

# History matching with an ensemble Kalman filter and discrete cosine parameterization

Behnam Jafarpour · Dennis B. McLaughlin

Received: 8 August 2007 / Accepted: 7 January 2008 / Published online: 29 January 2008  
© Springer Science + Business Media B.V. 2008

**Abstract** History matching of large hydrocarbon reservoirs is challenging because of several reasons including: (1) scarcity of available measurements relative to the number of unknowns, leading to an ill-posed inverse problem, (2) computational effort required for large reservoir problems, and (3) the need to insure that solutions are geologically realistic. All of these problems can be helped by using algorithms that rely on efficient and parsimonious descriptions (or parameterizations) of reservoir properties. This paper demonstrates the use of a novel parameterization approach, the discrete cosine transform, for history matching with a recently introduced sequential estimation technique, i.e., the ensemble Kalman filter. The proposed approach exploits the structure of the estimation and parameterization algorithms to reduce the size of reservoir states (pressures and saturations) as well as parameters (e.g., intrinsic permeability) with a marginal loss in accuracy. The introduced methodology eliminates redundancy in posing the estimation problem and results in additional computational savings. Application and generality of this approach are demonstrated using two waterflooding experiments characterized by different types of geological variability.

**Keywords** History matching · Ensemble Kalman filter · Discrete cosine transform · Reservoir description · Parameterization

## 1 Introduction

Predictive modeling of the production behavior of hydrocarbon reservoirs is of paramount importance for reservoir engineers. Accurate reservoir models can provide reliable production forecasts for guiding operational activities and development plans. Effective description of geological formations and their flow-related properties (such as permeability) is a significant part of any modeling exercise.

Several data acquisition techniques operating at various spatial scales provide information about the distribution of geological facies and rock properties within a petroleum reservoir. These include core characterizations, well logging, and seismic surveys. However, data acquisition costs limit the extent of such activities and the amount of data that can be collected in practice. Data limitations make it necessary to interpolate the available data to infer information about unobserved regions. Any interpolation scheme has to make assumptions (either implicit or explicit) about the structure of geological formations in the reservoirs. These assumptions can introduce a significant amount of uncertainty [1, 2].

To reduce the uncertainty in reservoir descriptions, reservoir model parameters are usually adjusted to improve the match between simulated and observed production data, a process known as history matching [3, 4]. Numerical reservoir simulators rely on discretized computational grids that require petrophysical properties such as intrinsic permeability and porosity to be defined for each grid block [5]. However, because geological media are connected

---

B. Jafarpour (✉)  
Department of Petroleum Engineering, Texas A&M University,  
401F Richardson Building,  
College Station, TX 77843, USA  
e-mail: behnam@pe.tamu.edu

D. B. McLaughlin  
Department of Civil and Environmental Engineering,  
Massachusetts Institute of Technology,  
Cambridge, MA, USA

layers of rocks with similar physical properties, a strong correlation (or redundancy) exists between property values at neighboring grid blocks. When a history-matching procedure independently adjusts property values in every grid block, many different combinations of block values may yield comparable matches to a given set of production data. If this occurs, the history-matching problem is ill posed (i.e., it is not possible to obtain a unique solution).

The problem of ill posedness can be mitigated if the spatially variable grid block-based property values are replaced by a smaller number of parameters that capture all (or most) of the variability. This process is called parameterization (or sometimes reparameterization). Parameterization often improves ill posedness because it constrains the set of possible grid block values. A parameterization approach should be selected so that it meets the needs of the reservoir simulator while retaining geological realism. An efficient and realistic parameterization can yield improved grid-based property estimates (by dealing with ill posedness) and can decrease computational effort (by reducing the number of variables to be estimated).

Many parameterization approaches have been proposed in the petroleum-engineering literature. These approaches obtain a reduced set of parameters by using methods that range from simple zonation [6] to more complex mathematical transforms [7–10]. Mathematical transforms that are based on spectral decomposition have the advantage of efficiently separating the large- and small-scale spatial features of a geological property such as permeability. If the finer details are omitted, an approximate description of the original field is achieved with fewer parameters.

In a recent publication [8], we proposed the discrete cosine transform (DCT) for parameterization of a permeability field (and other spatially variable quantities) and compared its performance and flexibility with the Karhunen–Loeve transform (KLT). These two linear transformations use different basis functions. In [8], we discuss and demonstrate through examples some of the important advantages of the DCT.

In this paper, we obtain an efficient and accurate history-matching algorithm by combining DCT parameterization with the ensemble Kalman filter (EnKF) [11], a recursive state estimation technique. In this algorithm, the DCT parameterization is applied to pressure and saturation as well as permeability, and the Kalman filter state vector includes DCT-transformed values of all three variables. This important extension of the parameterization concept is especially useful for Kalman filtering because it greatly reduces the size of the sample covariance matrices used to derive state estimates.

We begin by discussing how the EnKF can be applied to the reservoir history-matching problem. Then, we describe two example problems that may be used to test the EnKF with different parameterization options. We first consider

the implementation and performance of an EnKF history-matching algorithm that uses a conventional grid block-based description of spatial variability. Next, we present an overview of DCT parameterization and illustrate its advantages with a modified EnKF history-matching algorithm. We conclude with a brief discussion of implementation and efficiency issues.

## 2 History matching with the ensemble Kalman filter

The EnKF was first introduced by Evensen [11] in 1994 as a way to extend the classical Kalman filter to nonlinear problems [12]. A detailed discussion of the method and its evolution in the past decade as well as an efficient implementation of it can be found in [13]. Recent applications of the EnKF to history matching for petroleum engineering include [14–16].

The Kalman filter is designed to estimate the state vector  $x_t$  of a dynamic system from noisy measurements that are related to the states through a specified measurement equation. In petroleum reservoir applications, the state vector typically includes pressure and oil saturation (in the examples of interest here, capillary pressure is negligible, so oil and water pressures are approximately equal). The measurements in our waterflooding examples are well-bottom-hole pressures and flow rates for each phase.

Reservoir history matching is designed to estimate petrophysical properties such as intrinsic permeability, which are needed for predictive simulations. This exercise is often posed as a parameter estimation problem, rather than a state estimation problem. If a parameter estimation perspective is adopted, the desired estimates are typically obtained by minimizing a performance measure that quantifies the mismatch between measured and predicted states. However, it is also possible to formulate the history-matching problem as a state estimation problem. In this case, the properties of interest are simply added to the state vector and estimated together with the other states (this process is often called state augmentation). State augmentation can also include measurements that are nonlinearly related to the states (e.g., flow rates and bottom-hole pressures). A state estimation approach to history matching is particularly useful in real-time control applications, where new decisions are made as conditions change and new data become available.

The EnKF can be viewed as a recursive Bayesian procedure that approximates conditional probability densities of the state. These densities at time  $t$  are the forecast (or prior) density  $p[x_t|y_{0:t-1}]$ , conditioned on all measurements  $y_{0:t-1}$  taken before time  $t$  (subscript  $0:t-1$  refers to all time steps from initial to previous  $[t-1]$  time step), and the updated (or posterior) density  $p[x_t|y_{0:t}]$ , conditioned on all measurements through time  $t$ .

Various properties of these conditional densities, such as the mean (the minimum error variance value) or the mode (most likely value) may be adopted as state estimates. A history-matching procedure based on a state-augmented EnKF provides parameter and state estimates as well as distributions of possible values around these estimates. When a conventional grid block-based approach is adopted, the discrete states are composed of the values of the state variables at the grid blocks. When a DCT parameterization is used, the discrete states are the coefficients of the DCT basis vectors.

The EnKF approximates the forecast and updated densities with relatively small sets (or ensembles) of  $N$  random realizations ( $N=300$  was used in our experiments based on reported sensitivity analysis in [18]), denoted by  $x_{t|t-1}^j$  and  $x_{t|t}^j$ , respectively, where  $j=1,\dots,N$  denotes a particular replicate. An initial ensemble  $x_{0|0}^j$  drawn at random from a specified population of initial states (pressures, saturations, and permeabilities) is used to initialize the sequence of forecasts and updates. Subsequent forecasts are derived from the augmented state equation, which may be written as:

$$x_{t|t-1}^j = f_t(x_{t-1|t-1}^j, \mathbf{u}_{t-1}, \mathbf{w}_{t-1}^j); j = 1, \dots, N \quad (1)$$

where  $\mathbf{u}_{t-1}$  is a vector of known (non-random) time-dependent boundary conditions and controls and  $\mathbf{w}_{t-1}^j$  is a random vector that accounts for uncertain model errors. The elements of the vector function  $f_t(\cdot, \cdot, \cdot)$  corresponding to pressure and saturation represent solutions obtained from the reservoir simulator ECLIPSE [17] by propagating  $x_{t-1|t-1}^j$  and  $\mathbf{u}_{t-1}$  forward from  $t-1$  to  $t$ . The elements of  $f_t(\cdot, \cdot, \cdot)$  corresponding to augmented parameters are selected to insure that  $x_{t|t-1}^j = x_{t-1|t-1}^j$  when  $x_{t|t-1}^j$  is an augmented parameter (e.g., permeability). This requirement follows because the parameter estimate should not change between updates.

