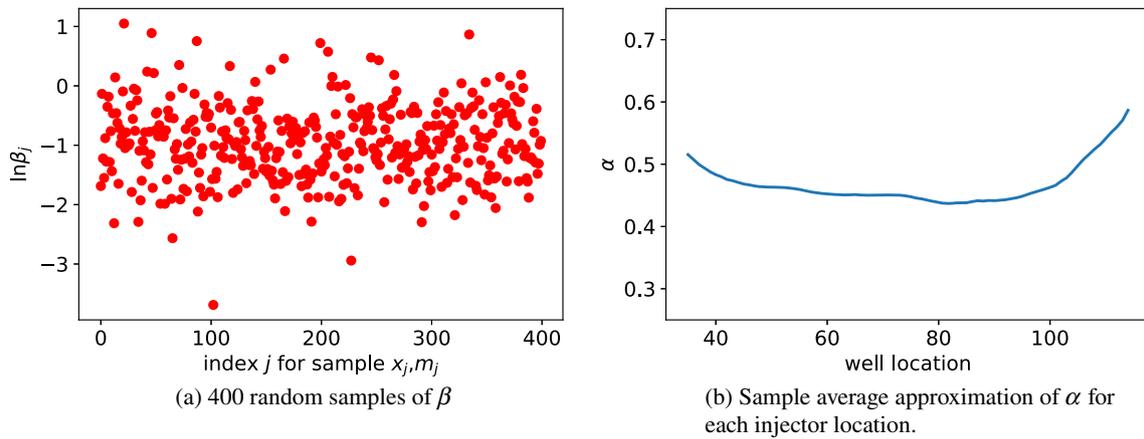


Fig. 4 a, b Variability of the partial bias correction β and the bias correction factor α for the 1D flow problem

location based on single realizations from the posterior distribution for model parameters gives very poor results (Fig. 6b).

4 Case study 2: Drilling-order problem in synthetic model

4.1 Reservoir model

In this example, we design a set of experiments to study the bias properties of the drilling-order problem applied to the synthetic REEK Field model [18, 29, 40], and the performances of various methods for estimating the bias for both complete and heuristic sequences by using the information from distinct controls and individual realizations. The objective of the robust optimization problem is maximization of the 10-year expected NPV with respect to the drilling schedule of wells. Learned heuristic search applied to the mean model with bias correction is used to optimize the drilling sequence under geological uncertainty. To illustrate the quality of the robust optimal

solution, we compute both optimal solution based on SAA and deterministic solutions of individual realizations. Moreover, we also investigate the possibility of optimizing the first few wells by using learned heuristic search with limited search depths.

The REEK model is a three-phase black-oil reservoir model with $40 \times 64 \times 14$ grid cells, of which 34,770 are active cells. An ensemble of 100 geologically consistent model realizations is used to empirically represent uncertainty in the porosity field, permeability field, and fault transmissibility multipliers. Table 1 shows the reservoir properties and control variables in REEK field. For the test problem, we assume that eight vertical fully penetrating wells (5 producers and 3 injectors) with fixed locations need to be drilled sequentially, and the first well is drilled at the beginning of simulation. The assumed drilling period is 6 months for all wells and wells begin operating immediately after drilling. Figure 7 shows the well locations and initial oil saturation of one randomly chosen model realization. The injectors are positioned around the oil-water contact and controlled by a maximum injection rate of 10,000 m³/day and a maximum BHP of 320 bars. The producers are distributed throughout the oil-containing area and controlled by a maximum production rate of 6000 m³/day and a minimum BHP of 250 bars.

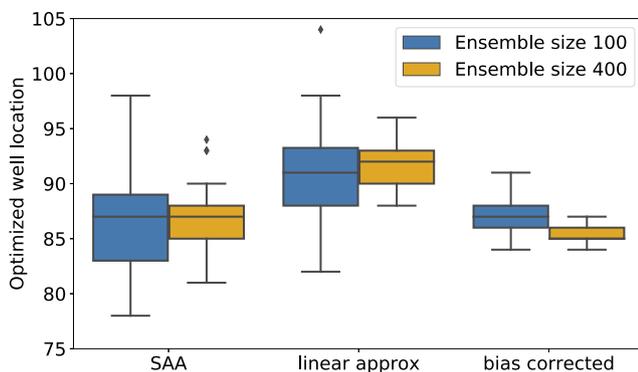


Fig. 5 Optimized well locations for three different methods and two ensemble sizes

4.2 Estimation of correction factor

Properties of α and β We randomly select 5 different complete drilling sequences and apply them to the entire ensemble of realizations and the reservoir mean model to obtain the actual values of the multiplicative bias correction factor α (7) and their partial correction factor β of individual realizations (8). Figure 8 shows the observed values in α and β from 5 random complete drilling sequences. Symbols that are in the same color indicate the partial correction factor $\beta(x_i, m_j, \bar{m})$ at a fixed drilling sequence x_i applied

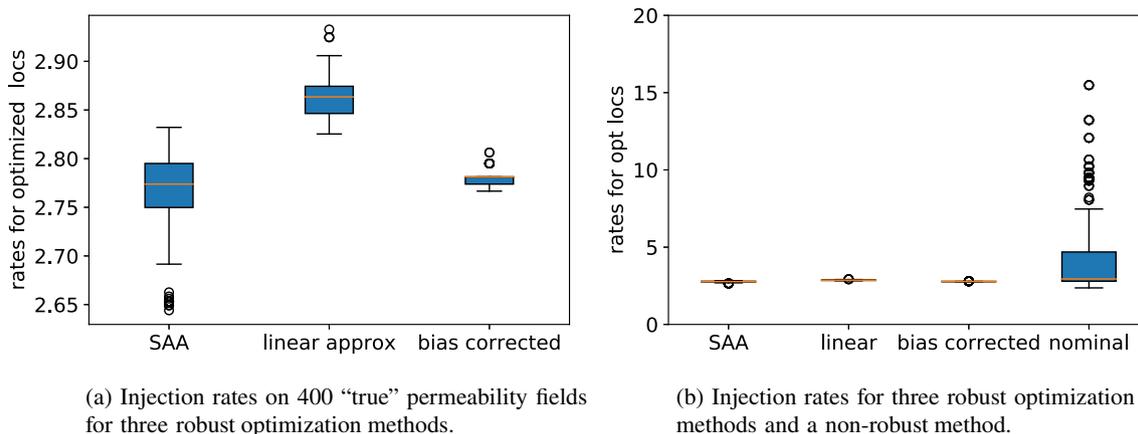


Fig. 6 a, b Values of the objective function from optimization of well location in the 1D flow problem (Small values are better)

to different reservoir models and the mean model. The x -axis displays the index j of a single reservoir model m_j . We observe that $\beta(x_i, m_j, \bar{m})$ at a fixed control changes significantly with geological uncertainty, but $\beta(x_i, m_j, \bar{m})$ at a fixed reservoir model changes slightly with different controls. The horizontal dashed lines represent the bias correction factors α_i for the 5 drilling sequences, computed by averaging over samples of β . The variability in α_i for these 5 drilling sequences is very small and the mean for each is close to 1.1.

To obtain more statistically reliable results, we sample 560 different drilling sequences for studying the bias properties. For 5 producers and 3 injectors, there are 56 possible combinations of drilling sequences based on the types of wells (e.g., [P, P, P, P, W, P, W, W], [P, P, W, P, P, P, W, W]). We randomly select 10 distinct drilling sequences from each combination and compute their actual values of α and β . Figure 9 shows the distributions of α

and β values obtained from these 560 random complete drilling sequences. α has a value near 1.108 for most of the drilling sequences (Fig. 9a), indicating that SAA of NPV is almost 10% higher than the NPV computed using the mean model. To accurately estimate the expected NPV, it is necessary to correct the bias in the initial approximation that is obtained in the reservoir mean model. Out of the 560 drilling sequences, only 22 (approximately 4%) sequences have α values lower than 1.09, whereas β changes between 0.6 and 1.5 with a larger variability (Fig. 9b). Interestingly, almost all of the drilling sequences with atypical bias correction factor α are from two extreme drilling sequences, i.e., either all producers or all injectors are drilled first. The results of these two combinations are plotted in yellow. This finding was unexpected and suggests that to more accurately estimate the bias for general drilling sequences, we should avoid sampling controls from these two extreme combinations since they seem to provide less useful information and vice versa.

Figure 10 compares the distributions of variance in β at a fixed reservoir model and a fixed control after eliminating two extreme combinations (i.e., [P, P, P, P, P, W, W, W] and [W, W, W, P, P, P, P, P]). As previously observed in Fig. 8, the variability of β among different drilling sequences at a fixed reservoir model is smaller than the variability of β among different reservoir models at a fixed drilling sequence. The average value of α can be estimated by averaging all observed values in β from distinct controls. In this example, the variance in α is only 0.000058 for general drilling sequences. The estimates of expected NPV with a bias correction factor $\bar{\alpha}$ could be accurate for most of the drilling sequences, but this will not always be the case in other problems. For the case of large variability in α , we should consider the similarity between the samples and control variables for estimating the bias for specific controls. In this case, the variability in β is much larger than

Table 1 Reservoir properties and control variables in REEK model

Field	REEK model
Number of grid blocks	40 × 64 × 14
Number of active cells	34770
Permeability (md)	0 to 3500 (average 733)
Porosity	0 to 0.45 (average 0.159)
Fault transmissibility multiplier	0 to 1 (average 0.105)
Number of geological realizations	100
Number of wells (all vertical wells)	8 (5 producers and 3 injectors)
Number of all possible drilling sequences	8! = 40320
Maximum production rate (m ³ /day)	6000
Minimum BHP of producers (bars)	250
Maximum injection rate (m ³ /day)	10000
Maximum BHP of injectors (bars)	320

Fig. 7 Well locations and initial oil saturation in one realization of the REEK model ensemble

