Clustering techniques for value-of-information assessment in closed-loop reservoir management

Barros, E.G.D.[1,*], Van den Hof, P.M.J.[2] and Jansen, J.D.[1]

[1] Delft University of Technology

[2] Eindhoven University of Technology

[*] Currently employed at TNO

Abstract

The application of closed-loop reservoir management (CLRM) to real-field cases can be computationally demanding. An even higher computational load results from procedures to assess the value of information (VOI) in CLRM. Such procedures, which are performed prior to field operation, i.e. during the field development planning (FDP) phase, require extreme amounts of simulations. Therefore, we look for alternatives to reduce this computational burden. In particular we study various clustering techniques to select a limited number of representative members from an ensemble of reservoir models. Using k-means clustering, multidimensional scaling and tensor decomposition techniques, we test the effectiveness of different dissimilarity measures such as distance in parameter space, distance in terms of flow patterns, and distance in optimal sets of controls. We apply several of these measures to a VOI-CLRM exercise using a simple 2D reservoir model which results in a reduction of the necessary number of forward reservoir simulations from millions to thousands. Finally, as a first step towards large-scale application, we assess the VOI in a larger benchmark case study.

Keywords: Representative model realizations, cluster analysis, clustering features, value of information, plausible truths, closed-loop reservoir management

E.G.D. Barros (corresponding author)
e.barros@tudelft.nl / eduardo.barros@tno.nl
Tel: +31 642318565

P.M.J. Van den Hof
p.m.j.vandenhof@tue.nl

J.D. Jansen
j.d.jansen@tudelft.nl

Clustering techniques for value-of-information assessment in closed-loop reservoir management

Barros, E.G.D.[1,*], Van den Hof, P.M.J.[2] and Jansen, J.D.[1]

[1] Delft University of Technology

[2] Eindhoven University of Technology

[*] Currently employed at TNO

Abstract

The application of closed-loop reservoir management (CLRM) to real-field cases can be computationally demanding. An even higher computational load results from procedures to assess the value of information (VOI) in CLRM. Such procedures, which are performed prior to field operation, i.e. during the field development planning (FDP) phase, require extreme amounts of simulations. Therefore, we look for alternatives to reduce this computational burden. In particular we study various clustering techniques to select a limited number of representative members from an ensemble of reservoir models. Using k-means clustering, multidimensional scaling and tensor decomposition techniques, we test the effectiveness of different dissimilarity measures such as distance in parameter space, distance in terms of flow patterns, and distance in optimal sets of controls. We apply several of these measures to a VOI-CLRM exercise using a simple 2D reservoir model which results in a reduction of the necessary number of forward reservoir simulations from millions to thousands. Finally, as a first step towards large-scale application, we assess the VOI in a larger benchmark case study.

1. Introduction

Modern reservoir management workflows include uncertainty quantification (UQ) based on reservoir simulation models. An increasingly popular UQ practice in the reservoir engineering community uses ensembles of reservoir model realizations to account for the geological uncertainties, which, however, contributes to increasing the computational costs of these workflows. Closed-loop reservoir management (CLRM) is a combination of life-cycle optimization and computer-assisted history matching, both accounting for uncertainties and demanding a significant amount of simulations. For this reason, the application of the CLRM framework in combination with UQ can be extremely computationally expensive. Workflows to assess the value of information (VOI) in CLRM during the field development planning (FDP) phase require even more simulations, which, at

the current level of hardware development, makes real-field applications unfeasible [9]. Therefore, we look for alternatives to reduce this computational cost.

The development of more practical ways of a-priori assessing the value of future measurements has been a topic of several studies recently. Some of these have focused on the use of proxy models to reduce the number of high-fidelity reservoir simulations required for the VOI analysis [15, 21]. Cardoso and Durlofsky [14], He et al. [20], Hewson [22], and Jansen and Durlofsky [29] investigated the use of reduced-order modeling to speed-up production optimization and history matching procedures. Others have proposed a more approximate definition of VOI which simplifies their procedure (Le and Reynolds [34-35]). Eidsvik et al. [18] have envisaged more sophisticated design of experiments to be a promising alternative to alleviate the computational costs of VOI assessment workflows. Recently, Shirangi and Durlofsky [48] presented a general framework that uses clustering techniques to determine representative models to accelerate computations for optimization under uncertainty. Insuasty et al. [24] also showed how clustering methods based on flow-relevant dissimilarity measures can be used to form reduced ensembles. Liu and Forouzanfar [38] showed the importance of flow-based dissimilarity measures for clustering geological realizations and forming reduced ensembles to optimize production in fractured reservoirs. This paper explores the use of clustering techniques to select subsets of representative model realizations to speed-up production optimization and other computational procedures present in the workflow for VOI assessment proposed by Barros et al. [9].

In the Background section we briefly recap our previously proposed methodology for VOI assessment in CLRM and review some previous work on cluster analysis. Next, in the Methodology section, we identify opportunities to apply clustering within the original procedure and we describe our approach to reduce the computational costs in different steps of the VOI assessment to come up with a more practical workflow. Thereafter, in the Examples section, we illustrate the application of the proposed measures to accelerate VOI calculations and we compare the results with those obtained with the original procedure. This work is a further extension of Barros et al. [10].

2. Background

2.1. Closed-loop reservoir management (CLRM)

Closed-loop reservoir management (CLRM) is a combination of frequent life-cycle production optimization and data assimilation (also known as computer-assisted history matching). Life-cycle optimization aims at maximizing a financial measure, typically net present value (NPV), over the producing life of the reservoir by optimizing the production strategy. This may involve well location optimization, or, in a more restricted setting, optimization of well rates and pressures for a given

configuration of wells, on the basis of one or more numerical reservoir models. Data assimilation involves modifying the parameters of one or more reservoir models, or the underlying geological models, with the aim to improve their predictive capacity, using measured data from a potentially wide variety of sources such as production data or time-lapse seismic. For further information on CLRM see, e.g., [16, 23, 26-28, 40, 44, 53].

2.2. VOI assessment in CLRM

Recently, we proposed a new methodology to assess the VOI of future measurements within the CLRM framework [9]. Our approach consists of "closing the loop" in the design phase to simulate how future information, to be obtained during the producing life-time of the reservoir, comes into play in the context of optimal reservoir management. By considering both data assimilation and optimization in the procedure, we are able "to not only quantify how information changes knowledge, but also how it influences the results of decision making" [9]. This is possible because a new production strategy is obtained every time the models are updated with new information, and the strategies with and without additional information can be compared in terms of the value of the optimization objective function (typically NPV) obtained when applying these strategies to a virtual asset (a synthetic truth).

One of the key aspects of this methodology is the idea of using an ensemble of N_{truth} "plausible truths" to account for the fact that in reality we do not know the true reservoir nor the outcome of the future measurements whose value we would like to assess. This requires extensive use of robust optimization and history matching procedures: for N_{truth} plausible truths we have N_{truth} robust optimizations under prior uncertainty, N_{truth} history matches to assimilate the future measurements and N_{truth} robust optimization given the posterior uncertainty. Note that in [9] we also propose an accelerated procedure where the number of prior robust optimizations is reduced from N_{truth} to 1. We consider this accelerated form of the workflow in this paper.

The VOI assessment considered here can be explained with the decision tree depicted in Figure 1. There are two types of decision here: the first one relates to the possibility of acquiring additional measurements, and the second one concerns the actions (i.e., production strategies) that have to be defined to operate the asset. The main premises in CLRM are that: (I) this second type of decisions (typically a sequence of actions) is done in an optimized way by taking into account the geological uncertainty (i.e., robust optimization; see section 2.3), and (II) any additional information can be assimilated to improve our model predictions (i.e., data assimilation; see section 2.4) and in turn allow us to update our actions (again through robust optimization) to improve performance. (These points (I) and (II) are indicated in the decision tree.) The decision tree in Figure 1 also shows how

the multiple plausible truths are used: depending on which plausible truth is considered, the outcome of the additional measurements is different, which results in a different optimized updated strategy to be implemented. We can then compute the improvement in performance ($\Delta J = J_{post} - J_{prior}$) for each plausible truth, which allows us to assess the (expected) VOI and its statistics. Note that in classical decision theory literature the term VOI is often used to refer to the expected VOI, but here our formulation with multiple plausible truths provides a framework where VOI is treated as a random variable.

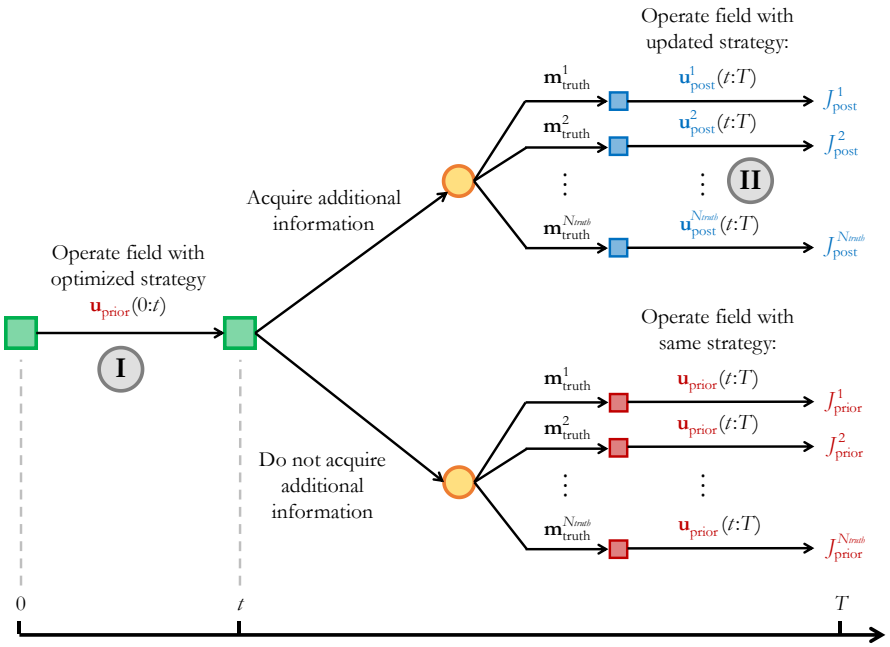

Figure 1: Decision tree for VOI assessment in closed-loop reservoir management. The production strategies, represented by vector **u,** consist of sequential operational decisions to be made. The option of acquiring additional information during the reservoir life-time (here at time *t*) provides an opportunity to update the strategy for the remaining time. The determination of production strategies $\mathbf{u}_{prior}$ and $\mathbf{u}_{post}$ is achieved with the closed-loop framework, by combining optimization and data assimilation (schematically indicated by steps I and II). Because the outcome of future measurements at time *t* is unknown a priori, multiple plausible truth scenarios are considered to account for this uncertainty and allow us to simulate plausible measurement outcomes.

2.3. Robust optimization

Robust life-cycle optimization uses one or more ensembles of geological realizations (reservoir models) to account for uncertainties and determine the production strategy that maximizes a given objective function over the ensemble $\mathbf{M} = \{\mathbf{m}_1, \mathbf{m}_2, \ldots, \mathbf{m}_N\}$; see, e.g., [50, 52]. Typically, the objective function optimized is the net present value (NPV). J_{NPV} is defined as

$$J_{NPV} = \mu_{NPV} = \frac{1}{N}\sum_{i=1}^{N} J_i \,, \tag{1}$$

where μ_{NPV} is the ensemble mean of the objective function values J_i of the individual realizations. The objective function J_i for a single realization *i* is defined as

$$J_i = \int_{t=0}^{T} \frac{q_o(t,\mathbf{m}_i)\, r_o - q_{wp}(t,\mathbf{m}_i)\, r_{wp} - q_{wi}(t,\mathbf{m}_i)\, r_{wi}}{(1+b)^{t/\tau}}\, dt\,, \tag{2}$$

where t is time, T is the producing life of the reservoir, q_o is the oil production rate, q_{wp} is the water production rate, q_{wi} is the water injection rate, r_o is the price of oil produced, r_{wp} is the cost of water produced, r_{wi} is the cost of water injected, b is the discount factor expressed as a fraction per year, and τ is the reference time for discounting (typically one year). The outcome of the optimization procedure is a vector **u** containing the settings of the control variables over the producing life of the reservoir. Typical elements of **u** are monthly or quarterly settings of well head pressures, water injection rates, valve openings etc. Note that, although the optimization is based on N models, only a single strategy **u** is obtained, because only one strategy can be implemented in reality. Note also that, despite being very disseminated among CLRM practitioners, the robust optimization approach presented in [50] is only one way of dealing with uncertainty in production optimization. An alternative approach is to balance risk and return within the optimization by including well-defined risk measures or other utility functions in the objective function; see, e.g., [13, 49].

2.4. Data assimilation

Efficient data assimilation algorithms are the second essential element of CLRM. Many methods for reservoir-focused data assimilation have been developed over the past years, and we refer to [2, 19, 41-42] for overviews. An essential component of data assimilation is accounting for uncertainties, and it is generally accepted that this is best done in a Bayesian framework:

$$p(\mathbf{m}\,|\,\mathbf{d}) = \frac{p(\mathbf{d}\,|\,\mathbf{m})\, p(\mathbf{m})}{p(\mathbf{d})}\,, \tag{3}$$

where p indicates the probability density, and **d** is a vector of measured data (e.g. oil and water flow rates or saturation estimates from time-lapse seismic). In equation (3) the terms $p(\mathbf{m})$ and $p(\mathbf{m}\,|\,\mathbf{d})$ represent the prior and posterior probabilities of the model parameters **m**, which are, in our setting, represented by prior and posterior ensembles respectively. The underlying assumption in data assimilation is that a reduced uncertainty in the model parameters leads to an improved predictive capacity of the models, which, in turn, leads to improved decisions. In our CLRM setting, decisions take the form of control vectors **u**, aimed at maximizing the objective function J.

2.5. Model selection

We use multiple ensembles of realizations to account for geological uncertainties. Typical ensembles are formed by tens or hundreds of realizations, making the procedures involved computationally intensive. The cost, in terms of the amount of simulations required, of robust optimization and

history matching algorithms tends to scale linearly with the size of the ensemble (i.e., O(N)), while the VOI assessment workflow described above scales with the square of the ensemble (i.e., O($N \cdot N_{truth}$) ≈ O(N^2)). Thus, a decrease in the number of realizations considered in the analysis may lead to significant reduction in the computational cost and make the VOI assessment problem more tractable. The challenge is how to cleverly select a subset of realizations which can represent the full ensemble to quantify the uncertainty. Others have worked on this problem; e.g., Armstrong et al. [5] use stochastic programming with recourse to reduce the number of scenarios to be considered and Sarma et al. [45] recommend the use of a minimax approach to efficiently select representative models from a large ensemble by matching target percentiles. Lu et al. [39] proposed an adaptive algorithm for robust optimization using approximate gradients which uses representative subsets of model realizations to avoid evaluating the objective function of the full set of models at every iteration. This work focusses on the use of clustering techniques to automate the selection of representative model realizations, along the lines of the works of Shirangi and Durlofsky [48] and Liu and Forouzanfar [38].

2.5.1. Clustering

Cluster analysis aims to group a set of N objects

$$\mathbf{\Theta} = \begin{bmatrix} \boldsymbol{\theta}_1 & \boldsymbol{\theta}_2 & \cdots & \boldsymbol{\theta}_N \end{bmatrix} = \begin{bmatrix} \theta_{11} & \theta_{12} & \cdots & \theta_{1N} \\ \theta_{21} & \theta_{22} & \cdots & \theta_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \theta_{M1} & \theta_{M2} & \cdots & \theta_{MN} \end{bmatrix}, \tag{4}$$

into N_{repr} clusters according to the similarity between the objects; see, e.g., [7]. Note that here the objects have been chosen as vectors $\boldsymbol{\theta}_i$, $i = 1,2,\ldots N$, in an M-dimensional space (e.g., N realizations of M grid block permeability values) but they could also be scalars, matrices or higher-order objects (tensors). Clustering has been widely used in pattern recognition, machine learning and statistics [6] and is broadly classified into partitional and hierarchical categories. As the name suggests, partitional clustering separates the objects into exclusive clusters such that the objects within a cluster are more similar to each other than to the objects in another cluster. On the other hand, hierarchical clustering, also known as connectivity-based clustering, connects objects to form clusters based on their distance. The connected objects in clusters can then be represented using a dendrogram (i.e., a diagram with a tree structure).