A version of the classical Kalman filter update [12] is used to update the forecast replicates obtained from Eq. 1:

$$x_{t|t}^j = x_{t|t-1}^j + \text{Cov}[x_{t|t-1}^j, y_{t|t-1}^j] \text{Cov}^{-1}[y_{t|t-1}^j, y_{t|t-1}^j] [y_t - y_{t|t-1}^j]; j = 1, \dots, N \quad (2)$$

This update reflects information obtained from the measurement  $y_t$  taken at time  $t$ . The notation  $\text{Cov}[\cdot, \cdot]$  represents the sample covariance between the ensembles associated with the two arguments. The perturbed measurement prediction  $x_{t|t-1}^j$  is obtained from the following measurement equation:

$$y_{t|t-1}^j = g_t(x_{t|t-1}^j) + \mathbf{v}_t^j; j = 1, \dots, N \quad (3)$$

where  $\mathbf{v}_t^j$  is a vector of measurement errors, drawn at random from a specified population, and  $g_t(\cdot, \cdot)$  is the

measurement operator that relates the measurements at  $t$  to the states at  $t$ . Equations 1, 2, and 3, together with the initial state  $x_{0|0}^j$ , define the EnKF recursion for the problem of interest here.

It is useful to note some of the distinctive attributes of the ensemble version of the Kalman filter. The EnKF can produce non-Gaussian sample distributions of the states  $x_{t|t-1}^j$  and  $x_{t|t}^j$ . There is no need to linearize or otherwise approximate the state transition function  $f_t(\cdot, \cdot, \cdot)$  or to assume that the random states and inputs are Gaussian. However, for large sample sizes, the EnKF replicates tend to converge to the true conditional densities  $p[x_t|y_{0:t-1}]$  and  $p[x_t|y_{0:t}]$  only when all states and measurements are jointly Gaussian. If this condition is not met, the filter's sample densities and moments are only approximations.

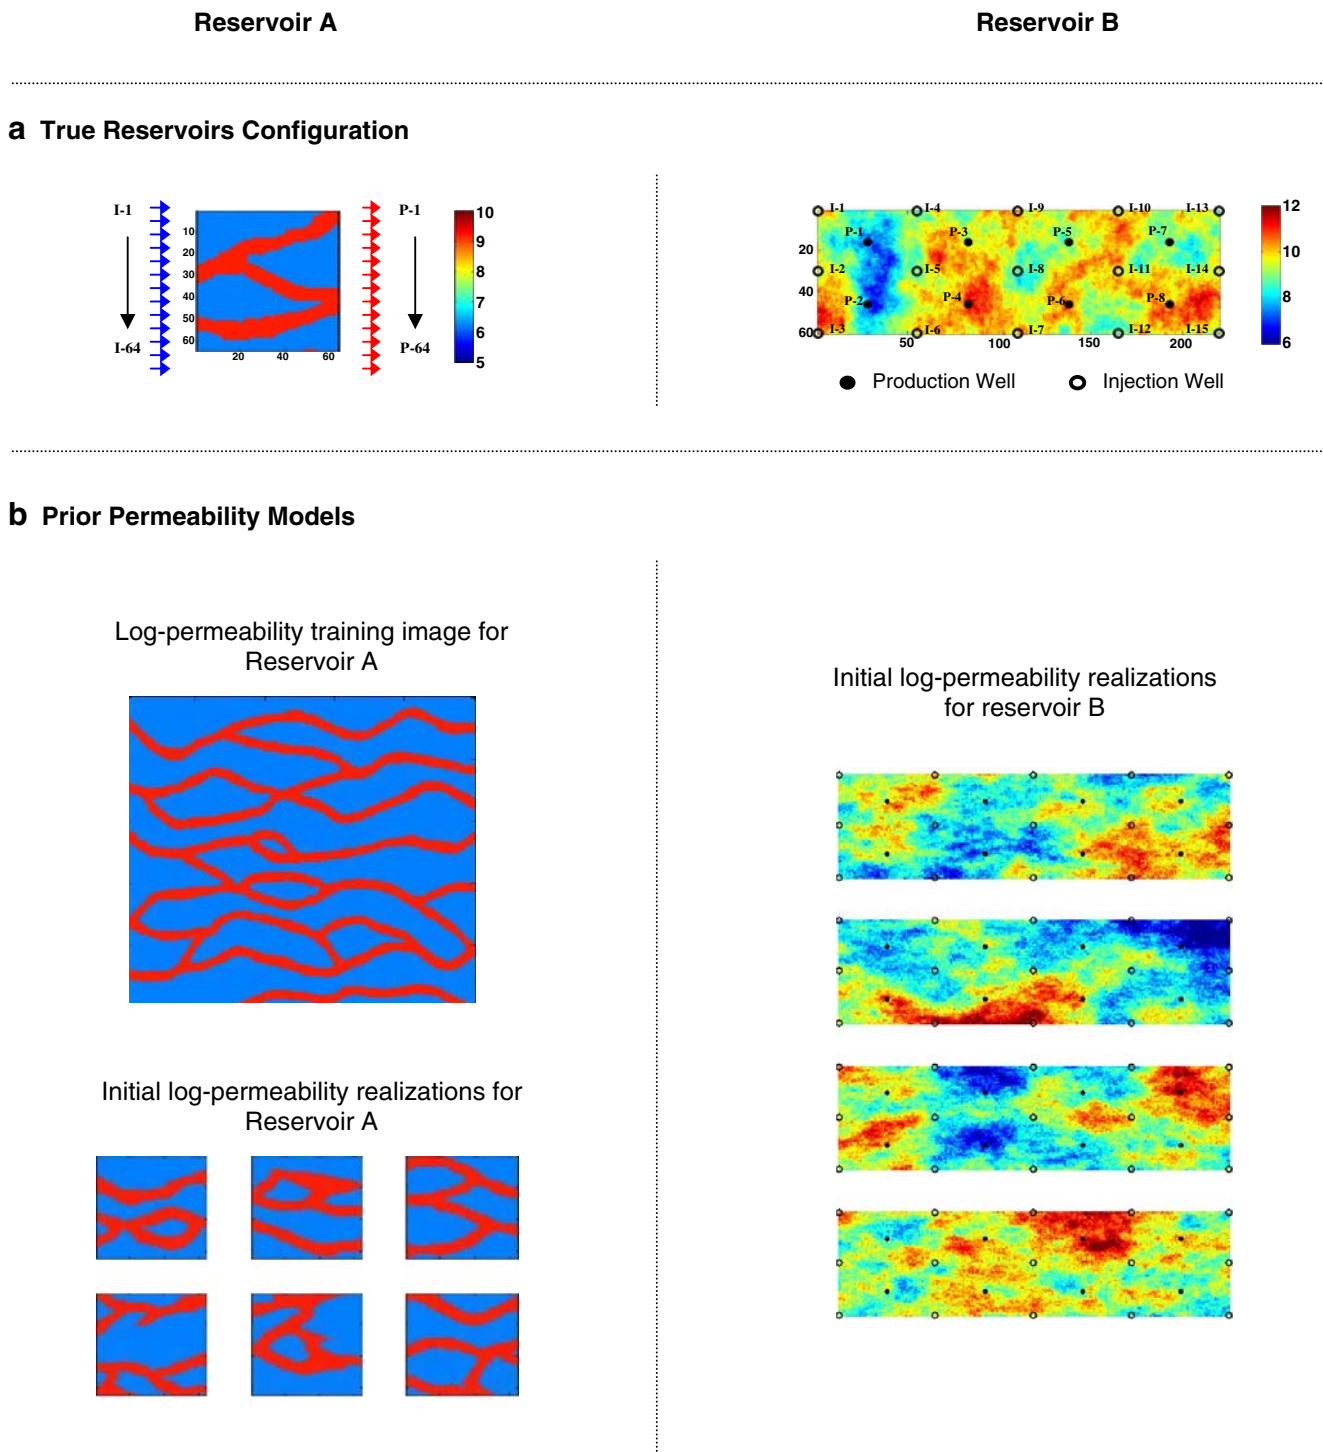
The covariance inversion operation in Eq. 2 is computationally expensive for large problems and can be ill conditioned. This can be avoided by applying a pseudo-inverse procedure based on a singular value decomposition, as proposed in [13]. However, for large problems, even the singular value decomposition approach can be very slow. In general, more efficient implementations of the update step may be necessary for large problems. In this paper, we achieve efficiency by using the smaller-state dimension provided by a DCT parameterization.

### 3 Formulation of the example problems

#### 3.1 Reservoir descriptions

We use two waterflooding examples to demonstrate the effectiveness of the approach presented in this paper. Figure 1a shows the reservoir configurations for these two experiments. Reservoir A has a line drive with 64 injectors (on the left) and 64 producers (on the right), while reservoir B is produced using a pattern drive as shown in Fig. 1a. General information about the simulations for each reservoir is provided in Table 1. The simulations are performed with the ECLIPSE reservoir simulator, which is set up for a two-phase (oil and water) black oil flow. The total simulation time is 1,080 days, divided into 12 intervals of 90 days.

The injection wells are operated with specified flow rates, while the production wells are operated under bottom-hole pressure control. The well settings are specified without any intent to maximize oil recovery. Rather, the dynamic well controls are used to excite the reservoir and increase the information content of its response. Each injection/production scenario lasts for 90 days before it is replaced with a new set of well settings. Details about production scenarios in each experiment are summarized in Tables 2 and 3.



**Fig. 1** Reservoirs configuration with the “true” permeability fields (a) and prior permeability models (b) for experiments A1, A2, B1, and B2

In this study, the only source of uncertainty considered is the permeability, which is treated as a spatially variable random field. Initial and boundary conditions are assumed to be known perfectly, and dynamic model errors are assumed to be negligible. In situations where these assumptions may not hold, additional error sources may be included in the ensemble-filtering process.