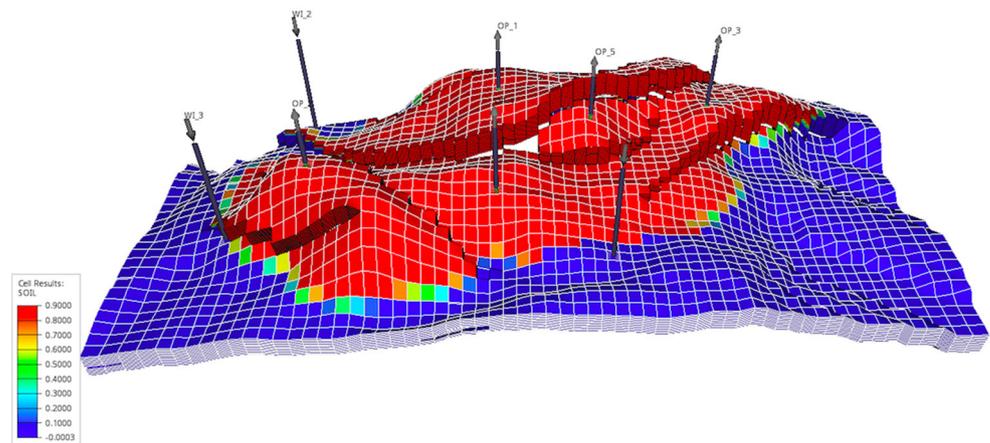
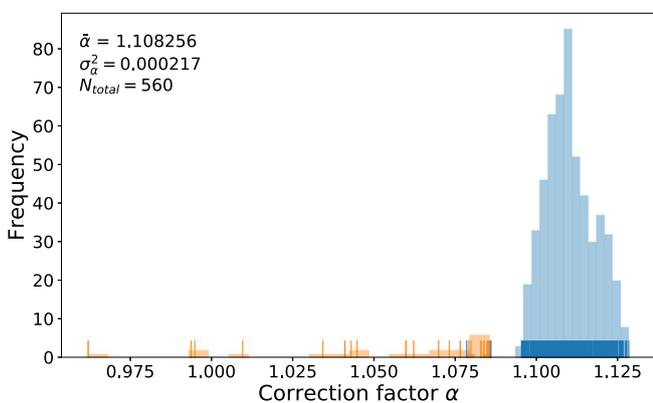
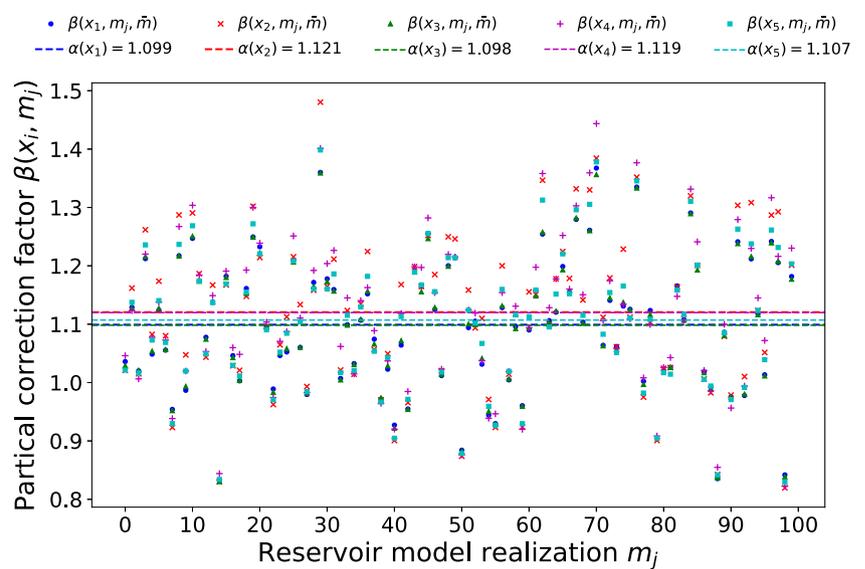
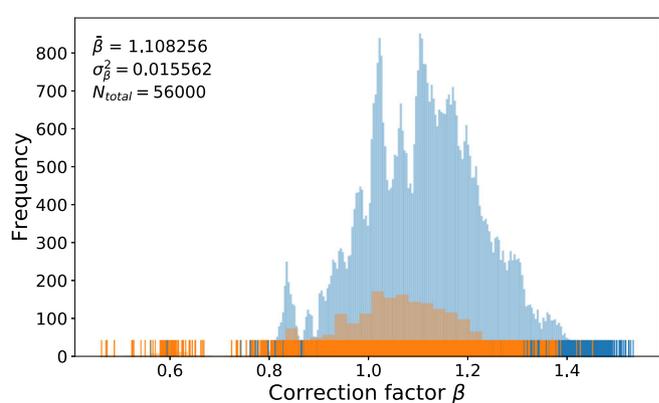


Fig. 8 Comparison of bias correction factor α and partial correction factor β of 5 different complete drilling sequences



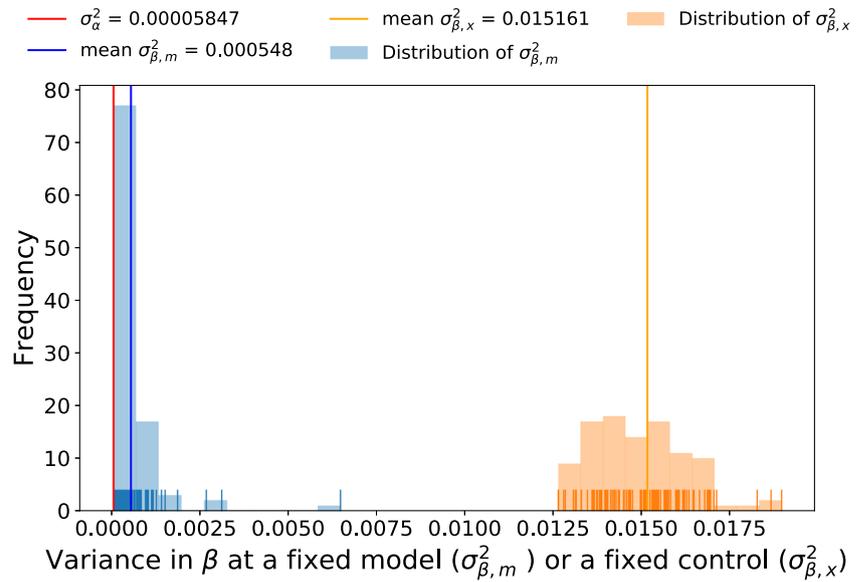
(a) Distribution of bias correction factor α



(b) Distribution of partial correction factor β

Fig. 9 a, b Comparison of the distributions of α and β of 560 different complete general drilling sequences

Fig. 10 Distribution of variance in β at a fixed model or a fixed control

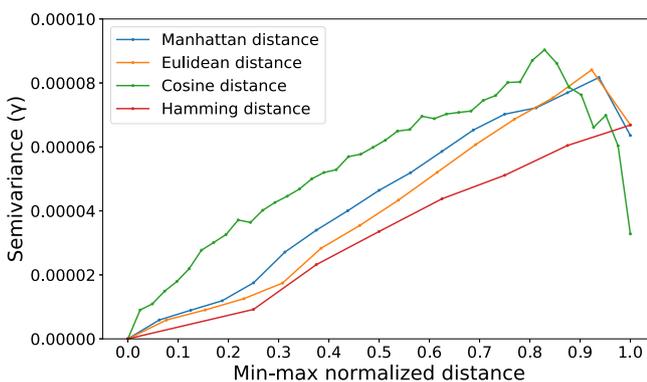


the variability in α for general complete drilling sequences, and the ratio $\sigma_\beta^2/\sigma_\alpha^2$ is almost 267. For a small number of samples, local estimates $\hat{\alpha}_{loc}$ based on β (12) might be away from the actual bias correction factor. The error of local estimates caused by sampling error due to small sample size could be as large as 40%. Hence, it is critical for this problem to use a taper length with a relatively large effective sample size for pure distance-based localization or apply regularized localization to avoid generating estimates that are far away from $\bar{\alpha}$.

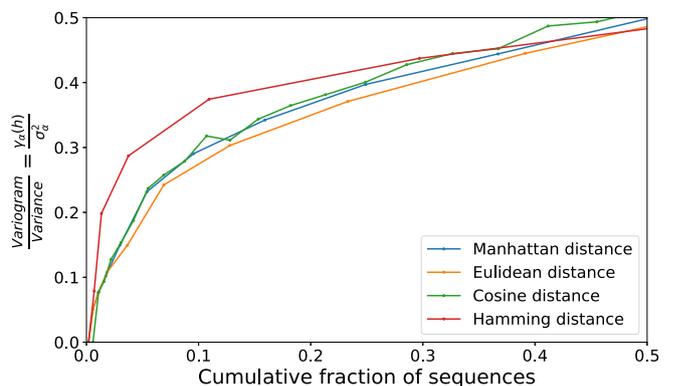
To measure the similarity of drilling sequences for the bias correction factor, we investigate the use of four different distance metrics (i.e., the Manhattan, Euclidean, cosine, and Hamming distances). Figure 11a shows the empirical variograms of α for these four distance metrics. We observe that all four variograms show an approximately linear increase with distance over most of their ranges.

Overall, each of these four distance metrics can measure the similarity of drilling sequences for expected NPV. For the purpose of identifying similar drilling sequences, a measure of distance is better than another if it correlates better over larger numbers of samples.

Figure 11b shows the cumulative fraction of drilling sequences within an upper bound on the variogram of α . Among a set of random controls, very few of them will have small Hamming distances. We note that, compared with the other three distance metrics, the Hamming distance yields fewer similar drilling sequences at a fixed threshold variance of α . In other words, when a fixed fraction of samples is considered for local estimation, drilling sequences extracted by the Hamming distance will potentially have more significant variability in α . By contrast, the Manhattan, Euclidean, and cosine distance metrics have similar behaviors and all perform better than



(a) Empirical variogram of the bias correction factor α as obtained using different distance metrics



(b) Cumulative fraction of sequences within a certain upper bound on the variogram of α

Fig. 11 a, b Performance comparison of four distance metrics applied in the drilling-order problem

the Hamming distance for local estimation based on the observations from randomly selected drilling sequences. In our subsequent experiments, we use the Manhattan distance to measure the distance between two drilling sequences for estimating α values for specific controls.

Accuracy of estimates of α In this experiment, we compare various methods for estimating the bias correction factor α of different drilling sequences. We first sample a number (N_x) of distinct general drilling sequences to obtain a set of observations of β , which requires $2 \times N_x$ simulations. To evaluate the quality of each of the three approaches (i.e., average bias correction factor, pure-distance localization, and regularized localization), we use each method to estimate α for 540 different drilling sequences. Estimation of α for various drilling sequences using these approaches does not require additional samples of β , but to estimate the expected NPV of a specific drilling sequence, we need to perform one additional simulation using the mean model to obtain the initial approximate value of expected NPV. In other words, the amount of computation required for computing the expected NPV of 540 sequences by using the information from 100 distinct controls and individual realizations is $2 \times 100 + 540 = 740$ simulations, which is much less than the cost for SAA (54,000 simulations). In this experiment, we use the accuracy of the estimated α as a criterion for judging the quality of the expected NPV based on the reservoir mean model with bias correction.