K-means clustering is one of the most used partitional clustering methods due to its simplicity [12]. The user predefines the number N_{repr} of sets C_j, $j = 1,2,\ldots N_{repr}$, that each contain a total of N_j indices corresponding to the objects belonging to each cluster, where the clusters are not necessarily of equal

size. The algorithm then attempts to iteratively improve the partitioning to achieve the lowest intra-cluster distance. This minimization problem can be formulated as follows:

$$C_{opt} = \arg\min_{C} \sum_{j=1}^{N_{repr}} \sum_{k \in C_j} d_{jk}\left(\boldsymbol{\theta}_k, \overline{\boldsymbol{\theta}}_j\right)^2 , \tag{5}$$

where $C = \{C_1, C_2, \ldots, C_{N_{repr}}\}$ is the set of N_{repr} clusters, i.e. a set of sets of indices, and $d_{jk}(\boldsymbol{\theta}_k, \overline{\boldsymbol{\theta}}_j)$ is the distance between one of the N_j data points within each cluster and the cluster centroid $\overline{\boldsymbol{\theta}}_j$ computed as

$$\overline{\boldsymbol{\theta}}_j = \frac{1}{N_j} \sum_{k \in C_j} \boldsymbol{\theta}_k . \tag{6}$$

The first step to use cluster analysis consists of choosing a feature operator F to compare the model realizations $\mathbf{m}_i$. The feature operator could just select a number of parameters (e.g., grid block permeability values) of the vectors of model parameters $\mathbf{m}_i$, or it could represent a more complex operation like a full simulation to compute the NPV or a sequence of saturation snapshots. Using this operator, the set of features $\Theta = [\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \ldots, \boldsymbol{\theta}_N]$ is formed, where $\boldsymbol{\theta}_i = F(\mathbf{m}_i)$. The clustering algorithm can then generate the distances required to determine C_{opt}. Caers [12] and Aydin and Caers [1] propose that uncertainty models should take into account the specific purpose of flow modeling. The selection of representative model realizations that we try to achieve here is nothing but an exercise of modeling uncertainty. Thus, the choice of the appropriate feature operator is also purpose-dependent; there is no one-method-fits-all solution when it comes to cluster model realizations.

Note that the sets $\boldsymbol{\theta}_i$ are elements of an M-dimensional space, where M can be very large. Unfortunately, most clustering algorithms do not work efficiently in higher dimensional spaces because of the inherent sparsity of the data, and as M grows, distance measures become increasingly meaningless ([31, 43]). A solution to this problem is to eliminate some of the dimensions of the feature space. However, if done wrongly, this may cause information loss and introduce wrong correlations between the model realizations. Aggarwal et al. [4] have shown that the projection of high-dimensional data spaces on reduced-order subspaces can lead to improved clustering results.

2.5.2. Projection methods

There is more than one method to project datasets on a reduced-order space. (Note that this is sometimes referred to as reducing the dimensionality of the "feature space".) One of them involves the use of tensor decomposition techniques. Tensor decomposition is strongly related to principal component analysis (PCA) or singular value decomposition (SVD). It enables the transformation of data into a compact representation while honoring their structure (e.g., spatial-temporal correlations)

which is usually degraded with the vectorization step in SVD approaches. These techniques can be used to compress large datasets stored as tensors by constructing low-rank approximations with minimal approximation or reconstruction error. For instance we may form a dataset $\mathbf{\Theta}$ in the form of a tensor representation of the data, $\mathbf{\Theta} = F(\mathbf{m}_1, \mathbf{m}_{2,\ldots}\ \mathbf{m}_N)$, which better preserves their structure than using a vector representation. E.g., we could construct a 3D tensor $\mathbf{\Theta}$ by stacking up the two-dimensional permeability fields (matrices) of an ensemble of 2D model realizations. We then perform the following minimization [24]:

$$\min_{\varphi_{1:I},\psi_{1:J},\chi_{1:K}} \left\| \mathbf{\Theta} - \sum_{i=1}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K} \sigma_{ijk}\left(\boldsymbol{\varphi}_i \otimes \boldsymbol{\psi}_j \otimes \boldsymbol{\chi}_k\right) \right\|_{\mathcal{F}} \tag{7}$$
$$\text{subject to } \boldsymbol{\varphi}_{i'}^T\boldsymbol{\varphi}_{i''} = \delta_{i'i''}, \boldsymbol{\psi}_{j'}^T\boldsymbol{\psi}_{j''} = \delta_{j'j''}, \boldsymbol{\chi}_{k'}^T\boldsymbol{\chi}_{k''} = \delta_{k'k''},$$

where σ_{ijk} is the core tensor (i.e., an all-orthogonal and ordered tensor which is analogous to the coefficients matrix in classical SVD), $\boldsymbol{\varphi}_i$, $\boldsymbol{\psi}_j$, $\boldsymbol{\chi}_k$ are orthonormal basis functions, $\|.\|_{\mathcal{F}}$ represents the Frobenius norm, and the symbol $\otimes$ denotes the tensor (outer) product over a vector space. This can be extended to tensors with more dimensions. For more information on tensor-based model-order reduction, we refer to Insuasty et al. [24] and Afra and Gildin [3], who also show that the solution to the minimization problem in (7) can be approximated by performing a higher-order SVD (HOSVD). In this case, the tensor is flattened (or unfolded) in a planar matrix structure where we can operate similarly to classical SVD. This allows us to determine the basis functions and the coefficients associated with them. Like in classical SVD, a truncation can then be applied to retain only those basis functions that explain the most dominant patterns in the data, thus resulting in the lower-dimensional representation we were aiming for. One of the modes of $\mathbf{\Theta}$ in our applications (here we assume mode k) typically refers to the model uncertainty, characterized by the N model realizations. We can apply a truncated SVD to the unfolded form of $\mathbf{\Theta}$ in this mode $(\mathbf{\Theta}_{(k)}\mathbf{\Theta}_{(k)}^T = \mathbf{U}_k\mathbf{\Sigma}_k\mathbf{V}_k^T)$ and use the obtained coefficients ($\mathbf{U}_k$) to derive the dissimilarity measure to cluster the realizations. Insuasty et al. [24] showed that this approach allows us to compare model realizations based on very rich datasets, such as the temporal evolution of the spatial distribution of pressures and saturations inside the reservoir. They were able to select a subset of realizations representative in terms of dynamic flow patterns and form reduced ensembles to perform robust production optimization more efficiently [24]. Insuasty [25] reported a comparison between the performance of a conventional SVD approach and a tensor decomposition for the compression of large spatial data and showed the benefits of using the latter. Liu and Forouzanfar [38] use saturation maps and well production data to cluster realizations of naturally fracture reservoirs: in this case the number of model parameters (i.e., for the discrete fracture network (DFN) models) may be different among the different realizations, making the parameter-based clustering less obvious. They also use SVD-type of projections to reduce the dimensionality of the feature space.

Another tool to represent model realizations $\mathbf{m}_i$ in a lower dimension is multidimensional scaling (MDS). It refers to techniques that use distance measures to produce a $\tilde{\boldsymbol{\theta}}_i$ representation of data points $\boldsymbol{\theta}_i$ in a reduced M_{MDS}-dimensional space with $M_{MDS} \ll M$. MDS was first introduced for the analysis of proximities in [47]. In recent years, the machine learning community has applied MDS for nonlinear dimension reduction. Kruskal [33] argues that MDS can be complementary to clustering techniques. Scheidt and Caers [46] introduced MDS in the reservoir simulation community, and since then many reservoir applications have been documented; see, e.g., Caers [12].

In this work, we use the Euclidean distance to derive the dissimilarities between the objects and, instead of following a classical MDS approach, which would be equivalent to PCA, we use non-metric MDS [36]. When applying non-metric MDS, a measure of fit, referred to as "stress" in the MDS literature, can be calculated to quantify the conformance of the representation $\tilde{\boldsymbol{\Theta}}$ to the original data $\boldsymbol{\Theta}$, which can be a criterion to define the number M_{MDS} of dimensions of the reduced space; see [32]. Low values of stress (i.e., below 5 %) indicate an excellent fit between dissimilarities and distances, and thus a good representation of the original samples in the reduced-dimensional space. We can then use this stress to determine the appropriate number of dimensions M_{MDS} to proceed with: we start with a small number of dimensions and increase this number until the stress value drops below an acceptable level. For further information on dimensionality reduction, including MDS and PCA approaches, we refer to [17] and [36]. Once we obtain the representation $\tilde{\boldsymbol{\Theta}}$, we perform clustering (e.g., also with the k-means method) based on the coordinates of the data points in the reduced-dimensional space.

3. Methodology

As discussed in the previous section, the VOI assessment workflow presented in [9] requires an excessive amount of simulations to be applied in practice. This is mainly due to the extensive use of robust optimization and history matching and to the fact that multiple plausible truths are considered. The most demanding steps constitute opportunities for considerable acceleration of the workflow. The focus of this work is on the use of model selection to achieve this goal. Thus, it is about looking for approximated results by compromising the rigor in UQ for the sake of computational speed-up.

In this section, we describe first how to select representative models to speed-up robust optimization and history matching, and then how to assess the quality of the results with the accelerated procedures. Afterwards, we explain how we can accelerate the VOI analysis by picking representative plausible truths. We also discuss the choice of the most appropriate feature to distinguish model

realizations in the different parts of the workflow. Finally, we combine all these measures to come up with a new and faster VOI assessment workflow.

3.1. Speeding-up robust optimization

The whole idea behind accelerating robust optimization is to reduce the number of reservoir simulations required. This is done by reducing the number of model realizations in the ensemble used in the optimization.

We start with a full ensemble $\mathbf{M}_{full}$ of N model realizations. The first step is deciding the number N_{repr} of representative realizations to form the reduced ensemble $\mathbf{M}_{repr}$. This number should reflect the speed-up factor we would like to achieve or the maximum ensemble size we can afford to use with the available computational resources.

The second step is choosing a feature F relevant to the problem to be optimized and building our dataset to apply the clustering algorithms. In our case, we are using reservoir simulation to evaluate the objective function in the water flooding optimization process. Therefore, it seems to be important to distinguish the model realizations on the basis of their simulated dynamical behavior. One option is to rely on the fact that model parameters (e.g., permeabilities and porosities) tend to correlate with the reservoir flow characteristics and use them as the feature to distinguish the realizations $\mathbf{m}_i$. An alternative is to perform reservoir simulations and work with features associated with the dynamics of the system, in which case we may consider model states or flow patterns (e.g., pressure/saturation snapshots and streamlines), or model outputs (e.g., well production data and NPV evolution over time), as it has been studied by Shirangi and Durlofsky [48], Insuasty et al. [24] and Liu and Forouzanfar [38].

The next step is preparing the feature data set for clustering, by applying projection methods to reduce the dimensionality of the feature space if necessary. Thereafter, k-means clustering is performed. Once the N_{repr} clusters are formed, one realization is selected as the representative of each cluster, forming the reduced ensemble $\mathbf{M}_{repr}$ of N_{repr} representative realizations. A common choice is to pick the closest realization of the cluster centroid as the representative of that cluster. Note that the derived clusters may have different sizes, as the clustering algorithms are not constrained to form groups containing the same number of realizations. Based on this observation, different weights proportional to the cluster sizes can be assigned to the representative realizations to reflect the number of realizations in their respective clusters. In this case, averages and other statistics of the reduced ensemble $\mathbf{M}_{repr}$ are computed with weights, unlike what is done for the full ensemble $\mathbf{M}_{full}$ where all the N realizations are usually assumed to be equiprobable and therefore equal weights. Finally, robust optimization is performed over $\mathbf{M}_{repr}$ resulting in an optimized strategy $\mathbf{u}_{repr}^{opt}$.

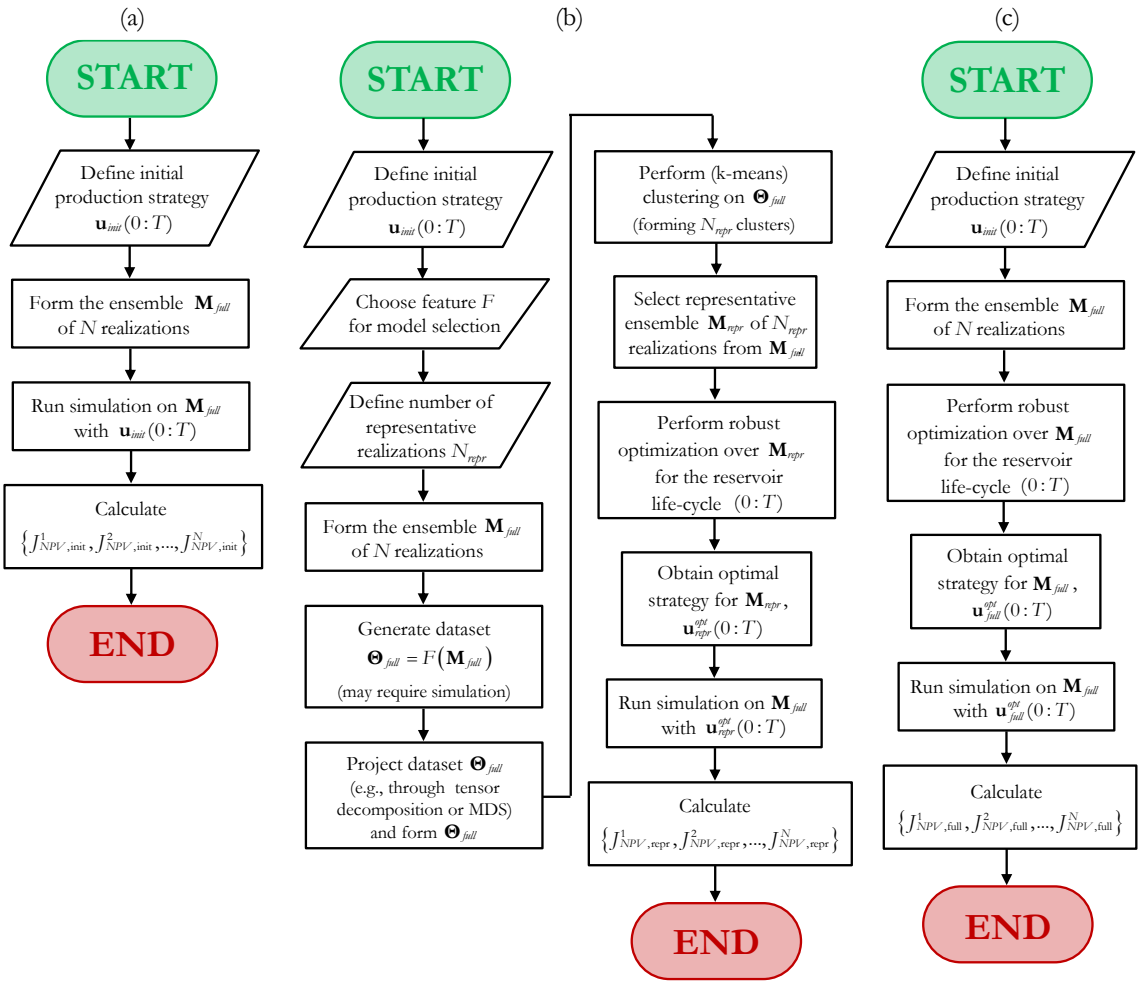

Figure 2: Workflow to evaluate the use of representative realizations for efficient robust optimization. (a) Computation of objective function values for unoptimized production, (b) robust optimization using representative realizations and (c) robust optimization using the full ensemble (reference).

In order to assess the ability of reduced ensembles to reproduce the results obtained by performing the optimization over the full ensemble $\mathbf{M}_{full}$, we compare the performance of the derived production strategies over $\mathbf{M}_{full}$ (Figure 2). Note that this is done only for validation purposes; once we start using the approach to accelerate the optimization, we proceed only with the steps in Figure 2 (b) and rely on optimization over the reduced ensemble $\mathbf{M}_{repr}$.

3.2. Speeding-up history matching

History matching procedures can also be accelerated by considering a reduced ensemble of realizations. Figure 3 illustrates how to compare the performance of representative ensemble to full ensemble based on the history-matched models. The first step is to create the representative ensemble using the techniques described in the *Model selection* section. This step is the same as what is done to form the reduced ensemble for robust optimization. Next, history matching is performed on both full and representative ensembles, resulting in two different posterior ensembles. We then

perform simulations on the realizations of the posterior ensembles to generate data for comparison and validation of the representative ensemble. Note that, when applying representative model selection to make the history matching procedure more efficient, only the workflow in Figure 3 (right) is carried out.

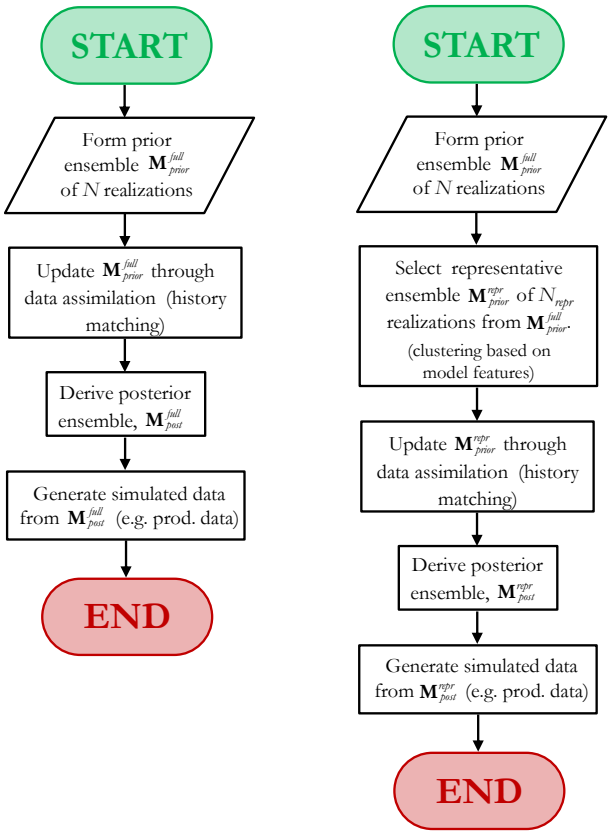

Figure 3: Workflow to evaluate the use of representative realizations for efficient history matching: procedure over the full ensemble (reference for validation) (left) and over the reduced ensemble (right).