In its update step, the EnKF uses measurements of bottom-hole pressure observations in injection wells and oil and water flow rate measurements in production wells. In each experiment, the “true” injection bottom-hole pressures and volumetric production quantities are generated by running the simulator with a specified “true” permeability field. Uncorrelated zero-mean random measurement errors



**Table 1** Simulation and assimilation information for reservoirs A and B

	Reservoir A	Reservoir B
Simulation information		
Simulation	Two-phase oil–water	Two-phase oil–water
Simulation time (days)	12×90=1,080	12×90=1,080
Grid blocks	64×64×1	60×220×1
Grid block dimensions (m)	10×10×10	10×10×10
Geometry	2D-Areal	2D-Areal
Porosity	0.2 (constant)	0.2 (constant)
Connate water saturation	0.1 (uniform)	0.1 (uniform)
Initial pressure	3,000 psi (uniform)	3,000 psi (uniform)
Total injected water	1 PV	1 PV
Number injection wells	64	15
Number production wells	64	8
Injection well control	Water Rate	Water Rate
Production well control	BHP	BHP
Observation information		
Observation intervals	90 days	90 days
Injection wells	bhp	bhp
Noise standard deviation	20 psi	100 psi
Production wells	Oil and water rates	Oil and water rates
Noise standard deviation	20 sbpd	100 sbpd

are added to these “true” pressures and flow rates. The standard deviations of the random measurement errors for each experiment are listed in Table 1.

The classical EnKF update can potentially give unphysical saturation values outside the range [0,1]. To avoid this problem, our Kalman filter updates the following inverse error function transform of saturation, which is distributed over  $(-\infty, +\infty)$ :

$$x^* = \text{erf}^{-1}(2x - 1) \quad (4)$$

After each filter update, the saturation is computed by applying the inverse transformation to obtain updated saturations for future simulation runs.

### 3.2 Reservoir permeabilities

Reasonable probabilistic models for the permeability field need to be constructed before a successful EnKF implementation can be obtained [18]. In an ensemble implementation, a probabilistic model is simply a procedure for generating realistic unconditional permeability replicates. In practice, the generation procedure should use as many sources of information as possible and should be based on sound geological and geostatistical principles.

In our synthetic experiments, unconditional permeability realizations were generated for two structurally different formations, identified as reservoirs A and B. These correspond to a predominantly channelized environment (fluvial deposition) and a prograding nearshore basin, respectively. Reservoir A is the subject of our experiments A1 and A2, while reservoir B is considered in experiments B1 and B2. The “true” permeability fields shown in Fig. 1a were used to generate noisy measurements and to generate “true” saturation and pressure values. The set of true permeability, pressure, and saturation states is used to assess history-matching performance. The true permeability field for reservoir A is a random realization drawn from the training image in Fig. 1. The “true” permeability field shown for reservoir B corresponds to the top layer of the SPE10-comparative model [19]. In the original SPE10 model, the contrast between the highest and lowest permeability values is about six orders of magnitude with the lowest permeability values in the order of  $10^{-3}$  mD. We rescaled the SPE10 permeability values to the range  $10^2$ – $10^5$  mD to avoid numerical stability issues of the forward simulations and speed the convergence of the linear solver in the reservoir simulator.

We use two different geostatistical modeling methods to generate prior permeability replicates that are structurally similar to the two true fields. For reservoir A, the geologic features common to all members of the training set are defined by the channelized training

**Table 2** Production scenarios for reservoir A

Production period	IG1	IG2	IG3	IG4	PG1	PG2	PG3	PG3
	Injection volume (PV)				Production BHP(psi)			
A. 1–180 days	1/6	1/6	1/6	1/6	2,950	2,950	2,950	2,950
B. 181–360 days	1/6	0	0	1/6	2,950	3,000	3,000	2,950
C. 361–540 days	0	1/6	1/6	0	3,000	2,950	2,950	3,000
D. 541–720 days	0	1/6	1/6	0	3,000	2,950	2,950	3,000
E. 721–900 days	1/6	0	0	1/6	2,950	3,000	3,000	2,950
F. 901–1,080 days	1/6	1/6	1/6	1/6	2,950	2,950	2,950	2,950

IG1 Wells I-1 to I-16, IG2 wells I-17 to I-32, IG3 wells I-33 to I-48, IG4 wells I-49 to I-64, PG1 wells P-1 to P-16, PG2 wells P-17 to P-32, PG3 wells P-33 to P-48, PG4 wells P-49 to P-64

**Table 3** Production scenarios for reservoir B

Production period	IG1	IG2	IG3	PG1	PG2
	Injection volume (PV)			Production BHP(psi)	
A. 1–180 days	1/6	1/6	1/6	2,500	2,500
B. 181–360 days	0	1/6	0	2,500	3,000
C. 361–540 days	1/6	0	0	2,500	3,000
D. 541–720 days	0	0	1/6	3,000	2,500
E. 721–900 days	1/6	1/6	1/6	2,500	2,500
F. 901–1,080 days	0	1/6	0	3,000	2,500

*IG1* Injection wells I-1, I-4, I-7, I-10, and I-13 all on the top row in Fig. 1, *IG2* injection wells I-2, I-5, I-8, I-11, and I-14 all in the middle in Fig. 1, *IG3* injection wells I-3, I-6, I-9, I-12, and I-15 all on the bottom row in Fig. 1, *PG1* production wells P-1, P-3, P-5, and P-7 in Fig. 1, *PG2* production wells P-1, P-4, P-6, and P-8 in Fig. 1

image shown in Fig. 1b. This training image has  $250 \times 250 \times 1$  grid blocks comprising two facies types: low-permeability background shale and high-permeability sandstone channels. The background shale permeability is assumed to be 500 mD, while the embedded high-permeability sandstone channels have a permeability value of 10,000 mD.

The permeability realizations shown in Fig. 1b are generated using the multiple-point geostatistical algorithm *SNESIM* [20, 21]. Each realization is discretized over a  $64 \times 64 \times 1$  ( $640 \times 640 \times 10$  m) grid block system. The varying shape and geometry of the channel facies in these realizations are the major source of uncertainty. The highly structured nature of the facies distribution in these realizations suggests a distinctive preferential flow displacement pattern. In addition, a high level of correlation (redundancy) is observed in description of these facies, suggesting that the field could be represented much more efficiently if instead of the block values a more appropriate description was adopted (this will be discussed in later sections).

Figure 1b also shows sample realizations for reservoir B, which lies in a different type of geological formation. The reservoir B samples are from a prior permeability ensemble generated with the two-point geostatistical algorithm *SGSIM* [2, 21]. Each realization is discretized over a grid of  $60 \times 220 \times 1$  blocks, with each block size of  $10 \times 10 \times 10$  ft. The reservoir B field is based on an exponential variogram model with assumed correlation lengths of 100 and 50 grid blocks in the  $x$  and  $y$  directions, respectively. This correlation length ratio is not based on the true SPE10 model. In fact, the correlation length in the top layer of the SPE10 model appears to be higher in the  $y$  direction than in the  $x$  direction. However, the permeability variance is selected to cover the expected range of permeability values in the SPE10 dataset.

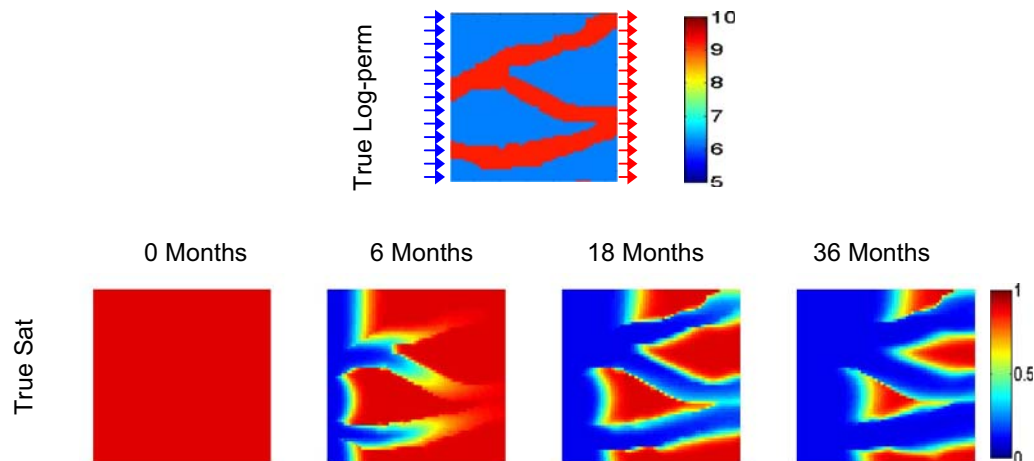

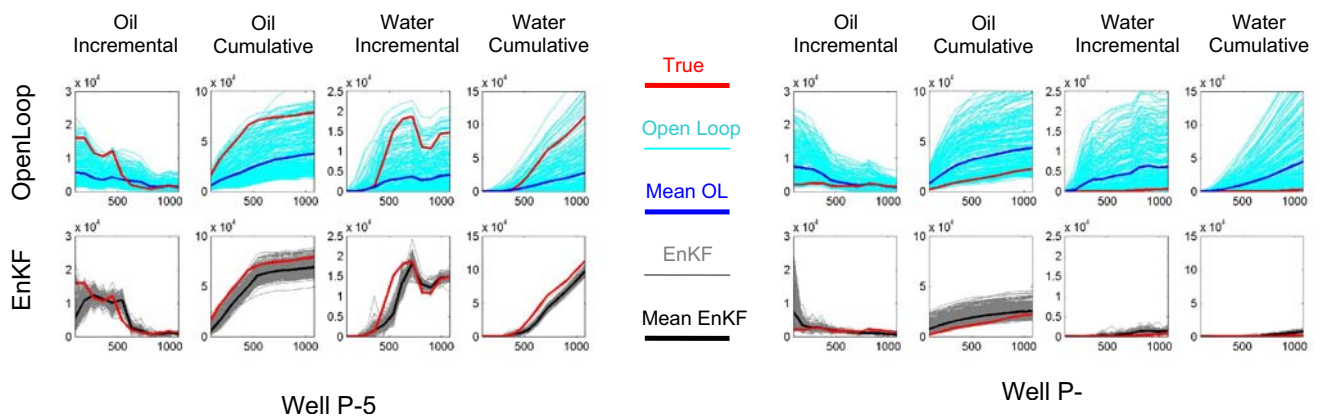
#### 4 History-matching results for grid block-oriented estimation

In this section, we present the EnKF history-matching results for estimating grid block values of pressure, saturation, and permeability. We consider two experiments based on reservoirs A and B, respectively.

