Journal of Petroleum Science and Engineering 217 (2022) 110858

Contents lists available at ScienceDirect

Journal of Petroleum Science and Engineering

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

An ensemble-based decision workflow for reservoir management

Yuqing Chang [a], Geir Evensen [a,b,*]

[a] NORCE–Norwegian Research Center, Bergen, Norway
[b] NERSC–Nansen Environmental and Remote Sensing Center, Bergen, Norway

ARTICLE INFO

Keywords:
EnRML
EnOPT
Decision making under uncertainty
Closed-loop reservoir management
History matching

ABSTRACT

It is challenging to make optimal field development and reservoir management decisions with diminishing resources and low-emission requirements. For an optimal exploitation of the reservoir fluids, it is necessary to introduce advanced digital tools and account for geological uncertainty when making decisions. This paper discusses an ensemble-based probabilistic decision-making workflow for closed-loop reservoir management. The reservoir model is a synthetic but realistic model with oil, gas, and water. We show how to use ensemble methods in a workflow for integrated uncertainty prediction, history matching, and robust optimization. The workflow uses advanced ensemble-based history-matching techniques to update a reservoir model in the annual maturation process. After the history-matching update, an optimization process decides the best wells to drill next to gain the highest net present value. An additional robust decision step evaluates the wells over the ensemble of reservoir models to ensure they lead to a positive ensemble-averaged net present value. The approach allows for up-to-date predictions, including uncertainty estimates, leading to improved decision support for field development, well-planning, production strategies, and reservoir management.

1. Introduction

Traditionally, we have made field-management decisions based on a single "best" reservoir-model prediction. Sometimes this best model is accompanied by an ad-hoc P90 and P10 model to represent the uncertainty in light of the two extreme cases. We can now do better by introducing ensemble methods, which use model ensembles to represent the uncertainty. The ensemble methods allow for a consistent representation of uncertainty that propagates from the geological inputs to the simulated outputs. Furthermore, by using ensemble history-matching (HM) methods, such as the iterative ensemble smoothers, we can use the ensemble correlations between the geological inputs and the simulated outputs to compute regression updates of the geological description. The iterative ensemble smoothers can handle weak nonlinearities and have worked well in many reservoir field cases. An advantage of using an ensemble of models is that we avoid defining a "best" model to base our decisions. Instead, we can make robust decisions based on the whole ensemble of model realizations. E.g., it is possible to compute which alternative control strategy results in the maximum net present value (NPV) when averaged over the ensemble of reservoir models. The existence of an ensemble of reservoir models and predictions also provides a basis for robust optimization of model controls that are the best in average over the whole ensemble.

The ensemble formulation facilitates a framework for integrated uncertainty analysis, history matching, optimization, and decision making. This paper demonstrates such a framework in a closed-loop formulation. We recursively update the reservoir model with new production data. Based on the updated models, we recompute the optimization of the control variables and re-evaluate the decision alternatives. This closed-loop ensemble framework is generally applicable to many prediction problems within and outside the geosciences.

Jansen et al. (2005) formulated the closed-loop reservoir management problem, later discussed by Jansen et al. (2009) and Skjervheim et al. (2015). The more prominent focus on developing methods for solving the history-matching problem may have delayed the development of closed-loop reservoir management systems. Additionally, closed-loop reservoir management adds another level of complexity as it involves both HM and optimization loops, which we need to solve iteratively, and the approach becomes less practical to apply and time consuming to run. history-matching is a prerequisite for a working reservoir optimization and further integration into the closed-loop framework. This paper differs from some other recent publications on optimization of drilling order and ensemble-based closed-loop reservoir management, e.g., Silva et al. (2017), Lu and Reynolds (2020), Barros et al. (2020), Leeuwenburgh et al. (2010, 2016). We also include a decision process in the workflow, and we point out some weaknesses of the EnOPT method when used with significant geological uncertainty.

* Corresponding author at: NORCE–Norwegian Research Center, Bergen, Norway.
E-mail addresses: yuch@norceresearch.no (Y. Chang), geev@norceresearch.no (G. Evensen).

https://doi.org/10.1016/j.petrol.2022.110858
Received 16 March 2022; Received in revised form 13 June 2022; Accepted 10 July 2022
Available online 20 July 2022

The ensemble-history-matching formulation has evolved from early attempts of using the ensemble Kalman filter (EnKF) developed by Evensen (1994, 2003, 2009) for recursive updating of reservoir models. The initial approach was similar to the recursive model updating in weather forecasting, except the focus was on the joint updating of the model parameters and the model state (Nævdal et al., 2003; Haugen et al., 2008). In weather prediction, one only updates the model state before making a new prediction of the future. A breakthrough paper by Skjervheim et al. (2011) introduced the ensemble smoother from Van Leeuwen and Evensen (1996) for parameter estimation, using the approach from Evensen (2009, Chap. 10). With further extensions of the initial smoother method to use iterations for reducing the impacts of nonlinearity (Chen and Oliver, 2012, 2013; Emerick and Reynolds, 2012), iterative ensemble smoothers have become the method of choice in ensemble history matching. These methods simplify the history-matching problem by only updating the model parameters. Furthermore, they compute the solution by integrating the model from the initial time using the updated parameters, as in traditional history matching. In this paper, we have used the ensemble-subspace-EnRML method by Evensen et al. (2019) and Raanes et al. (2019).

Van Essen et al. (2009) introduced the concept of robust optimization. He showed that optimizing controls over an ensemble of reservoir models gave superior results to optimization over a single model realization. A popular robust optimization method in the reservoir community is ensemble-based optimization (EnOPT) (see, e.g., Lorentzen et al., 2006; Chen et al., 2009). Characteristic of (EnOPT) is a statistical approximation to the objective function's gradient computed from ensemble statistics. Stordal et al. (2016) and Fonseca (2016) have further developed EnOPT to the version used in the current paper. Although we have used EnOPT in this paper, we will claim that the method still has issues handling problems with significant geological uncertainty as this leads to additional stochasticity in the gradient calculations. Recent alternative methods may be the learned-heuristics approach by Wang and Oliver (2019) or the mean-model method introduced by Wang and Oliver (2021), which both avoid using a gradient. These methods are promising for robust optimization, leading to reduced computational requirements and avoiding using an ensemble gradient.

Finally, we should implement the decision alternative if its implementation leads to a positive NPV. The ensemble approach allows estimating the resulting NPV over the ensemble of updated reservoir models and provides a foundation for making a robust decision. The following sections will explain and demonstrate our closed-loop decision workflow on a simplified reservoir model. The workflow is readily applicable to realistic reservoir models and can include the geological modeling following the approach by Zachariassen et al. (2011), Skjervheim et al. (2012), and Hanea et al. (2015). Our example below focuses on optimizing the drilling order of pre-planned wells given a specific drilling schedule. Thus, we consider a discrete optimization problem.

2. Methodology

This section introduces the closed-loop reservoir management workflow and briefly introduces the ensemble methods used. We define a model state $\mathbf{x}$ that characterizes the reservoir model and its uncertain geological variables. In our example, we assume uncertain porosity and permeability fields, and fault multipliers. These are reservoir parameters that are important for characterizing the volume of oil and its distribution in the reservoir. Additionally, the permeability together with the fault transmissibilities determine the connectivity and fluid flow in the reservoir. We will use history matching to improve our prior knowledge of $\mathbf{x}$. The model controls $\mathbf{u}$ can be continuous variables such as imposed production and injection rates, but in our example, they represent the discrete drilling order of the wells. We assume a reservoir model

$$\mathbf{y} = \mathbf{g}(\mathbf{x}, \mathbf{u}), \tag{1}$$

where $\mathbf{x} \in \Re^n$ is the vector of n poorly known model parameters, $\mathbf{u} \in \Re^{n_c}$ is a vector of n_c model controls, and $\mathbf{y} \in \Re^m$ is a set of m predicted measurements.

2.1. The ensemble-based decision workflow

The closed-loop decision workflow in Fig. 1 comprises various decision steps for reservoir management during the field development. The workflow is independent of the methods used for history matching and optimization, although it is ideally suited to ensemble methods. In other words, closed-loop reservoir management is the standard way of working. However, traditionally one has used less sophisticated manual methods for model updating and history matching. Furthermore, as an alternative to advanced optimization methods, one has used intuition, experience, and model simulations to decide the next reservoir-development steps and the drainage strategy.

Our robust ensemble workflow uses ensemble history-matching and optimization methods recursively in time. We condition model parameters to all current and past observations to create the best possible ensemble of geologically consistent models that fits the measurements within their estimated uncertainty. We robustly optimize decision alternatives over the ensemble of history-matched geological models and evaluate each alternative's expected objective function at decision points.

2.2. Ensemble-based optimization

Given a reservoir, where we use the variable $\mathbf{x}$ to represent the geological model state, and a set of control variables $\mathbf{u}$ defining its development and operation, we can define the objective function (e.g., evaluating the net present value) as $\mathcal{O}(\mathbf{u}, \mathbf{x})$. Thus, the objective function is a function of the control parameters and model state. The optimization goal is to find the controls in $\mathbf{u}$ that maximize the expected value of $\mathcal{O}$ (Chang et al., 2019).

We will use EnOPT to maximize the objective function and start from the pre-conditioned steepest-ascent method

$$\mathbf{u}^{i+1} = \mathbf{u}^i + \eta^i \, \mathbf{C}_{uu} \, \nabla_{\mathbf{u}} \, \mathcal{O}(\mathbf{u}^i, \mathbf{x}), \tag{2}$$

where the superscript i is the iteration index, η^i is the step size, $\mathbf{C}_{uu}$ is a pre-determined symmetric and positive definite matrix, and $\nabla_{\mathbf{u}} \mathcal{O}(\mathbf{u}, \mathbf{x}) \in \Re^{N \times 1}$ is the gradient of the objective function to the control variables evaluated at $(\mathbf{u}, \mathbf{x})$.