Note also that, in this section, we discuss the selection of representative model realizations only as a principle to make history matching computationally more efficient. This is a mechanism that will fit in the VOI workflow as it will be described in a later section, where important choices not specified in Figure 3 will be determined according to the application considered (e.g., whether to assimilate data measured at a single time or during a time interval; which production strategy **u** is used to generate the data). It is also important to highlight that there is a variety of data assimilation methods and that this approach of deriving a reduced ensemble is not always suitable. Ensemble-based methods (e.g., EnKF) rely on a sufficiently large number of samples to derive the directions to update the ensemble and will not work well if N_{repr} is too small. As we will see later, for the examples in this work, we skip this ensemble reduction whenever the EnKF method is used.

3.3. Representative plausible truths

The selection of fewer plausible truths for the VOI analysis can help reducing the computational cost of the workflow at a different level than the acceleration of robust optimization and history matching. The plausible truths are model realizations which we pick to play the role of truth in the CLRM framework. Thus, the goal remains the same: to select representative model realizations.

The challenge is to find relevant features to distinguish these realizations considering their role in the workflow. Although we are still interested in the reservoir management problem (i.e., the water flooding process in our case), the plausible truths are not directly involved in the optimization procedure; we perform the optimizations on the realizations of the prior and posterior ensembles. Due to this difference in roles, the features which are relevant to select representative realizations for the robust optimization and history matching may be not the most appropriate to distinguish plausible truths. As we mentioned before, literature suggests there is no one-method-fits-all solution for choosing the selection features and, therefore, we look for fit-for-purpose solutions.

The methodology for VOI assessment presented in [9] and briefly summarized in the *Background* section accounts for the decision making process, which in CLRM takes the form of optimized production strategies. When we perform the VOI assessment following the workflow presented in [9] (considering N_{truth} plausible truths), we obtain the solution schematically depicted in Figure 4. There is a different production strategy $\mathbf{u}^i_{post}$ corresponding to each plausible truth $\mathbf{m}^i_{truth}$. Also, each plausible truth has its own pair $\{J^i_{NPV,prior}, J^i_{NPV,post}\}$ that is directly related to the VOI calculation. These (possibly unique) associations unveil a mechanism to distinguish plausible truths according to the decisions made (or their consequences) in each scenario. This suggests that we should look for features that carry a similar structure, attributing a different production strategy (or model input) to each plausible truth. We refer to these features as decision-based features.

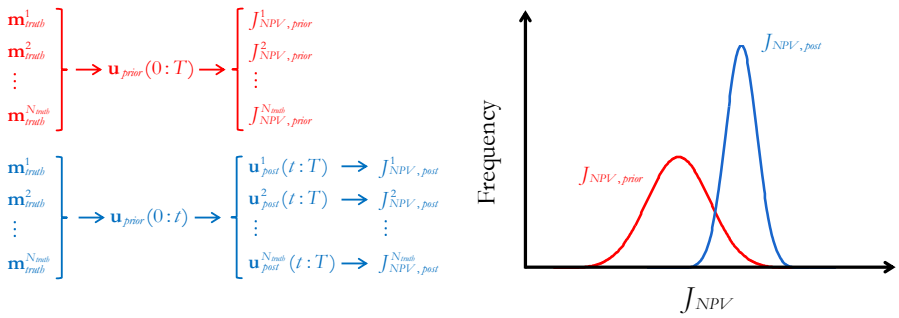

Figure 4: Typical solution obtained with the VOI workflow presented in [9]: a different decision for each plausible truth (left) resulting in improved performance. The NPV plots show prior (red) and posterior (blue) distributions sampled by the plausible truths (right).

In previous sections, we discussed possible features to distinguish model realizations for robust optimization and history matching purposes. It is important to highlight that in both applications all

the model realizations are submitted to the same production strategy. In this context, the features used to select representative realizations rely on the fact that we can compare them through their inherent characteristics (i.e., model parameters) or response to a given strategy (i.e., model states and model outputs). From here on, we refer to these features as model-based features. Figure 5 summarizes the main characteristics of both model-based and decision-based features. Model-based features do not account for the fact that the different scenarios imply different decisions, while the decision-based ones do, connecting them to the VOI assessment setting introduced in [9].

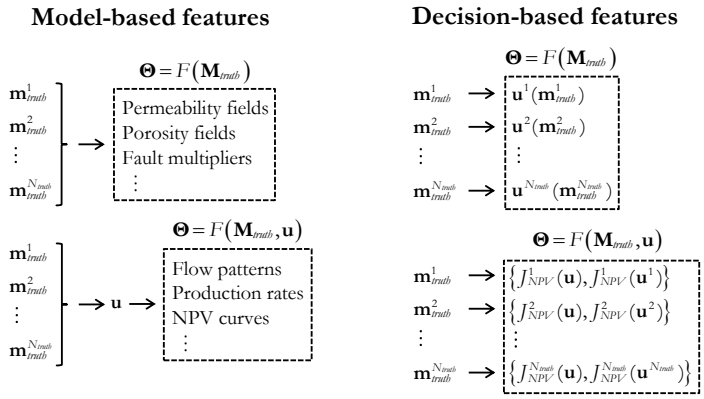

Figure 5: Schematic comparison between model-based and decision-based features. Model-based features rely exclusively on the characteristics of the models and their response to a fixed input, while decision-based features rely on the distinction of models through the decisions associated with each scenario.

The problem becomes then how to determine a set of decisions (or production strategies) $\mathbf{u}^i, i = 1, \ldots, N_{truth}$ that can uniquely identify each plausible truth $\mathbf{m}^i_{truth}$. Very importantly, in addition to that, we seek ways of obtaining these decisions $\mathbf{u}^i(\mathbf{m}^i_{truth})$ in a computationally attractive way so that we can accelerate the VOI workflow when selecting representative plausible truths based on them.

As mentioned before, in the CLRM context the decisions take the form of optimized production strategies. Thus, a good way of identifying characteristics of the optimal configurations related to each of the plausible truths is to perform separate optimizations on them. If we consider an initial ensemble of N_{truth} plausible truths, this means we need to perform N_{truth} optimizations. (Note that the computational cost associated with these N_{truth} optimizations is significantly lower than the cost of the full VOI workflow, which makes this approach suitable for accelerating the VOI assessment.) As a result, we obtain a set of N_{truth} optimal production strategies $\mathbf{u}^i_{opt}$ and N_{truth} optimal objective function values $J^i_{NPV}(\mathbf{u}^i_{opt})$.

Typically, the optimal production strategies $\mathbf{u}^i_{opt}$ tend to be very different from each other because the plausible truths are different model realizations. Following this reasoning, the optimal production strategies seem to be appropriate to support the clustering of plausible truths. One of the potential problems with this approach concerns the limitations of the optimization methods used. With the

chances of reaching local optima, there is a risk that we can never derive truly optimal production strategies. On the other hand, one may argue that the optimization methods used to derive the optimized strategies for clustering are the same as the ones used in the CLRM and VOI workflows considered here. Thus, in principle, even if not truly optimal, these optimized strategies could still be used for the purpose of identifying representative plausible truths. Another problem of this approach is the risk of non-uniqueness of optimal solutions for the production optimization problem, due to the possible presence of redundant degrees of freedom in the high-dimensional space of control variables [50]. This may result in multiple production strategies that are equally optimal, which could put the validity of this approach at stake. A possible way to avoid the redundancy is to perform an a-priori tensor- or SVD-based decomposition of a large set of possible controls and perform the optimization in a reduced control space. Alternatively, one could impose temporal and/or spatial correlations on the controls which also reduces the degrees of freedom in the control space. A downside of such a-priori measures is that they may lead to lower NPV values. As an alternative, we therefore apply an a-posteriori tensor decomposition of the set of optimal production strategies and retain a fraction of the basis functions by truncation based on their energy. By doing so, we intend to capture only the main trends of the data and reduce the effects of the non-uniqueness of the optimal production strategies, although we note that this a-posteriori decomposition of the controls does not guarantee an improvement of the situation. As a final remark on this issue, we note that for these N_{truth} nominal optimizations would in principle be performed with the same optimization methods that will be used for the VOI and CLRM exercises and that the risk of non-uniqueness of optimal production strategies will also be present there. Ultimately, the selection of representative plausible truths based on these strategies would still be consistent with the VOI workflow, even in its limitations.

On the other hand, the optimal objective function values $J^i_{NPV}(\mathbf{u}^i_{opt})$ typically tend to be close to each other because the different optimal production strategies compensate for the differences in the model realizations (i.e., given the same well locations, the optimal sweep of the reservoir tends to be similar for all the realizations); see, e.g., the schematic NPV distributions in Figure 4. Thus, the dissimilarities between the realizations in terms of NPV are less pronounced for their optimal configurations, and, because of that, they are less suitable to help in the selection of representative plausible truths. However, in combination with the objective function values $J^i_{NPV}(\mathbf{u}_{prior})$ obtained with a robust strategy $\mathbf{u}_{prior}$ (i.e., optimized to maximize the mean objective function given the initial uncertainty), these data reveal how much we may benefit if we learn or observe the truth for each one of the plausible truths. We can then distinguish plausible truths according to the gains associated with their optimal configurations, which are directly related to the VOI, and this can be useful for

our purposes. One of the advantages of using these $\{J^i_{NPV}(\mathbf{u}_{prior}), J^i_{NPV}(\mathbf{u}^i_{opt})\}$ data features is avoiding the problem of non-unique optimal solutions discussed in the previous paragraph.

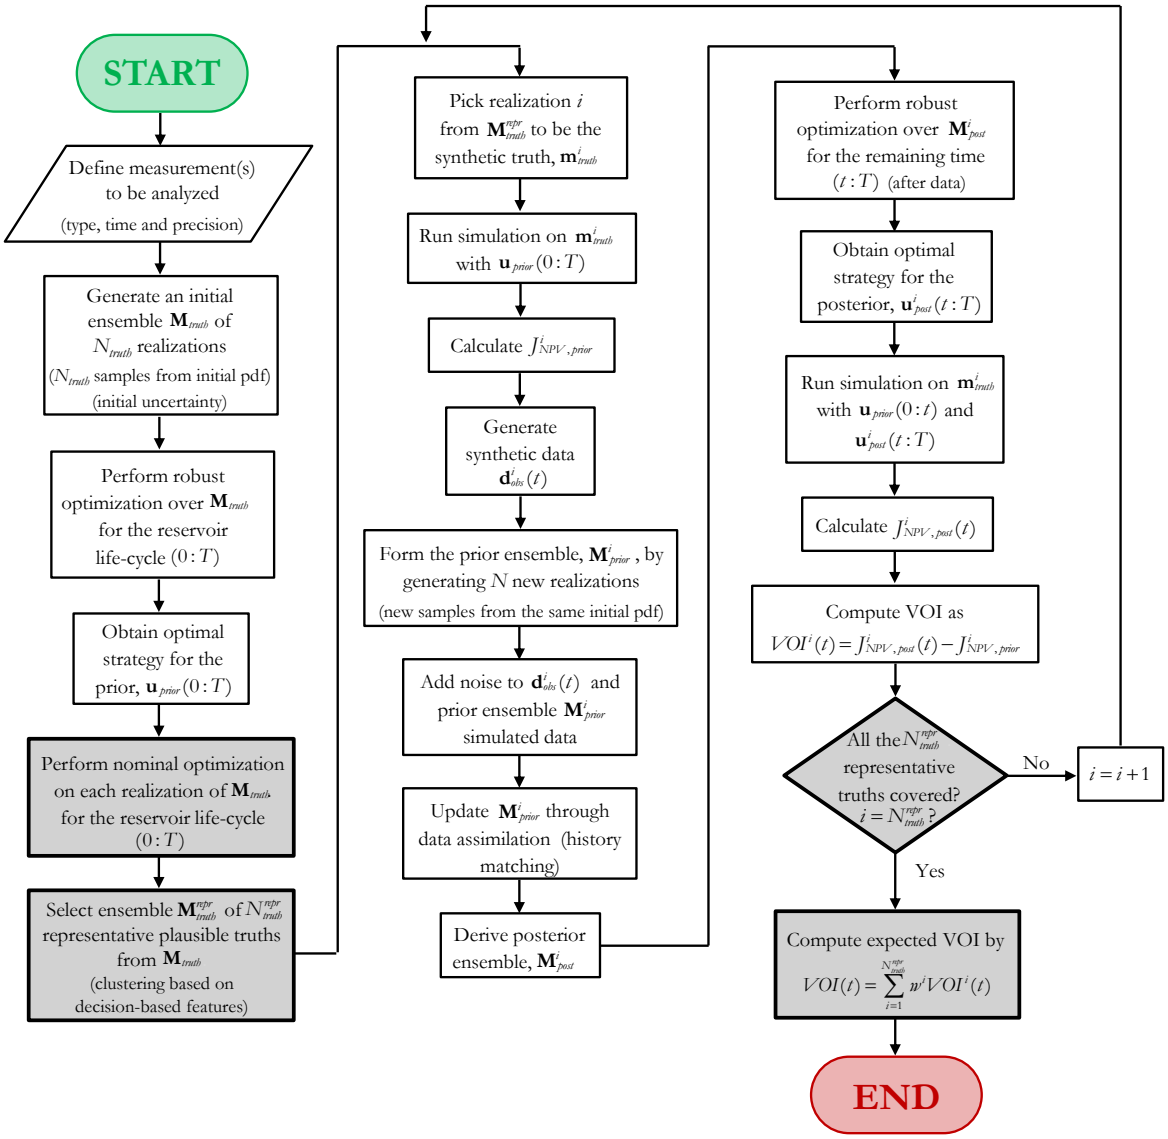

Figure 6: Workflow to compute the expected VOI for a single observation time using representative plausible truths and full ensembles for robust optimization and history matching. The main modifications in the workflow with respect to the original procedure presented in [9] are highlighted in grey bold boxes.

We introduce the selection of representative plausible truths to our original VOI assessment workflow [9] and we obtain the procedure depicted in Figure 6 for cases with a single observation time. The main difference compared to the original workflow is that, before entering the loop where each one of the N_{truth} realizations of the initial ensemble $\mathbf{M}_{truth}$ is picked to be the truth $\mathbf{m}^i_{truth}$, we have a few more pre-processing steps. First, a step where we optimize each one of the realizations and then a step where we perform the clustering to select N^{repr}_{truth} representative plausible truths based on the decision-based features as explained above. Another minor change in the workflow refers to the computation of the statistics of VOI: before, the plausible truths were considered (for simplicity) to be equiprobable, but, now, the selected plausible truths in $\mathbf{M}^{repr}_{truth}$ may have different weights w_i

assigned by our selection procedure (i.e., weights proportional to the number of realizations in each cluster). Note that there is a computational cost associated with the additional N_{truth} nominal optimizations required, but that this extra cost is minor when compared to the cost of the full workflow. Another point to realize is that these N_{truth} nominal optimizations would be performed anyway if we carry out a value of clairvoyance analysis (VOC; see [9]) to determine the upper bound for the VOI assessment.

3.4. Accelerated VOI assessment

We combine the ideas of the three previous sections to the original VOI workflow of Barros et al. [9]. The result is a new workflow for accelerated VOI assessment, shown in Figure 7.