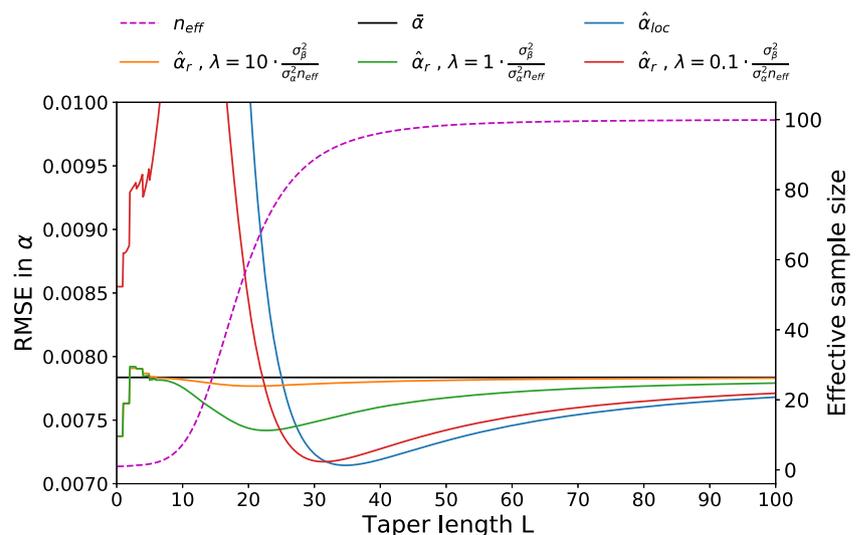
Figure 12 shows RMSE of the estimates of the 540 values of α based on a set of 100 observed values in β . The black horizontal line indicates the result from a fixed estimated α obtained by averaging samples of β (11). Pure distance-based localization and regularized localization are applied with different taper lengths between 1 and 200 (for reference, the maximum Manhattan distance between

complete sequences with eight wells is 32). We observe that purely local estimates $\hat{\alpha}_{loc}$ are potentially more accurate than the estimates computed by averaging over samples of β when a taper length of $L > 26$ is used. In this case, the effective sample size is relatively large ($n_{eff} > 80$), and the best value of the taper length in this test is approximately 34. When the taper length extends over 170, all samples are assigned high weights and the performance is close to the approximation of $\bar{\alpha}$. On the contrary, the quality of local estimates is very sensitive to the taper length because of sampling error when the taper length is small (e.g., $n_{eff} \approx 10$ at $L = 15$). The error of estimated values resulting from the small number of samples is significant.

Since our estimate of optimal weighting in regularized estimation of α depends on the value of σ_α^2 , which we are unlikely to know accurately, we investigated the sensitivity of results to variation in the magnitude of the weighting term. Three different levels of the ratio of the variance of α and β (i.e., $\lambda = 0.1 \cdot \sigma_\beta^2 / \sigma_\alpha^2 n_{eff}$, $\sigma_\beta^2 / \sigma_\alpha^2 n_{eff}$, and $10 \cdot \sigma_\beta^2 / \sigma_\alpha^2 n_{eff}$) were evaluated. For $\lambda = 10 \cdot \sigma_\beta^2 / \sigma_\alpha^2 n_{eff}$, the regularized estimates $\hat{\alpha}_r$ are close to $\bar{\alpha}$ and the RMSE is similar. For $\lambda = 0.1 \cdot \sigma_\beta^2 / \sigma_\alpha^2 n_{eff}$, the effect of regularization is small; consequently, the estimates at a small taper length are inaccurate due to sampling errors. However, the results show that $\lambda = \sigma_\beta^2 / \sigma_\alpha^2 n_{eff}$ reduces the sensitivity of the estimates to the taper length, while still generating more accurate values even for cases with a small number of samples. As a result, the corresponding regularized estimates are potentially more accurate for a wider range of taper lengths.

To reduce the influence of sampling error on the conclusions, we repeat each experiment 100 times, each time using a fresh sample of controls to obtain observations of β . For every single test, we apply three different sample sizes ($N_x = 100, 200, \text{ and } 400$) to investigate the effect of

Fig. 12 Comparison of the RMSEs of estimates of α obtained using different approaches based on 100 observed values of β



the sample size on the accuracy of the estimated α and the best value of the taper length for distance-based localization. Figure 13 shows the RMSEs of α estimates obtained using different approaches with the three different sample sizes after repeating each test 100 times. In general, the estimates are improved by using more samples. The improvement, however, is small for estimates obtained without localization (horizontal dashed lines) as the estimate of the mean for β is already quite accurate for $N_x = 200$.

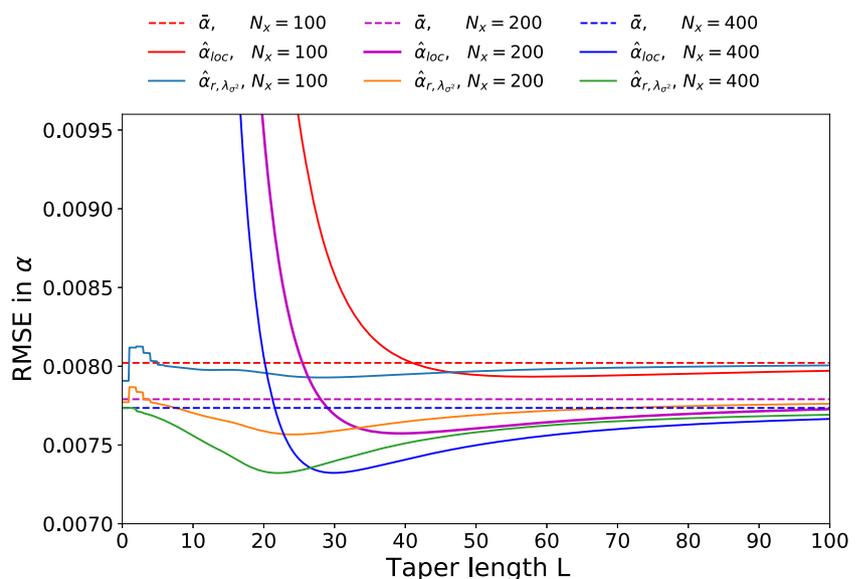
Distance-based localization offers a further increase in accuracy, and the improvement is more significant with a larger sample size. In this example, we observe that the RMSE of the local estimates at the optimal taper length is close to that of the estimates computed from the optimal weights based on the covariance of the partial correction factors (26). For the drilling-order problem, the well-position based distance distribution obtained from a fixed set of random controls changes slightly when the control variable for estimating the bias correction factor is varied. In such a case, the distributions of elements in covariance matrix $\text{cov}(\mathbf{b}, \mathbf{b}_0)$ (42) in terms of the distance are stable. The variance of the correction factor increases almost linearly with distance, making the optimal weight based on the covariance of partial correction factors decreases with the distance of the samples. As a result, the optimal weights would be close to the local weights at the optimal taper length for the drilling-order problem.

As observed previously, regularization with $\lambda = \sigma_\beta^2 / \sigma_\alpha^2 n_{\text{eff}}$ appears to always produce more accurate estimates of α than those obtained by averaging of all values of β . When the taper length is greater than a certain threshold, it seems that pure distance-based localization slightly outperforms regularized localization. The reason for this is that local estimates are potentially more

accurate than $\bar{\alpha}$ with large effective sample sizes; however, the regularization term reduces the effect of localization and results in estimates that are closer to $\bar{\alpha}$. Both the variance of α and the best value of the taper length are generally unknown. Therefore, instead of using regularized localization based on an uncertain value of $\sigma_\alpha^2 / \sigma_\beta^2$, for this problem, we use pure distance-based localization with a relatively large effective sample size to estimate α values for different controls.

We investigated the effect of taper length on the error in the estimate of bias correction factor, α , for three different sample sizes. For each sample size, we repeated the experiment 100 times to obtain reliable estimates. The results show that the optimal taper length for local estimates of α decreases as the sample size increases (Fig. 13). As the optimal taper length is also subject to sampling error, we investigated the variability. Figure 14 compares the frequency distributions of the optimal taper lengths for each ensemble size. When only 100 samples are used to estimate α , the best value of taper length is shown to be highly sensitive to the initial samples (histogram in blue). In approximately 1/3 of the cases, the optimal taper length is larger than 200, which means that the average value of all β will provide a better estimate than a purely distance-based weighting with a small taper length. For some controls, the actual α values might be very close to $\bar{\alpha}$, such that more samples will be required for local estimation to provide any further improvement in accuracy. When N_x is increased from 100 to 400, the optimal taper length is less variable, varying between 25 and 30 in most cases (histogram in green). It seems that for cases with sample sizes larger than the ratio $\sigma_\beta^2 / \sigma_\alpha^2$, the best value of the taper length is likely to lie near a point where the value of $\sigma_\beta^2 / \sigma_\alpha^2 n_{\text{eff}}$ is approximately 1. In that case, the variability in the local

Fig. 13 RMSEs of estimates of α with three different sample sizes



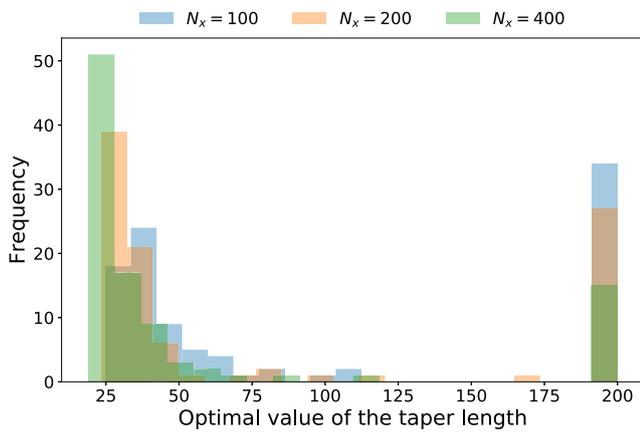
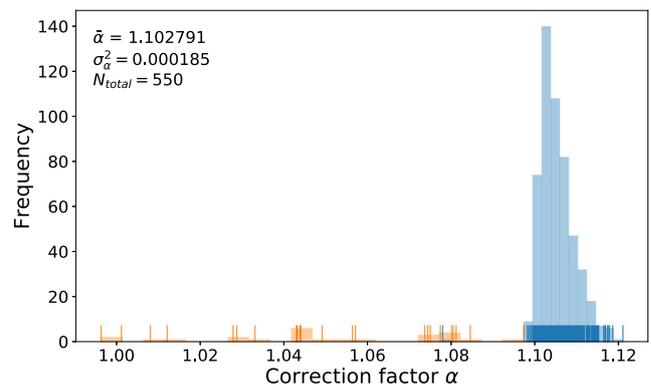


Fig. 14 Distribution of the optimal taper length for the pure distance-based localization

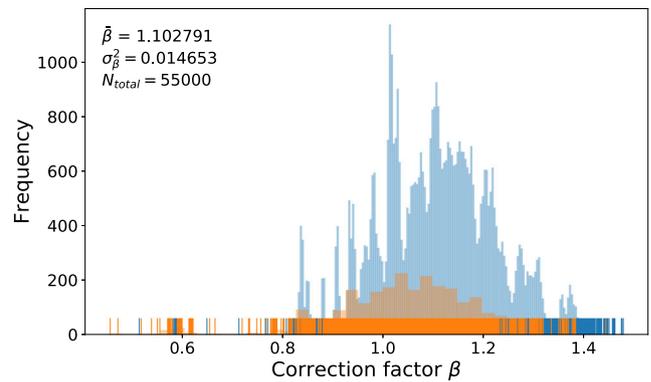
estimates will be close to the variability in the actual α values. When the variance of α is unknown, it might be difficult to select the best value of the taper length. However, the results show that the local estimates are not highly sensitive to the taper length when the effective sample size is relatively large, in which case local estimates are still more accurate than estimates obtained using the average value of β .