Figure 2 summarizes the EnKF history-matching results for grid block-oriented estimation with reservoir A (Experiment A1). Figure 2a shows the true log permeability and its corresponding saturation profiles after 0, 3, 18, and 36 months. The prior permeability for reservoir A follows a strictly bimodal distribution. Moreover, saturation profiles also tend to show a bimodal behavior except near the moving waterfront. This bimodality can make the estimation of facies distribution a challenging task for the linear update in Eq. 2, which is known to be optimal only under jointly Gaussian priors and observations [12]. Nevertheless, the strong correlations between states and observations help to insure that the covariance-based Kalman updates are in the right direction, provided that correlation information is correctly specified.

Figure 2b shows the mean over the ensemble of the updated permeability and saturation profiles for reservoir A at the EnKF measurement (update) times. While the estimated permeability field has captured the main large-scale permeability trends present in the true field, the channels are somewhat disconnected in the interior of the field. The EnKF update appears to have difficulty capturing permeability connectivity (channel structure), although the prior replicates are all connected. A contributing factor to this loss of connectivity is, at least in part, the absence of observations away from the boundaries where the wells are located. Covariances derived from the prior permeability replicates are used to interpolate between the dynamic observations at the ends of the domain where wells (and measurements) are located. Estimates in the interior far from the wells tend to have higher uncertainties (not shown).

A close inspection reveals that the locations of the “channel” boundaries produced by the history-matching procedure are slightly offset. For example, in the final permeability estimate (rightmost plot), the top channel is closer to the top boundary than it is in the true field. As a result, the estimated saturation profile in the top right region shows no residual oil, in contrast to the true saturation profile that reveals that a considerable amount of oil remains. A similar trend is observed in the low-permeability region in the middle of the reservoir between the bottom and middle channels. Because the estimated low-permeability feature is too narrow compared to the true permeability field, the leftover oil in the middle of the reservoir is underestimated. While this visual analysis of estimation results is not a

**a Experiment A1 “True” log-permeability and saturations****b Experiment A1 mean log-permeability and saturation estimates****c Experiment A1 production forecast**

**Fig. 2** Results for experiment A1 (estimating grid block properties with the EnKF for reservoir A): **a** the true log-permeability field and snapshots of its corresponding synthetic true saturation profiles generated by ECLIPSE reservoir simulator; **b** pixel-based EnKF mean

log permeability and saturation estimates at selected update times; **c** production observation (*red*), open loop forecast and its mean (*cyan* and *blue*), and EnKF forecast and its mean (*gray* and *black*)

precise performance measure, it provides a useful qualitative assessment, particularly when considered in conjunction with more quantitative measures.

Particularly relevant quantitative performance measures for history matching can be derived from the differences between the predicted and measured oil and water production (prediction errors) at different time steps. When the predicted values are open looped (i.e., when they are obtained without history matching), the root-mean-squared prediction error defines a reference performance level. If the history-matching procedure is helpful, then predictions produced by the EnKF should give a smaller root-mean-squared error. Moreover, the filter performance should improve over time, as more measurements are collected. This improvement should be accompanied by a decreased ensemble spread around an ensemble mean that is close to the true value.

The production forecast curves for reservoir A are shown in Fig. 2c. Plots for incremental and cumulative oil (first and second columns) and water (third and fourth columns) productions are shown for wells P-5 and P-30. These wells correspond to two separate regions (inside and outside the channel facies) in the production end of the field (on the right). Well P-30 is located in a low-permeability area and does not show any sign of water breakthrough. This is in contrast to well P-5, which is inside a high-permeability region and experiences early water breakthrough. The first and second rows in each plot show the open loop and EnKF forecast results, respectively. Figure 2c indicates that oil production predictions improve with time, suggesting that the updates continuously improve the predictive power of the reservoir model after each analysis step. This can also be observed from the improvement in the permeability field (Fig. 2b).

While the rate of oil production generally decreases in time, water production is initially close to 0 until breakthrough occurs and increases thereafter. As a result, initial deviations between predicted and true production curves are larger for oil and smaller for water. The third and fourth columns of Fig. 2c indicate that the water production has less ensemble spread and better prediction quality than oil production even at later times. This can be explained by observing that the water breakthrough occurs after a few EnKF updates have already been carried out, when a more accurate model is available for generating water production predictions.

It is important to note that although the later updates improve forecast quality, the cumulative production plots (second row) do not converge to the observed values because of the initial forecast errors. This highlights an important implication in production optimization. In general, early injection and production control decisions are based on poorer knowledge of reservoir properties and are therefore

suboptimal. This implies that more frequent monitoring and model updating should be performed at early stages of the reservoir's lifetime.

Figure 3 summarizes the EnKF history-matching results for grid block-oriented estimation with reservoir B (experiment B1). The true log permeability and saturation profiles after 0, 3, 18, and 36 months are shown in Fig. 3a, while the ensemble mean permeability and saturation updates are shown in Fig. 3b. The third panel in the figure (Fig. 3c) shows the production plots for wells P-1 and P-8 (see Fig. 1). P-1 is a production well located in a low-permeability region (refer to Fig. 1), while P-8 is in a high-permeability zone.

From an estimation viewpoint, experiment B1 differs in several ways from experiment A1. The ratio of observation points to the total number of grid blocks for reservoir B is much less than for reservoir A, i.e., 23 of 13,200 compared with 128 of 4,096. On the other hand, the reservoir B observation points are distributed more evenly throughout the field and cover a larger area. This observing strategy can be expected to provide more information content (better observability) than the one used for reservoir A. Another major difference between the two experiments is the more nearly Gaussian character of the prior permeability distribution in experiment B1 (although the saturation profile still shows bimodal characteristics). This nearly Gaussian behavior is more compatible with the implicit assumptions of the EnKF.

It is clear from Fig. 3b that the filter is able to distinguish high- and low-permeability regions in the field. The updates at earlier times appear to have only found the two major high- and low-permeability areas. However, further details of the permeability field are retrieved at later update times. As expected, permeability trends in the vicinity of observation points are estimated more accurately than in regions between the wells.

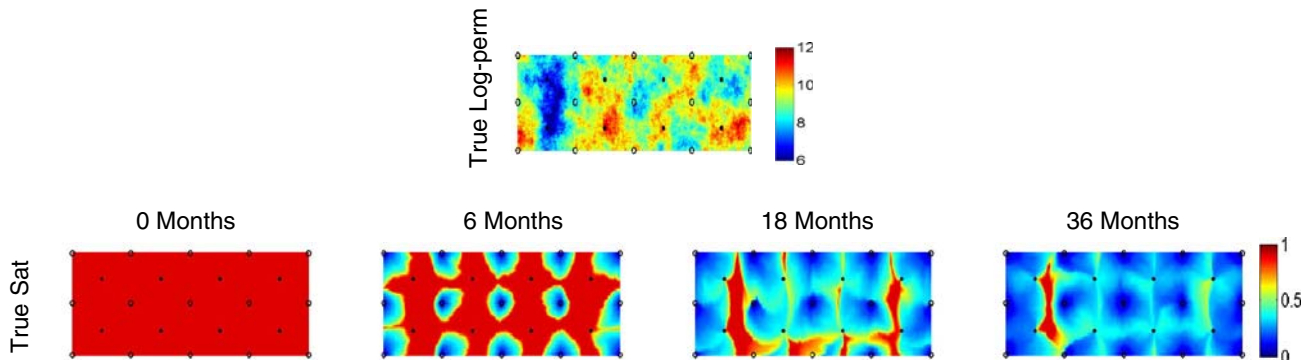
The estimated saturation results after 18 and 36 months appear to be very close to the true saturation profiles at those times. However, the saturation profile at 6 months is very different from true saturation. This is explained by the inaccurate permeability estimates at the initial update steps.

Production plots in the second row of Fig. 3c show the incremental improvement in the quality of the EnKF estimates (first and third columns) compared to open loop results (first row). It is clearly seen that the EnKF ensemble spread is substantially reduced in comparison with the open loop production ensemble. The mean EnKF production forecast approaches the true production curve, indicating an improvement in estimation accuracy.