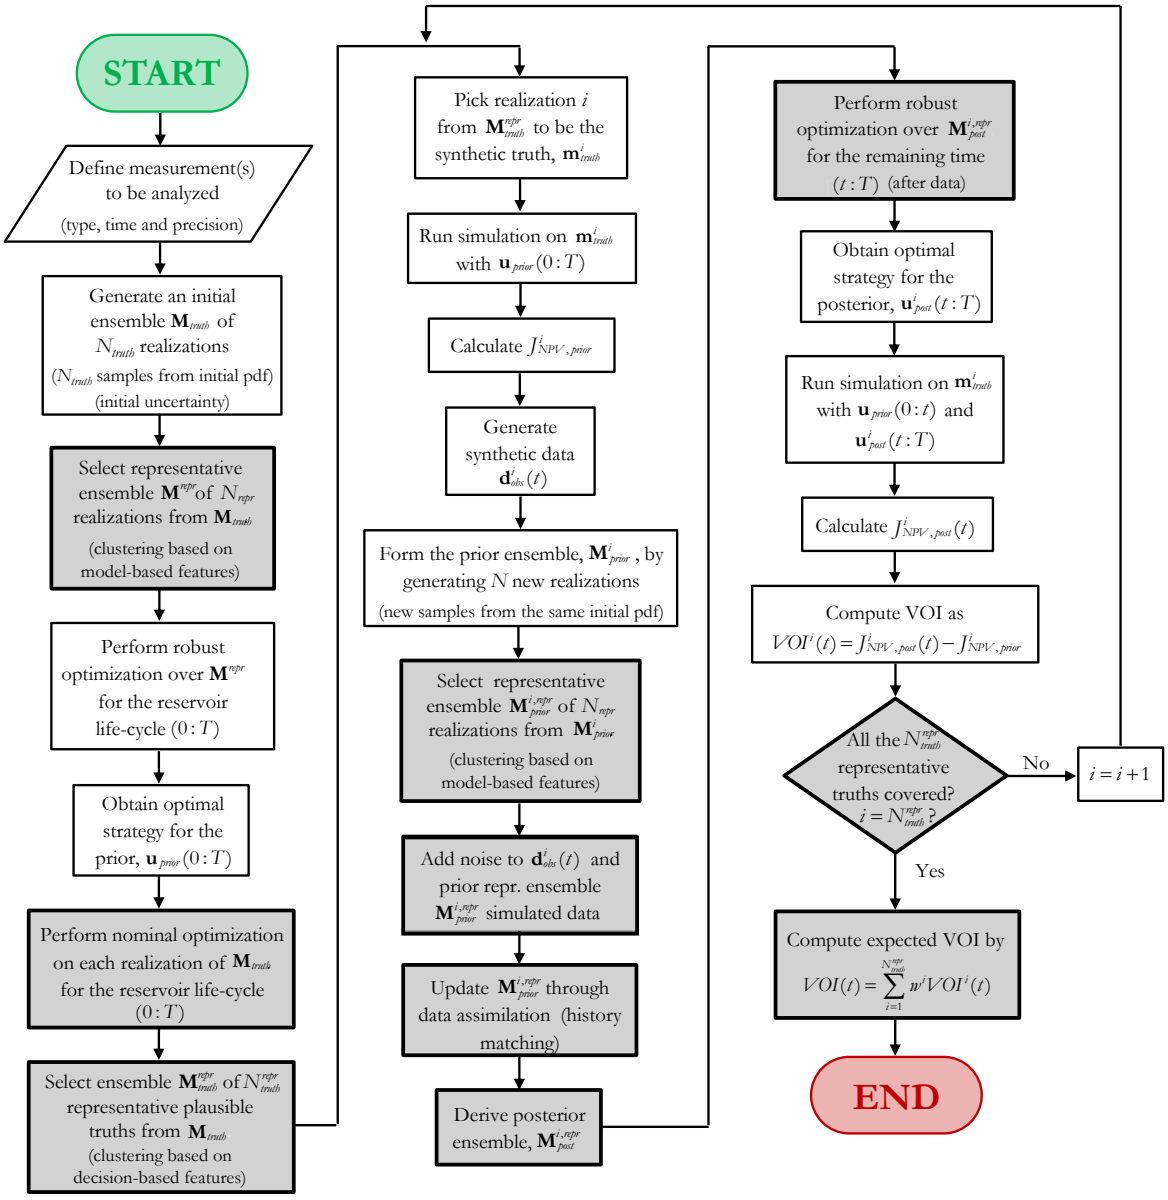

Figure 7: Workflow to compute the expected VOI for a single observation time using representative plausible truths and reduced ensembles of representative model realizations for robust optimization and history matching. The main modifications in the workflow with respect to the original procedure [9] are highlighted in grey bold boxes.

We note that in this flowchart we consider the size of the ensemble of plausible truths and the prior ensembles to be N_{truth} and N respectively. Note that, these ensembles may be chosen to have the same size, but that this is not a necessity. The same holds for the number of representative realizations for robust optimization (and history matching) and the number of representative plausible truths, which here are chosen as N_{repr} and N_{truth}^{repr} respectively.

Given these considerations, we can expect a speed-up factor proportional to $(N_{truth} / N_{truth}^{repr}) \cdot (N / N_{repr})$ by using this accelerated procedure. This means that, if we are able to select reduced ensembles with 10 times fewer realizations, we can reduce the number of required reservoir simulations by a factor of 100.

4. Examples

4.1. Robust optimization with reduced ensembles

4.1.1. 2D five-spot model

As a first step to illustrate our approach, we used the same small 2D five-spot example (441 grid blocks) from our previous paper; see [9] for the complete description of the example. Originally, we had ensembles of $N = 50$ model realizations to characterize the geological uncertainties. To accelerate the robust optimizations, we considered a reduced number of representative realizations, $N_{repr} = \{3, 5, 10\}$, representing approximately 5 %, 10 % and 20 % of the number of realizations in the full ensembles. We evaluated the performance of different features for clustering: permeability field, oil saturation snapshots at every control time interval (every 150 days) and NPV time-series. We also studied the effect of different projection methods by using both MDS and tensor decomposition to perform the projection of the feature space before clustering. For the non-metric MDS we used the standard implementation available in Matlab ("mdscale" function) with the previously mentioned stress criterion (here 5 %) to determine the dimension of the reduced space. For the tensor decomposition we used the implementation of the HOSVD described in [24], with a cut-off of 95 % in terms of the energy of the eigenvalues to determine the number of basis functions to be retained for the uncertainty dimension. The production strategy used to generate the features for clustering was fixed as the initial production strategy chosen as the starting point of the optimization (here mid in-between bounds: p_{prod} = 250 bar for the producers and p_{inj} = 400 bar for the injector).

By applying the steps as in Figure 2 to compare the performance of the robust optimization with reduced and full ensembles, we obtained results in terms of NPV. Figure 8 depicts the results for the two examples expressed as probability distribution function (pdf) plots. The NPV values for the

unoptimized production (Figure 2 (a)) are shown in grey and the reference results (Figure 2 (c)) in black. The results obtained using the reduced ensembles (Figure 2 (b)) are represented by the colored lines according to the number of representative realizations. We repeated the procedure with several ensembles $\mathbf{M}_{full}$ and we obtained results similar to the ones shown here.

Generally, all the production strategies $\mathbf{u}_{repr}^{opt}$ optimized with the representative ensemble performed very well when compared to the strategy $\mathbf{u}_{full}^{opt}$ for the full ensemble. We notice that the selection with three representative realizations (5 %) from the ensemble performs poorer in most cases, and fails badly in the case depicted in Figure 8 (e). This shows that, for this example, taking an amount of representative realizations corresponding to only 5 % of the original ensemble size is insufficient whereas taking 10 % and 20 % give good results in all cases. Also, we notice that the MDS transformation helps the selection of representative realizations better than the tensor decomposition, especially when using the permeability field as selection feature (Figure 8 (a) and (d)) but also when using the oil saturation snapshots (Figure 8 (b) and (e)). We believe the better performance of MDS here to be related to the small size of the model (441 grid blocks) and the (relatively) smooth character of the geological realizations (see [9]). For larger models and cases with more complex geological features, in which there is more spatial-temporal structure to be preserved, it is expected that tensor decomposition should perform the best [24]. Overall, based on the results from the 2D five-spot model, oil saturation snapshots and NPV time-series seem to be the most suitable selection features based on the results from 2D five-spot model, which is in alignment with the conclusions drawn by Shirangi and Durlofsky [48] after comparing several model-based features within their general framework.

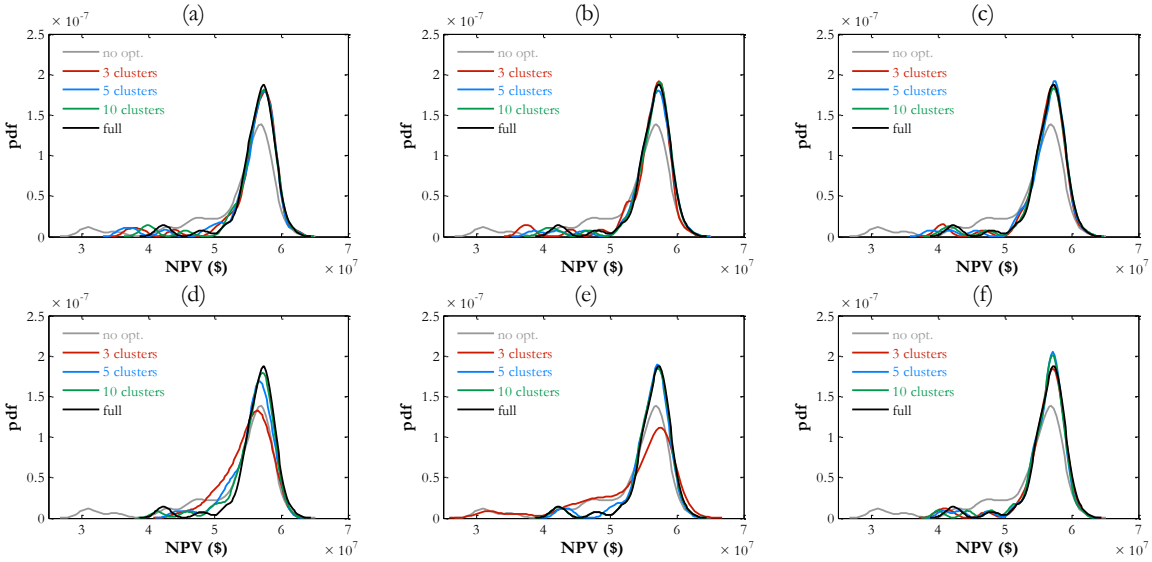

Figure 8: Results of robust optimization over an ensemble of the 2D five-spot model expressed in terms of NPV pdf plots. (a) Permeability as feature and MDS as projection method, (b) oil saturation snapshots as feature and MDS as projection method, (c) NPV time-series as feature and MDS as projection method, (d) permeability as feature and tensor decomposition as projection method, (e) oil saturation snapshots as feature and tensor decomposition as projection method, and (f) NPV time-series as feature and tensor decomposition as projection method.

4.1.2. Egg model

The Egg model is a synthetic reservoir model created to serve as a benchmark for water flooding optimization, computer-assisted history matching and CLRM applications. The model consists of $N = 100$ realizations of a channelized reservoir with $60 \times 60 \times 7$ grid cells. Its 18,553 active cells give it the shape of an egg (Figure 9). The field is produced through water flooding, with four producers and eight injectors in defined locations, and has been used for several studies; see, e.g., [50]. For further information on a standardized version of this model, we refer to [30].

For this case study, we considered the optimization of the water injection rates for a production period of 3,600 days. The rates can be adjusted every 360 days (i.e., $M = 10$ control time intervals) in the range of $0 \leq q_{inj} \leq 79.5$ m³/day, and the maximum injection pressure allowed is $p_{inj,max} = 420$ bar. The bottom-hole pressure of the producers is kept constant at $p_{prod} = 395$ bar. The robust optimization experiments were carried out with the help of the AD-GPRS simulator to obtain the required predictions and gradients [11, 54] for each model realization which were used in a simple implementation of the steepest ascent algorithm.

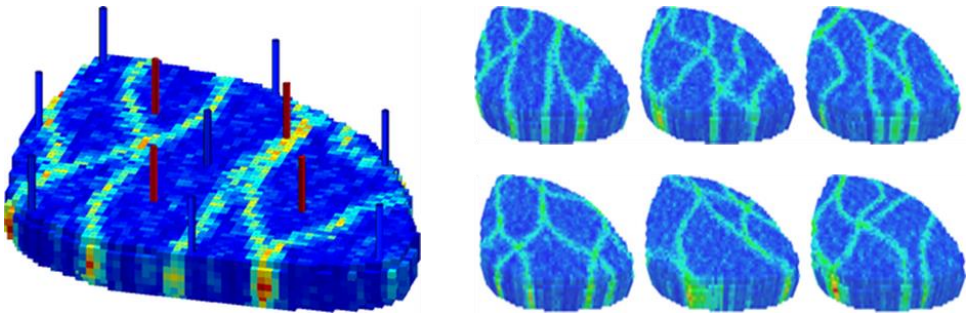

Figure 9: The Egg model: a channelized reservoir with 8 injectors (blue) and 4 producers (red) (left). Six randomly chosen realizations of the permeability distribution (right). (after [30, 50])

In this example, to illustrate the use of representative ensembles, we considered a reduced number of representative realizations, $N_{repr} = \{5, 10, 20\}$, representing again 5 %, 10 % and 20 % of the number of realizations in the full ensemble. Because this is a larger model, we could not afford repeating as many optimization experiments as the ones done for the 2D five-spot model. We chose to do experiments considering: two clustering features only – one model-parameter feature (permeability distributions) and one model-response feature (oil saturation snapshots) – and the two projection methods as used before (MDS and tensor decomposition) – which are expected to behave more different from each other in larger cases.

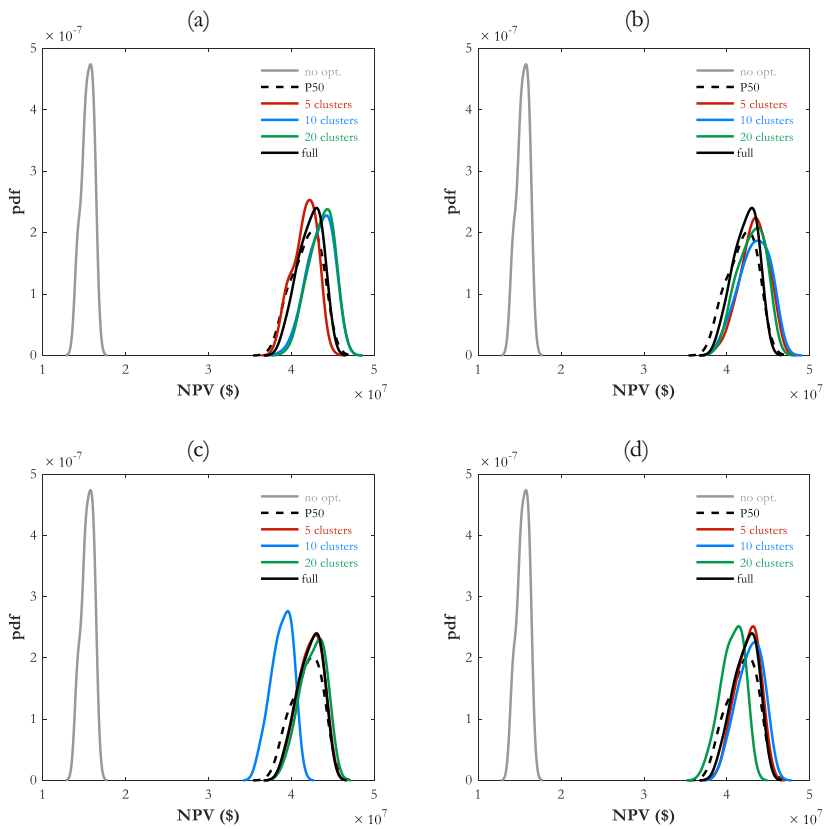

Figure 10: NPV pdf plot of an ensemble from the Egg model. (a) Permeability as feature and MDS as projection method, (b) permeability as feature and tensor decomposition as projection method, (c) oil saturation snapshots as feature and MDS as projection method, and (d) oil saturation snapshots as feature and tensor decomposition as projection method.

Figure 10 depicts the results obtained. The color scheme is the same as in Figure 8, but here we also include the case in which the optimization is performed over the P_{50} model only for comparison. First, we observe that, in this case, robust optimization results in significantly higher NPV values compared to those obtained with the unoptimized strategy. We can also see that the optimizations over the reduced ensembles were able, in most of the cases, to achieve a similar performance compared to the optimization performed over the full ensemble. In some cases, the optimization with the representative realizations managed to outperform the reference results, which reminds us that the optimization techniques used here are not perfect and cannot guarantee globally optimal solutions for the production optimization problem. Moreover, we observe some inconsistent results: in Figure 10 (c), the robust optimization based on $N_{repr} = 10$ selected models performed significantly worse than the one using $N_{repr} = \{5, 20\}$ realizations, and in Figure 10 (d) the optimization based on $N_{repr} = 20$ realizations was the one that performed the worst. We expected to see an increase in the performance of the robust optimization as N_{repr} gets closer to N. Although these unexpected results could be related to limitations of our implementation (e.g., imperfections within the optimization algorithms and simulations), we have not yet been able to find a conclusive explanation for it.

4.2. History matching with reduced ensembles

4.2.1. 2D five-spot model

To illustrate the workflow from Figure 3, we first tested it in a history matching twin experiment on the 2D five-spot example. For this exercise, we used an adjoint-based RML method (see [41-42]) and AD-GPRS to obtain the required gradients to update the uncertain parameters. As the version of AD-GPRS made available to us only provided the gradients with respect to permeability multipliers, this example was slightly modified with respect to the original model used in previous sections: here we assumed the porosity field to be homogeneous and known (ϕ = 0.2), and thus not to be updated throughout the history matching. We considered the availability of measurements of water and oil field production rates at a single observation time, and measurement errors of ε_{prod} = 5 m^3/day. We repeated the history matching for different observation times. The production strategy was fixed at p_{prod} = 250 bar for the producers and p_{inj} = 400 bar for the injector.

First, we performed the history matching over the full ensemble of N = 49 realizations, with one additional realization used to generate the synthetic measurements. Next, we considered only the representative realizations. Like in the previous section, we repeated the procedure for several choices, by varying the number of representative models, the selection feature and the projection method. Figure 11 displays the results obtained with N_{repr} = 10 representative models selected based on oil saturation snapshots projected by tensor decomposition for measurements at t_{data} = 750 days. The results for other clustering settings, which are reported in [51], were comparable to the ones showed here. From these results, we could not make any conclusions about which projection method (MDS or tensor decomposition) performs the best.