Correction factors for heuristic sequences During the heuristic search process, a set of heuristic sequences with different numbers of selected wells is used in combination with online learning techniques to estimate the maximum value among the possible complete drilling sequences constrained by previous wells. Hence, we need to compute the expected NPVs of some specific heuristic sequences to optimize either an entire drilling sequence or only the first few wells when applying a learned heuristic search. In this experiment, we study the bias properties of heuristic sequences and the possibility of using mean model bias correction to accurately estimate their expected NPVs.

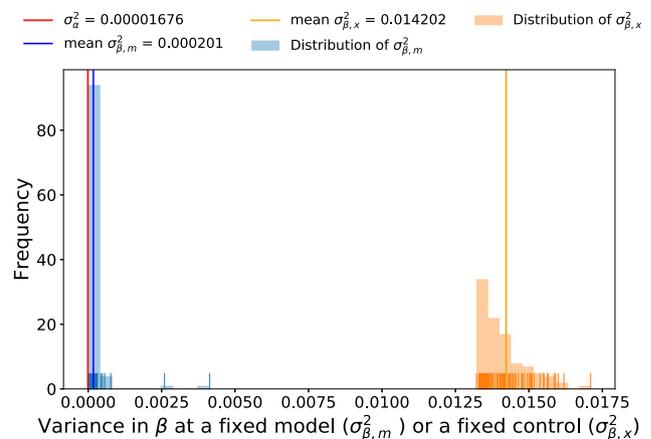
Figure 15a and b show the distributions of α and β from 550 random heuristic sequences, 400 of which have fewer than four selected wells (i.e., all possible heuristic sequences with $N_s \leq 3$ are considered). The histogram in orange represents the bias from the heuristic sequences in which either all producers or all injectors are drilled sequentially first. Figure 15c compares the distributions of the variance of β for a fixed reservoir model and a fixed heuristic sequence after the elimination of controls from those two extreme combinations. Relative to the observations for complete drilling sequences, the bias for heuristic sequences has some similar properties: (1) the $\bar{\alpha}$ for heuristic sequences is close to that for complete sequences, but the variability in α is smaller; (2) heuristic sequences in which either all producers or all injectors



(a) Distribution of the bias correction factor α



(b) Distribution of the partial correction factor β



(c) Distributions of the variance of β for a fixed model and a fixed heuristic sequence

Fig. 15 a–c Distributions of α , β , and σ_β^2 of heuristic drilling sequences

are drilled first have atypical and relatively smaller values of α ; (3) β values for heuristic and complete sequences have similar distributions, and the variability in β is much larger than the variability in α ; and (4) the variance of β among different heuristic sequences for a fixed model is much smaller than that among different realizations for a fixed heuristic sequence. We observe that the variance of

α for heuristic sequences is only 0.0000168. In this case, the expected NPV that is estimated from the mean model with $\bar{\alpha}$ will have a high level of accuracy. Distance-based localization might require a large number of samples to improve the estimates, and even then, the improvement will be small.

Figure 16 shows the empirical variogram of α for heuristic sequences based on the Manhattan distance. The number of selected wells constrains the distance between heuristic sequences. For eight wells, the maximum Manhattan distance between heuristic sequences with fewer than four selected wells is only 12. For the sequences in which at least half of the wells are drilled simultaneously ($N_s \leq 4$), while the variance of α increases with distance, the correlation of α is still strong even for the sequences with the largest distance. For $N_s \leq 5$ and $N_s \leq 6$, however, the variance of α decreases beyond a certain distance. The main reason for this is that heuristic sequences with a significant gap in N_s generally have a large distance, but they might also have similar values of α . If we sample heuristic sequences with a large N_s (i.e., in which most of the wells are drilled sequentially) to estimate the bias for sequences with a small N_s (i.e., in which most of the wells are drilled simultaneously), distance-based localization might not significantly improve the estimates even with a large number of samples, since most of the samples will be at long distances, but may still have similar α values. Therefore, to ensure the effectiveness of distance-based localization, we should use samples with a small gap in N_s to estimate the bias for specific heuristic sequences with N_s selected wells.

Figure 17 shows the RMSEs of the estimated α values for heuristic sequences with different numbers of selected wells obtained by using the information of distinct heuristic controls. As in the previous section, to reduce the sampling error, we repeat each experiment 100 times. As expected, the results show that the average value of β is a good

approximation of α for most heuristic sequences due to the small variability in α . The error in the estimates can be only slightly reduced from 0.0034 to 0.0032 by increasing the sample size from 100 to 400, but at the cost of 600 additional simulations. In this case, it is not necessary to use a larger sample size to improve the estimates. The estimated α based on 100 observed values of β is already very close to the actual α values, although this might not always be the case in other problems. The effect of distance-based localization is small even if we use $N_x = 400$, which means that more samples would be needed for distance-based localization to improve the estimates further. As observed in the case of complete drilling sequences, the regularized estimates obtained from heuristic sequences also prove that a regularization term based on $\bar{\alpha}$ with a parameter of $\sigma_\beta^2/\sigma_\alpha^2 n_{\text{eff}}$ reduces the sensitivity of the estimates to the taper length, such that the estimates are potentially more accurate for a wider range of taper lengths. The results for both complete and heuristic sequences illustrate the effectiveness of using MMBC to estimate the expected NPV. Although the effect of localization on improving the estimates is relatively small in this example, this might not always be the case in other problems. For cases with considerable variability in α , the improvement in accuracy achieved by using distance-based localization with an appropriate taper length or regularized localization will be more significant.

4.3 Robust optimization under geological uncertainty

Optimization of complete sequences In this experiment, we apply learned heuristic search with MMBC to compute the RO drilling sequence under geological uncertainty. A set of 100 random heuristic drilling sequences with different numbers of selected wells is used to generate 100 observations of the partial correction factor β . This requires

Fig. 16 Empirical variogram of α for heuristic sequences based on the Manhattan distance

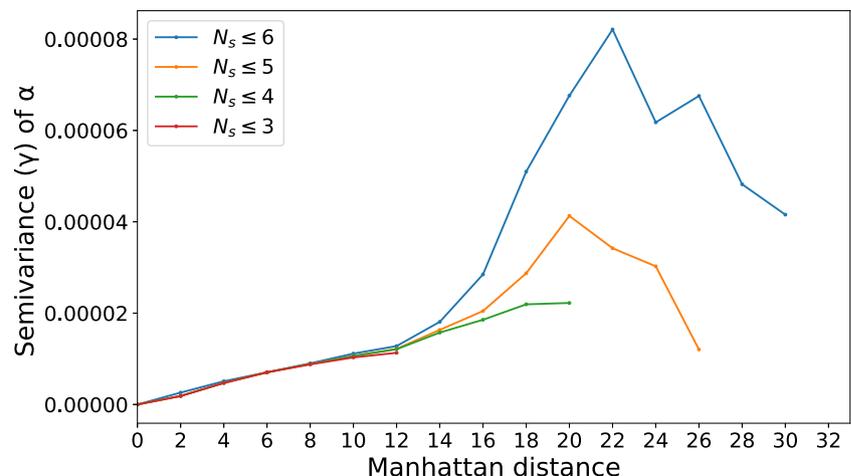
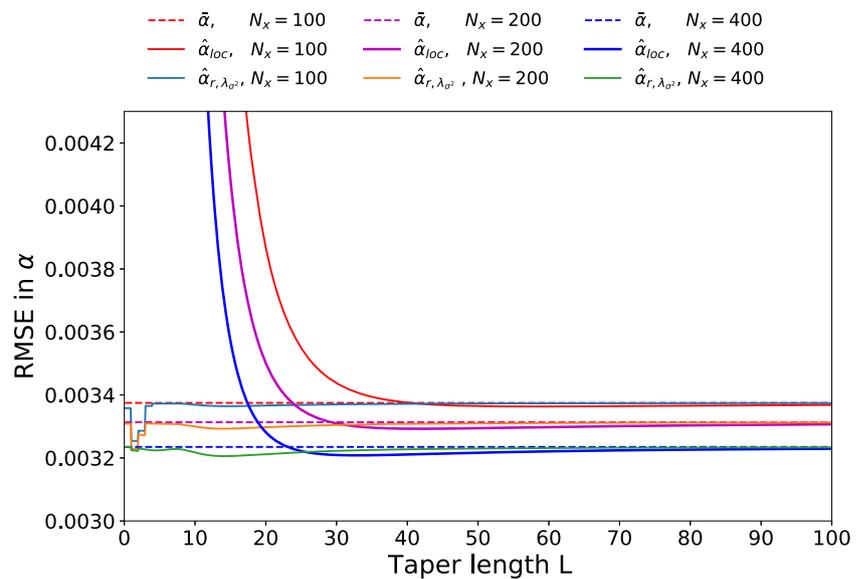


Fig. 17 RMSEs of estimates of α for heuristic sequences obtained using different approaches