The results presented in Figs. 2 and 3 confirm the effectiveness of the EnKF for history matching. However, while updating grid block values of states and parameters may be computationally reasonable for small problems, there are a few drawbacks to this approach. First, the



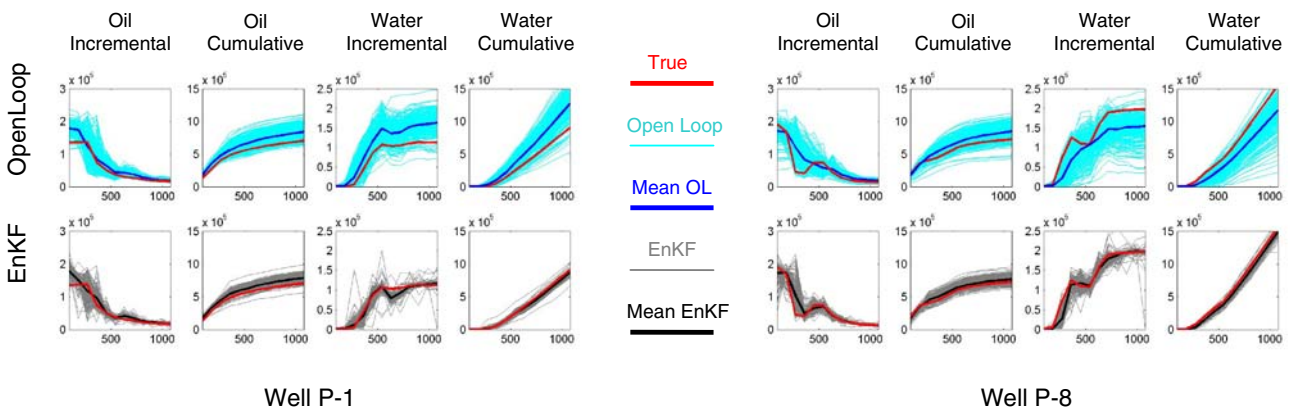
### a Experiments B1 “True” log-permeability and saturations



### b Experiment B1 mean log-permeability and saturation estimates



### c Experiment B1 production forecast



**Fig. 3** Results for experiment B1 (estimating grid block properties with the EnKF for reservoir B): **a** the true log permeability field and snapshots of its corresponding synthetic true saturation profiles generated by ECLIPSE reservoir simulator; **b** pixel-based EnKF mean

log-permeability and saturation estimates at selected update steps; **c** production observation (*red*), openloop forecast and its mean (*cyan* and *blue*), and EnKF forecast and its mean (*gray* and *black*)

computational cost of applying the update to individual grid block values can become considerable for large problems ( $\sim 10^6$ – $10^7$  grid blocks). Second, neighboring blocks have strongly correlated values both within and across states. In

this case, grid block-oriented estimation can be both wasteful and geologically unrealistic. For example, much work can be expended to obtain locally disconnected estimates for problems with continuous facies (Fig. 2b). It

is possible to obtain a better posed and less computationally demanding history-matching problem by changing the way that spatial variability is represented. This is the topic of the next section.

## 5 Parameterization with the discrete cosine transform

This section provides a brief overview of the DCT, with an emphasis on its use for parameterizing spatially variable reservoir properties such as permeability. The DCT was first introduced in [22] for signal decorrelation and has been widely used for image compression after its introduction [23, 24]. A detailed discussion of the use of DCT for parameterization and its properties is given in [8].

The discrete one-dimensional DCT of a signal  $u(n)$  of length  $N$  can be written as [23] (using notation in [8]):

$$v(\mu) = \alpha(\mu) \sum_{n=0}^{N-1} u(n) \cos \left[ \frac{\pi(2n+1)\mu}{2N} \right] \quad 0 \leq \mu \leq N-1$$

$$\alpha(\mu) \equiv \begin{cases} \sqrt{\frac{2}{N}} & \mu = 0 \\ \sqrt{\frac{1}{N}} & 1 \leq \mu \leq N-1 \end{cases} \quad (5)$$

The inverse DCT is then given by:

$$u(n) = \sum_{\mu=0}^{N-1} \alpha(\mu) v(\mu) \cos \left[ \frac{\pi(2n+1)\mu}{2N} \right] \quad 0 \leq \mu \leq N-1 \quad (6)$$

Extension of the above equations to higher dimensions is given in [23, 24]. However, the separability property of DCT can be exploited to achieve computational savings by applying the one-dimensional transform in each direction (see Appendix). For a two-dimensional image, the basis functions (images) are arranged according to their orientation and level of detail (frequency content) in a descending order from upper left to lower right (see Fig. 4a). Depending on the desired level of details in the approximation, more high frequency components (lower right modes) are included.

Suppose that the original image is represented by a grid of  $N$  equal size blocks (or pixels), with a distinct variable value (e.g., grayscale intensity or permeability) assigned to each block. When a DCT parameterization is used, the variable value in each block is represented as the sum of  $K$  of the cosine basis functions. Because the basis functions are fixed, the  $K$  DCT-weighting coefficients constitute an alternative way to describe the image (i.e., once these coefficients are specified the block values can be reconstructed). If  $K=N$ , the DCT coefficient representation can reproduce the image perfectly. When  $K < N$ , some information is usually lost, but the DCT coefficient representation is more compact. However, in many cases, the compressed

DCT representation of the image is still quite good, even for  $K$  much less than  $N$ .

The DCT can be concisely summarized with a pair of linear equations. The inverse transformation from  $K$  DCT basis function coefficients ( $v_t$ ) to  $N$  grid block values ( $u_t$ ) is given by:

$$u_t = \mathbf{A}^T v_t \quad (7)$$

The forward transformation from  $N$  grid block values ( $u_t$ ) to  $r$  DCT basis function coefficients ( $v_t$ ) is given by:




$$v_t = \mathbf{A}^T u_t \quad (8)$$

Where  $\mathbf{A}^T$  is an  $N$  by  $K$  matrix with columns consisting of the  $K$ -retained DCT basis functions. Equation 7 follows from Eq. 8 and the fact that the DCT basis functions are orthonormal. In practice, the DCT coefficients are computed with a sequence of one-dimensional fast Fourier transforms [25].

Figure 4b shows an example to illustrate the compression capabilities of the DCT. An image of MIT's symbolic dome is shown with a truck on top of it. The logarithm of the magnitude of the DCT transform coefficients of this image is shown next to it. It is clear that only a few of the DCT coefficients in the upper left corner have a large magnitude (red areas), implying that the corresponding basis functions are responsible for most of the variability in the image. Note that the upper left corner DCT basis functions are those with the largest periods (or lowest wave numbers), as indicated in Fig. 4a. Figure 4c shows a series of approximations for the original image, obtained by keeping 0.1, 1, 5, and 10% of the total number of DCT basis function coefficients. A reasonably good approximation of the image is obtained by including only 5% of the original coefficients (the truck on the top of the dome is clearly visible when only 5% of the coefficients are kept). This powerful compression property of the DCT transformation is responsible for its widespread use in JPEG image compression standards [24].

While size reduction (efficiency) is a common theme for image compression and parameterization of spatially distributed reservoir properties, the two problems are fundamentally different. In image compression, the weighting coefficient assigned to each basis function is known because the DCT coefficients can be computed directly from the known image. Consequently, it is easy to rank all  $N$  basis coefficient values and select the set of  $K$  basis vectors that make the greatest contribution (i.e., have the largest coefficients).

The parameterization problem is different because the image is not known and the  $K$ -leading weighting coefficients cannot be identified in advance. In this case, it is common to retain the  $K$  basis function terms that are most important in the aggregate, over some set of specified

**a Low-frequency discrete Cosine transform modes****b An example image and its log-DCT coefficients****c Low-rank approximations with increasing number of DCT modes**

**Fig. 4** Compression power of the discrete cosine transform: **a** sample discrete cosine transform modes (bases); **b** an example image (the famous MIT dome) with its log-DCT coefficients; **c** approximate representations of the example image with increasing number of included modes

images (called a training set) or, alternatively, to simply specify, a priori, the  $K$  basis function terms that are believed to be the most important. In the first case, the DCT is said to be “trained,” while in the second case, it is “untrained” [8]. In either case, the unknown basis function coefficients must be estimated from a history-matching procedure.

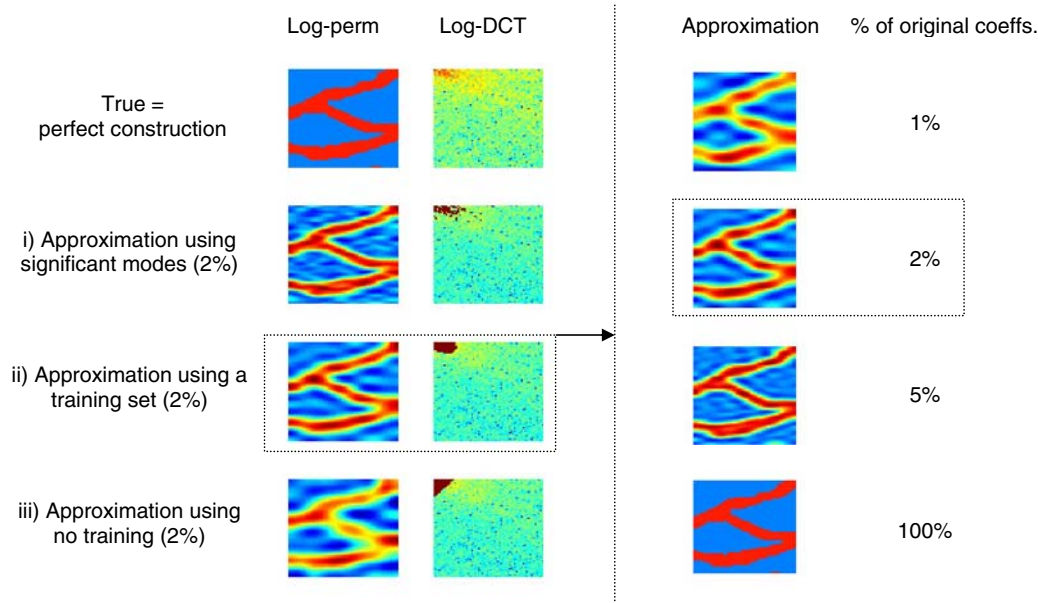
The trained parameterization is performed by applying the DCT on each ensemble member to obtain an ensemble of the DCT coefficients. The training is then performed by averaging the magnitude of the DCT coefficients over the

ensemble and retaining the largest  $K$  coefficients, which in turn identify the corresponding basis images. The computational complexity of the DCT is very favorable because of the separability of the transform that allows for low-dimensional computations.