The EnOPT algorithm (Lorentzen et al., 2006; Chen et al., 2009) introduces an approximate stochastic representation for the gradient. The method computes the gradient from an ensemble of controls, $\mathbf{u}_j \sim \mathcal{N}(\mathbf{u}, \mathbf{C}_{uu})$ as a linear regression between the controls and the resulting NPVs. We use the subscript j as the ensemble realization index. From here and onwards, we only discuss the gradient approximation at the current iteration i. For the simplicity of notation, we, therefore, skip the iteration superscript on $\mathbf{u}_j^i$ and note that $\mathbf{x}$ is independent of i. For a given geological model $\mathbf{x}$, we can approximate the gradient as

$$\begin{aligned} \nabla_{\mathbf{u}} \, \mathcal{O}(\mathbf{u}, \mathbf{x}) &\triangleq \mathbf{C}_{uu}^{-1} \mathbf{C}_{uo} \\ &\approx \mathbf{C}_{uu}^{-1} \, \frac{1}{N} \sum_{j=1}^{N} \Big(\mathbf{u}_j - \overline{\mathbf{u}} \Big) \Big(\mathcal{O}(\mathbf{u}_j, \mathbf{x}) - \overline{\mathcal{O}} \Big). \end{aligned} \tag{3}$$

Here $\overline{\mathbf{u}}$ is the arithmetic average of the controls at iteration i, and

$$\overline{\mathcal{O}} = \frac{1}{N} \sum_{j=1}^{N} \mathcal{O}(\mathbf{u}_j, \mathbf{x}), \tag{4}$$

is the arithmetic average of NPVs for all controls. Thus, we only need to compute the sample covariance, $\mathbf{C}_{uo}$, between the controls $\mathbf{u}_j$ and

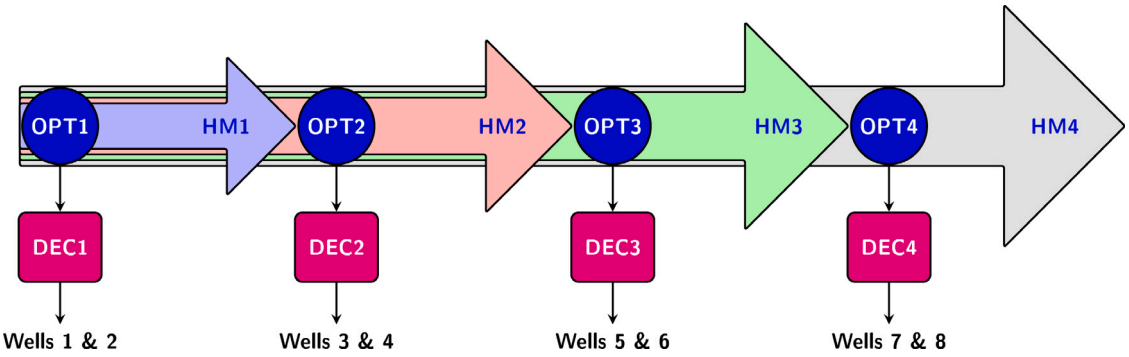

Fig. 1. DIGIRES decision workflow for field development. Drilling two wells per year, we use the workflow to decide on the optimal drilling schedule and how many wells to drill.

the resulting ensemble of NPVs, $\mathcal{O}(\mathbf{u}_j, \mathbf{x})$. The previous formulation defines the original EnOPT algorithm (Lorentzen et al., 2006; Chen et al., 2009).

Chen et al. (2009) also used EnOPT in the case with uncertain geology. They introduced uncertainty in the geological reservoir characterization and represented it by an ensemble of N model realizations. Formally, introducing an ensemble of reservoir models leads to an ensemble of simultaneous optimization problems, one for each geologic realization. We can then find one optimal control vector for each geologic realization. An alternative by Van Essen et al. (2009) is to define an optimization problem where we sum the individual optimization problems and solve for one common optimal control vector.

It is common to associate only one realization $\mathbf{u}_j$ to each geologic realizations $\mathbf{x}_j$ as this simplifies the formulation and avoids running N^2 simulations. Stordal et al. (2016) showed that, for a sufficiently large ensemble size, this approximation does not introduce any bias in the gradient approximation. Thus, in the case of geological uncertainty, we can use the following definition for the gradient

$$\begin{aligned}\nabla_{\mathbf{u}} \mathcal{O}(\mathbf{u}, \mathbf{x}) &\triangleq \mathbf{C}_{uu}^{-1}\mathbf{C}_{uo} \\ &\approx \mathbf{C}_{uu}^{-1}\frac{1}{N}\sum_{j=1}^{N}\left(\mathbf{u}_j - \overline{\mathbf{u}}\right)\left(\mathcal{O}(\mathbf{u}_j, \mathbf{x}_j) - \overline{\mathcal{O}}\right).\end{aligned} \tag{5}$$

Here we have redefined $\overline{\mathcal{O}}$ from Eq. (4) to use the ensemble of geologic realizations as

$$\overline{\mathcal{O}} = \frac{1}{N}\sum_{j=1}^{N}\mathcal{O}(\mathbf{u}_j, \mathbf{x}_j). \tag{6}$$

A further modification of the EnOPT algorithm with geological uncertainty by Fonseca (2016) modifies the averaging of the NPVs in Eq. (6) with the following definition for the gradient

$$\begin{aligned}&\nabla_{\mathbf{u}} \mathcal{O}(\mathbf{u}, \mathbf{x}) \approx \\ &\mathbf{C}_{uu}^{-1}\frac{1}{N}\sum_{j=1}^{N}\left(\mathbf{u}_j - \overline{\mathbf{u}}\right)\left(\mathcal{O}(\mathbf{u}_j, \mathbf{x}_j) - \mathcal{O}(\overline{\mathbf{u}}, \mathbf{x}_j)\right).\end{aligned} \tag{7}$$

In cases where the geological uncertainty is high, Eq. (7) gives a lower variance for the gradient estimate than Eq. (5), see Stordal et al. (2016). Therefore, we will use the EnOPT method with Eq. (7) in this work. In this version of EnOPT, we must also compute the NPV using $\overline{\mathbf{u}}$ for all geological realizations in each iteration to evaluate the quality of the updated control and the stopping criteria. We have used a stopping criteria based on the improvement of objective function between two iterations. If the objective function fails to improve after three backtracking trial steps, the optimization algorithm stops. Thus, each optimization iteration requires $2N$ simulations.

2.3. Subspace EnRML

In history matching, given measurements $\mathbf{d}$ of $\mathbf{y}$, we estimate the posterior and conditional probability density function (PDF) of $\mathbf{x}$ from Bayes' formula written as

$$f(\mathbf{x} \mid \mathbf{d}) \propto f(\mathbf{x}) f(\mathbf{d} \mid \mathbf{g}(\mathbf{x})). \tag{8}$$

Here $f(\mathbf{x} \mid \mathbf{d})$ is the conditional PDF of the parameters $\mathbf{x}$ given the observations $\mathbf{d}$. This posterior PDF is proportional to the product of the prior PDF $f(\mathbf{x})$ and the likelihood $f(\mathbf{d} \mid \mathbf{g}(\mathbf{x}))$. The prior PDF describes the prior knowledge of the parameters, while the likelihood describes the uncertainty in the observations. The proportionality constant is just one over the integral of the right-hand side of Eq. (8). It ensures that the posterior PDF integrates to one, but for ensemble methods, it turns out that we do not need to compute it. For additional information on this formulation, see the discussion in Evensen et al. (2022, Chap. 2). Note also that for the history-matching problem, we either treat the controls $\mathbf{u}$ as given, or if they are uncertain, we can include them in $\mathbf{x}$ as explained by Evensen (2021). Note again that in the current paper, we do not optimize the controls but rather the drilling sequence, making it possible to add the controls to the HM state vector.

It is common to introduce the normally-distributed priors

$$f(\mathbf{x}) = \mathcal{N}(\mathbf{x}^{\mathrm{f}}, \mathbf{C}_{xx}), \tag{9}$$

$$f(\mathbf{d} \mid \mathbf{g}(\mathbf{x})) = f(\mathbf{e}) = \mathcal{N}(\mathbf{0}, \mathbf{C}_{dd}), \tag{10}$$

where $\mathbf{x}^{\mathrm{f}} \in \Re^n$ is the prior estimate of $\mathbf{x}$ with error-covariance matrix $\mathbf{C}_{xx} \in \Re^{n\times n}$, and $\mathbf{C}_{dd} \in \Re^{m\times m}$ is the error-covariance matrix for the measurements. We can then write Eq. (8) as

$$\begin{aligned}f(\mathbf{x} \mid \mathbf{d}) \propto &\exp\left\{-\frac{1}{2}\left(\mathbf{x} - \mathbf{x}^{\mathrm{f}}\right)^{\mathrm{T}}\mathbf{C}_{xx}^{-1}\left(\mathbf{x} - \mathbf{x}^{\mathrm{f}}\right)\right\} \\ &\times \exp\left\{-\frac{1}{2}\left(\mathbf{g}(\mathbf{x}) - \mathbf{d}\right)^{\mathrm{T}}\mathbf{C}_{dd}^{-1}\left(\mathbf{g}(\mathbf{x}) - \mathbf{d}\right)\right\}.\end{aligned} \tag{11}$$

Maximizing $f(\mathbf{x} \mid \mathbf{d})$ is equivalent to minimizing the cost function

$$\begin{aligned}\mathcal{J}(\mathbf{x}) = &\frac{1}{2}\left(\mathbf{x} - \mathbf{x}^{\mathrm{f}}\right)^{\mathrm{T}}\mathbf{C}_{xx}^{-1}\left(\mathbf{x} - \mathbf{x}^{\mathrm{f}}\right) \\ &+ \frac{1}{2}\left(\mathbf{g}(\mathbf{x}) - \mathbf{d}\right)^{\mathrm{T}}\mathbf{C}_{dd}^{-1}\left(\mathbf{g}(\mathbf{x}) - \mathbf{d}\right).\end{aligned} \tag{12}$$

Ensemble methods for history matching tries to sample the posterior PDF in Eq. (11). The Randomized Maximum Likelihood (RML) method discussed by Oliver et al. (1996) and Kitanidis (1995) samples Eq. (11) approximately by minimizing an ensemble of cost functions written for each realization as

$$\begin{aligned}\mathcal{J}(\mathbf{x}_j) = &\frac{1}{2}\left(\mathbf{x}_j - \mathbf{x}_j^{\mathrm{f}}\right)^{\mathrm{T}}\mathbf{C}_{xx}^{-1}\left(\mathbf{x}_j - \mathbf{x}_j^{\mathrm{f}}\right) \\ &+ \frac{1}{2}\left(\mathbf{g}(\mathbf{x}_j) - \mathbf{d}_j\right)^{\mathrm{T}}\mathbf{C}_{dd}^{-1}\left(\mathbf{g}(\mathbf{x}_j) - \mathbf{d}_j\right).\end{aligned} \tag{13}$$

Here $\mathbf{x}_j^{\mathrm{f}} \sim \mathcal{N}(\mathbf{x}^{\mathrm{f}}, \mathbf{C}_{xx})$ and $\mathbf{d}_j \sim \mathcal{N}(\mathbf{d}, \mathbf{C}_{dd})$ are realizations of the parameters and measurements sampled from their prior distributions. *Note that, in the nonlinear case, the minimizing solutions will not precisely sample the posterior non-Gaussian distribution, but only provides an approximate sampling of it.*

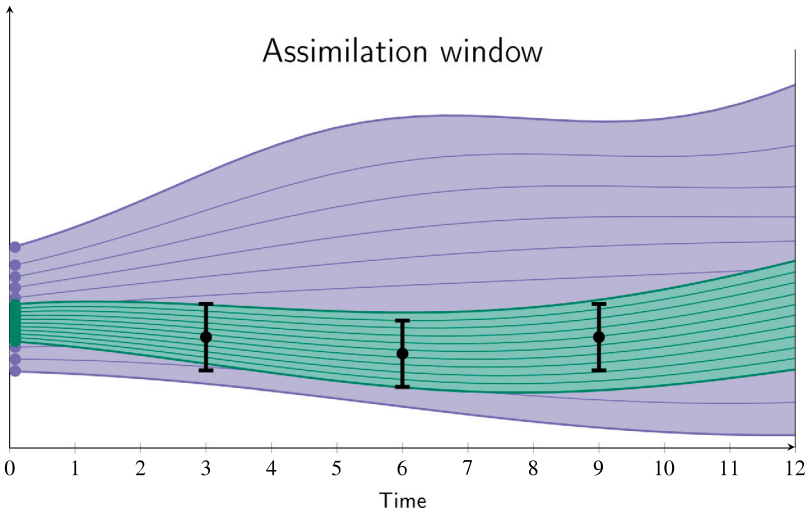

Fig. 2. The figure illustrates the history-matching problem assuming a perfect model solved using an ensemble approach. One defines an assimilation time window and updates the prior model input parameters (blue bullets) such that the posterior ensemble of model realizations (green) better fits the observations (black dots). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)
Source: The figure is taken from Evensen et al. (2022).