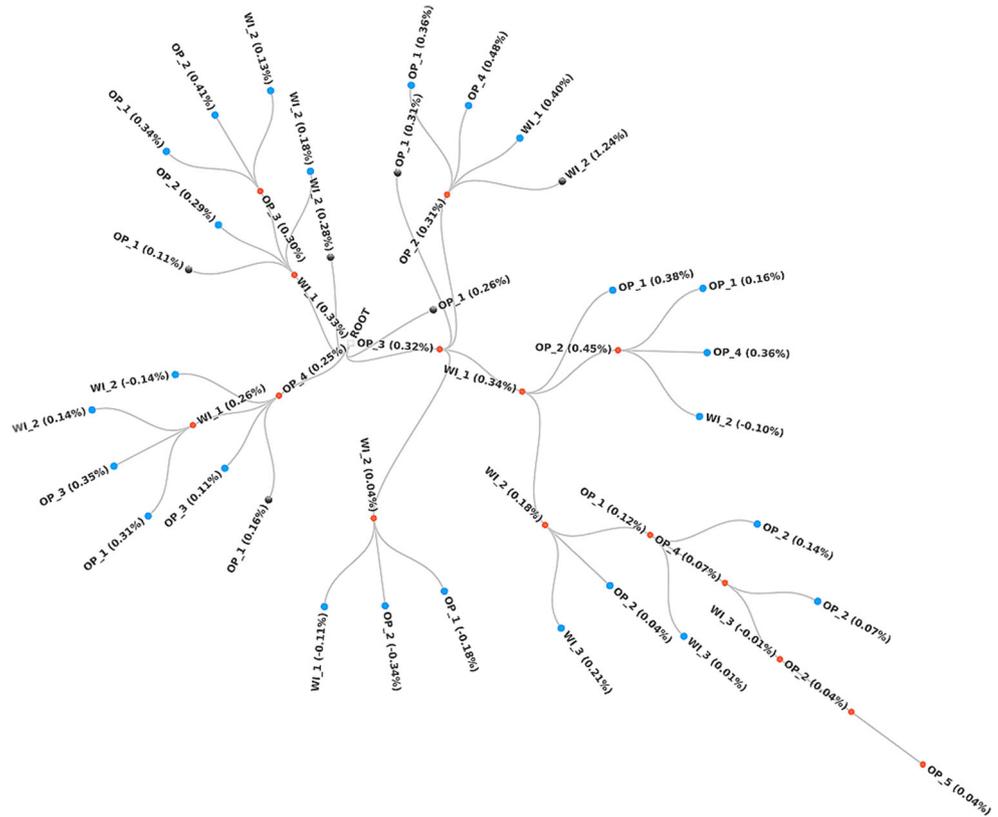
200 simulations, as the NPV must be computed for the drilling sequence run using the mean model, and for a model realization. Because the variance of α is unknown, we apply pure distance-based localization to estimate the bias for different drilling sequences. To avoid inaccuracy in the local estimates caused by sampling error due to a small sample size, we choose a taper length near 50, which is large enough to ensure that the effective sample size is relatively large. To evaluate the effectiveness of the various approaches, we compute the “gold standard” RO strategy based on SAA and the NO strategies of individual realizations. The performance of RO/NO optimal strategies is investigated by comparing their NPV distributions in the entire ensemble of 100 realizations with that of 560 random complete drilling sequences.

Figure 18 shows MLHS-SR search strategy applied to the mean model with bias correction. To speed up the search process while maintaining the highest possible solution quality after space reduction, we consider at least half of the remaining wells with high economic indicator values as the next possible actions (i.e., prior space reduction) and preserve three partial paths with the highest estimated expected NPVs for the remaining search process (i.e., posterior space reduction). For nodes marked in orange, the paths through them have the highest estimated expected NPV for complete sequences and have been expanded in the most promising directions. For each selected direction, only the partial paths go through blue nodes (or orange nodes if the paths have been extended), and black nodes are evaluated after prior space reduction. However, the paths with black nodes are pruned in the process of posterior space reduction due to their lower estimated values and will not be considered during the search process. In this case, the search evaluates 46 paths through 13 decision steps to optimize

an entire drilling sequence: [OP_3, WL_1, WL_2, OP_1, OP_4, WL_3, OP_2, OP_5]. Considering the cost of obtaining the partial correction factor β (200 simulations), only 246 simulations are required to compute the RO drilling sequence for eight wells. For comparison, when MLHS-SR is used with SAA, it is necessary to perform more than 3000 simulations to obtain a solution with the same quality. By more extensively reducing the search space, we can find the same solution faster (235 simulations). However, excessive space reduction (e.g., searching along only one direction) will lead to a drilling sequence with a less-than-optimal value (– 0.71%). Because a large fraction of the cost is a result of estimation of the bias correction factor, the cost reduction achieved by evaluating fewer controls might not be very significant, especially for small fields with large ensemble sizes. In such cases, it is not advisable to prune the space significantly to reduce the cost, because the computational cost might be reduced only slightly while yielding a solution with a suboptimal value.

The percentage presented near each node in Fig. 18 represents the error of the estimated expected NPV of each partial path (i.e., heuristic sequences) compared with SAA. The average error of the estimated expected NPV from the mean model is reduced from – 9.20 to 0.21% by estimation of the bias correction factor. The maximum error on the estimated expected NPV is only 0.54%. In this example, by using the information from only 100 samples of the drilling sequence and individual model realizations, the expected NPVs of different drilling sequences can be accurately estimated. The optimal drilling order from the learned heuristic search is computed by repeatedly expanding the search in the directions with the highest estimated values. Because the bias correction factor for the drilling-order problem in the REEK Field is relatively stable, adding such

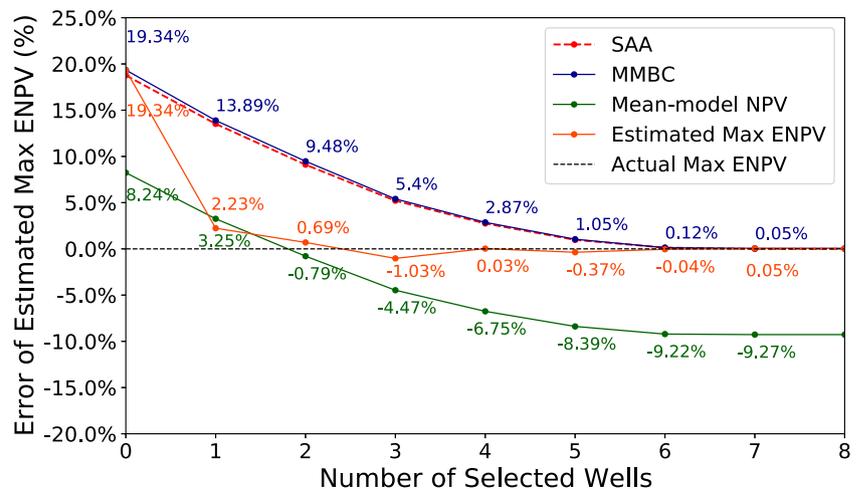
Fig. 18 Search strategy of MLHS-SR



a multiplicative factor does not significantly change the initial ranking of the heuristic sequences in the mean model. Hence, we observe that MLHS-SR performed directly in the mean model (i.e., linear approximation) finds the same solution obtained using the bias-corrected mean model. However, for the cases with considerable variability in α , a learned heuristic search with MMBC can potentially find a better drilling sequence than that obtained using the linear approximation.

Figure 19 shows the error on the maximum expected NPV of the complete drilling sequence along the optimal path that is estimated by using heuristic sequences with MMBC and online learning techniques. We first remove the bias in the initial approximations of the expected NPV obtained from the mean model by estimating a bias correction factor α (green curve \rightarrow blue curve), and we then correct the estimates of the maximum expected NPV obtained from heuristic sequences through online learning

Fig. 19 Error of estimates of maximum NPV along the optimal path



mechanisms (blue curve → orange curve). We notice that the initial estimates of the maximum expected NPV from heuristic sequences (blue curve) always overestimate the actual maximum value for the complete sequence and show an exponentially decaying trend with an increasing decay factor throughout the search process. However, the accuracy of the estimated maximum expected NPV is significantly improved by learning the errors on approximate values obtained from previous drilling steps (orange curve). In this case, the average error along the optimal path is only 0.22%. With such accurate estimates, the final optimized sequence is very likely to have a high expected NPV and to be near the true optimal sequence. When MLHS is applied without any space reduction, the average NPV of the optimized sequence is 0.60% lower than that of the solution obtained with space reduction. The main reason for this is that some estimated values underestimate the actual maximum value (e.g., the error for the partial path with three selected wells is – 1.03%), meaning that if the expected NPV of the current solution is higher than all of these underestimated values, the search will terminate with the current solution. In this experiment, an appreciable amount of space reduction (e.g., at least half of the remaining wells are considered as the next possible action) is shown to effectively prune paths that are less likely to correspond to the optimal drilling order, leading to a better solution with a lower cost.

Figure 20 compares the NPV distributions for RO/NO sequences and a set of 560 random drilling sequences. The results show that the average NPV is significantly increased by optimizing the drilling sequence with a learned heuristic search. RO sequence has a higher average NPV than either the NO solutions or the set of random drilling sequences, being as much as 25.86% higher than the average NPVs of some of the random drilling sequences. Table 2 summarizes the optimization results for each of the approaches. At the

cost of approximately 200 additional simulations, the RO strategy obtained from the mean model with bias correction is shown to be significantly better than the deterministic solutions for individual realizations. We observe many of the NO strategies have similar NPV distributions over a set of 100 realizations. It seems that in the REEK model, the optimal drilling sequence does not vary significantly with the geological uncertainty.

Although the geological properties vary significantly from realization to realization, the optimized deterministic well sequences for 100 model realizations have common characteristics. In particular, the optimized sequence for each single realization always starts with one producer and one injector, in that order. For 92 of the 100 individual realizations, the last well drilled in the optimal sequence is a producer. In most cases, Producer OP_3 and Injector WI_1 are drilled first, and Producer OP_5 is the last well drilled. In the REEK Field case, with five producers and three injectors, the optimal wells for these three drilling positions appear to be nearly independent of the geological properties. As a result, the RO solution is likely to be near several deterministic optimal solutions for this field. It is clear that this will not always be the case for other fields.