Figure 11 (left) shows the simulated forecasts of field production rates for the prior ensemble. The thin dashed lines represent each one of the N = 49 realizations of the full prior ensemble while the thicker dashed lines correspond to the N_{repr} = 10 representative ones. The thick solid lines show the forecast generated with the synthetic truth, and the yellow circles indicate the measurements available. We observe that, although the synthetic truth seems to be captured by the ensemble, there is a large spread in the predictions. Figure 11 (right) shows the predictions a posteriori, after the history matching was performed over the full and reduced prior ensembles. We note that the assimilation of the available measurements contributed to a significant reduction in the spread of the curves, but, more importantly, we observe a reasonably good agreement between the uncertainty characterized by the full and reduced posterior ensembles. We observe a slight reduction in the spread of the curves, which was expected because smaller ensembles tend to underestimate uncertainty in the forecasts. There are measures such as ensemble inflation that could possibly minimize this undesirable effect, but they were not considered in this work. We also emphasize that

different weights have been assigned to the representative realizations forming the reduced ensemble while the realizations of the full ensemble are equiprobable, and for this reason the visual comparison may not be the most appropriate way of assessing the quality of the approximation here. Despite these remarks, the results suggest that representative realizations can be used to make history matching procedures more efficient without compromising the uncertainty quantification to an unacceptable level.

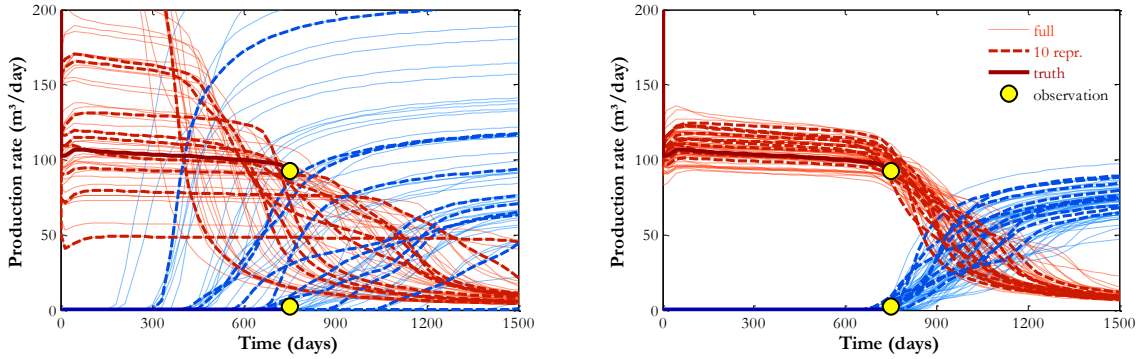

Figure 11: History matching results for the 2D five-spot model. Solid lines represent the prediction from the synthetic truth, the dots correspond to the synthetic data to be matched and the dashed lines represent the predictions of the realizations to be updated. Red lines correspond to oil production and blue lines to water production.

4.2.2. Egg model

Based on the learning from the 2D five-spot example, we repeated the same procedure for the Egg model to confirm that it is possible to make the history matching more efficient also in larger case studies. Once again, we used the RML method and AD-GPRS to obtain the required gradients to update the uncertain permeability field. The settings were the same as the ones described before, but here the injection rates were fixed to q_{inj} = 79.5 m³/day. The observations considered were field production rates and bottom-hole pressures measured in the injectors, available at t_{data} = 1,800 days, with measurement errors of ε_{prod} = 5 m³/day and ε_{BHP} = 10 bar. Here we used the full ensemble with N = 99 realizations, plus one synthetic truth. As before, we considered several choices for the selection of representative realizations, but we display only one set of the results. The results for other clustering settings can be found in [51].

Figure 12 displays the results obtained with N_{repr} = 10 representative models selected based on oil saturation snapshots projected by tensor decomposition. Figure 12 (left) shows the simulated forecasts of field production rates for the prior ensembles, and Figure 12 (right) for the posterior ensembles. The color scheme is the same as in Figure 11. The differences between the prior and posterior predictions of production rates are minor, indicating that the main contribution to the mismatches is probably related to the bottom-hole pressure measurements in the injectors, displayed in Figure 13. The black dashed lines represent the realizations of the full prior ensemble while the

blue lines correspond to the representative ones. The red lines show the forecast generated with the synthetic truth, and the yellow circles indicate the measurements available. Like in the 2D five-spot example, we observe that the reduced ensemble constitutes a good approximation of the full ensemble. As before, we repeated the exercise considering other choices for the selection of representative realizations, but we displayed only one set of the results here. The results for other clustering settings can be found in [51]. From those results, we observed that N_{repr} = 10 seemed to be close to the limit as the minimum number of representative realizations to approximate the uncertainty characterized by the full ensemble. We also noticed that MDS and tensor decomposition lead to similar performance, but that the latter technique resulted in reduced ensembles which provided slightly larger spreads in the production forecasts.

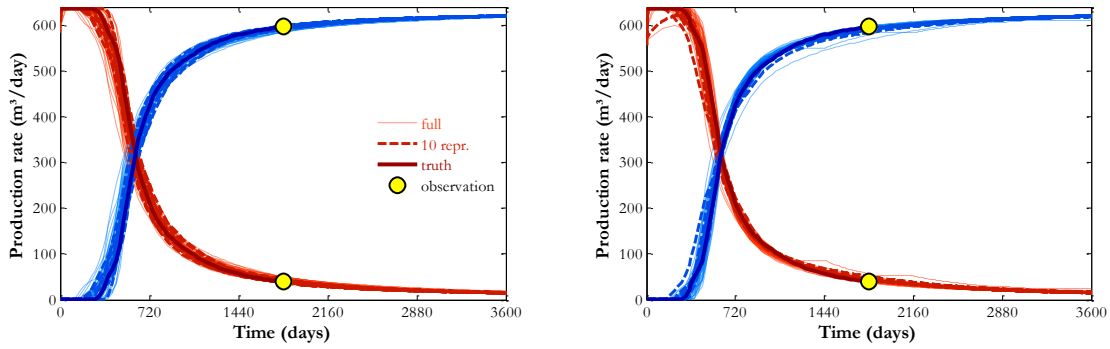

Figure 12: History matching results for the Egg model in terms of the predictions of field production rates: prior (left) and posterior (right). Red lines correspond to oil production and blue lines to water production. Solid lines represent the prediction from the synthetic truth, the yellow dots correspond to the synthetic data to be matched and the dashed lines represent the predictions of the realizations to be updated (full and reduced ensembles; see legend).

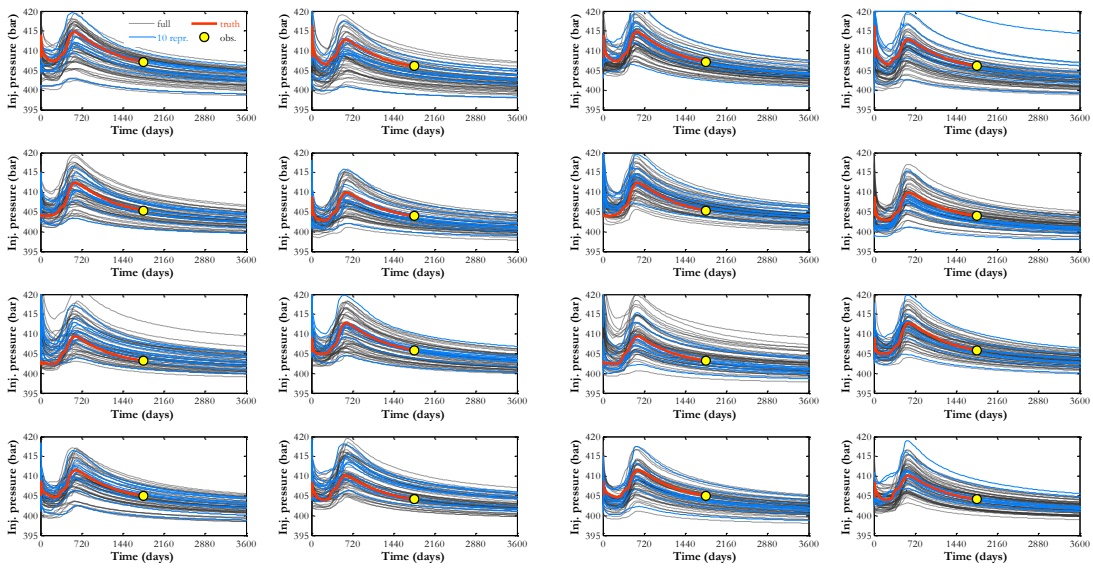

Figure 13: History matching results for the Egg model in terms of the predictions of BHP at the 8 injectors: prior (left) and posterior (right). Solid red lines represent the prediction from the synthetic truth, the yellow dots correspond to the synthetic data to be matched and the dashed lines represent the predictions of the realizations to be updated (full ensemble in black and reduced ensemble in blue; see legend).

4.3. Accelerating VOI assessment

4.3.1. Representative plausible truths

4.3.1.1. 2D five-spot model

To illustrate the selection of representative plausible truths, we applied to the 2D five-spot model the workflow to assess the VOI for a single observation time with a reduced number of plausible truths (Figure 6). The workflow was repeated for different observation times, t_{data} = {150, 300, … , 1,350} days. The history matching step was performed with the EnKF module of MRST [37]. We assessed the VOI of the production data (total flow rates and water-cuts) with absolute measurement errors (ε_{flux} = 5 m³/day and ε_{wct} = 0.1). The VOC and the VOI were computed for each of the nine observation times, and we compared the results against those obtained using the original workflow [9] with the full ensemble $\mathbf{M}_{truth}$ of plausible truths (N_{truth} = 50), which serves as a reference. Note that in this subsection the focus is on the acceleration by reducing the number of plausible truths only, thus the robust optimization and history matching steps are performed using full ensembles.

First, we checked our hypothesis that the model-based features are not suitable for selecting plausible truths. For that, we selected N_{truth}^{repr} = 5 plausible truths through clustering based on: permeability fields, NPV evolution profiles and flow patterns (i.e., pressure and saturation snapshots). No projection methods (e.g., MDS) were used. We also tested a random selection. The results are shown in Figure 14, including the reference results. The different lines represent percentiles and mean of the VOC and VOI distributions as a function of the time when the additional information (or clairvoyance) becomes available. We do not interpret or explain the VOI and VOC results here; we refer to [9] for this purpose. Here we assess the ability to obtain similar results with fewer plausible truths.

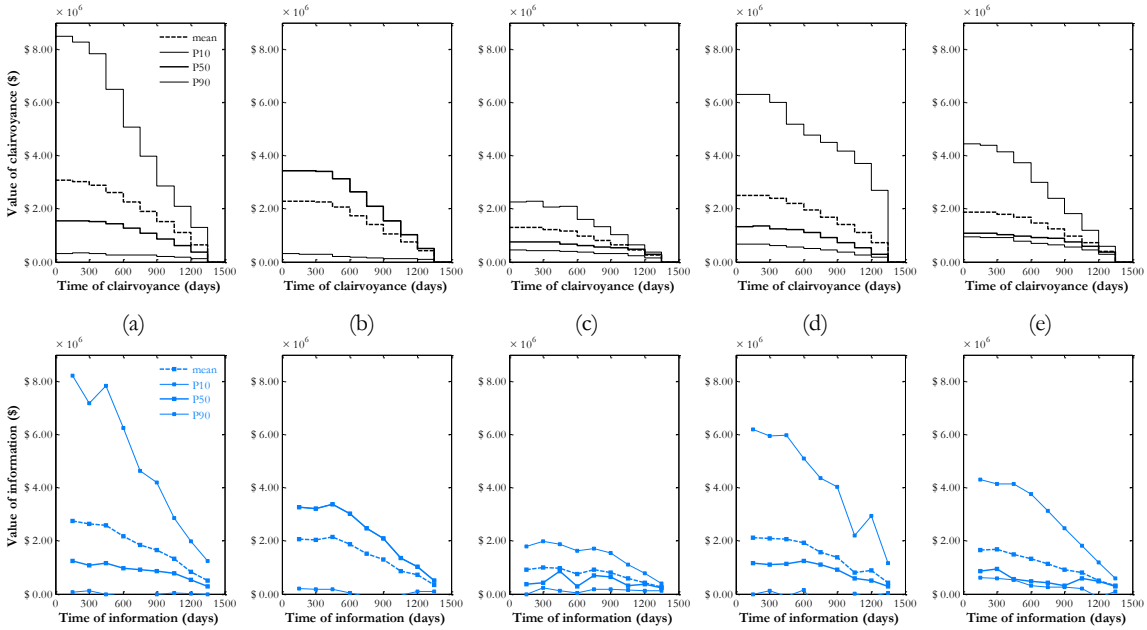

Figure 14: Results of the VOI (and VOC) assessment for the 2D five-spot model using selection of 5 plausible truths based on model-based features. (a) Reference obtained using the original workflow, (b) random selection, (c) selection based on permeability field, (d) selection based on NPV time series and (e) selection based on flow patterns.

Overall, none of the selections in Figure 14 is able to satisfactorily reproduce the reference results, shown in Figure 14 (a). Although a selection based on some feature (Figure 14 (c), (d) and (e)) is clearly better than a random selection (Figure 14 (b)), these results seem to support our hypothesis that the model-based features are not the most appropriate means to select representative plausible truths.

Next, we repeated the same procedure with our proposed decision-based features. This required a nominal optimization on each of the N_{truth} = 50 plausible truths initially considered. Figure 15 shows the data we obtain from these optimizations, which can be used for clustering and model selection. It becomes clear that these features create a space in which we can distinguish the samples and select those instances of $\mathbf{M}_{truth}^{repr}$ that can better represent the entire population $\mathbf{M}_{truth}$. Figure 15 (a) displays each of the plausible truths plotted in a two-dimensional space: the first dimension corresponds to the objective function values (i.e., NPV) before the nominal optimizations and the second to the objective function values after the optimizations. The clusters are shown in different colors and the selected plausible truths are marked as squares. Figure 15 (d) exhibits the optimal production strategy for each plausible truth: the plots show the BHP controls at each of the five wells of the 2D five-spot model; the colors indicate the clusters formed and the thicker lines highlight the representative of each cluster. Figure 15 (a) and (b) display the clusters obtained using the objective-function feature (data displayed in Figure 15 (a)), and Figure 15 (c) and (d) show the clusters derived from the optimal-strategy feature (data displayed in Figure 15 (d)). The first point to note is that, as one would expect, both features lead to different clustering results. Besides that, even though we observe a more

clear separation between the colors in Figure 15 (a) and in Figure 15 (d), the colors in their respective pairs (b) and (c) still show some degree of clustering.

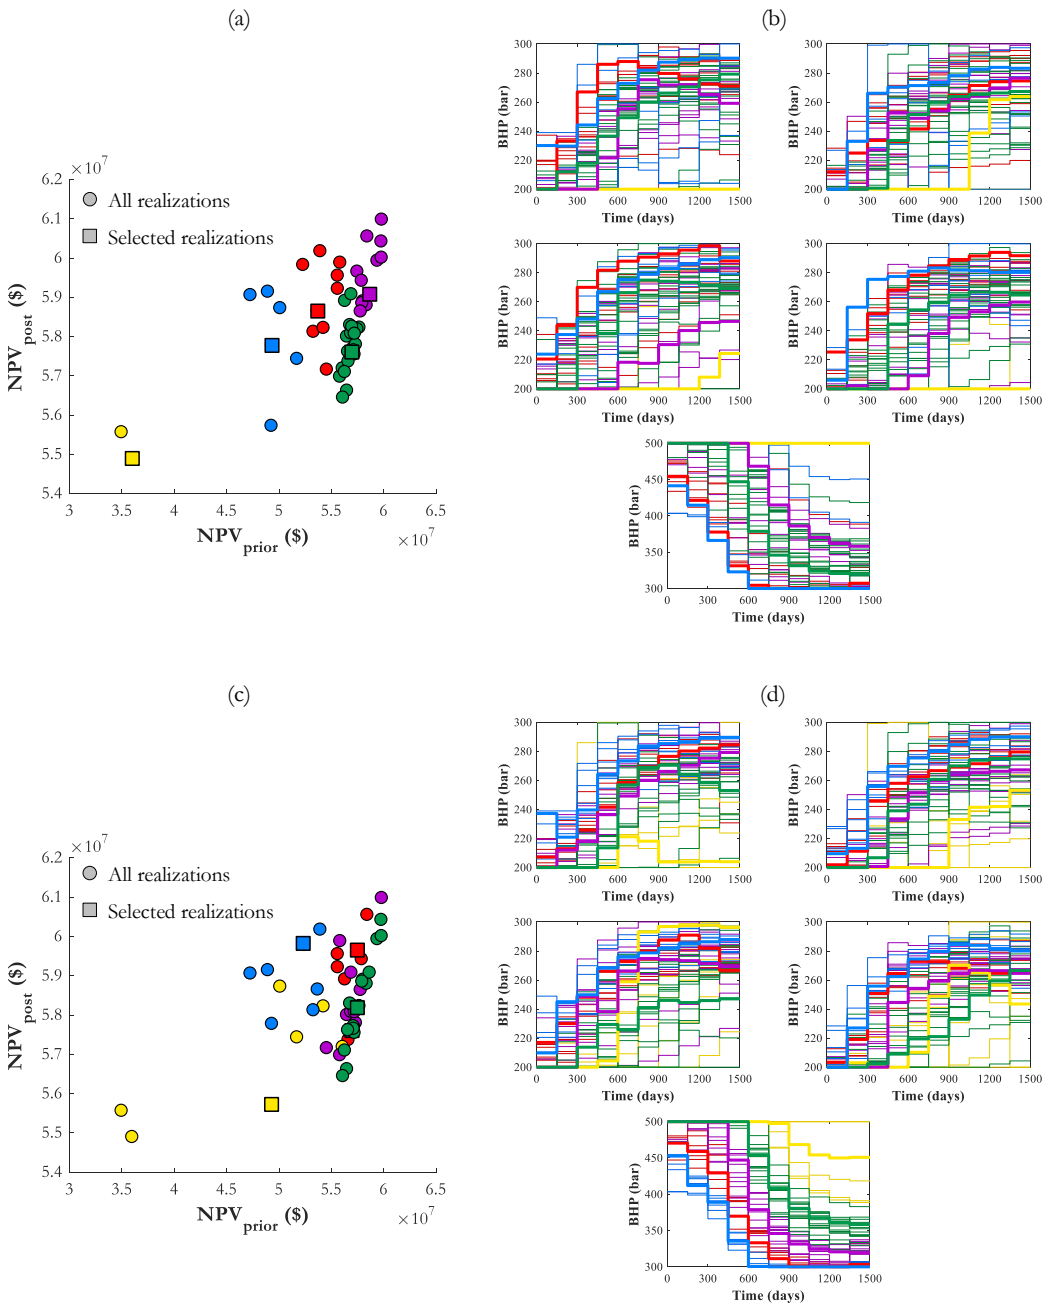

Figure 15: Proposed decision-based features for the selection of representative plausible truths. Objective function values, before and after nominal optimizations ((a) and (c)). Optimal production strategies (bottom-hole pressures at all the wells for every control interval) ((b) and (d)). Selection based on objective function values ((a) and (b)) and selection based on optimal production strategies ((c) and (d)).