In the standard EnRML method by Chen and Oliver (2013), we minimize the cost functions in Eq. (13) by using a gradient descent method. Characteristic for ensemble methods is the use of a representation for the model sensitivity (i.e., the tangent-linear operator $\mathbf{G}$ of $\mathbf{g}(\mathbf{x})$) defined by the linear regression

$$\mathbf{G} = \mathbf{C}_{yx}\mathbf{C}_{xx}^{-1} \tag{14}$$

By using this regression, we are replacing the exact model sensitivities with an average best least-squares fit defined by (14). In the ensemble methods, we replace the covariances with their ensemble representations

$$\overline{\mathbf{G}} = \overline{\mathbf{C}}_{yx}\overline{\mathbf{C}}_{xx}^{-1}. \tag{15}$$

This formulation requires the pseudo inversion of large matrices and is expensive to compute.

The subspace EnRML was introduced by Evensen et al. (2019), Raanes et al. (2019), and they discuss its theoretical foundation as well as the practical implementation. Since, in the ensemble formulation, the solution is confined to an initially defined ensemble subspace, they showed that it is possible to formulate the EnRML method to search for the solution in the ensemble subspace. This formulation leads to a significantly more efficient and simpler algorithm where the number of unknowns is the number of ensemble members N rather than the number of parameters n. While $N = \mathcal{O}(100)$, we often have $n = \mathcal{O}(10^6)$. Theoretically, the ensemble subspace RML and the EnRML methods give identical results, but the subspace variant is more stable when n is large in practice.

Fig. 2 illustrates the ensemble HM concept. The prior input parameters (blue bullets) lead to an ensemble prediction with considerable uncertainty that does not well represent the observations. Then, the iterative ensemble smoothers converge to a posterior set of parameter values (green bullets), resulting in an ensemble prediction better agreeing with the observations (black dots). For the details of the ensemble subspace RML algorithm and its implementation we refer to the two original papers by Evensen et al. (2019) and Raanes et al. (2019), but see also the discussion in Evensen (2021) and Evensen et al. (2022).

3. Workflow and problem description

This section explains the closed-loop decision workflow illustrated in Fig. 1 and Table 1 applied to a reservoir model. We will use it to select the drilling schedule optimally. In the next section, we briefly introduce the reservoir model before demonstrating the decision workflow on a reservoir management case.

3.1. Reservoir model

We use the REEK reservoir model provided by Equinor and described by Hanea et al. (2017). The model simulates three fluid phases consisting of oil, water, and gas. The model grid has dimensions $40 \times 64 \times 14$ and consists of 35 840 grid cells. On average, the original oil in place (OOIP) is $4\,831 \times 10^7$ sm^3 (Oguntola and Lorentzen, 2021). The reservoir has three zones (UpperReek, MidReek, and LowerReek) with six faults of different transmissibilities, and it has five producers and three injectors (see Fig. 3). Thus, we use a relatively small but realistic model to demonstrate the ensemble workflows. The REEK model was also used in Evensen et al. (2019), Evensen (2021) for testing iterative ensemble smoothers for reservoir history matching. We consider porosity, permeability and fault transmissibilities as uncertain model parameters. A model simulation requires one CPU minute on a single CPU.

3.2. Problem description

We have defined a demonstration case with eight preplanned wells, five producers, and three injectors. Given the constraint of drilling two wells a year, the goal is to determine the optimal drilling schedule for the preplanned wells as part of an annual maturation process. We represent the geological reservoir uncertainty for porosity, permeability and fault transmissibilities by using 100 geological realizations. Given a control strategy $\mathbf{u}_j$ on a geological realization $\mathbf{x}_j$, with $Q(t_i)$ represents the function $Q(t_i) = Q(\mathbf{u}_j, \mathbf{x}_j, t_i)$, we define the objective function of NPV as

$$\begin{aligned}\mathcal{O}(\mathbf{u}_j, \mathbf{x}_j) = &\sum_{k=1}^{N_t} \frac{Q_{op}(t_k) r_{op} + Q_{gp}(t_k) r_{gp}}{(1+d)^{t_k/\tau}} \\ &- \sum_{k=1}^{N_t} \frac{Q_{wp}(t_k) r_{wp} + Q_{wi}(t_k) r_{wi} + D(t_k)}{(1+d)^{t_k/\tau}},\end{aligned} \tag{16}$$

where $Q_{op}(t_k), Q_{gp}(t_k), Q_{wp}(t_k)$ and $Q_{wi}(t_k)$ are the total oil production, gas production, water production and water injection volumes (in SM3) over the time interval t_k, respectively. r_{op}, r_{gp}, r_{wp} and r_{wi} are the corresponding oil price, gas price, water disposal cost and water injection cost (in \$/SM3), respectively. $D(t_i)$ is the total well drilling and completion costs during time interval t_i. d is the discount rate per year and τ is the number of days per year. We use constant values for d and τ in this case, where $d = 0.08, \tau = 365$.

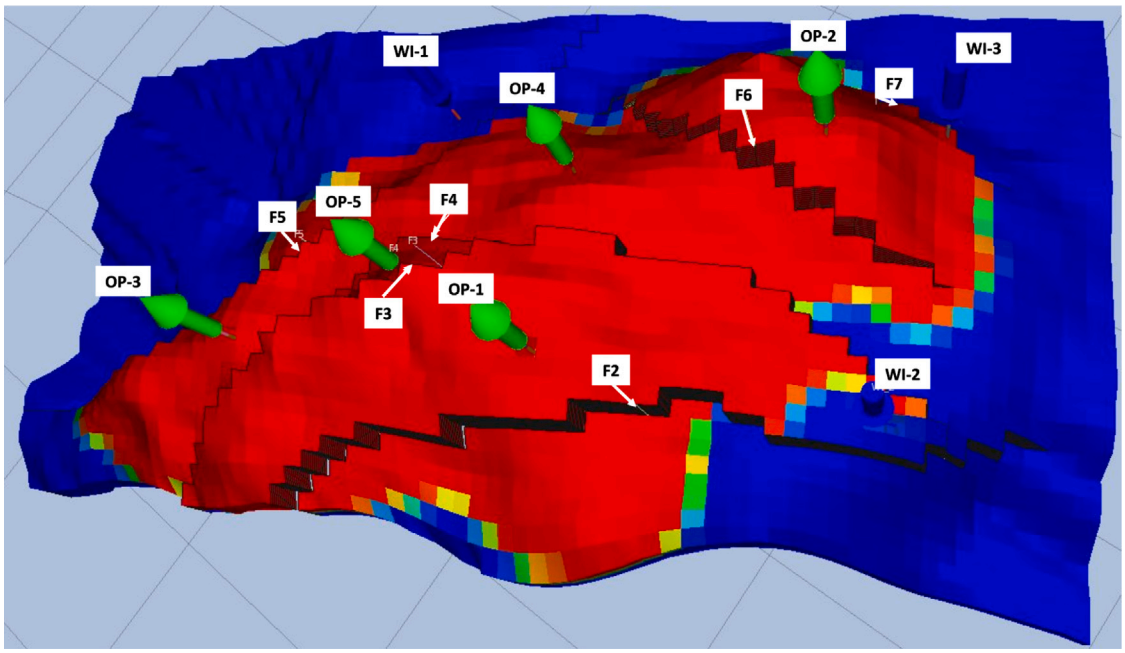

Fig. 3. The positions of wells and faults on the initial oil saturation map of the Reek field.

Table 1
Example workflow process. We discuss the results from the various workflow steps in Section 4.

INIENS	Generate initial ensemble based on all prior information from well logs, 3D-seismic, and geological interpretation.
ENSPRED0	Initial ensemble prediction used to assess well production potentials
OPT1	Robust optimization to decide which two wells to drill the first year
DEC1	Robust assessment of NPV for drilling zero, one, or two optimal wells
PROD1	Produce reservoir with two first wells for one year
HM1	Ensemble history matching using data for year one
ENSPRED1	Ensemble prediction using history-matched models
OPT2	Robust optimization to decide wells three and four to drill the second year
DEC2	Robust assessment of NPV for drilling zero, one, or two optimal wells
PROD2	Produce reservoir including eventual additional wells for one more year
HM2	Ensemble history matching using data for year two
ENSPRED2	Ensemble prediction using history-matched models
OPT3	Robust optimization to decide wells five and six to drill the third year
DEC3	Robust assessment of NPV for drilling zero, one, or two optimal wells
PROD3	Produce reservoir including eventual additional wells for one more year
HM3	Ensemble history matching using data for year three
ENSPRED3	Ensemble prediction using history-matched models
DEC4	Robust assessment of NPV for drilling zero, one, or both of the final two wells

In the NPV calculation in Eq. (16) we have decided to ignore possible additional costs from carbon taxes. These taxes vary between countries, but we can easily add the additional costs from a carbon tax to the objective function. A carbon tax will reduce the profitability of each well, and it might change the decisions. As the energy companies now have increased emphasis on reducing CO_2 releases, almost regardless of the carbon taxation, it is becoming important to address multi-objective optimization where we maximize NPV while minimizing the CO_2 emissions. The current workflow will allow for such future applications.