Partial optimization Wang and Oliver [40] have shown that learned heuristic search is a viable means of solving the sequential drilling-order problem for complete sequences. In many cases, however, only the first few wells in the drilling sequence may be needed. As shown in Fig. 18, the optimal partial path will not always extend in only one direction during the search. In such cases, it is not guaranteed that the first few wells can always be optimized along the final optimal complete sequence by limiting the depth of the search path. In this experiment, we investigate the possibility of optimizing the first few wells without

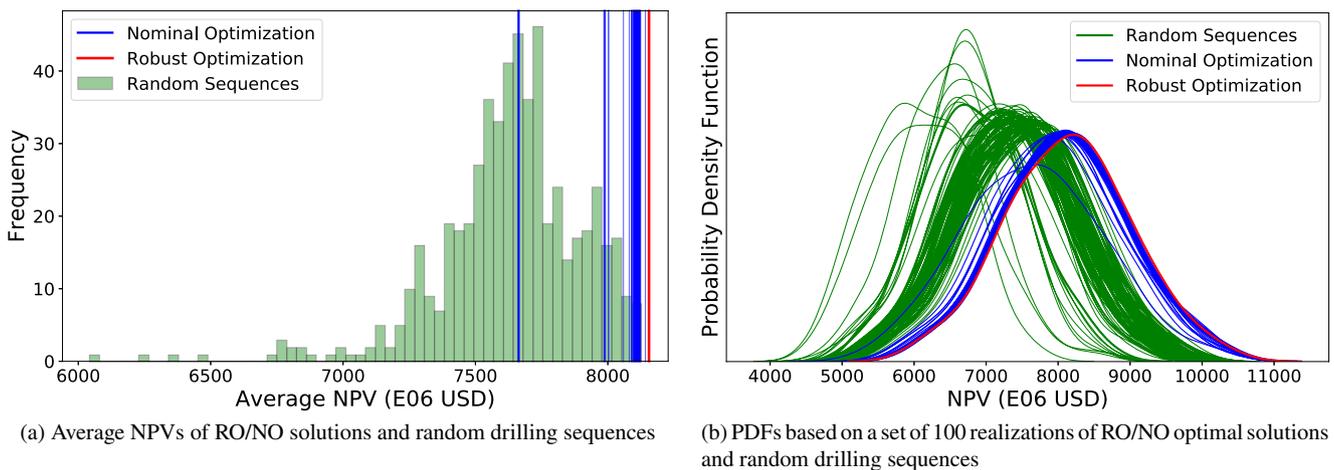


Fig. 20 a, b Results obtained through learned heuristic search strategies compared with random drilling sequences

Table 2 Comparison of robust and nominal optimization results

	Average NPV (E09 USD)	Simulations	Percentage change
Robust optimization	8.15	246	Ref
Nominal optimization	7.66 to 8.15	22 to 49	– 6.04 to 0%
Random sequences	6.04 to 8.12	–	– 25.86 to – 0.36%

finding the entire optimal sequence by performing learned heuristic search with a limited depth (i.e., a depth-limited search and an iterative depth-limited search).

Figure 21 shows the search strategies for the first two and three wells using MLHS-SR with depth-limited search (DLS). To optimize the first few wells, we terminate the search with the first selected partial path that includes one additional well. Figure 21a shows that the first optimal partial path with three selected wells is [OP_3, WI_1, OP_2], which indicates that the optimized sequence of the first two wells is [OP_3, WI_1]. The path through Producer OP_3 and Injector WI_1 has been previously expanded as the most promising direction in the second decision step; however, the search changes direction after the evaluation of the extended partial paths due to their underestimated values. Seven additional decision steps are required to guide the search back to the path along the final optimized complete sequence. In this case, if the search were to be stopped immediately with the first selected partial path with two wells (i.e., [OP_3, WI_1]), then the same solution for the first two wells could be obtained faster (i.e., 26 simulations would be eliminated). However, this will not always be the case for other search depths. Figure 21b shows that

although the first expanded direction with three selected wells is [OP_3, WI_1, OP_2], the path with four wells [OP_3, WI_1, WI_2, OP_1] provides a better solution for the first three wells (i.e., [OP_3, WI_1, WI_2]) since with more observations from longer paths, online learning mechanisms can improve the estimated values and guide the search closer to the true optimal solution. A total of 234 and 240 simulations are required to optimize the first two and three wells, respectively, under geological uncertainty, of which 34 and 40 simulations, respectively, are performed during the search process.

Figure 22 shows the search strategies for the first two and three wells using MLHS-SR with iterative depth-limited search (IDLs). In each decision step, we consider only the paths extended from the previous decisions and optimize the next well based on the first selected partial path with two more wells. We optimize the first well based on the first-visited best partial path with two selected wells (i.e., [OP_3, WI_1]), then we compute the second well based on the paths starting with OP_3. In contrast to the search strategies obtained using DLS, as shown in Fig. 21, the paths starting with Injector WI_1 and Producer OP_4 are pruned, thus avoiding 14 simulations. Figure 23 shows the numbers of

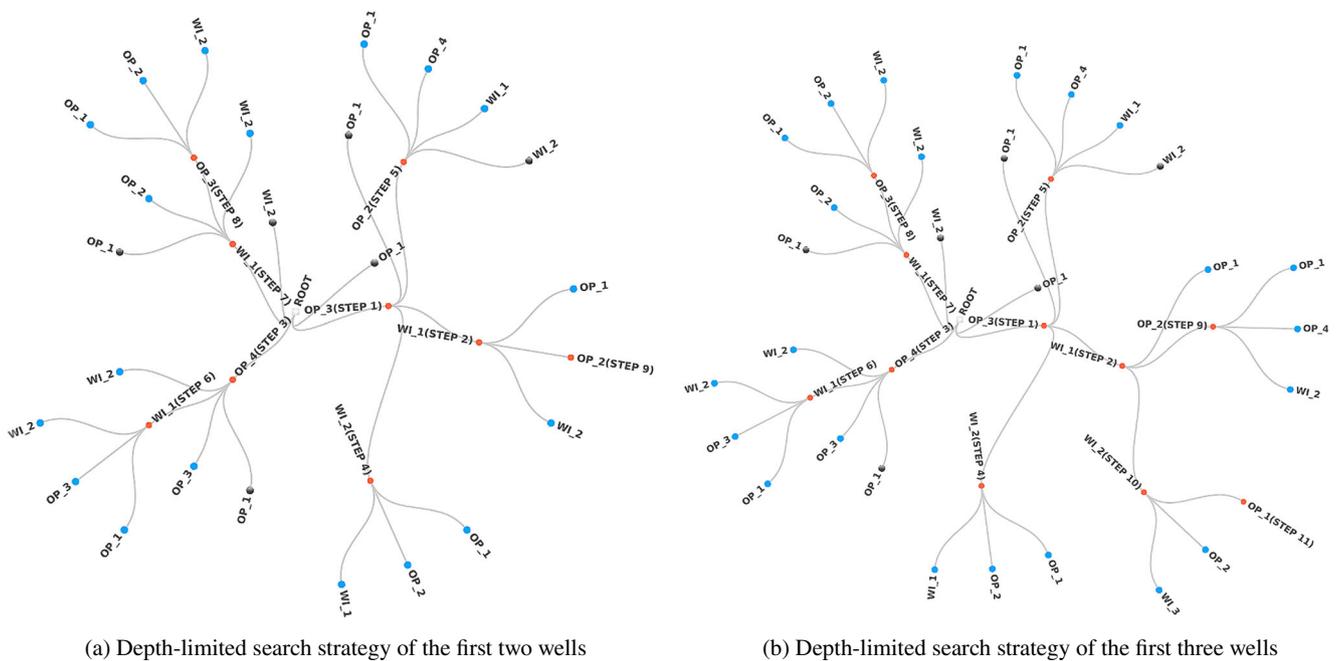


Fig. 21 a, b Optimization of the first two and three wells using MLHS-SR with DLS

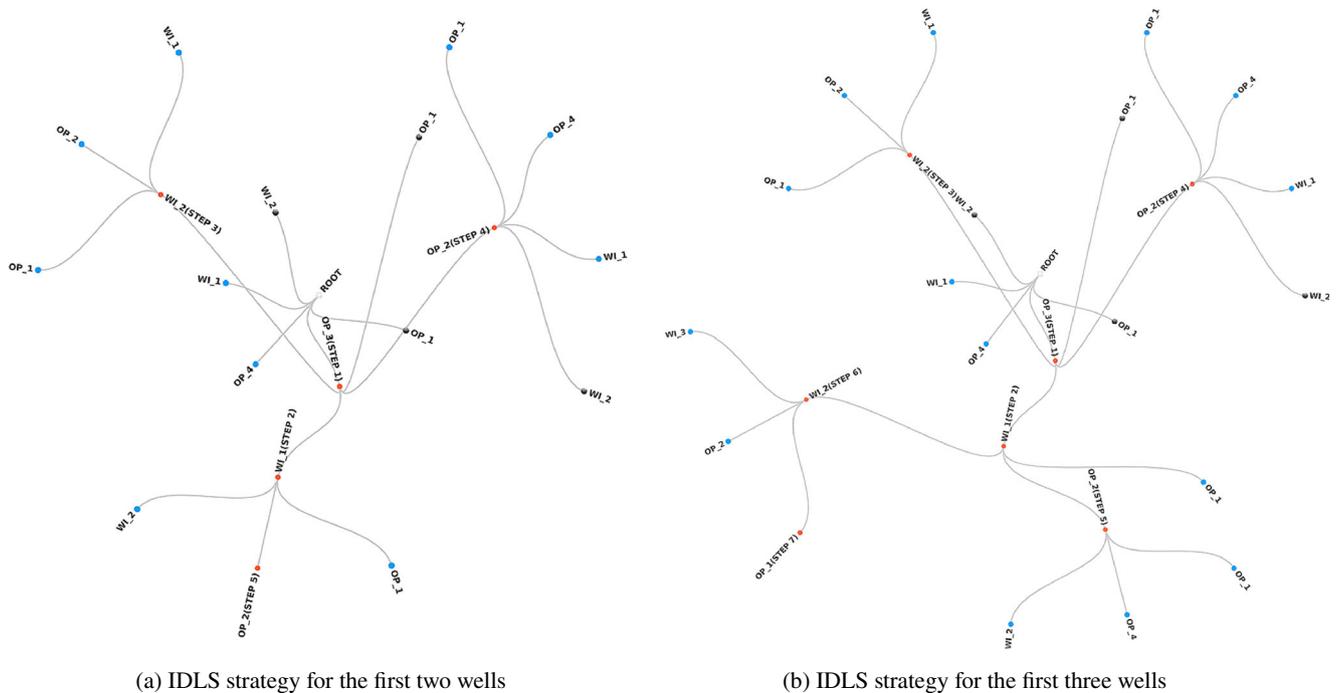


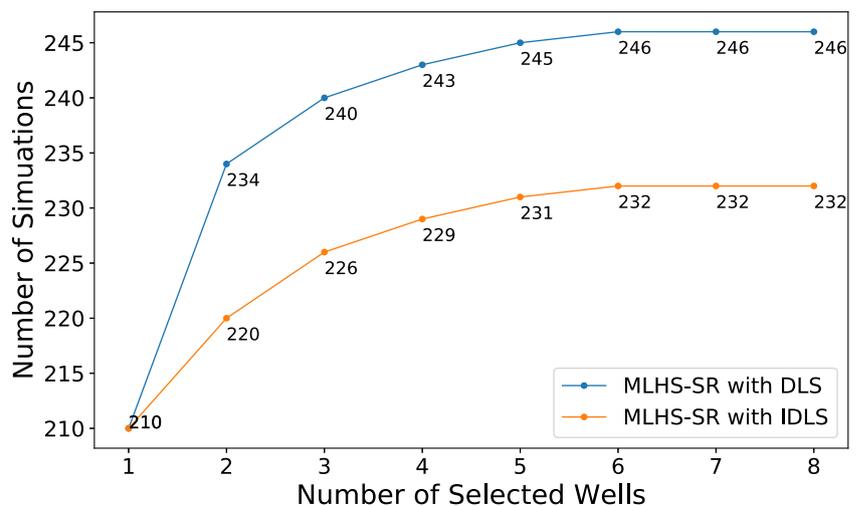
Fig. 22 a, b Optimization of the first two and three wells using MLHS-SR with IDLS