Figure 16 depicts the VOI results obtained by picking $N_{truth}^{repr} = 5$ representative plausible truths according to these new features. This time we observe better selections, which are able to repeat the reference results (Figure 16 (a)) almost perfectly. We note that the selection based on the objective

function values (Figure 16 (b)) succeeds in reproducing the reference results for VOC, but less for VOI. In contrast, the selection based on optimal production strategies (Figure 16 (c)) performs fairly well for both VOC and VOI. Therefore, the optimal production strategy is the feature we chose for selecting representative plausible truths for this application.

After that, we investigated the impact of the non-uniqueness of optimal production strategies in the model selection. For that, we carried out nominal optimizations on each of the plausible truths starting from three different initial solutions. Figure 17 shows the results: Figure 17 (b), (c) and (d) correspond to the selection obtained with the three different starting points and Figure 17 (e) to the selection based on the data of the three optimizations altogether. We observe that the results are not identical, which confirms that the optimal production strategies may be nonunique and, thus, not suitable for model selection purposes.

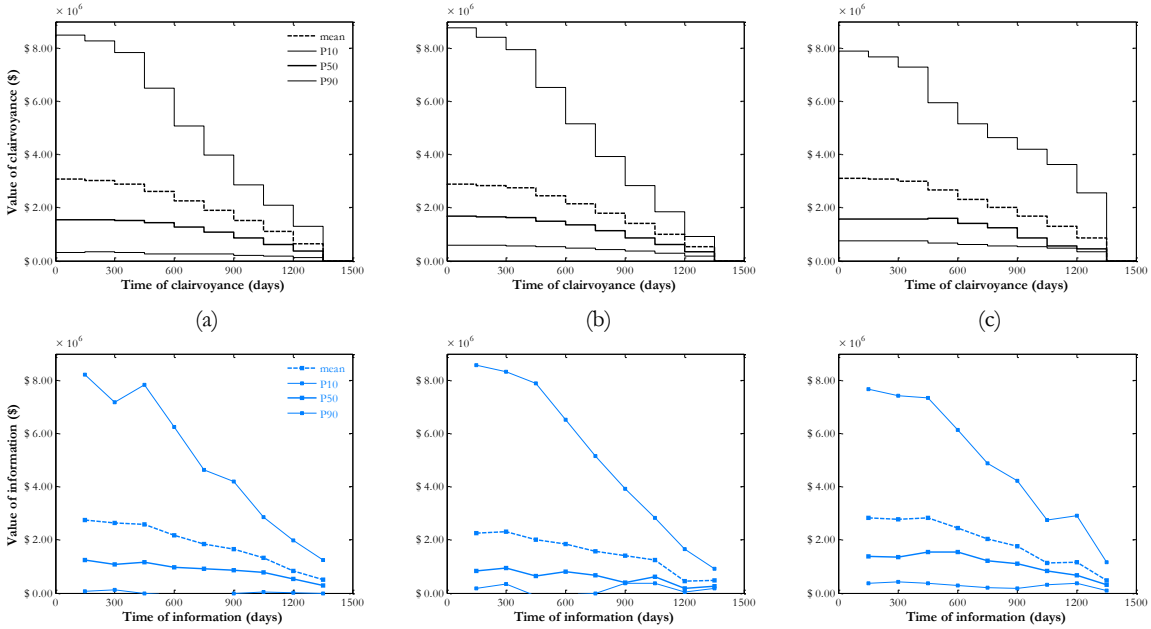

Figure 16: Results of the VOI (and VOC) assessment for the 2D five-spot model using selection of 5 plausible truths based on decision-based features. (a) Reference obtained using the original workflow, (b) selection based on objective function values and (c) selection based on optimal production strategies.

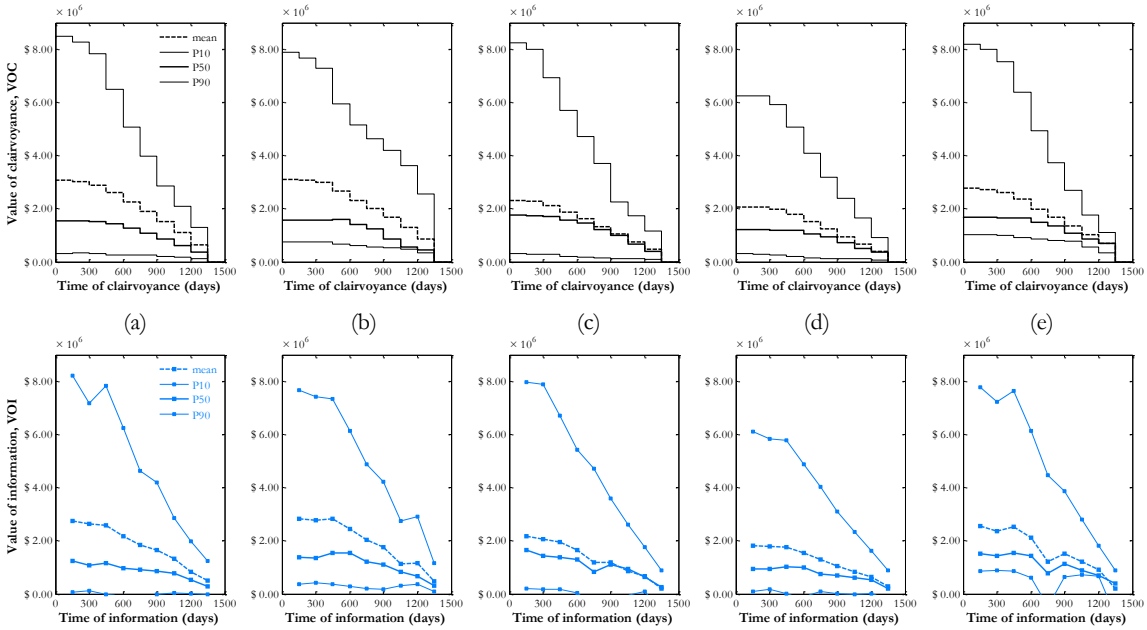

Figure 17: Results of the VOI (and VOC) assessment for the 2D five-spot model selecting representative plausible truths based on optimal production strategies obtained with different starting solutions. (a) Reference results, (b) starting from robust optimal solution, (c) starting from greedy controls (maximum injection and maximum drawdown in the producers), (d) starting from mid in-between bounds and (e) selection based on data from all the three optimizations.

As a measure to remediate this problem, we applied a projection based on a truncated tensor decomposition to the optimal production strategies dataset. By retaining the fraction of the basis functions that preserves 90% of the energy of the singular values, we hope to capture only the main trends and reduce the effect of non-uniqueness of the optimal solutions. Figure 18 shows the results obtained with such a projection for the same optimal production strategies. The difference between the results for the three different optimization starting points (plots (b)-(e)) is smaller than in Figure 18. Note that in this case we are not trying to show which figure (Figure 17 or Figure 18) depicts results closer to the reference like we did when comparing previous figures. The point here is to check whether the adopted measure reduces the impact of having different starting points, which indicates that this additional step could be helping us to minimize the effect of non-uniqueness of the optimal solutions.

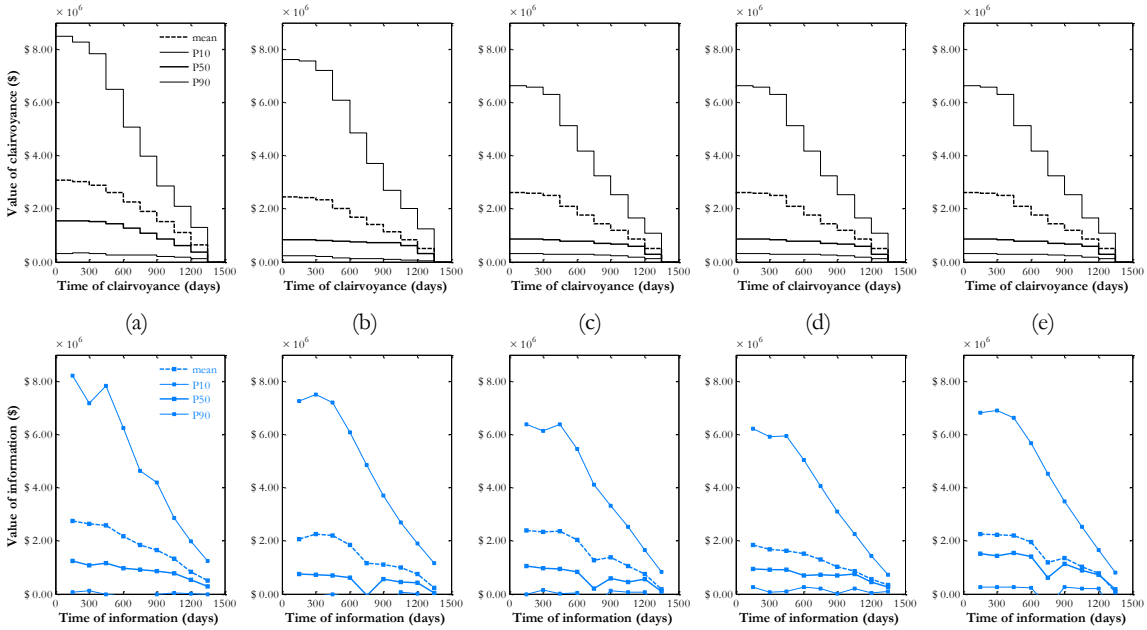

Figure 18: Results of the VOI (and VOC) assessment for the 2D five-spot model selecting representative plausible truths based on optimal production strategies projected by tensor decomposition. (a) Reference results, (b) starting from robust optimal solution, (c) starting from greedy controls (maximum injection and maximum drawdown in the producers), (d) starting from mid in-between bounds and (e) selection based on data from all the three optimizations.

4.3.2. Fully accelerated workflow

4.3.2.1. 2D five-spot model

Finally, we applied all the measures we discussed so far (i.e., selection of representative ensembles for robust optimization, history matching and plausible truths) following the accelerated procedure for VOI assessment depicted in Figure 7. Note that here the history matching was performed on the full ensembles because our implementation with MRST used the EnKF method, which is not reliable for very small ensembles. After the history matching steps, we selected representative realizations to accelerate the robust optimization over the posterior ensembles. Figure 19 presents the results. Again, Figure 19 (a) exhibits the reference results with $N_{truth} = 50$ plausible truths and robust optimization over ensembles of $N = 49$ realizations. Figure 19 (b) corresponds to the results obtained by using $N_{truth}^{repr} = 5$ representative plausible truths and full ensembles with $N = 49$ for the robust optimizations. Figure 19 (c) shows the results when considering all the $N_{truth} = 50$ plausible truths and reduced ensembles for robust optimizations. And Figure 19 (d) displays the results obtained with $N_{truth}^{repr} = 5$ plausible truths and ensembles of $N_{repr} = 5$ realizations for the optimizations.

We see that the acceleration measures allow us to obtain similar results by considering only 10 % of the original number of realizations. We also observe that even the combination of the two acceleration measures described is still able to correctly approximate the main trend of the reference

results. Note that the lines plotted in Figure 19 (b) and (d) represent percentiles and mean values of VOI based on N_{truth}^{repr} = 5 samples, while in Figure 19 (a) and (c) these values are computed with N_{truth} = 50 samples, and this should be taken into account when interpreting the quality of the results.

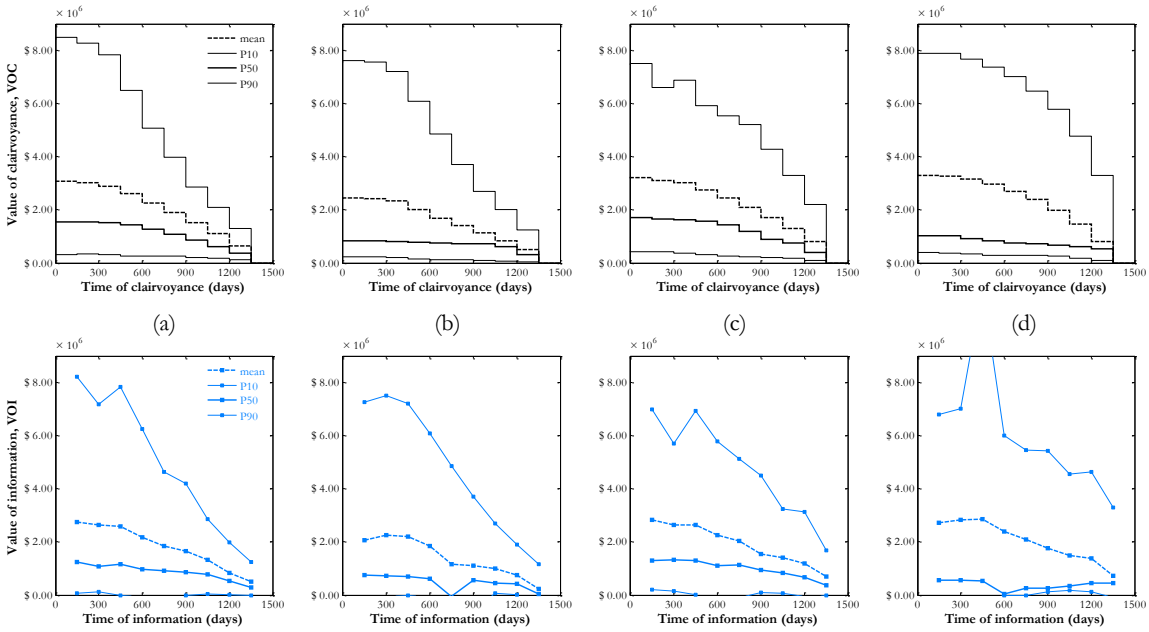

Figure 19: Results of the new accelerated VOI (and VOC) assessment for the 2D five-spot model. (a) Reference results, (b) only selecting representative plausible truths, (c) only reducing the ensembles for robust optimization, and (d) combining both measures.

In terms of computational cost, the results in Figure 19 (a) require approximately 1.5 million simulations. This number includes forward and backward simulations. The results shown in Figure 19 (b) and (c) require 150,000 and 170,000 simulations, respectively. And the results in Figure 19 (d) need 17,000 simulations to be computed. Thus, by applying all measures to reduce the number of model realizations considered in the assessment, we were able to alleviate the computational cost of the workflow by a factor of 88, which is quite significant.

4.3.2.2. Egg model

As a final test, we applied all the model selection measures discussed in this paper (Figure 7) to make the VOI assessment possible for the Egg model. We considered N_{truth}^{repr} = 10 representative plausible truths and ensembles of N_{repr} = 10 representative realizations for the history matching and robust optimization. The workflow to assess the VOI of a single observation time was repeated for different observation times, t_{data} = {360, 720, … , 3,240} days. The entire exercise was performed with AD-GPRS to evaluate the objective functions and obtain the required gradients for the production optimization and history matching steps. The history matching was performed using the RML method. We assessed the VOI of the production data (field production rates and pressures in the

injectors) with absolute measurement errors (ε_{prod} = 5 m³/day and ε_{BHP} = 10 bar). The VOC and the VOI were computed for each of the nine observation times.