We are using a twin experiment approach for the geology, but we do not know the optimal drilling strategy. We could have evaluated the optimal drilling strategy by a brute force simulation on the reference model, although with eight wells that would require 40320 simulations, which is more than we can afford with the current model and computer resources. Also, the purpose of this work is not to find the absolutely optimal drilling strategy but to demonstrate a consistent and affordable workflow that leads to a near-optimal drilling strategy within the uncertainty in the reservoir knowledge.

Initially, the reservoir uncertainty is significant as we have limited information to determine the reservoir properties. Thus, when we compute the two initial wells to drill, we do this under considerable reservoir uncertainty. However, as soon as we start producing the reservoir, we will obtain dynamic data and use these in the history matching to improve the reservoir characterization. The subsequent optimization of the wells' drilling schedule will be more accurate due to the reduced reservoir uncertainty. In other words, we propose a workflow where we recursively introduce new optimal wells to drill next. We condition the model ensemble on the new dynamic data as soon as the wells provide new dynamic data. The reservoir model and the drilling order optimization will improve from one year to another. Additionally, the accuracy of the ensemble predictions will improve with each model update and lead to decision-making with reduced uncertainty.

We start by dividing the reservoir lifetime into different phases:

Initialization: We define the initial reservoir model and its uncertainty and represent it using the ensemble of model realizations (**INIENS**). We assigned uncertainty to the porosity, permeability, and fault multipliers (see the detailed specification below).

We then use this model ensemble in ensemble predictions **ENSPRED1** to define well targets and propose a drainage strategy for the reservoir.

First year: We must decide which wells to drill first (i.e., assuming two wells drilled in parallel). Although there are no observation data at the initial period, we can still robustly optimize on drilling sequence using the initial ensemble. After deciding on the two first wells from the optimization in **OPT1**, we assess their NPV over the ensemble in **DEC1** before drilling them and producing the reservoir for the first year (**PROD1**). We acquire dynamic data from the first year during production and use these data to history match the model in **HM1**. From the history-matched model, we can run ensemble predictions **ENSPRED1** to assess future production or even plan new well-targets to evaluate during the decision process for the following year.

Second year: We repeat the procedure from year one for the second year. However, we start by using the now improved history-matched models with reduced geological uncertainty from the first year to make predictions and optimize for the two subsequent wells to drill in **OPT2**. After that, we evaluate if the two optimal wells give a positive NPV in **DEC2**, and if they are economical, we decide to drill them. Following the reservoir production, we condition the model on all the newly acquired data in **HM2**, and we are ready to plan new wells and optimize their drilling order for the third year.

Third year: We repeat the procedure from the second year to decide on the subsequent two wells.

Fourth year: In the current example, there are only two wells left to decide on for the fourth year, so we do not need the optimization process; we can perform the decision-making **DEC4** directly through ensemble scenario studies. However, we could have used reservoir ensemble from the third year to plan additional well-targets to select from in an optimization process **OPT4**.

The above division into an annual process is quite common in reservoir units. It is called an annual maturation process (AMAP) and follows a natural business process with yearly budgeting and economic predictions. The AMAP includes model updating, the planning of new wells and their evaluation, and the decision to drill or not to drill. Our ensemble workflow is a complete workflow supporting the AMAP and includes advanced history-matching and optimization tools in the process. And the workflow is robust, accounting for the total uncertainty.

Our experiment follows the steps given in Table 1. We start with an ensemble of 100 model realizations with different porosities, permeabilities, and fault multipliers.

The porosity fields have an average value equal to the average porosity from core data, and the variance should reflect the variability between the cores. We sampled the porosity fields from a distribution with mean equal to 0.18 and standard deviation of 0.1, restricted to the interval $[0.00001, 2.7]$. For the logarithm of the permeability we sampled the random fields from a distribution with mean equal to 2.4 and standard deviation of 0.8, with permeability values restricted to the interval $[0, 2000]$. For both porosity and permeability, we used an non-isotropic decorrelation scale of 1.0 km in the x-direction and 2.0 km in the y-direction. The vertical correlation between the sampled fields is 0.6, while the correlation between porosity and permeability is 0.75.

In addition to having uncertainty in the porosity and permeability, we assume the fault transmissibilities to be uncertain, and we sample them as log-uniform between zero and one.

We also define one model realization to represent the true reservoir geology. Given a drilling order, we simulate this model to obtain the "observed data" on which we condition the ensemble. Thus, we use a twin-experiment approach.

4. Results

Following the workflow shown in Fig. 1 and Table 1, we will focus on the results of the optimization, history matching, and the final decision. To have a more straightforward comparison, we discuss the optimization steps **OPT1**, **OPT2**, and **OPT3** together. After that, we will discuss the results from the history-matching steps **HM1**, **HM2**, and **HM3**, and finally we discuss the final decision **DEC4**. As the ensemble methods used for robust optimization are less mature than the ensemble history matching, we will discuss the drilling order optimization in more detail.

4.1. Drilling order optimization

To use EnOPT with discrete variables, we define an underlying continuous variable, drilling priority that varies between zero and one, to estimate. Then the well with the highest drilling priority is drilled first. Initially, we generate a sequence of drilling priorities between zero and one for all the wells, defining our starting drilling sequence. After that, we generate an ensemble of drilling priorities by adding random numbers drawn from a normal distribution $\mathcal{N}(0, \sigma^2)$ to generate perturbed drilling sequences. The variance parameter allows for tuning the magnitude of the perturbations in the ensemble of drilling priorities to ensure that the resulting drilling sequences differ from each other, but not too much. In our experiments we used $\sigma = 0.1$.

To determine which wells we should drill first, we optimized the drilling priorities over the reservoir lifetime. Given an ensemble of drilling priorities, the mean priorities define the current iteration's best drilling sequence. We then compute the NPV values over the geological ensemble from the production of the reservoir lifetime for all geological realizations. The average NPV is used to evaluate which drilling sequence is the best.

We could use EnOPT on a case without geological uncertainty, i.e., only one realization of the reservoir model. We would then have a gradient evaluated as the regression between the ensemble of predicted NPVs and the ensemble of controls. When adding a stochastic component to the reservoir model, we reduce the correlation between the control variables and the predicted NPV, and this causes problems for the EnOPT method. We can imagine cases where the geological uncertainty is so considerable that it completely masks any correlation between the NPVs and controls, and in this case, the method will fail. Thus, we restrict using EnOPT to cases of modest geological uncertainty. In the current example, we also expect the impact of uncertain geology to be the largest in the first iteration before we have started to condition the geological ensemble to dynamical measurements.

We have plotted the results from the three optimization steps **OPT1**, **OPT2**, and **OPT3** in Fig. 4 from the upper to the lower rows of plots. We show the mean NPV values versus the optimization iterations in the left column and the NPVs for the best run in the boxplots in the right column. In Table 2 we list the starting and resulting drilling sequences for all the optimization experiments. We also note that we might run several trials in each iteration until we obtain a new sequence that leads to a higher NPV.

For the first optimization step, OPT1 shown in the upper row of Fig. 4, we ran five optimization experiments using two different starting points. As the initial geological uncertainty is significant, performing several experiments searching for the best possible initial drilling sequence makes sense. The runs initialized with different random seeds for generating the control perturbations result in different optimization results, even when the geological ensemble is the same. In addition to solving a problem with multiple local maxima, the EnOPT method uses a stochastic gradient. Hence, there is a possibility of finding different results when rerunning the optimization using other random seeds or starting points. The gray line in the upper-left plot shows a trend of declined NPVs from iteration two because this experiment had a slightly different stopping criterion compared with other runs.

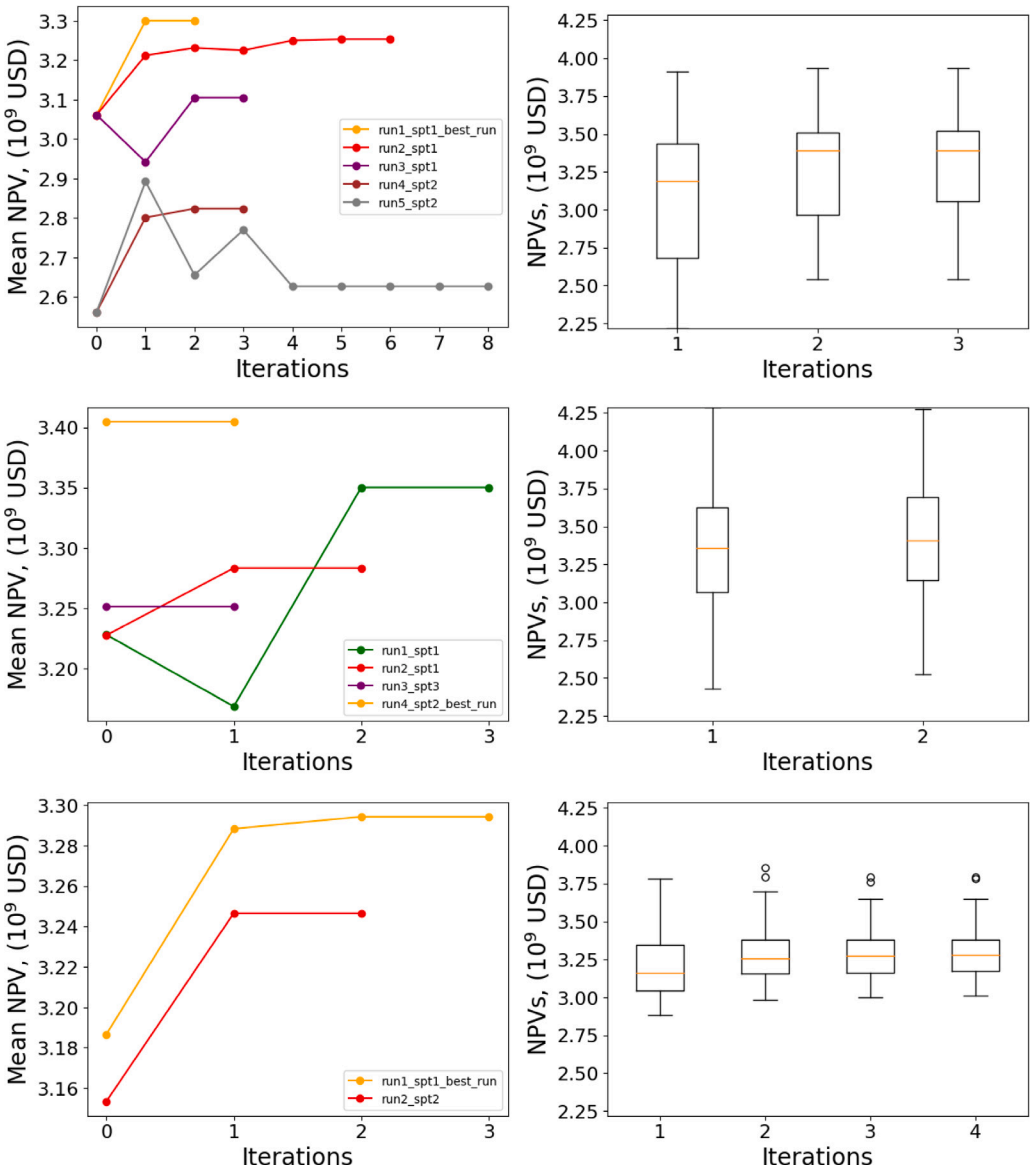

Fig. 4. The plots show the optimization steps **OPT1**, **OPT2**, and **OPT3** from the upper to the lower plots. In the left column we show the ensemble-mean NPV as a function of the iterations with the orange lines representing the results of the best runs. The labels, stp1, stp2, and stp3, denote starting point 1–3, for the optimization, and run1_stp1 and run2_stp1 denote two runs from the same starting point but with different random seeds for the control perturbations. The boxplots in the right column show the NPVs for the best runs's mean strategy evaluated on all geological realizations. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