simulations needed to optimize the first few wells by using MLHS-SR with DLS and IDLS. The results for this example show that both methods can be used to optimize the first few wells along the final optimized complete sequence by controlling the search depth, but IDLS finds a solution faster by eliminating some unnecessary node evaluations caused by underestimated values.

The performance of DLS and IDLS depends on the online learning techniques applied. When the estimated values are sufficiently accurate such that the search process will expand the path in only one direction, both DLS and IDLS can optimize the first few wells along the final optimal

complete drilling sequence at the same cost. However, the estimated values are generally less accurate earlier in the search process due to the lack of previous sequential observations, meaning that the search might later change direction. Compared with DLS, IDLS is more likely to optimize the first few wells faster without any loss of the solution quality since paths with low NPV generally will not continue to be extended when appropriate online learning mechanisms are used. In this case, because only 246 simulations are needed to optimize an entire sequence for eight wells via MLHS-SR, the cost reduction achieved by optimizing only the first few wells is not very significant.

Fig. 23 Comparison of the costs of optimizing the first few wells using MLHS-SR with DLS and IDLS



For larger fields, however, MLHS-SR with IDLS can be used to efficiently optimize only the first few wells in the sequence.

5 Conclusions

In this paper, we presented a methodology for robust optimization of the expected value of an expensive-to-compute objective function for which the uncertainty is characterized by an ensemble of model realizations. A straightforward approach using sample averaging to approximate the expected value is prohibitively expensive in most real reservoir applications. In our approach, we compute a multiplicative bias correction to the value of the objective function computed from the first term of the Taylor series expansion. Computation of the bias correction requires two function evaluations for each random sample of the control variables: one evaluation using the mean reservoir model and a second evaluation using a random model realization. Our results show that the information from distinct controls and individual realizations can be used to accurately estimate the bias in the approximation of expected NPV obtained from the mean model and that robust optimization can be performed with good approximations of the expected NPV while requiring many fewer simulation runs than SAA. Based on our experiments, we arrive at the following conclusions:

- Approximations of the expected NPV obtained using the mean reservoir model are generally poor compared with estimates obtained using SAA, but approximations from the mean model can be significantly improved by estimating a multiplicative bias correction factor, which is estimated by simulating a modest number of randomly selected controls on both reservoir realizations and the mean reservoir model.
- Distance-based localization can improve the estimated values of the correction factor. The performance of this approach the estimation of the bias-correction factor depends on the distance measure and the taper length, both of which must be estimated.
- Regularized localization with an appropriate parameter based on the variance of the bias correction factor reduces the sensitivity of the estimates to the taper length; consequently, regularized estimates are accurate for a wider range of taper lengths. When the variance of the bias correction factor is known, regularized localization is the preferred method.
- The RO solutions obtained using MMBC and SAA are of similar quality and are both superior to the deterministic solutions for individual realizations and the mean reservoir model (i.e., linear approximation).

- A learned heuristic search can be performed either to optimize complete drilling sequences or to optimize only the first few wells in the sequence at a reduced cost by limiting the search depth. Compared with DLS approach, a solution of the same quality can be found faster by performing an IDLS with effective online learning techniques.

The major advantage of using mean model bias correction is that we only need to perform simulation using the mean model to obtain an initial approximate value of expected NPV, for which the bias will be corrected by a multiplicative correction factor estimated based on the observations of partial corrections from similar controls. Although the methodology was applied to the problem of determining optimal drilling order, the methodology is fairly general and does not require that the objective function be general or that the control variables be continuous. The effect of localization of the estimate could be significant in the case of large variability in the bias correction factor, especially when the ensemble size is large such that more observations from distinct controls and individual realizations could be used for bias estimation. However, to ensure the effectiveness of distance-based localization for bias correction, we need to design appropriate distance measures for identification of sets of similar controls. For the drilling-order problem, a distance metric based on the well position can effectively measure the similarity between drilling sequences. In this work, we presented a method for efficiently estimating the expected value of an objective function by applying a multiplicative bias correction to the value obtained from the mean model. It would not be difficult to modify the bias-correction method for use in multiple objective optimizations under uncertainty (e.g., standard deviation, percentiles, expected value). For example, it would be possible to compute an estimate of the variance of the objective-function value, by squaring the value of the objective function from the mean model and multiplying it by the variance of partial corrections, which can also be estimated by using information from distinct controls and individual model realizations.

Funding Open Access funding provided by NORCE Norwegian Research Centre AS. This research was supported through the DIGIRES project by the Research Council of Norway and industry partners Equinor, Petrobras, Aker BP, Neptune Energy, WinterhallIDEA, Lundin Norway, and Vår Energy. Access to Eclipse licenses was provided by Schlumberger.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material

is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Nomenclature f^* , Actual expected objective function value of optimal control; α , Bias correction factor; $\bar{\alpha}$, Average value of bias correction factor; \bar{m} , Expected value of model parameter; $\bar{\mu}_{n_s}$, Mean single-step ratio of variability in estimated value; β , Partial correction factor for individual realizations; Δt , Time interval; Δx , Length interval; δ , Distance between controls; γ , Single-step ratio of estimated value; $\hat{\alpha}_{loc}$, Local estimate; $\hat{\alpha}_r$, Regularized estimate; $\hat{\gamma}$, Forecast error of initial evaluation function value; \hat{f} , Learned evaluation function; λ , Regularization parameter/Lagrangian parameter (depending on context); \mathbf{b}_0 , Vector of β at a fixed control x_0 ; \mathbf{b} , Vector of β from random controls and realizations; \mathbf{P} , Vector of drilling sequence; \mathbf{w} , Weight vector of β ; μ , Viscosity of fluid; ω , Weight; ρ , Distance-based weight; σ_α^2 , Variance of bias correction factor α ; σ_β^2 , Variance of partial correction factor β ; θ , Permeability; A , Cross-sectional area; b , Discount rate; C , Covariance of model parameter; d , Number of remaining actions; E , Expected value of objective function; f , Objective function/evaluation function (depending on context); f_m , First derivatives of objective function with respect to model parameter; f_{mm} , Second derivatives of objective function with respect to model parameter; g , Actual economic value from previous actions; h , Heuristic function/estimated maximum future value; k , Log-permeability; L , Taper length; m , Model parameter; n_{eff} , Effective sample size; n_s , Environment state at decision stage; N_e , Total number of model realizations; N_s , Number of selected wells; N_w , Total number of wells; N_x , Total number of random controls; NPV , Net present value; P , Pressure; q , Production/injection rate; r , Reward/cost; T , Total number of time steps ; t , Time; W , Cost of drilling well; x , Control variable.

Subscripts Δt^Φ , Learning time period; Φ , Learning technique index; f_n , Evaluation function at a specific state; h_n , Heuristic function at a specific state; i , Control index/learning period index/cell index (depending on context); j , Model realization index/time step index(depending on context); k , Well index; o , Oil ; s , Decision stage; w , Water/well (depending on context); wi , Water injection; \cos , Cosine distance; L_1 , Manhattan distance; L_2 , Euclidean distance; eff , Effective sample size.

Superscripts \bar{m} , Reservoir mean model; Φ , Learning; i , Injection; p , Production; lin , Linear approximation; $quad$, Quadratic approximation.

Appendix 1

To derive Eq. 24, recall that:

$$E[\mathbf{bb}^T] = \text{cov}(\mathbf{b}, \mathbf{b}) + E[\mathbf{b}]E[\mathbf{b}^T] \tag{33}$$

Assuming stationarity, this simplifies to:

$$E[\mathbf{bb}^T] = \text{cov}(\mathbf{b}, \mathbf{b}) + \bar{\alpha}^2 \mathbf{1}\mathbf{1}^T \tag{34}$$

Similarly:

$$E[\mathbf{bb}_0^T] = \text{cov}(\mathbf{b}, \mathbf{b}_0) + \bar{\alpha}^2 \mathbf{1}\mathbf{1}^T \tag{35}$$

The unbiasedness condition gives that:

$$\bar{\alpha}^2 \mathbf{1}\mathbf{1}^T \mathbf{w} = \frac{1}{N} \bar{\alpha}^2 \mathbf{1}\mathbf{1}^T \mathbf{1}. \tag{36}$$

Appendix 2

The covariance function C is related to the variance and the semivariogram:

$$C(h) = \sigma^2 - \gamma(h), \tag{37}$$

where σ^2 is the variance, h is the distance between two observations, and $\gamma(h)$ is the semivariogram at distance h .

Since the observation β is obtained from random controls of individual realizations, we need to model the covariance function of β with two terms, i.e., variability in β at a fixed model realization and variability in β at a fixed control variable:

$$C_\beta = C_\beta(h_x) + C_\beta(h_m), \tag{38}$$

where $C_\beta(h_x)$ is the covariance function for β at fixed model realization; h_x is distance between observations; $C_\beta(h_m)$ is the covariance function for β at fixed control variable; h_m is the distance of realizations corresponding to the observed controls, i.e., $h_m = 0$ if observed values are from the same realization, $h_m = 1$ if observed values are from different realizations.