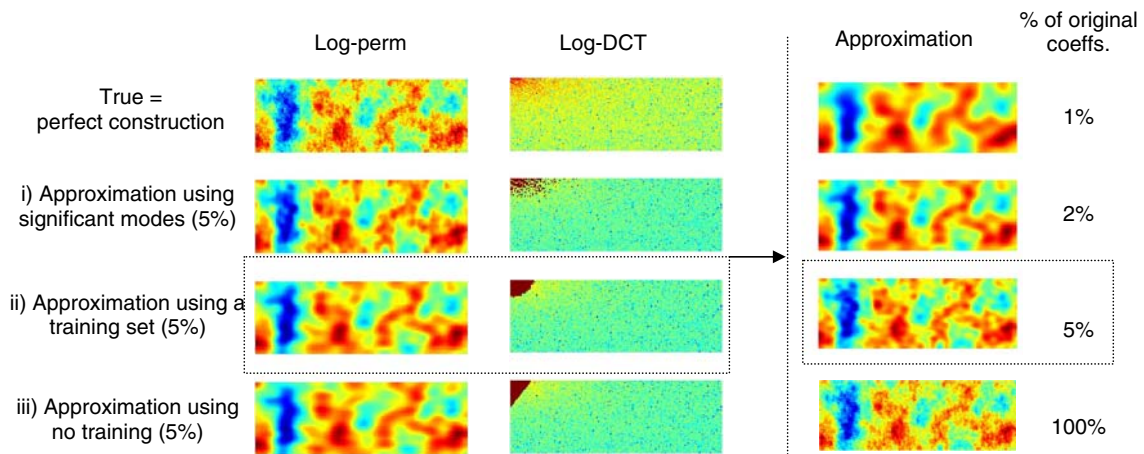
We can use the two permeability models discussed earlier to demonstrate alternative methods for selecting dominant DCT basis functions. Figure 5a and b show in row 1 the true permeabilities along with their DCT coefficients for reservoirs A and B, respectively.



### a DCT parameterization for Reservoir A



### b DCT parameterization for Reservoir B



**Fig. 5** Mode selection and approximation with the DCT: **a** flexible strategies for mode selection with or without prior training; **b** approximate representation of a sample permeability with increasing number of modes

Rows 2, 3, and 4 of Fig. 5a and b show approximations using, respectively, 2 and 5% of the original coefficients for reservoirs A and B. The retained basis functions are associated with (1) the most significant coefficients of the true permeability field (row 2), (2) a training set of prior permeability fields (row 3), or (3) an a priori selection that does not rely on a training set (row 4). In our subsequent history-matching examples, we use a training set to select the  $r$  retained basis function terms (option 2). This is a reasonable approximation to the approach that would likely be used in practice. However, we show in [8] that an untrained approach for selecting basis function terms may

be more robust than a trained approach because it is not misled by incorrect training images. The most appropriate choice depends on the application.

Figure 5a and b also show how image approximation quality changes with an increasing number of retained basis functions. It should be noted that perfect reconstruction (synthesis) of the original permeability field is possible if no truncation is carried out ( $r=n$ ). For our history-matching examples, we keep 2% of the original set of trained basis function coefficients for the channelized model of reservoir A ( $K=0.02N$ ) and 5% for the SPE model of reservoir B ( $K=0.05N$ ).



## 6 History-matching results for the discrete cosine transform parameterization

In this section, we revisit the block-oriented history-matching experiments described in Section 4 to illustrate how the DCT method can provide low-dimensional parameterizations suitable for history matching. Various history-matching parameterization methods have been investigated in the past [6–10]. Typically, these methods have only been used to reduce the number of unknown geological properties (e.g., permeabilities and/or porosities) adjusted in the history-matching procedure. In this case, the grid block values of pressure and saturation states are obtained from the reservoir simulator, using grid block property values derived from the parameterization.

When the EnKF is used for history matching, the forecast pressure and saturation states are derived from the simulator, but the updated pressure and saturation are derived from the Kalman update of Eq. 2. Therefore, history matching with the EnKF lends itself to more parameterization options. In particular, an efficient parameterization can be used to reduce the number of updated pressure and saturation states as well as the number of unknown permeability or porosity states.

The DCT parameterization of the EnKF update replaces the forecast grid block values of pressure, saturation, and permeability (a total of  $3n$  states) contained in  $x_{t|t-1}$  by their corresponding DCT coefficients (a total of  $3r$  states). These DCT coefficients are contained in  $v_{t|t-1}$  and computed from Eq. 8. The forecast DCT coefficients are updated with an appropriate version of Eq. 2 (expressed in terms of the DCT state vectors  $v_{t|t-1}$  and  $v_{t|t}$  rather than the grid block-oriented state vectors  $u_{t|t-1}$  and  $u_{t|t}$ ). The updated grid block values of pressure, saturation, and permeability are then constructed from the updated DCT coefficients and the specified DCT basis functions, using Eq. 5.

The implications of the reduction in state vector dimension provided by the DCT parameterization can be significant for large problems. If the state dimension is large the EnKF update scheme can be very time consuming. With our proposed parameterization of all the states (pressure, saturation, and geological properties), the size of the EnKF update is considerably reduced, making the filtering approach much more practical for realistic large-scale applications.

In Section 4 (experiments A1 and B1), we estimated the values of the permeability, pressure, and saturation fields for each individual grid block. In general, this approach can result in discontinuous estimates that are geologically inconsistent with prior knowledge. The loss of structural continuity observed in the block-oriented approach can be mitigated if we use a truncated parameterization such as DCT that emphasizes large-scale (low wave number)

features over small-scale (high wave number) features. In this case, the estimated permeabilities are more correlated over space, and continuous channels are more likely to be properly identified.

These properties of the DCT are illustrated in Figs. 6 and 7, which show the history-matching results from the EnKF with DCT parameterization for reservoirs A and B (experiments A2 and B2), respectively. Figures 6a and 7a show the true log permeability and its corresponding saturation profiles after 0, 3, 18, and 36 months for the two reservoirs. Figures 6b and 7b show the updated ensemble mean permeability and saturation for each case. Finally, Figs. 6c and 7c compare the incremental and cumulative oil and water production forecast with the open loop and true values for the same production wells as in Figs. 2c and 3c.

It is clear from a comparison of Figs. 2b and 6b that the estimated permeability and saturation in Fig. 6b are smoother than those in Fig. 2b. The smoothness and continuity introduced by the DCT parameterization appear to have improved the connectivity between the channel facies. Another qualitative observation is that the saturation estimates in Fig. 6b seem to be slightly better than those in Fig. 2b. In particular, the residual oil values observed in the top right and bottom right regions of the reservoir after 36 months are closer to the true saturation in Fig. 6b than they are in Fig. 2b.

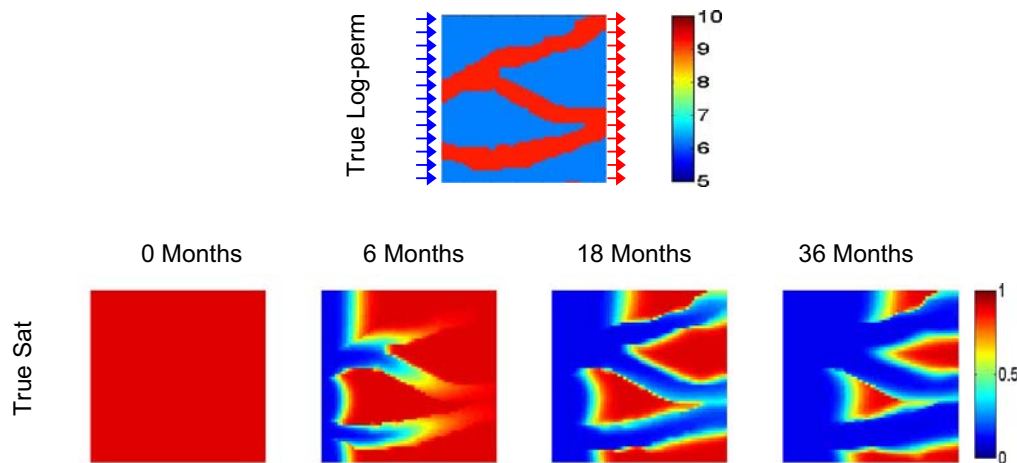
The production plots in Fig. 6c for the DCT parameterization are almost identical to those in Fig. 2c for the block-oriented alternative. This is not surprising as the permeability and saturation profiles are only marginally different in the two experiments, apart from the smoothness in the case of parameterized estimates. Our analysis of the production plots in Fig. 6c also applies to Fig. 2c as there is no visible difference between the results.

Figure 7 shows the results for reservoir B when DCT parameterization is used to concisely describe the saturation, pressure, and permeability fields. In this case, the results are smoother and more continuous than the grid block-oriented history-matching results shown in Section 4 (see Fig. 3). The pressure and saturation plots are in better agreement because of the continuity and smoothness of them even in the grid-based estimation.


These results imply that estimation of the fine details in the parameterization does not contribute significantly to the overall estimation results. In other words, the observations do not provide the information needed to resolve fine-scale details, and there is no added value in including those details in the estimation problem. An attempt to estimate grid block state and parameter values increases the computational cost of the EnKF update while possibly giving spurious and geologically unexpected outcomes.