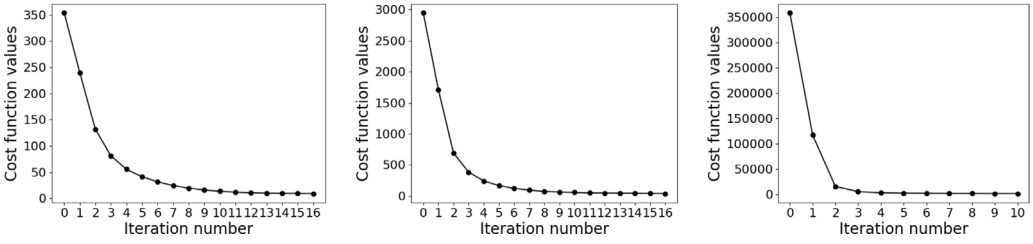

Fig. 5. Cost function values as a function of the iteration number for the history-matching steps **HM1**, **HM2**, and **HM3**.

To avoid the EnOPT algorithm getting trapped at a local optimum, we perform backtracking, i.e., we recompute the optimization with a smaller step size but using the same gradient, which sometimes allows us to find a drilling order with higher NPV. The stopping criteria for optimization is when the objective function fails to improve after several trial steps of backtracking. For the experiment represented by the gray line in the upper left plot of Fig. 4, if the backtracking fails five times, we re-sample the ensemble of drilling priorities with reduced variance. If this procedure fails three times, the algorithm stops. This approach is a time-consuming process, even for minor problems. For other experiments, with limited computational resources, we assume

Table 2
The drilling sequences obtained from the different optimization runs. Blue color marks the experiments and results that give the highest mean NPV values. Red color marks the starting points.

Runs	Year 1		Year 2		Year 3		Year 4		NPV ($\$10^6$)
OPT1									
spt1	*WI-2*	*OP-5*	*OP-1*	*WI-3*	*OP-4*	*WI-1*	*OP-2*	*OP-3*	*3 061*
run1-spt1	OP-1	WI-2	OP-4	WI-3	OP-5	WI-1	OP-2	OP-3	3 299
run2-spt1	OP-1	WI-2	OP-2	OP-5	WI-1	WI-3	OP-4	OP-3	3 253
run3-spt1	OP-1	WI-3	OP-5	WI-2	OP-4	WI-1	OP-2	OP-3	3 104
spt2	*WI-1*	*WI-3*	*OP-1*	*OP-5*	*OP-3*	*OP-4*	*OP-2*	*WI-2*	*2 561*
run4-spt2	OP-3	WI-3	OP-1	WI-1	OP-4	OP-5	OP-2	WI-2	2 823
run5-spt2	OP-5	WI-2	WI-1	WI-3	OP-3	OP-2	OP-1	OP-4	2 626
OPT2									
spt1	OP-1	WI-2	*OP-5*	*WI-1*	*OP-2*	*WI-3*	*OP-4*	*OP-3*	*3 227*
run1-spt1	OP-1	WI-2	OP-2	WI-3	OP-5	WI-1	OP-4	OP-3	3 350
run2-spt1	OP-1	WI-2	OP-3	WI-3	OP-4	OP-5	OP-2	WI-1	3 283
spt2	OP-1	WI-2	*OP-4*	*WI-3*	*OP-3*	*WI-1*	*OP-2*	*OP-5*	*3 405*
run4-spt2	OP-1	WI-2	OP-4	WI-3	OP-3	WI-1	OP-2	OP-5	3 405
spt3	OP-1	WI-2	*OP-5*	*WI-3*	*OP-2*	*WI-1*	*OP-3*	*OP-4*	*3 251*
run3-spt3	OP-1	WI-2	OP-5	WI-3	OP-2	WI-1	OP-4	OP-3	3 251
OPT3									
spt1	OP-1	WI-2	OP-4	WI-3	*WI-1*	*OP-3*	*OP-2*	*OP-5*	*3 187*
run1-st1	OP-1	WI-2	OP-4	WI-3	OP-2	OP-3	OP-5	WI-1	3 294
spt2	OP-1	WI-2	OP-4	WI-3	*OP-3*	*OP-5*	*WI-1*	*OP-2*	*3 153*
run2-st2	OP-1	WI-2	OP-4	WI-3	OP-2	WI-1	OP-3	OP-5	3 246
DEC4									
p5w1	OP-1	WI-2	OP-4	WI-3	OP-2	OP-3	OP-5	WI-1	3 231
p5	OP-1	WI-2	OP-4	WI-3	OP-2	OP-3	OP-5		3 228
w1	OP-1	WI-2	OP-4	WI-3	OP-2	OP-3		WI-1	3 255
nowells	OP-1	WI-2	OP-4	WI-3	OP-2	OP-3			3 257

the optimization converges if the backtracking trial steps fail three times.

For **OPT2** and **OPT3**, we used only three backtracking steps when the algorithm stops at a local optimum. After three trial backtracking steps, we assume that the algorithm converges to a local optimum if the algorithm cannot find a better solution. For **OPT2** we used three different starting points, with SPT2 giving the best result. We generated the starting sequences for the optimization randomly.

The boxplots in the right column show a significant improvement with a reduced spread of the NPVs from **OPT1** to **OPT3**. In the boxplots, we evaluate the NPV of each iteration's best drilling strategy over the geological ensemble, and we illustrate the minimum and maximum values together with the first quartile, the third quartile and median represented by the lower and upper edge of the box, and the orange line inside the box, respectively.

We see that EnOPT benefits from a reduced geological uncertainty induced by the history-matching steps. In **OPT3** the stochastic gradient becomes more stable, and the NPV estimation more accurate, leading to improved robustness of the method. We have a sound improvement from one iteration to the next for the two cases with different starting points. Still, the two starting points results in different drilling sequences, and we might improve the results further by running from additional starting points.

From the results displayed in Fig. 4, we notice that for **OPT1**, the three different starting sequences result in different solutions. Changing the initial well-priority ensembles result in other solutions even with the same starting point. We improved the drilling sequence compared to the random starting sequences, but we would assume that the best solution we found is still sub-optimal solution or near-global-optimal solution.

In our experiment, the first two wells drilled contribute most strongly to the total NPV. The reason is that these wells gain benefit first, and we produce them for the longest time. When we have multiple wells in the optimization sequence, we will have many local optima. From **OPT1**, all the solutions have a producer and injector pair for the first year. Thus, the first optimization **OPT1** is about finding the best injector–producer pair for year one. Depending on the starting point, the optimization locks on to a producer–injector pair and then continues exploring the wells for the subsequent years. Hence, if EnOPT locks on to a sub-optimal producer–injector pair for the first year, this corresponds to a local optimum. Therefore, it is important to run the optimization from different starting points to reduce the risk of converging to a local optimum.

4.2. History matching

We drill the two new wells having the highest priority and produce them along with the existing wells for a year after each optimization step. After a year's production, we use the produced and injected rates to history-match our ensemble of models. The additional observations obtained during the last year's production provide new information about the reservoir and lead to better reservoir characterization and lower geological uncertainty. For the history matching, we use the subspace EnRML method introduced above. Note that we initialize the history matching with the same initial ensemble and conditions to all observations available from the start of production. And for the reservoir simulations, we integrate the model ensemble from the initial time. A sequential updating of the reservoir using only the newly acquired observations is possible using the ensemble Kalman filter but introduces additional computational issues as discussed by Skjervheim et al. (2011).

In Fig. 5, the average cost-function values as a function of the iteration number illustrates the convergence of the subspace EnRML method. The additional observations introduced every year lead to a higher initial cost-function value. However, after a few iterations, the updated model parameters result in a cost-function value reduced to a level similar to the number of measurements.

4.3. Case analysis

The initial ensemble allows us to run ensemble simulations (**ENSPRED0**) with different drilling-order and well configurations, and we can assess the production uncertainty consistently with our prior-uncertainty assumptions. The optimization step **OPT1** takes this assessment one step further and computes the optimal wells to drill first. The optimization (see the upper plots in Fig. 4 and Table 2)