$C_\beta(h_x)$ can be obtained from the average variogram of β of all ensemble realizations:

$$C_\beta(h_x) = \frac{1}{N} \sum_{j=1}^N \sigma_\beta^2(m_j) - \frac{1}{N} \sum_{j=1}^N \gamma(h_x, m_j), \tag{39}$$

where $\sigma_\beta^2(m_j)$ and $\gamma(h_x, m_j)$ are the variance and variogram of β at a fixed realization m_j , respectively.

For fixed control variable, the values of β_{ij} and $\beta_{ij'}$ will be correlated for model realizations m_j and $m_{j'}$ at a small distance. Here, we assume that the model realizations are far enough apart such that the observed values of β from different realizations are independent. Then the covariance function for β for fixed control variable can be described as:

$$C_\beta(h_m) = \begin{cases} \frac{1}{N_x} \sum_{k=1}^{N_x} \sigma_\beta^2(x_k) & \text{if } h_m = 0 \\ 0 & \text{if } h_m = 1, \end{cases} \tag{40}$$

where $\sigma_\beta^2(x_k)$ is the variance of observed β from different realizations at a fixed control x_k .

$$\text{cov}(\mathbf{bb}^T) = \begin{bmatrix} \text{cov}(\beta_{11}, \beta_{11}) & \text{cov}(\beta_{11}, \beta_{22}) & \cdots & \text{cov}(\beta_{11}, \beta_{NN}) \\ \text{cov}(\beta_{22}, \beta_{11}) & \text{cov}(\beta_{22}, \beta_{22}) & & \text{cov}(\beta_{22}, \beta_{NN}) \\ \vdots & & \ddots & \vdots \\ \text{cov}(\beta_{NN}, \beta_{11}) & \text{cov}(\beta_{NN}, \beta_{22}) & \cdots & \text{cov}(\beta_{NN}, \beta_{NN}) \end{bmatrix} \tag{41}$$

$$\text{cov}(\mathbf{bb}_0^T) = \begin{bmatrix} \text{cov}(\beta_{11}, \beta_{01}) & \text{cov}(\beta_{11}, \beta_{02}) & \cdots & \text{cov}(\beta_{11}, \beta_{0N}) \\ \text{cov}(\beta_{22}, \beta_{01}) & \text{cov}(\beta_{22}, \beta_{02}) & & \text{cov}(\beta_{22}, \beta_{0N}) \\ \vdots & & \ddots & \vdots \\ \text{cov}(\beta_{NN}, \beta_{01}) & \text{cov}(\beta_{NN}, \beta_{02}) & \cdots & \text{cov}(\beta_{NN}, \beta_{0N}) \end{bmatrix} \quad (42)$$

References

- Alexanderian, A., Petra, N., Stadler, G., Ghattas, O.: Mean-variance risk-averse optimal control of systems governed by PDEs with random parameter fields using quadratic approximations. *SIAM/ASA J. Uncertain. Quantif.* **5**(1), 1166–1192 (2017)
- Barros, E.G.D., Maciel, S., Moraes, R.J., Fonseca, R.M.: Automated clustering based scenario reduction to accelerate robust life-cycle optimization. In: *ECMOR XVI-16th European Conference on the Mathematics of Oil Recovery*, Barcelona. European Association of Geoscientists & Engineers (2018)
- Beyer, H.-G., Sendhoff, B.: Robust optimization—a comprehensive survey. *Comput. Methods Appl. Mech. Eng.* **196**(33), 3190–3218 (2007)
- Cardoso, M.A., Durlofsky, L.J.: Linearized reduced-order models for subsurface flow simulation. *J. Comput. Phys.* **229**(3), 681–700 (2010a)
- Cardoso, M.A., Durlofsky, L.J.: Use of reduced-order modeling procedures for production optimization. *SPE J.* **15**(2), 426–435 (2010b)
- Chen, Y., Oliver, D.S., Zhang, D.: Efficient ensemble-based closed-loop production optimization. *SPE J.* **14**(4), 634–645 (2009)
- Chen, C., Wang, Y., Li, G., Reynolds, A.C.: Closed-loop reservoir management on the Brugge test case. *Comput. Geosci.* **14**(4), 691–703 (2010)
- Chen, P., Villa, U., Ghattas, O.: Taylor approximation and variance reduction for PDE-constrained optimal control under uncertainty. *J. Comput. Phys.* **385**, 163–186 (2019)
- Darlington, J., Pantelides, C.C., Rustem, B., Tanyi, B.A.: An algorithm for constrained nonlinear optimization under uncertainty. *Automatica* **35**(2), 217–228 (1999)
- Darlington, J., Pantelides, C.C., Rustem, B., Tanyi, B.A.: Decreasing the sensitivity of open-loop optimal solutions in decision making under uncertainty. *Eur. J. Oper. Res.* **121**(2), 343–362 (2000)
- Denney, D.: Pros and cons of applying a proxy model as a substitute for full reservoir simulations. *J. Petrol. Technol.* **62**(7), 634–645 (2010)
- Doucet, A., Godsill, S., Andrieu, C.: On sequential Monte Carlo sampling methods for Bayesian filtering. *Stat. Comput.* **10**(3), 197–208 (2000)
- Durlofsky, L.J.: Upscaling and gridding of fine scale geological models for flow simulation. In: *8th International Forum on Reservoir Simulation Iles Borromees* (2005)
- Fonseca, R.M., Chen, B., Jansen, J.D., Reynolds, A.C.: A stochastic simplex approximate gradient (StoSAG) for optimization under uncertainty. *Int. J. Numer. Methods Eng.* **109**(13), 1756–1776 (2017)
- Gaspar Ravagnani, A.T.F.D.S., Muñoz Mazo, E.O., Schiozer, D.J.: A case study of the structure of the process for production strategy selection. *Int. J. Model. Simul. Petrol. Ind.* **4**(1), 1–14 (2011)
- Gaspari, G., Cohn, S.E.: Construction of correlation functions in two and three dimensions. *Q. J. R. Meteorol. Soc.* **125**(554), 723–757 (1999)
- Hamming, R.: *Coding and Information Theory*. Prentice-Hall, New York (1980)
- Hanea, R.G., Casanova, P., Hustoft, L., Bratvold, R.B., Nair, R., Hewson, C., Leeuwenburgh, O., Fonseca, R.M.: Drill and learn: a decision making workflow to quantify value of learning. In: *SPE Reservoir Simulation Conference*. Society of Petroleum Engineers (2017)
- Hart, P.E., Nilsson, N.J., Raphael, B.: A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans. Syst. Sci. Cybern.* **4**(2), 100–107 (1968)
- Jansen, J.D., Durlofsky, L.J.: Use of reduced-order models in well control optimization. *Optim. Eng.* **18**(1, SI), 105–132 (2017)
- Jaro, M.A.: Advances in record-linkage methodology as applied to matching the 1985 census of Tampa, Florida. *J. Am. Stat. Assoc.* **84**(406), 414–420 (1989)
- Jesmani, M., Jafarpour, B., Bellout, M.C., Foss, B.: A reduced random sampling strategy for fast robust well placement optimization. *J. Petrol. Sci. Eng.* **184**, 106414 (2020)
- Kim, S., Pasupathy, R., Henderson, S.G.: A guide to sample average approximation. In: *Handbook of Simulation Optimization*, pp. 207–243. Springer (2015)
- Kish, L.: *Survey Sampling*. Wiley, New York (1965)
- Kleywegt, A.J., Shapiro, A., Homem-de Mello, T.: The sample average approximation method for stochastic discrete optimization. *SIAM J. Optim.* **12**(2), 479–502 (2002)
- Korenius, T., Laurikkala, J., Juhola, M.: On principal component analysis, cosine and euclidean measures in information retrieval. *Inf. Sci.* **177**, 4893–4905 (2007)
- Korf, R.E.: Depth-limited search for real-time problem-solving. *Real-Time Syst.* **2**(1–2), 7–24 (1990)
- Lamas, L.F., Botechia, V.E., Schiozer, D.J., Delshad, M.: Optimization for drilling schedule of wells in the development of heavy oil reservoirs. *Braz. J. Petrol. Gas* **11**(3), 165–173 (2017)
- Leeuwenburgh, O., Chitu, A.G., Nair, R., Egberts, P.J.P., Ghazaryan, L., Feng, T., Hustoft, L.: Ensemble-based methods for well drilling sequence and time optimization under uncertainty. In: *ECMOR XV-15th European Conference on the Mathematics of Oil Recovery* (2016)
- Levenshtein, V.I.: Binary codes capable of correcting deletions, insertions and reversals. *Sov. Phys. Dokl.* **10**(8), 707–710 (1966)
- Mathai, A.M., Provost, S.B.: *Quadratic Forms in Random Variables: Theory and Applications*. Marcel Dekker, New York (1992)
- Rahim, S., Li, Z.: Well placement optimization with geological uncertainty reduction. *IFAC-PapersOnLine* **48**(8), 57–62 (2015). 9th IFAC Symposium on Advanced Control of Chemical Processes ADCHEM 2015
- Robert, C.P., Casella, G.: *Introducing Monte Carlo Methods with R*. Springer, New York (2010)
- Ronald, S.: More distance functions for order-based encodings. In: *1998 IEEE International Conference on Evolutionary Computation Proceedings*. IEEE World Congress on Computational Intelligence, pp. 558–563 (1998)
- Schiavinotto, T., Stütze, T.: A review of metrics on permutations for search landscape analysis. *Comput. Oper. Res.* **34**(10), 3143–3153 (2007)
- Sharma, S.: *Applied Multivariate Techniques*. Wiley, New York (1996)
- Silva, V.L.S., Emerick, A.A., Couto, P., Alves, J.L.D.: History matching and production optimization under uncertainties—application of closed-loop reservoir management. *J. Petrol. Sci. Eng.* **157**, 860–874 (2017)

38. van Doren, J.F.M., Markovinović, R., Jansen, J.-D.: Reduced-order optimal control of water flooding using proper orthogonal decomposition. *Comput. Geosci.* **10**(1), 137–158 (2006)
39. Van Essen, G., Zandvliet, M., Van den Hof, P., Bosgra, O., Jansen, J.-D.: Robust waterflooding optimization of multiple geological scenarios. *SPE J.* **14**(1), 202–210 (2009)
40. Wang, L., Oliver, D.S.: Efficient optimization of well drilling sequence with learned heuristics. *SPE J.* **24**(5), 2111–2134 (2019)

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.