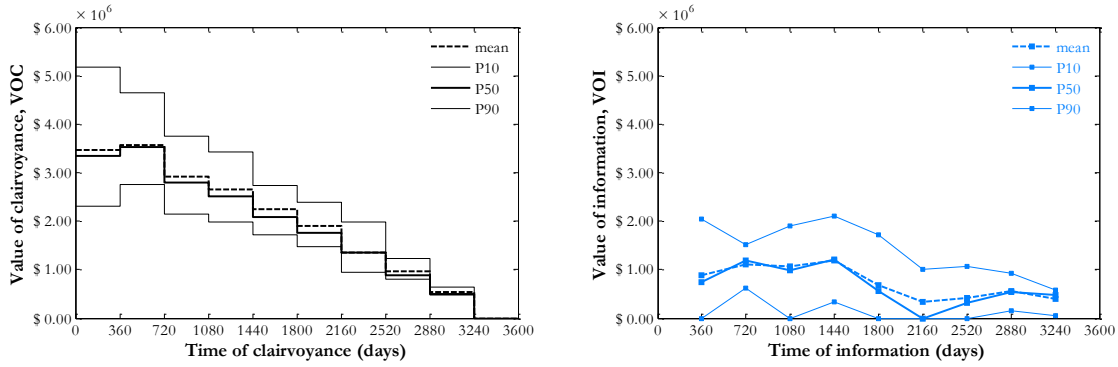

Figure 20: Results for the Egg model with a single observation time: VOC and VOI.

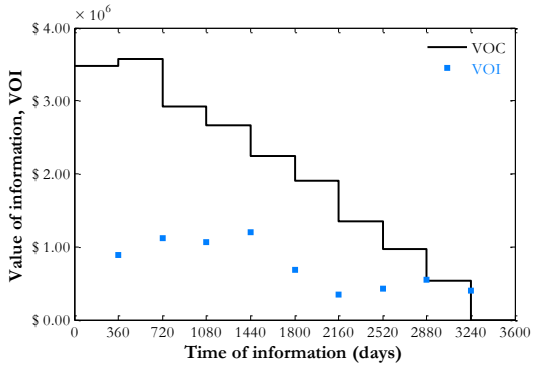

Figure 21: Results for the Egg model with a single observation time: expected VOC and VOI.

Figure 20 and Figure 21 display the results obtained. We note that the type and trend of these results share similarities with those observed for the 2D five-spot model. As before, the VOC decreases in a stepwise fashion and serves as an upper-bound for the VOI. On the other hand, this time the VOI values are significant lower than the VOC values, most likely because the field production rate measurements are not as informative as well-per-well measurements and also because of the larger number of uncertain parameters.

We also observe that, unlike the results for the 2D five-spot model, in Figure 20 and Figure 21 the VOI values fluctuate more around the overall trend. This could be related to this particular case study as VOI does not need necessarily to show a monotonic behavior over time. This could also be associated with imperfections of the optimization procedures involved in the workflow, i.e. the risk of reaching local optima in the production optimization and history matching steps. Finally, this fluctuation of VOI values could be explained by the number of truths considered being insufficient to derive proper VOI statistics. Even though there are no theoretical guarantees, we believe that increasing N_{truth}^{repr} could eliminate this behavior.

We also repeated the VOI assessment for the case with multiple observation times, according to the extension of the methodology presented in [8]. Figure 22 and Figure 23 display the results obtained,

which, once again, confirm the conclusions drawn from the previous case studies. However, unlike the results for the 2D example in [8], here the increase in VOI after the first observation time is more important.

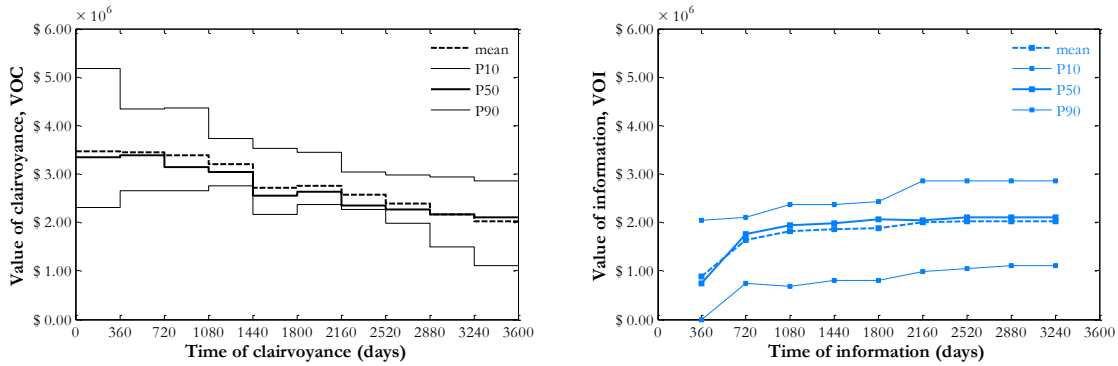

Figure 22: Results for the Egg model with multiple observation times: VOC and VOI.

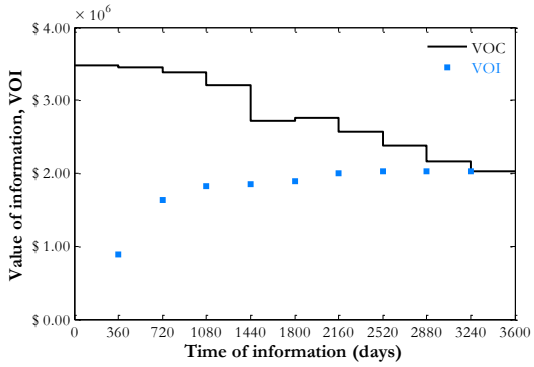

Figure 23: Results for the Egg model with multiple observation times: expected VOC and VOI.

Besides assessing the VOI of future measurements, our approach of "closing the loop" in the design phase produces a great amount of data which could be used to gain additional insight into the (closed-loop) reservoir management problem. Because we consider an ensemble of plausible truths, multiple CLRM sequences are obtained, including optimal production strategies (Figure 24), history matched ensembles and their predictions (Figure 25 and Figure 26). An in-depth analysis of these data may lead to a better understanding on how to manage the reservoir, serving as a sort of pre-computed operation manual.

In Figure 24, we note that the production strategies optimized with perfect geological knowledge already from $t = 0$ make a much better use of the operational flexibility available, allowing a better control over the reservoir system to improve its performance. The strategies derived through CLRM manage to use some of this flexibility (i.e., the strategies are somewhat spread over the allowed bounds for the control variables), but not as efficiently as in the case with clairvoyance (i.e., the strategies exhibit an even larger spread) because the learning obtained from production measurements is more limited. The robust strategy determined under prior uncertainty is not able to capture the possibility of reacting to future measurements.

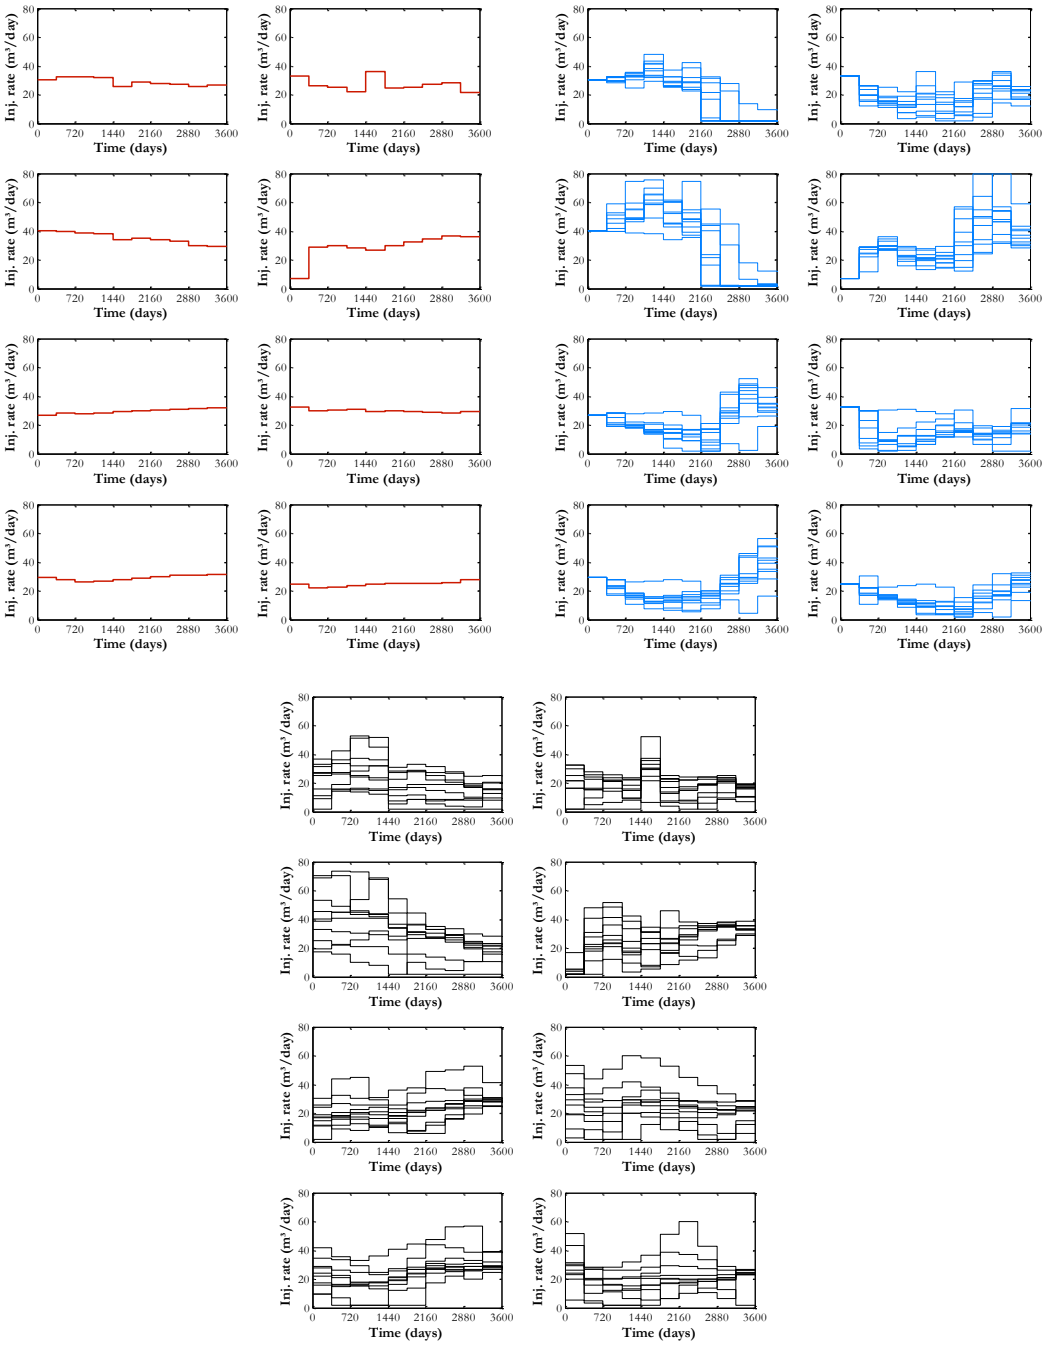

Figure 24: Optimal production strategies for the plausible truths considered Egg model: optimized under prior uncertainty (top left); obtained through CLRM with additional production measurements (top right); optimized under the assumption of clairvoyance available at $t = 0$.

Figure 25 depicts the performance of the production strategies discussed above in terms of the NPV predictions for the $N_{truth}^{repr} = 10$ representative plausible truths considered. First, we note that the strategy derived under prior uncertainty leads to cumulative NPV curves that descend after reaching their maximum, suggesting that production continued even after reaching uneconomical levels. This occurs because the optimization under prior uncertainty was performed over an ensemble of realizations different from the ensemble of $N_{truth}^{repr} = 10$ representative plausible truths considered, which causes the determined strategy to be suboptimal for the plausible truths.

We also observe that the CLRM strategies result in higher NPV values, but also in a larger spread. Finally, we notice a distinct behavior of the NPV curves derived with CLRM and under clairvoyance: the curves of the different plausible truths cross each other, while the curves derived under prior uncertainty never do that. The reason for this behavior is that we have a single production strategy optimized with the prior knowledge and one optimal strategy for each plausible truth in the other cases. These multiple optimal strategies have been tailored to their respective plausible truths, which allows them to achieve higher NPV values. This also makes them achieve their maximum at different times, causing the curves to cross.

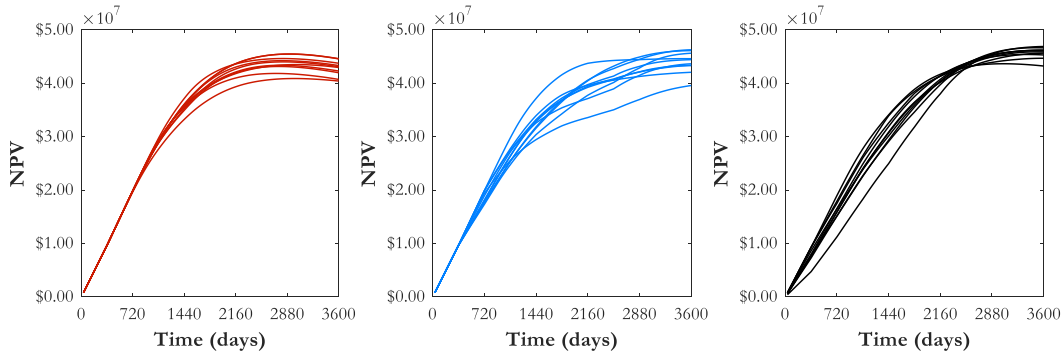

Figure 25: Cumulative NPV curves for the 10 representative plausible truths (Egg model) for the production strategies from Figure 24.

As discussed above, the optimized production strategies may be suboptimal for the plausible truths. In practice, when operating a field, the production would be interrupted if it reaches uneconomical levels (i.e., reactive control would be exercised). This consideration may have an impact on the NPV curves derived and on the outcome of our VOI assessment. Therefore, we repeated the exercise including reactive control as follows. Figure 26 shows the cumulative NPV curves obtained by interrupting the production simulated for the plausible truths at the time they reach their maximum. As a result, we notice that the curves that before would go down (in Figure 25) now stay flat after they reach their maximum value.

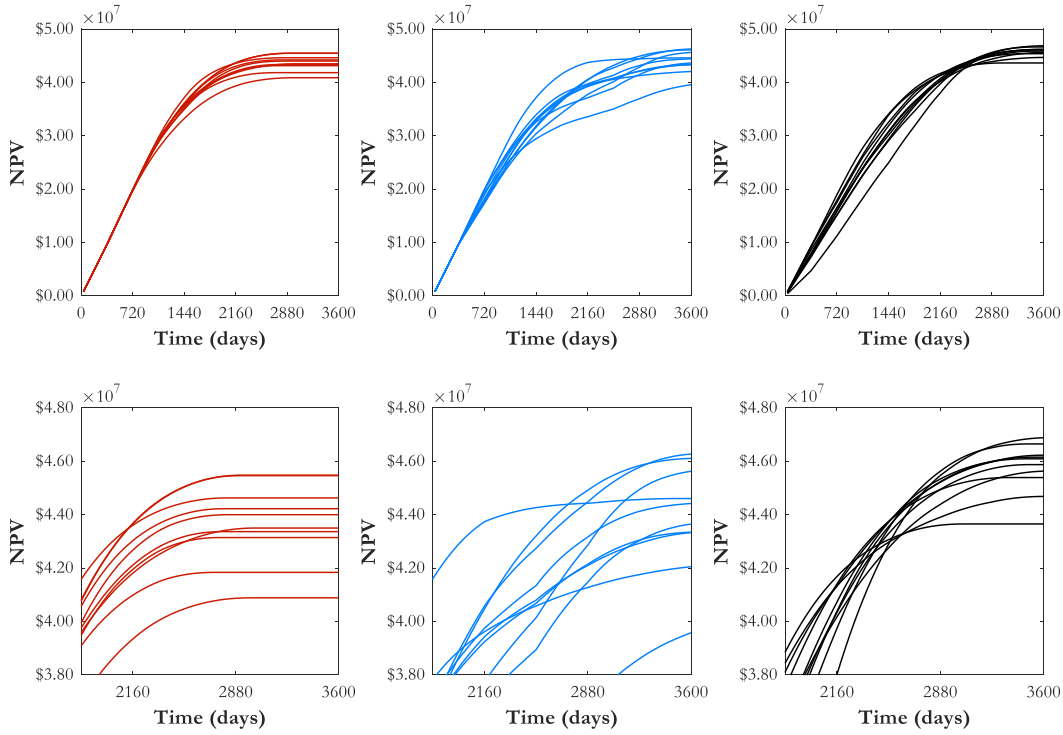

Figure 26: Cumulative NPV curves for the 10 representative plausible truths disregarding uneconomic production (Egg model) for the production strategies from Figure 24. The plots at the bottom show a zoomed-in view of the curves for the second half of the reservoir life-cycle.

Figure 27 and Figure 28 show the results of the VOI assessment obtained when disregarding the uneconomical production from the plausible truths, and we note that this causes the VOI and VOC to decrease in comparison with Figure 22 and Figure 23. This can be understood by comparing Figure 25 and Figure 26: when disregarding the uneconomical production, the baseline values for the VOI assessment (i.e., the final NPV achieved with the strategies optimized under prior uncertainty) increase, while this consideration does not have a major impact on the values obtained with additional knowledge. The VOI and VOC depicted Figure 27 and Figure 28 can also be verified in Figure 26 (bottom): from the zoomed-in plots we can see more clearly the incremental gain with respect to the prior strategy for the $N_{truth}^{repr} = 10$ plausible truths.