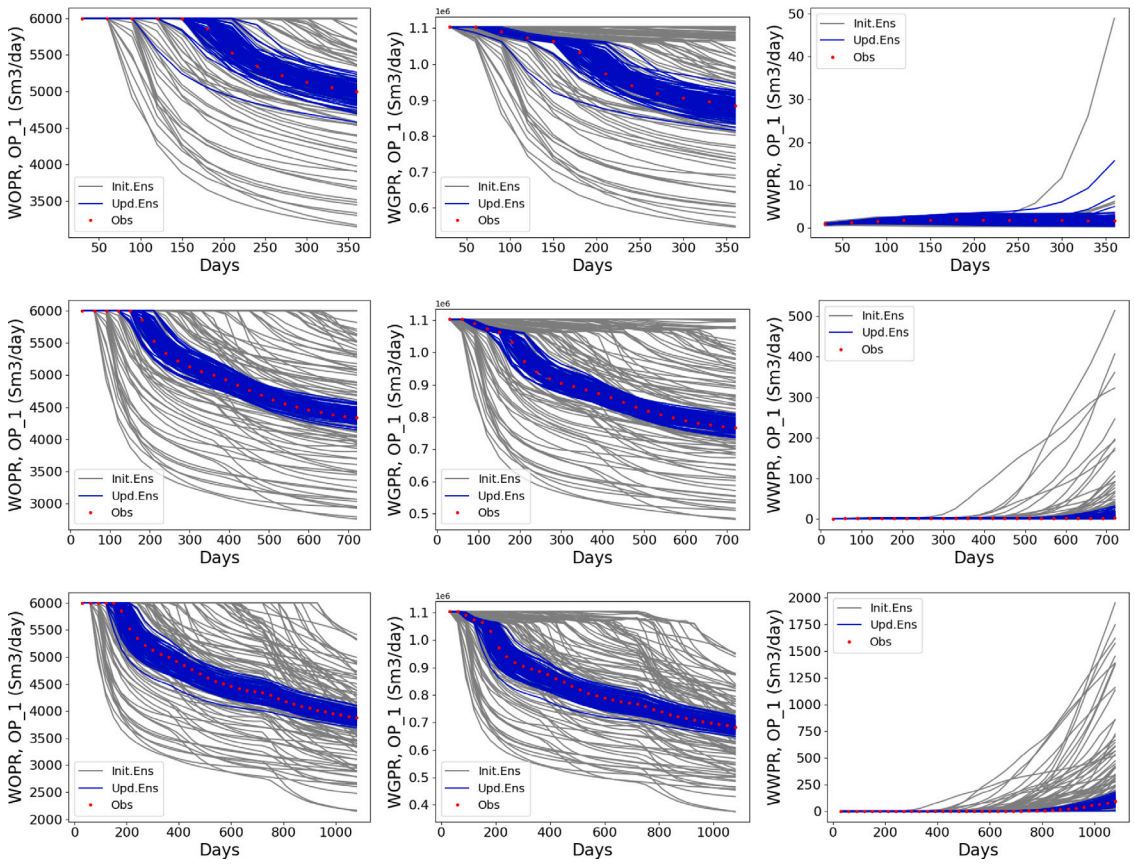

Fig. 6. Production profiles for OP-1 following the different history-matching steps **HM1** (upper), **HM2** (middle), and **HM3** (lower). Gray curves, blue curves and red dots represent production data for the prior ensemble, posterior ensemble and the observation data respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

tells us that the producer OP-1 and the injector WI-2 have the highest drilling priorities. As they also have a positive NPV, the decision **DEC1** is to drill these wells. We then produce the wells for one year (**PROD1**) and acquire the oil, water, and gas production rates for each well (WOPR, WGPR, and WWPR), and the injection rates (WWIR) as monthly observation data. Hence, we have 48 observations from the first year to use in the history-matching step **HM1**.

Fig. 6 shows the results from **HM1** for the firstly drilled producer in the upper plots (the middle and lower plots show the results after **HM2** and **HM3**, discussed below). The plots present the prior and posterior ensembles of production profiles and the observations. After the history matching, there is a clear improvement in the ensemble prediction's match to the observed production rates. Similarly, in Fig. 7, the upper left plot shows the improvement in the results for the water injector WI-2. Thus, the history-matching step has updated the reservoir porosity, permeability, and fault transmissibilities, such that OP-1 and WI-2 can produce and inject the imposed observed rates. We also expect that the updated ensemble of reservoir models will lead to a more accurate optimization and decision, with less uncertainty on the subsequent wells to drill due to the conditioning on observations by the HM.

Next, we use the history matched ensemble from **HM1** to optimize the subsequent two wells to drill. After the optimization step **OPT2**, the producer OP-4 and the injector WI-3 have the highest drilling priorities, and as these wells lead to a positive NPV, we decide (**DEC2**) to drill them next.

In **PROD2**, we produce the reservoir while including the new wells and the ones drilled initially. Hence, for the second year, we have the two producers, OP-1 and OP-4, together with the two injectors, WI-2 and WI-3. We produce the reservoir until the end of the second year while acquiring the monthly production and injection data, leading to a total of 144 observation data that we use in the following history-matching step, **HM2**.

We present the results from **HM2** in the second row of Fig. 6, the second row of Fig. 7 and the first row of Fig. 8. In Fig. 6, we see that the history matching further reduces the uncertainty of the prediction for OP-1 and that we obtain an excellent match to the observations. In particular, the history matching eliminates some problematic realizations that produce too much gas in the initial ensemble. We also obtained good results for the water injectors shown in Fig. 7 and the producer OP-4 presented in Fig. 8.

Again we follow the history-matching step with another optimization step **OPT3**, and we decide (**DEC3**) to drill the producers OP-2 and OP-3 with the highest drilling priorities next, as their production also leads to a positive NPV.

We then produce all drilled wells until the end of the third year while acquiring the monthly observation data. We now have 312 observation data on which we condition the model. In the lower row of plots in Figs. 6–9, we see the results from the final history-matching step **HM3**. For all wells, we obtain an excellent match to the data. We would, of course, expect this in a twin experiment where the ensemble statistics and model used is consistent with the reference case.

We do not need to optimize for drilling priorities when there are only two candidate drilling locations left for the final decision step **DEC4**. However, we could have included additional planned well alternatives in an optimization computation.

Thus, we have a final decision step **DEC4**, of determining whether to drill or not, the remaining wells OP-5 and WI-1. In the decision step **DEC4**, we evaluate the NPVs or different drilling scenarios. I.e., drill none of the wells (nowells), drill only OP-5 (p5), drill only WI-1 (w1), or drill both wells (p5w1). The procedure simulates the four different scenarios over the ensemble of history-matched models and evaluates the NPV of each realization for each scenario. We can then estimate the mean and uncertainty of the NPV for each scenario. We subtract

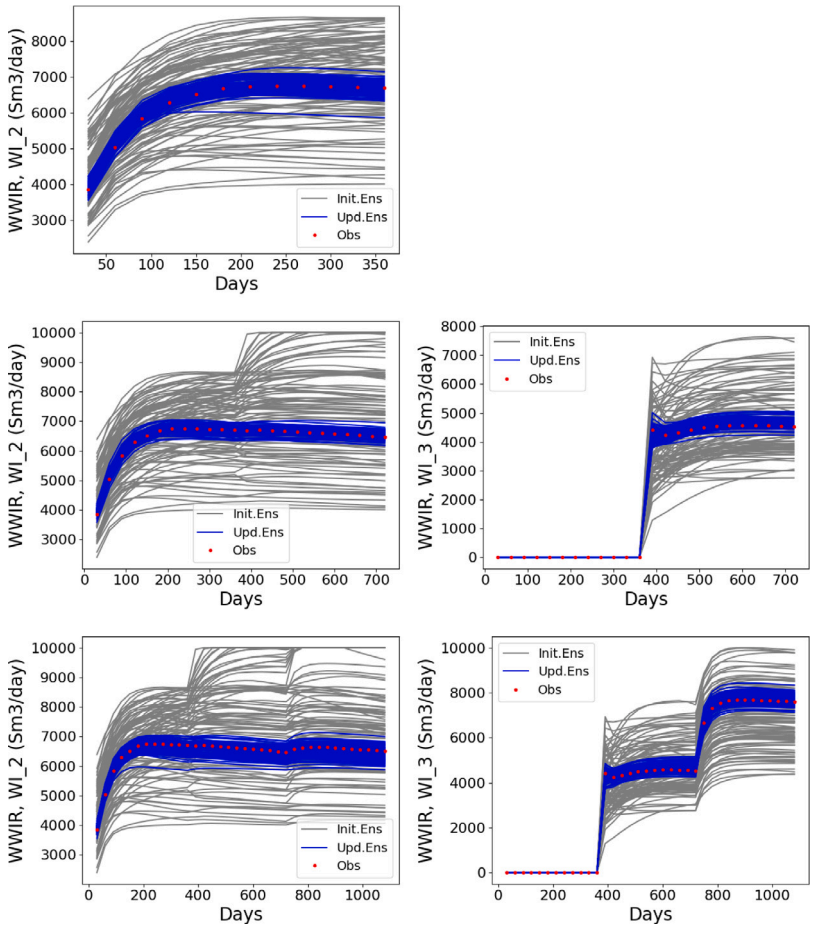

Fig. 7. Injections rates of WI-2 (left) and WI-3 (right) following the three history-matching steps **HM1** (upper), **HM2** (middle), and **HM3** (lower). Gray curves, blue curves and red dots represent data for the prior ensemble, posterior ensemble and the observation data respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

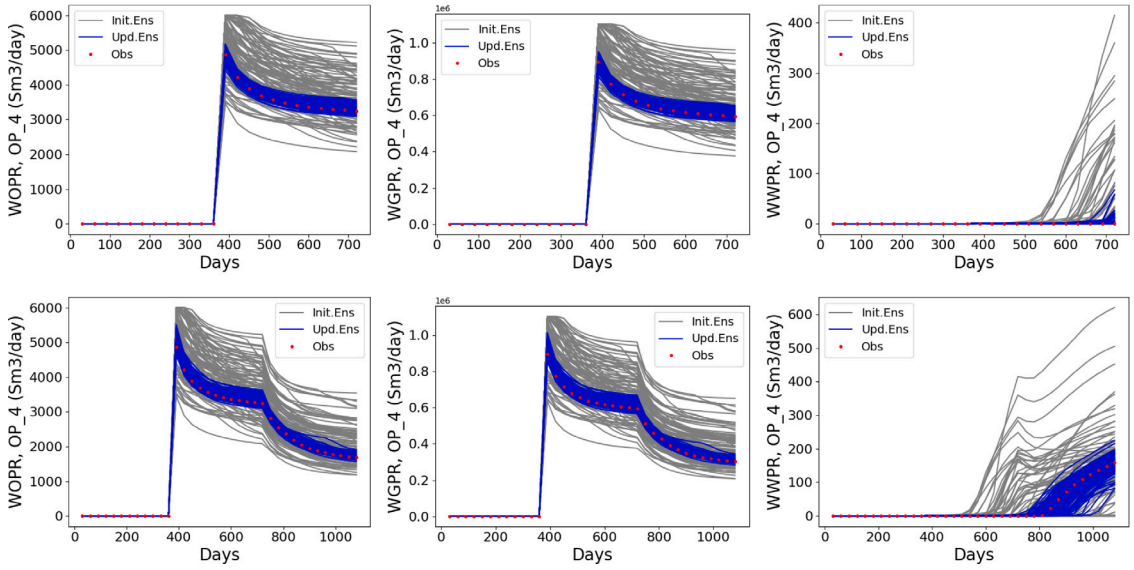

Fig. 8. Production profiles for OP-4 following the different history-matching steps **HM2** (upper) and **HM3** (lower). Gray curves, blue curves and red dots represent production data for the prior ensemble, posterior ensemble and the observation data respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the drilling costs of an assumed 60 million USD per well in the final evaluation. In a realistic setting, we should also include the operating costs of the wells.

We illustrate the NPV evaluation results in Fig. 10. In the left plot, we observe that the scenario of not drilling any wells yields the highest mean NPV and has the least uncertainty, as shown in the right plot.