The fact that our reduced state DCT parameterization gives comparable or even better results than a grid block-

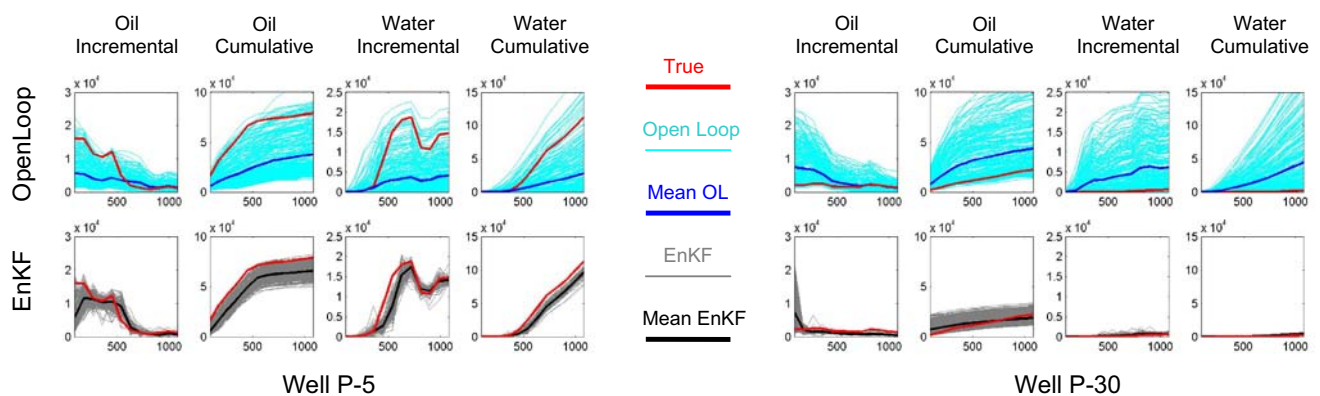
### a Experiment A2 “True” log-permeability and saturations



### b Experiment A2 mean log-permeability and saturation estimates



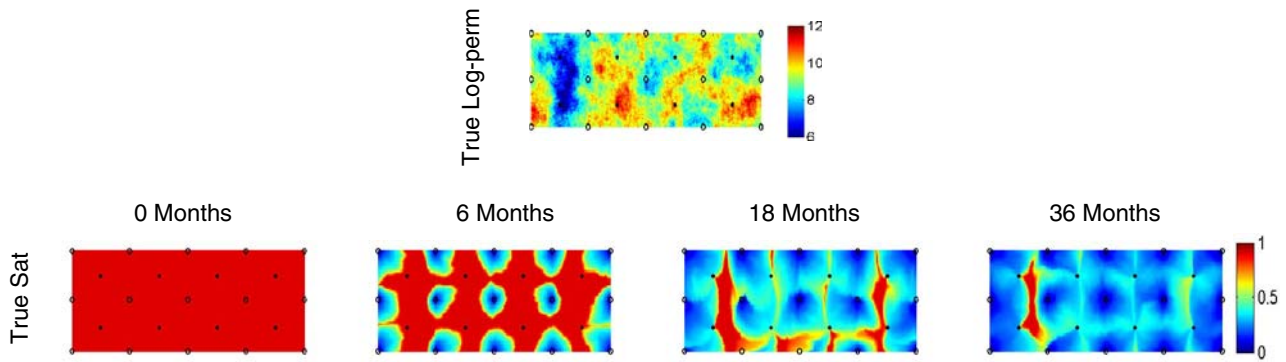
### c Experiment A2 production forecast




**Fig. 6** Results for experiment A2 (estimating DCT coefficient with the EnKF for reservoir A): **a** the true log permeability field and snapshots of its corresponding synthetic true saturation profiles generated by *ECLIPSE* reservoir simulator; **b** the mean EnKF log permeability and

saturation estimates at selected time-steps using the DCT parameterization; **c** production observation (red), open loop forecast and its mean (cyan and blue), and EnKF forecast and its mean (gray and black)

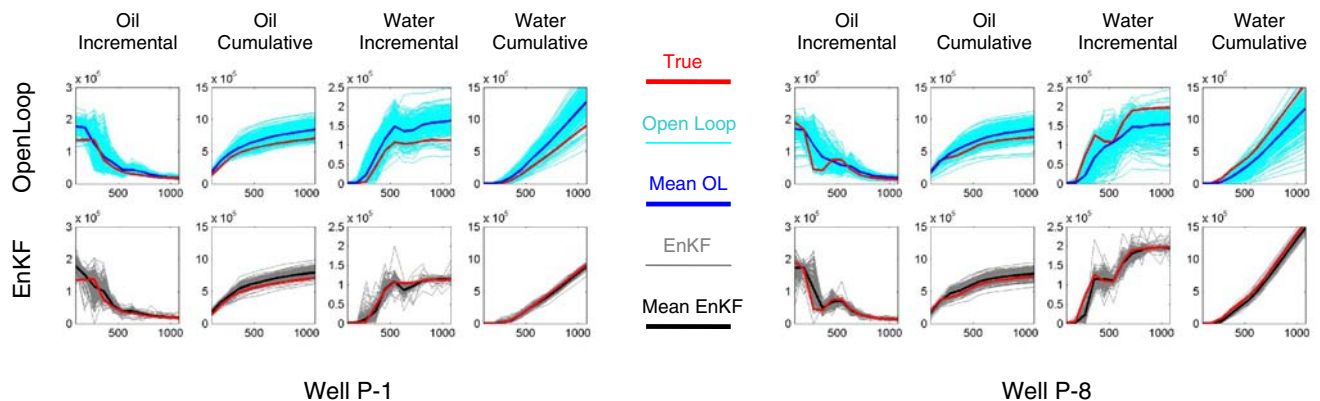
### a Experiments B2 “True” permeability and saturations



### b Experiment B2 mean permeability and saturation estimates



### c Experiment B2 production forecast



**Fig. 7** Results for experiment B2 (estimating DCT coefficient with the EnKF for reservoir A): **a** the true log permeability field and snapshots of its corresponding synthetic true saturation profiles generated by *ECLIPSE* reservoir simulator; **b** EnKF mean logperme-

ability and saturation estimates using the DCT parameterization; **c** production observation (*red*), open loop forecast and its mean (*cyan* and *blue*), and EnKF forecast and its mean (*gray* and *black*)

oriented approach has important computational implications. The dimensions of the DCT state vectors, which are  $82 \times 3 + 192 = 438$  and  $660 \times 3 + 31 = 2,011$  for reservoirs A and B, respectively, are only 3.5 and 5.1% of the

dimensions of the grid block-oriented states, which are  $64 \times 64 \times 3 + 192 = 12,480$  and  $60 \times 220 \times 3 + 31 = 39,631$ . A complete complexity analysis would, however, require that we account for the DCT forward and inverse transforma-

tions before and after the EnKF update, respectively. The DCT basis has several desirable properties that are discussed in [8]. Among these are the fixed (data independent) basis vectors that can be calculated and stored only once without any need for image statistics. More importantly, by virtue of the separability property of the DCT basis, the transformation can be implemented efficiently through one-dimensional row/column-wise operations (see Appendix).

The implementation of the DCT with the EnKF and its overall update complexity in our two-dimensional examples can be summarized as follows:

1. Compute and store one-dimensional basis vectors  $A_{N_x}^{(N_x, N_x)}$  and  $A_{N_y}^{(N_y, N_y)}$ , see Appendix (negligible complexity).
2. Run forecast simulations with the most current states and parameters.
3. Compute the forward DCT transform for each ensemble member using separability (see Appendix):  $(N_x^2 + N_y^2)N_{\text{ens}}$ .
4. Apply EnKF update using  $K$ -truncated coefficients [13]:  $\sim KN_{\text{ens}}^2$ .
5. Compute inverse DCT transform for each sample using EnKF updated DCT coefficients:  $(N_x^2 + N_y^2)N_{\text{ens}}$ .
6. Repeat steps 2–5 until all data are assimilated.

The total computational complexity for the update step using the DCT with above algorithm is approximately  $\left[2(N_x^2 + N_y^2)N_{\text{ens}} + KN_{\text{ens}}^2\right]$ . If the DCT parameterization is not used, the major computational cost of the grid block-based EnKF update is proportional to  $[(N_x \times N_y)N_{\text{ens}}^2]$  (see [13]). For values of  $N_x=N_y=100$ ,  $N_{\text{ens}}=100$ , and  $K=0.05 \times 100 \times 100 = 500$ , the computational complexity of the DCT-based EnKF update ( $=9 \times 10^6$ ) is at least an order of magnitude less than its grid block-oriented implementation ( $=1 \times 10^8$ ). It is important to note that above analysis did not take advantage of the additional savings associated with the sparsity of the truncated DCT. In addition, for constant parameters (such as permeability), the DCT transformation before each update step (except the first update) is not required as the values of the DCT coefficients from previous updates remain unchanged. As a result, the DCT complexity analysis presented above is a conservative estimate.

Finally, we note that while the above analysis suggests that the DCT-based implementation of the EnKF update is computationally more efficient to its grid block update alternative, when accurate grid-based flow simulation is used, the computational cost of the overall history-matching framework is dominated by forward simulation runs. The savings in the computation of the EnKF update is expected to become more important when faster/approximate simulation methods are used.