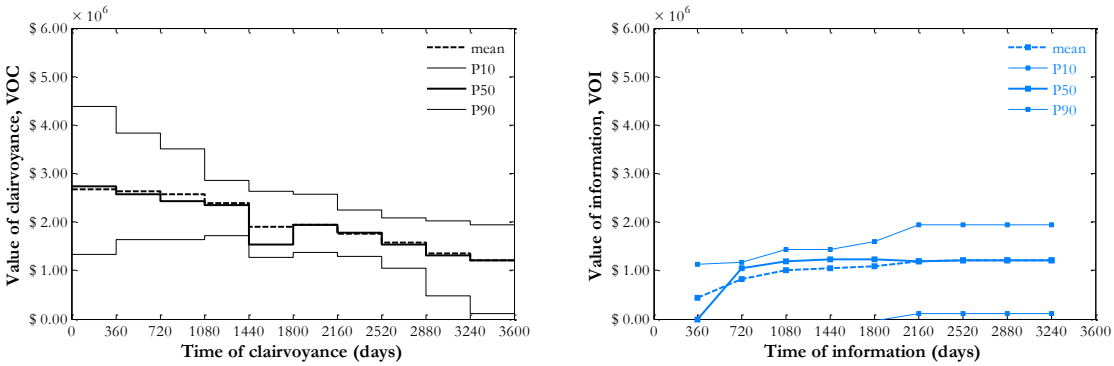

Figure 27: Results for the Egg model with multiple observation times: VOC and VOI disregarding the uneconomic production of the plausible truths.

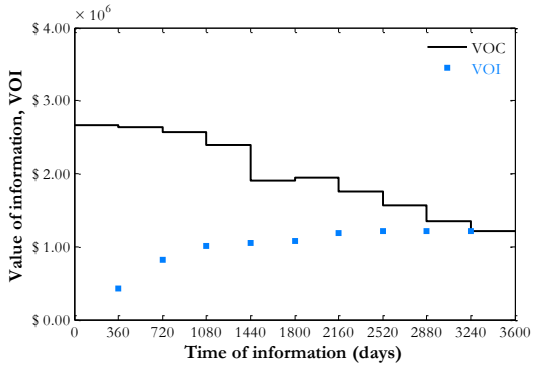

Figure 28: Results for the Egg model with multiple observation times: expected VOC and VOI disregarding the uneconomic production of the plausible truths.

5. Discussion and conclusions

We applied two different measures to decrease the amount of simulations required in VOI assessment workflows. First, we showed how to make the robust optimization and history matching procedures more efficient by reducing the size of the ensembles considered. We used several model-based features to select representative realizations through clustering and compared them against reference results. For the small 2D five-spot example (441 grid blocks), we concluded that oil saturation snapshots data transformed by tensor decomposition or MDS provide a good basis for the selection of reduced ensembles to speed-up robust optimization and history matching. Second, we introduced two new decision-based features to support the selection of representative plausible truths as another alternative to speed-up the VOI assessment procedure. We confirmed that optimal production strategies are most suitable for this purpose and that the choice of the selection feature is case-dependent even within the same workflow. A disadvantage of this approach is the challenge in deriving meaningful statistics of the VOI given the reduced number of plausible truths. Finally, we combined both aforementioned acceleration measures to design a new procedure for faster VOI assessment. For the 2D example, we were able to reduce the amount of required reservoir simulations from millions to tens of thousands. This significant reduction in computational costs represents an important step for the use of the VOI assessment in larger examples.

Based on the learning from the experiments with the small 2D model, we tested the acceleration measures on the medium-sized Egg model (18,553 grid blocks). The results confirmed the main conclusions obtained for the small model, reassuring us that clustering based on the appropriate features can support the selection of representative models to approximate the uncertainty characterized by the full ensembles of realizations and make our workflows much more efficient also in larger examples. In the end, by combining all the acceleration measures, we were able to apply our methodology for VOI assessment to the Egg model, which would otherwise have been

computationally intractable. The results obtained supported our main conclusions from previous case studies, which reassures the consistency of our approach to VOI assessment.

However, there is still scope for future research to further accelerate these workflows. In particular, the combination of the ideas presented in this paper with the use of surrogate models (e.g., proxies), reduced-physics models or multiscale methods may be necessary to develop practical tools for VOI assessment in realistic, large-scale applications. One of the results of the Egg models made us question whether the $N_{truth}^{repr} = 10$ plausible truths considered were enough to derive meaningful statistics of VOI. Most likely the arbitrary choice of the number of representative realizations also impacts the robust optimization and history matching steps. A more systematic approach could improve the choice of N_{repr} and N_{truth}^{repr}.

Another point that deserves attention concerns the generalization of this approach to field development problems. In this work we focused on reservoir management applications where the well configuration remains fixed, which is determinant for the flow behavior of the different model realizations. In workflows where the well configuration changes, the optimization results will likely be more sensitive to the underlying geology. This means that more research is needed to confirm whether the conclusions of this work on reservoir management can be extended to the field development context. Yet we believe that some of the ideas discussed here may be adapted or inspire the development of acceleration measures to VOI workflows in more general problems.

Finally, it is important to reflect on the use of such VOI workflows in practice. We have proposed some solutions to make these workflows more tractable computationally, which is an important step to enable their application to real-life problems. Another necessary step to make their use more attractive concerns the interpretability of the quantified VOI, a topic that still has to be addressed. Any physical interpretation of these results implies actually multiple interpretations, because the additional information might be useful in a different way for each truth scenario considered. To achieve interpretability in an automated way, we need to develop approaches that leverage all the data produced by the workflow to identify relationships between the various variables of problem with the computed VOI.

Acknowledgments

This research was carried out within the context of the ISAPP Knowledge Centre. ISAPP (Integrated Systems Approach to Petroleum Production) is a joint project of TNO, Delft University of Technology (TU Delft), ENI, Statoil and Petrobras. The EnKF module for MRST was developed by Olwijn Leeuwenburgh (TNO) and is available at http://www.isapp2.com/data-sharepoint/enkfmodule-for-mrst. The authors thank Lou Durlofsky and Oleg Volkov of Stanford

University for sharing a working version of the AD-GPRS reservoir simulator and its optimization module. The authors thank Fu Kai Yap for the contributions of his MSc thesis work to this research. The authors thank Siep Weiland and Edwin Insuasty of Eindhoven University of Technology for insightful discussions on tensor decomposition.

References

1. Aydin, O. and Caers, J. (2013). Image transforms for determining fit-for-purpose complexity of geostatistical models in flow modeling. *Computational Geosciences*, **17** (2), 417-429.
2. Aanonsen, S.I., Naevdal, G., Oliver, D.S, Reynolds, A.C. and Valles, B. (2009). The ensemble Kalman filter in reservoir engineering – a review. *SPE Journal*, **14** (3) 393-412.
3. Afra, S. and Gildin, E. (2016). Tensor based geology preserving reservoir parameterization with Higher Order Singular Value Decomposition (HOSVD). *Computers & Geosciences* **94**, 110–120.
4. Aggarwal, C.C., Wolf, J.L., Yu, P.S., Procopiuc, C., Park, J.S. (1999). Fast algorithms for projected clustering. *SIGMOD Record*, **28** (2), 61-72.
5. Arabie, P. and Hubert, L.J. (1996). An overview of combinatorial data analysis. *In* Arabie, P., Hubert, L.J., and Soete, G.D., Eds., *Clustering and Classification*, 5-63. World Scientific, River Edge, NJ.
6. Armstrong, M., Ndiaye, A., Razanatsimba, R. and Galli, A. (2013). Scenario reduction applied to geostatistical simulations. *Mathematical Geosciences*, **45**, 165-182.
7. Baker, M. (2015). Use of cluster analysis to improve representative model selection: a case study. Paper SPE 176408-MS presented at SPE/IATMI Asia Pacific Oil & Gas Conference and Exhibition, Nusa Dua, Bali, Indonesia, 20-22 October.
8. Barros, E.G.D., Leeuwenburgh, O., Van den Hof, P.M.J. and Jansen, J.D. (2015). Value of multiple production measurements and water front tracking in closed-loop reservoir management. Paper SPE-175608-MS presented at the SPE Reservoir Characterization and Simulation Conference and Exhibition, Abu Dhabi, UAE, 14-16 September.
9. Barros, E.G.D., Van den Hof, P.M.J. and Jansen, J.D. (2016). Value of information in closed-loop reservoir management. *Computational Geosciences*, **20** (3), 737-749.
10. Barros, E.G.D., Yap, F.K., Insuasty Moreno, E., Van den Hof, P.M.J. and Jansen, J.D. (2016). Clustering techniques for value-of-information assessment in closed-loop reservoir management. Paper presented at the 15th European Conference on Mathematics in Oil Recovery, Amsterdam, The Netherlands, 29 August - 1 September.
11. Bukshtynov, V., Volkov, O., Durlofsky, L.J., Aziz, K. (2015). Comprehensive framework for gradient-based optimization in closed-loop reservoir management. *Computational Geosciences*, **19** (4), 877-897.
12. Caers, J. (2011). *Modeling uncertainty in the earth sciences*. John Wiley & Sons.
13. Capolei, A., Suwartadi, E., Foss, B. and Jørgensen, J.B. (2015). A mean-variance objective for robust production optimization in uncertain geological scenarios. *Journal of Petroleum Science and Engineering*, **125**, 23-37.
14. Cardoso, M.A. and Durlofsky, L.J. (2010). Use of reduced-order modeling procedures for production optimization. *SPE Journal*, **15** (2), 426-435.
15. Chen, B., He, J., Xie, J., Sarma, P., Wen, X.-H., Chen, W.H. and Reynolds, A.C. (2016). Pilot design analysis using proxies and Markov Chain Monte Carlo method. Paper presented at the 15th European Conference on Mathematics in Oil Recovery, Amsterdam, The Netherlands, 29 August - 1 September.

16. Chen, Y., Oliver, D.S. and Zhang, D. (2009). Efficient ensemble-based closed-loop production optimization. *SPE Journal*, **14** (4), 634-645.

17. Cunningham, J.P. and Ghahramani, Z. (2015). Linear dimensionality reduction: survey, insights, and generalizations. *Journal of Machine Learning Research*, **16**, 2859-2900.

18. Eidsvik, J., Mukerji, T. and Bhattacharjya, D. (2015). Value of information in the earth sciences – Integrating spatial modeling and decision analysis, 1st ed, Cambridge University Press.

19. Evensen, G. (2009). *Data assimilation – The ensemble Kalman filter*, 2nd ed., Springer, Berlin.

20. He, J., Sarma, P. and Durlofsky, L.J. (2013). Reduced-order flow modeling and geological parameterization for ensemble-based data assimilation. *Computers and Geosciences*, **55**, 54-69.

21. He, J., Xie, J., Sarma, P., Wen, X.H., Chen, W.H. and Kamath, J. (2016). Proxy-based work flow for a priori evaluation of data-acquisition programs. *SPE Journal*, **21** (4), 1400-1412.

22. Hewson, C.W. (2015). Reduced-order modelling for production optimisation. MSc thesis, Delft University of Technology.

23. Hou, J., Zhou, K., Zhang, X.S. Kang, X.D. and Xie, H. (2015). A review of closed-loop reservoir management. *Petroleum Science*, **12**, 114-128.

24. Insuasty, E., Van den Hof, P.M.J., Weiland, S. and Jansen, J.D. (2017). Flow-based dissimilarity measures for reservoir models: a spatial-temporal tensor approach. *Computational Geosciences*, **21** (4), 645-663.

25. Insuasty, E. (2018). *A spatial-temporal approach to reduced complexity modelling for hydrocarbon reservoir optimization*. PhD thesis. Eindhoven University of Technology.

26. Jansen, J.D., Brouwer, D.R., Nævdal, G. and van Kruijsdijk, C.P.J.W. (2005). Closed-loop reservoir management. *First Break*, January, **23**, 43-48.

27. Jansen, J.D., Bosgra, O.H. and van den Hof, P.M.J. (2008). Model-based control of multiphase flow in subsurface oil reservoirs. *Journal of Process Control* **18**, 846-855.

28. Jansen, J.D., Douma, S.G., Brouwer, D.R., Van den Hof, P.M.J., Bosgra, O.H. and Heemink, A.W. (2009). Closed-loop reservoir management. Paper SPE 119098 presented at the SPE Reservoir Simulation Symposium, The Woodlands, USA, 2-4 February.

29. Jansen, J.D. and Durlofsky, L.J. (2016). Use of reduced-order models in well control optimization. *Optimization and Engineering*. Published online.

30. Jansen, J.D., Fonseca, R.M., Kahrobaei, S., Siraj, M.M., Van Essen, G.M. and Van den Hof, P.M.J. (2014). The egg model – a geological ensemble for reservoir simulation. *Geoscience Data Journal*, **1** (2), 192-195.

31. Keim, D., Berchtold, S., Böhm, C., Kriegel., H.P. (1997). A cost model for nearest neighbor search in high dimensional data space. *In Proc.* of the 16th Symposium on Principles of Database Systems, 78-86.

32. Kruskal, J.B. (1964). Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, **29** (1), 1-27.

33. Kruskal, J.B. (1976). The relationship between multidimensional scaling and clustering. *In Proc.* of Advanced Seminar Conducted by the Mathematics Research Center, the University of Wisconsin-Madison, 3-5 May, 17-44.

34. Le, D.H. and Reynolds, A.C. (2014). Optimal choice of a surveillance operation using information theory. *Computational Geosciences*, **18**, 505-518.

35. Le, D.H. and Reynolds, A.C. (2014). Estimation of mutual information and conditional entropy for surveillance optimization. *SPE Journal*, **19** (4), 648-661.

36. Lee, J.A. and Verleysen, M. (2007). *Nonlinear dimensionality reduction*, Springer, Berlin.

37. Lie, K.-A., Krogstad, S., Ligaarden, I. S., Natvig, J. R., Nilsen, H. M. and Skalestad, B. (2012). Open source MATLAB implementation of consistent discretisations on complex grids. *Computational Geosciences*, **16** (2), 297-322.

38. Liu, Z. and Forouzanfar, F. (2017). Ensemble clustering for efficient robust optimization of naturally fractured reservoirs. *Computational Geosciences*. Published online.

39. Lu, R., Forouzanfar, F. and Reynolds, A.C. (2017). An efficient adaptive algorithm for robust control optimization using StoSAG. *Journal of Petroleum Science and Engineering*, **159**, 314-330.

40. Naevdal, G., Brouwer, D.R. and Jansen, J.D. (2006). Waterflooding using closed-loop control. *Computational Geosciences*, **10** (1), 37-60.

41. Oliver, D.S., Reynolds, A.C. and Liu, N. (2008). *Inverse theory for petroleum reservoir characterization and history matching*, Cambridge University Press, Cambridge.

42. Oliver, D.S. and Chen, Y. (2011). Recent progress on reservoir history matching: a review. *Computational Geosciences*, **15** (1) 185-221.

43. Parsons, L., Haque, E. and Liu, H. (2004). Subspace clustering for high dimensional data: a review. *ACM SIGKDD Explorations Newsletter*, **6** (1) 90-105.

44. Sarma, P., Durlofsky, L.J. and Aziz, K. (2008). Computational techniques for closed-loop reservoir modeling with application to a realistic reservoir. *Petroleum Science and Technology*, **26** (10 and 11), 1120-1140.

45. Sarma, P., Chen, W.H. and Xie, J. (2013). Selecting representative models from a large set of models. Paper SPE 163671 presented at the SPE Reservoir Simulation Symposium, The Woodlands, USA, 18-20 February.

46. Scheidt, C. and Caers, J. (2009). Representing spatial uncertainty using distances and kernels. *Mathematical Geosciences*, **41** (4), 397-419.

47. Shepard, R.N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function. I. *Psychometrika*, **27** (2):125-140.

48. Shirangi, M.G. and Durlofsky, L.J. (2016). A general method to select representative models for decision making and optimization under uncertainty. *Computers and Geosciences*, **96**, 109-123.

49. Siraj, M.M., Van den Hof, P.M.J. and Jansen, J.D. (2016). Robust optimization of water-flooding in oil reservoirs using risk management tools. Paper presented at 11th IFAC Symposium on Dynamics and Control of Process Systems, Trondheim, Norway, 6-8 June.

50. Van Essen, G.M., Zandvliet, M.J., Van den Hof, P.M.J., Bosgra, O.H. and Jansen, J.D. (2009). Robust waterflooding optimization of multiple geological scenarios. *SPE Journal*, **14** (1), 202-210.

51. Yap, F.K. (2016). Representative models for history matching and robust optimization. MSc thesis, Delft University of Technology.

52. Yeten, B., Durlofsky, L.J. and Aziz, K. (2003). Optimization of nonconventional well type, location and trajectory. *SPE Journal*, **8** (3), 200-210.

53. Wang, C., Li, G. and Reynolds, A.C. (2009). Production optimization in closed-loop reservoir management. *SPE Journal*, **14** (3), 506-523.

54. Zhou, Y. (2012). Parallel General-Purpose Reservoir Simulation with Coupled Reservoir Models and Multi-Segment Wells. PhD Thesis, Stanford University.