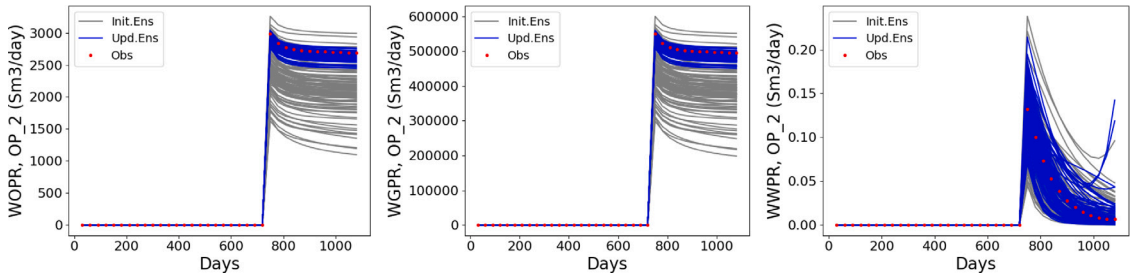

Fig. 9. Production profiles for OP-2 following the history-matching step **HM3**. Gray curves, blue curves and red dots represent production data for the prior ensemble, posterior ensemble and the observation data respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

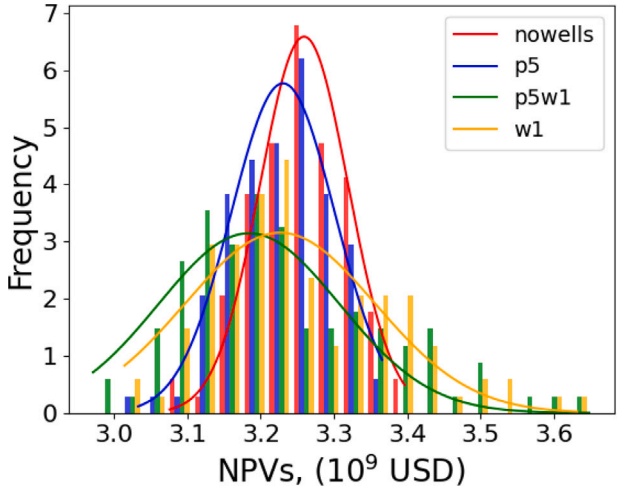

Fig. 10. The figure compares the four decision scenarios, i.e., drilling no wells (nowells), drilling only producer five (p5), drilling only water injector one (w1), or drilling both wells (p5w1). We plot the number of realizations falling in a particular NPV interval as histograms for the four decision scenarios as a function of the NPV given in 10^9 USD. The NPV values are after subtracting the drilling costs of USD 60M per well.

Therefore, the decision of the final step is not to drill any more of the planned wells and rather continue producing the reservoir with the existing six wells. We also note that the distributions are highly overlapping, and there are many realizations that, if used alone, would give a different conclusion on whether to drill none, one, or two wells. However, the robust decision is to select the alternative with the highest mean NPV when averaged over the ensemble. Note that the ensemble histograms are close to Gaussian in the current case, which simplifies the decision process.

Finally, we illustrate the updating of fault multipliers in Fig. 11. We show the prior, posterior, and actual reference values of the six fault multipliers and, from top to bottom, the results after **HM1**, **HM2**, and **HM3**. The positions of the faults and wells are in Fig. 3. We observe that the fault multipliers are more sensitive to the production data from wells near the faults. We denote the six faults as F2, F3, F4, F5, F6, and F7. The data are in the logarithm scale. We see that during **HM1**, the fault multiplier of F3 is more sensitive to the production of the open wells OP-1 and WI-2 because these two wells locate between faults F2 and F3. Similarly, we see a clear update for F3, F5, and F6 after **HM2**. The update results from the new observation data from the newly drilled wells OP-4 and WI-3. The well OP-4 locates between faults F4 and F5, while WI-3 is between the faults F6 and F7. Furthermore, **HM3** updates the faults F5 and F7. Because we drill wells OP-2 and OP-3 before **HM3**. OP-2 is close to F7, and OP-3 close to F5. The observations from these two wells help update faults F5 and F7. We also note that we get close to the reference values used in the "truth" model realization for all significantly updated fault multipliers. We have similar updates of the three-dimensional porosity and permeability fields, where data from different wells update different spatial regions. However, these updates are more elaborate, and we have left them out of this discussion.

5. Discussion

This work demonstrates a workflow of ensemble-based decision-making for closed-loop reservoir management. The decision workflow integrates well-drilling-sequence optimization and history-matching using ensemble methods that correctly account for uncertainty. Furthermore, using an ensemble of models allows for robust decision-making with decision alternatives evaluated over the ensemble of models. This decision workflow is independent of the history-matching and optimization methods used, but it requires an ensemble of model realizations to represent and predict the uncertainty. Hence, we only implement an option if it results in a positive NPV on average over the ensemble of models.

We have demonstrated the workflow on a simple reservoir case, REEK, provided by Equinor. We used the EnOPT method for the optimization and the subspace EnRML for history matching. While the history matching is well developed and its application well understood, we observe that the optimization methods and strategies are less mature and define an important area for further research. With significant reservoir uncertainty, it is not clear that EnOPT is the best method since the reservoir uncertainty lead to less correlation between the optimization controls and the predicted objective. The approach taken here uses multiple starting points, which helps the algorithm find a better solution and avoid getting trapped in local optima. We generated the starting sequences for the optimization randomly. An alternative strategy could be to start from the best sequence from the previous optimization step or perturbations of it if we wish to explore multiple starting points.

An argument against optimizing the drilling sequence over the reservoir's lifetime is that we initially have significant uncertainty in the reservoir description. The resulting optimal well sequence will likely change after the next history-matching steps, which reduce the reservoir uncertainty and improve the ensemble of models. So maybe EnOPT would work better and more easily find the optimal solution if we only optimize for one year at a time. We could even constrain the optimization to search for the optimal injector–producer pair, which reduces the number of alternative solutions to 15, i.e., we match each of the five producers to each of the three injectors.

Another issue is that EnOPT is a gradient method, and here we use it to optimize a discrete sequence of wells. We have assigned a continuous variable named drilling priority to each well, and we drill the wells in the order of their drilling priority. This procedure works well, as also shown by Hanea et al. (2017).

The history matching updates the model and reduces the reservoir uncertainty. Thus, after some model updates, we expect the reduced reservoir uncertainty to improve EnOPT performance. Still, we believe

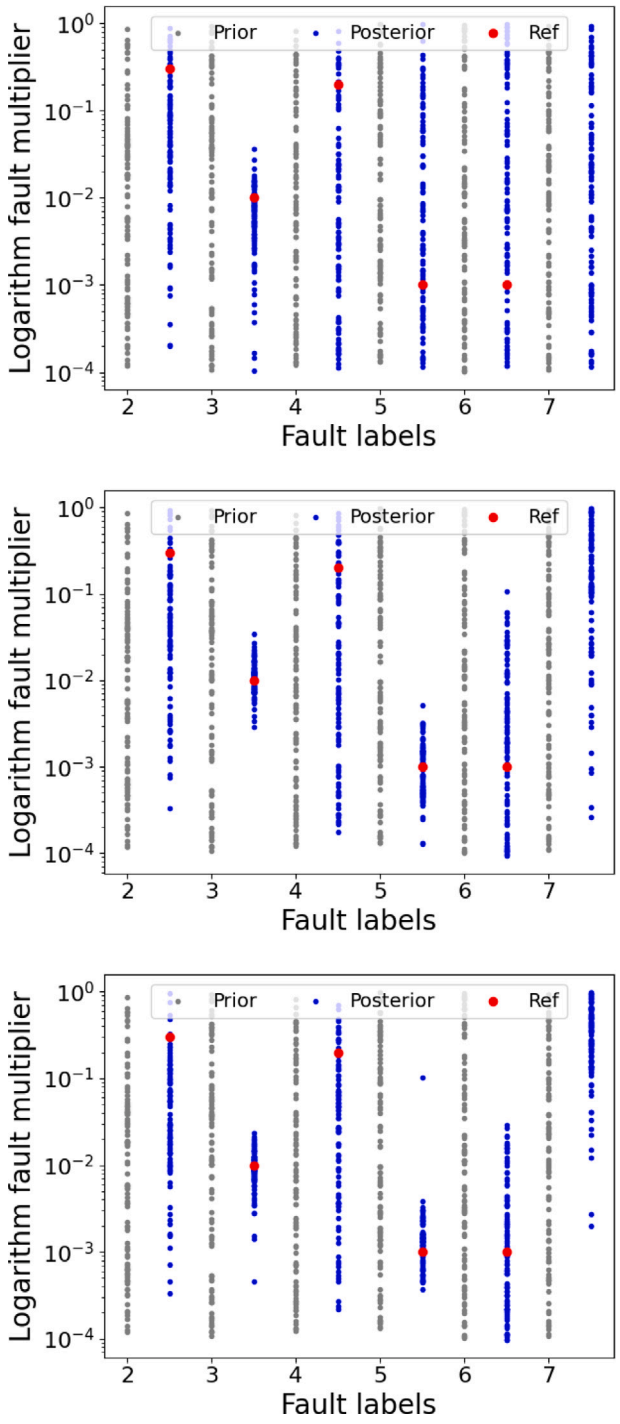

Fig. 11. Fault-multiplier updates from the history-matching steps **HM1** (upper), **HM2** (middle), and **HM3** (lower). Gray, blue and red dots represent prior, posterior and the true reference values. Note that we are using a log-scale for the fault multipliers. The zero reference value for F7 falls outside the range of the *y*-axis. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

that we should develop and evaluate other robust optimization methods that are not based on ensemble gradients. We are currently working with some of such methods, e.g., Wang and Oliver (2019, 2021) but it is too early to conclude on the performance of these methods compared to using EnOPT.

We have had less emphasis on the decision steps and assumed Gaussian ensemble predictions, which means we should base our decision on the ensemble mean. I.e., we select the alternative that gives the highest NPV value in average over the ensemble of reservoir models. In the case of non-Gaussian predictions, we might want to use more elaborate methods based on utility theory and include risk behaviors. At this point, we emphasize that the introduction of reservoir uncertainty into decision-making helps the decision-makers make more robust decisions.

A final exciting aspect of this work is that we now also have an efficient algorithm for "well-placement optimization". The traditional method for optimizing the placement of a well is to move it around in the reservoir until one finds the best well location. However, this is a highly nonlinear and challenging optimization problem, particularly for complex drilling trajectories and horizontal wells. Using the workflow outlined above, we can plan several additional wells based on the geologic understanding and optimize and select among these. Whenever it is uneconomical to drill any more wells, we continue producing the reservoir until the end of the reservoir life time. Every time we have updated the reservoir model with new information, we can further revise the drilling plans and targets.