## 7 Conclusions

In this paper, we have presented an efficient history-matching approach by combining the EnKF with a flexible and effective parameterization method, the DCT. The parameterization with DCT is shown to provide several important benefits including better preservation of geological continuity (in particular for channelized facies distribution), more relevant and smaller number of unknowns to estimate (improved ill-posedness), and computational efficiency. Two waterflooding examples with structurally different geological formations show that estimation results with the DCT parameterization retrieve the main large-scale features of the underlying true permeability field at lower computational cost than when grid block values are estimated. This implies that the computational effort spent in estimating individual grid block properties is unnecessary and undesirable if an acceptable approximation of the states and variables can be obtained with a smaller number of parameterized coefficients.

The approximation introduced by truncating the DCT representation of permeability seems to provide a smoother and more continuous estimate of the original permeability field, as might be expected given the fact that smaller-scale details are omitted in the truncated DCT. However, our examples suggest that the estimated saturations are less sensitive to approximations introduced by the truncated DCT, at least in determining the global pattern of oil displacement. This has important implications for the optimization and control of oil production operations, where knowledge of the movement of the saturation profile is used to adjust reservoir control variables. Overall, it looks like a DCT version of the EnKF provides an accurate, efficient, robust, and flexible option for reservoir history matching, particularly for real-time operations.

**Acknowledgments** Funding for this research has been provided by Shell International Exploration and Production. The ECLIPSE flow simulator code was provided by Schlumberger-Geoscientific.

## Appendix

### Separability of DCT Basis

In this appendix, the separability of DCT basis and its computational advantages over nonseparable compression transforms such as KLT (also known as principle component analysis) is presented. While two-dimensional DCT is discussed here, the approach is readily extended to higher-dimensional DCT, and its advantage is even more significant for higher-dimensional datasets.



The two-dimensional forward DCT of an image  $u(n_{xy})$  can be written as [23]:

$$v(\mu, \omega) = \frac{\alpha(\mu)\alpha(\omega)}{\sqrt{N_x N_y}} \sum_{m=0}^{N_x-1} \sum_{n=0}^{N_y-1} u(n_x, n_y) \cos \left[ \frac{\pi(2n_x+1)\mu}{2N_x} \right] \cos \left[ \frac{\pi(2n_y+1)\omega}{2N_y} \right] \quad (9)$$

$$\alpha(j) \equiv \begin{cases} \frac{1}{\sqrt{2}} & j = 0 \\ 1 & \text{otherwise} \end{cases}$$

where  $\mu=0, \dots, N_x-1$  and  $\omega=0, \dots, N_y-1$ .

Similarly, the inverse two-dimensional DCT transform is then given as [23, 24]:

$$u(n_x, n_y) = \frac{2}{\sqrt{N_x N_y}} \sum_{\mu=0}^{N_x-1} \sum_{\omega=0}^{N_y-1} \alpha(\mu)\alpha(\omega)v(\mu, \omega) \cos \left[ \frac{\pi(2n_x+1)\mu}{2N_x} \right] \cos \left[ \frac{\pi(2n_y+1)\omega}{2N_y} \right] \quad (10)$$

where  $n_x=0, \dots, N_x-1$  and  $n_y=0, \dots, N_y-1$ .

The separability property of the two-dimensional DCT can be illustrated as:

$$v(\mu, \omega) = \sqrt{\frac{2}{N_x}} \alpha(\mu) \sum_{n_x=0}^{N_x-1} \left\{ \sqrt{\frac{2}{N_y}} \alpha(\omega) \sum_{n_y=0}^{N_y-1} u(n_x, n_y) \cos \left[ \frac{\pi(2n_y+1)\omega}{2N_y} \right] \right\} \times \cos \left[ \frac{\pi(2n_x+1)\mu}{2N_x} \right] \quad (11)$$

The inner summation is an  $N_y$ -point one-dimensional DCT of the rows of  $u(n_{xy})$ , while the outer summation represents the  $N_x$ -point one-dimensional DCT of the columns of the semitransformed matrix. This implies that a two dimensional DCT transform can be implemented as a series of one-dimensional DCT transforms along the rows and then the columns or vice versa (the order in which this is done is theoretically immaterial).

The matrix form of the above one-dimensional forward and inverse DCT implementation can be represented, respectively, as:

$$[v] = [A_{N_x}]_{N_x \times N_x}^T [u]_{N_x \times N_y} [A_{N_y}]_{N_y \times N_y} \quad (12)$$

$$[u] = [A_{N_x}]_{N_x \times N_y}^T [v]_{N_x \times N_y} [A_{N_y}]_{N_y \times N_y}^T \quad (13)$$

where  $A_{N_x}$  and  $A_{N_y}$  are  $(N_x \times N_x)$  and  $(N_y \times N_y)$  transformation matrices, respectively.

If a direct transformation on the two-dimensional image is used through row (or column) concatenation to form a vector  $U$  with  $N_{xy}$  components, the one-dimensional transformation can be written as:

$$V = [A]_{N_x N_y \times N_x N_y} [U]_{N_x N_y \times 1} \quad (14)$$

where  $V$  is the DCT transformation coefficients.

The computational complexity of the two-dimensional concatenated DCT transformation in Eq. 14 is proportional to  $(N_x N_y \times N_x N_y)$ , while the complexity of the computations in the one-dimensional row/column operations of Eq. 12 is of order  $(N_x N_x + N_y N_y)$ . This clearly indicates the computational superiority of the separable basis used in the DCT. It is important to note that the reported complexity analysis in this section does not take advantage of the sparsity of the truncated DCT and are therefore upper bounds. More efficient implementation algorithms can be used in practice.

## References

1. Deutsch, C.V.: Geostatistical Reservoir Modeling. Oxford University Press, New York (2002)
2. Deutsch, C.V., Journel, A.G.: GSLIB: Geostatistical Software Library. Oxford University Press, New York (2002)
3. Gavalas, G.R., Shah, P.C., Seinfeld, J.H.: Reservoir history matching by Bayesian estimation. Soc. Petrol. Eng. J. **16**(6), 337–350 (1976)
4. Oliver, D.S., Reynolds, A.C., Bi, Z., Abacioglu, Y.: Integration of production data into reservoir models. Petrol. Geosci. **7**(9), 65–73 (2001)
5. Aziz, K., Settari, A.: Petroleum Reservoir Simulation. Applied Science, London (1979)
6. Reynolds, A.C., He, N., Chu, L., Oliver, D.S.: Reparameterization techniques for generating reservoir descriptions conditioned to variograms and well-test pressure data. Soc. Petrol. Eng. J. **1**(4), 413–426 (1996)
7. Jacquard, P., Jain, C.: Permeability distribution from field pressure data. Soc. Petrol. Eng. J. **5**, 281–294 (1965)
8. Jafarpour, B., McLaughlin, D.: Efficient permeability parameterization with discrete cosine transform. In: The Proceedings of the SPE Reservoir Simulation Symposium, Houston, TX, Paper SPE 106453 (2007)
9. Sahni, I., Horne, R.N.: Multiresolution wavelet analysis for improved reservoir description. SPERE **8**(1), 53–69 (2005)
10. Sarma, P., Durlofsky, L., Aziz, K.: A new approach to automatic history matching using Kernel PCA. In: The Proceedings of the SPE Reservoir Simulation Symposium, Houston, TX, Paper SPE 106176 (2007)
11. Evensen, G.: Sequential data assimilation with a non-linear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. J. Geophys. Res. **99**(C5), 10143–10162 (1994)
12. Kalman, R.E.: A new approach to linear filtering and prediction problems. Trans. ASME J. Basic Eng. **82**, 35–45 (1960)
13. Evensen, G.: The ensemble Kalman filter: theoretical formulation and practical implementation. Ocean Dyn. **53**, 343–367 (2003)
14. Nævdal, G., Johnsen, L.M., Aanonsen, S.I., Vefring, E.H.: Reservoir monitoring and continuous model updating using ensemble Kalman filter. Paper SPE-84372 (2003)
15. Gu, Y., Oliver, D.S.: History matching of the PUNQ-S3 reservoir model using the ensemble Kalman filter. Paper SPE 89942 (2005)
16. Wen, X.H., Chen, W.: Real-time reservoir model updating using ensemble Kalman filter. Paper SPE-92991 in the proceedings of the SPE Reservoir Simulation Symposium, pp 1–14 (2005)
17. ECLIPSE 100 (Black Oil): Reference Manual and Technical Description. Schlumberger GeoQuest, Houston (2006)
18. Jafarpour, B., McLaughlin, D.: Estimating reservoir permeabilities with the ensemble Kalman filter: the importance of the ensemble design. Paper SPE 108941 (2008)

19. SPE International: SPE Comparative Solution Project: General Information. <http://www.spe.org/csp/>
20. Strebelle, S., Journel, A.G.: Reservoir modeling using multiple-point statistics. In: The Proceedings of the SPE Annual Technical Conference and Exhibition, New Orleans, Paper SPE 71324 (2001)
21. Remy, N.: S-GeMS: A geostatistical earth modeling library and software. Ph.D. thesis, Stanford University (2004)
22. Ahmed, N., Natarajan, T., Rao, K.R.: Discrete cosine transform. *IEEE Trans. Comput.* **C-23**, 90–93 (1974)
23. Rao, K.R., Yip, P.: Discrete Cosine Transform: Algorithms, Advantages, Applications. Academic, Boston (1990)
24. Gonzalez, R.C., Woods, R.E.: Digital Image Processing, 2nd edn. Prentice Hall, Upper Saddle River, NJ (2002)
25. Brigham, E.O.: The Fast Fourier Transform and its Applications, p. 448. Prentice-Hall, Englewood Cliffs, NJ (1988)