A final remark is that there may be many other fields and applications outside petroleum where the closed-loop concept is sound. We anticipate that any observed time-recursive system where one specifies the controls variables that determine the system's future evolution could benefit from a closed-loop approach involving recursive use of robust control optimization, followed by a risk-based decision process to determine the controls, and finally, the use of data assimilation or history matching of newly acquired information to update the dynamic system variables and parameters. Potential application areas in the geosciences can be geothermal energy and groundwater management (as pointed out by a reviewer). Other areas, such as managing traffic flow by controlling traffic lights and redirecting traffic to less busy roads, or managing the medicine dosage for cancer treatment or in diabetic patients, would also benefit from the ensemble-based closed-loop concept.

CRediT authorship contribution statement

Yuqing Chang: Coding of methods, Running of simulations, Analysis of results, Plotting of results. **Geir Evensen:** Original idea and outline of work, Supervision, Main writer of paper, Funding acquisition, Project manager.

Declaration of competing interest

One or more of the authors of this paper have disclosed potential or pertinent conflicts of interest, which may include receipt of payment, either direct or indirect, institutional support, or association with an entity in the biomedical field which may be perceived to have potential conflict of interest with this work. For full disclosure statements refer to https://doi.org/10.1016/j.petrol.2022.110858. Geir Evensen reports financial support was provided by Research Council of Norway.

Acknowledgments

The authors thank Kristian Fossum, Rolf Johan Lorentzen and Andreas Størksen Stordal at NORCE for valuable discussions during this work. This work received support from the Research Council of Norway and the companies Equinor Energy AS, Aker BP ASA, Wintershall Dea Norge AS, Vår Energy AS, Petrobras, Lundin Energy Norway AS, and Neptune Energy Norge AS, through the Petromaks–2 DIGIRES project (280473) (http://digires.no).

References

Barros, E., den Hof, P.V., Jansen, J., 2020. Informed production optimization in hydrocarbon reservoirs. Opt. Eng. 21, 25–48. http://dx.doi.org/10.1007/s11081-019-09432-7.

Chang, Y., Lorentzen, R.J., Nævdal, G., Feng, T., 2019. OLYMPUS optimization under geological uncertainty. Comput. Geosci. http://dx.doi.org/10.1007/s10596-019-09892-x.

Chen, Y., Oliver, D.S., 2012. Ensemble randomized maximum likelihood method as an iterative ensemble smoother. Math. Geosci. 44, 1–26. http://dx.doi.org/10.1007/s11004-011-9376-z.

Chen, Y., Oliver, D.S., 2013. Levenberg-marquardt forms of the iterative ensemble smoother for efficient history matching and uncertainty quantification. Comput. Geosci. 17, 689–703. http://dx.doi.org/10.1007/s10596-013-9351-5.

Chen, Y., Oliver, D.S., Zhang, D., 2009. Efficient ensemble-based closed-loop production optimization. SPE J. 14 (4), 634–645. http://dx.doi.org/10.2118/112873-PA.

Emerick, A.A., Reynolds, A.C., 2012. History matching time-lapse seismic data using the ensemble Kalman filter with multiple data assimilations. Comput. Geosci. 16 (3), 639–659. http://dx.doi.org/10.1007/S10596-012-9275-5.

Evensen, G., 1994. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. J. Geophys. Res. 99 (C5), 10,143–10,162. http://dx.doi.org/10.1029/94JC00572.

Evensen, G., 2003. The ensemble Kalman filter: Theoretical formulation and practical implementation. Ocean Dyn. 53, 343–367. http://dx.doi.org/10.1007/s10236-003-0036-9.

Evensen, G., 2009. Data Assimilation: The Ensemble Kalman Filter, 2nd Springer, p. 320. http://dx.doi.org/10.1007/978-3-642-03711-5.

Evensen, G., 2021. Formulating the history matching problem with consistent error statistics. Comput. Geosci. 25, 945–970. http://dx.doi.org/10.1007/s10596-021-10032-7.

Evensen, G., Raanes, P.N., Stordal, A.S., Hove, J., 2019. Efficient implementation of an iterative ensemble smoother for data assimilation and reservoir history matching. Front. Appl. Math. Stat. 5, 47. http://dx.doi.org/10.3389/fams.2019.00047.

Evensen, G., Vossepoel, F.C., Van Leeuwen, P.J., 2022. Data Assimilation Fundamentals: A Unified Formulation for State and Parameter Estimation. Springer, ISBN: 978-3-030-96708-6, p. 254. http://dx.doi.org/10.1007/978-3-030-96709-3, Open access.

Fonseca, R., 2016. A Modified Gradient Formulation for Ensemble Optimization under Geological Uncertainty (Ph.D. thesis). Delft University of Technology, http://dx.doi.org/10.4233/uuid:e66b1e00-b4c2-43b8-91fa-c57773fcf24b.

Hanea, R.G., Casanova, P., Wilschut, F., Hustoft, L., Fonseca, R.M., 2017. Well trajectory optimization constrained to structural uncertainties. In: SPE Conference Paper. http://dx.doi.org/10.2118/182680-MS.

Hanea, R., Evensen, G., Hustoft, L., Ek, T., Chitu, A., Wilschut, F., 2015. Reservoir management under geological uncertainty using fast model update. In: SPE Conference Paper. p. 12, SPE-173305-MS.

Haugen, V.E., Nævdal, G., Natvik, L.-J., Evensen, G., Berg, A., Flornes, K., 2008. History matching using the ensemble Kalman filter on a North Sea field case. SPE J. 13, 382–391. http://dx.doi.org/10.2118/102430-PA.

Jansen, J.-D., Brouwer, R., Douma, S.G., 2009. Closed loop reservoir management. In: SPE Conference Paper. http://dx.doi.org/10.2118/119098-MS.

Jansen, J.D., Brouwer, D.R., Naevdal, G., Van Kruijsdijk, C.P.J.W., 2005. Closed-loop reservoir management. First Break 43–48. http://dx.doi.org/10.3997/1365-2397.2005002.

Kitanidis, P., 1995. Quasi-linear geostatistical therory for inversing. Water Resour. Res. 31 (10), 2411–2419. http://dx.doi.org/10.1029/95WR01945.

Leeuwenburgh, O., Chitu, A., Nair, R., Egberts, P., Ghazaryan, L., Feng, T., Hustoft, L., 2016. Ensemble-based Methods for Well Drilling Sequence and Time Optimization under Uncertainty. European Association of Geoscientists & Engineers, http://dx.doi.org/10.3997/2214-4609.201601871.

Leeuwenburgh, O., Egberts, P.J., Abbink, O.A., 2010. Ensemble Methods for Reservoir Life-Cycle Optimization and Well Placement. In: SPE Kingdom of Saudi Arabia Annual Technical Symposium and Exhibition, http://dx.doi.org/10.2118/136916-MS, SPE-136916-MS.

Lorentzen, R.J., Berg, A.M., Nævdal, G., Vefring, E.H., 2006. A new approach for dynamic optimization of water flooding problems. In: SPE Conference Paper. http://dx.doi.org/10.2118/99690-MS.

Lu, R., Reynolds, A., 2020. Joint optimization of well locations, types, drilling order, and controls given a set of potential drilling paths. SPE J. http://dx.doi.org/10.2118/193885-PA.

Nævdal, G., Johnsen, L.M., Aanonsen, S.I., Vefring, E., 2003. Reservoir monitoring and continuous model updating using the ensemble Kalman filter. SPE J. 10 (1), 66–74. http://dx.doi.org/10.2118/84372-PA.

Oguntola, M.B., Lorentzen, R.J., 2021. Ensemble-based constrained optimization using an exterior penalty method. J. Pet. Sci. Eng. 207, 109165. http://dx.doi.org/10.1016/j.petrol.2021.109165.

Oliver, D.S., He, N., Reynolds, A.C., 1996. Conditioning permeability fields to pressure data. In: ECMOR V-5th European Conference on the Mathematics of Oil Recovery. p. 11. http://dx.doi.org/10.3997/2214-4609.201406884.

Raanes, P.N., Stordal, A.S., Evensen, G., 2019. Revising the stochastic iterative ensemble smoother. Nonlin. Processes Geophys. 26, 325–338. http://dx.doi.org/10.5194/npg-2019-10.

Silva, V.L.S., Emerick, A.A., Couto, P., Alves, J.L.D., 2017. History matching and production optimization under uncertainties — Application of closed-loop reservoir management. J. Pet. Sci. Eng. 157, 860–874. http://dx.doi.org/10.1016/j.petrol.2017.07.037.

Skjervheim, J.-A., Evensen, G., Hove, J., Vabø, J.G., 2011. An ensemble smoother for assisted history matching. In: SPE Conference Paper. http://dx.doi.org/10.2118/141929-MS.

Skjervheim, J.A., Hanea, R.G., Evensen, G., 2015. Fast model update coupled to an ensemble based closed loop reservoir management. In: EAGE: Petroleum Geostatistics. EAGE, p. 5. http://dx.doi.org/10.3997/2214-4609.201413629.

Skjervheim, J.-A., van Lanen, X., Hulme, D., Stenerud, V.R., Zachariassen, E., Liu, S., Hove, J., Evensen, G., 2012. Integrated workflow for consistent model building from depth conversion to flow simulation - north sea field case. In: EAGE/SPE Conference Paper. http://dx.doi.org/10.3997/2214-4609.20148221.

Stordal, A.S., Szklarz, S.P., Leeuwenburgh, O., 2016. A theoretical look at ensemble-based optimization in reservoir management. Math. Geosci. 48 (4), 399–417. http://dx.doi.org/10.1007/s11004-015-9598-6.

Van Essen, G.M., Zandvlit, M.J., den Hof, P.M.J.V., Bosgra, O.H., Jansen, J.D., 2009. Robust waterflooding optimization of multiple geological scenarios. SPE J. 14 (1), 202–210. http://dx.doi.org/10.2118/102913-PA.

Van Leeuwen, P.J., Evensen, G., 1996. Data assimilation and inverse methods in terms of a probabilistic formulation. Mon. Weather Rev. 124, 2898–2913. http://dx.doi.org/10.1175/1520-0493(1996)124<2898:DAAIMI>2.0.CO;2.

Wang, L., Oliver, D.S., 2019. Efficient optimization of well drilling sequence with learned heuristics. SPE J. 24 (5), 2111–2134. http://dx.doi.org/10.2118/195640-PA.

Wang, L., Oliver, D.S., 2021. Fast robust optimization using bias correction applied to the mean model. Comput. Geosci. 25, 475–501. http://dx.doi.org/10.1007/s10596-020-10017-y.

Zachariassen, E., Skjervheim, J.A., Vabø, J.G., Lunt, I., Hove, J., Evensen, G., 2011. Integrated work flow for model update using geophysical monitoring data. In: EAGE/SPE Conference Paper. http://dx.doi.org/10.3997/2214-4609.201410330.