Comput Geosci (2017) 21:315–333
DOI 10.1007/s10596-017-9614-7

ORIGINAL PAPER

Direct forecasting of reservoir performance using production data without history matching

Addy Satija[1] · **Celine Scheidt**[2] · **Lewis Li**[2] · **Jef Caers**[3]

Received: 31 October 2016 / Accepted: 6 January 2017 / Published online: 25 January 2017

Abstract The conventional paradigm for predicting future reservoir performance from existing production data involves the construction of reservoir models that match the historical data through iterative history matching. This is generally an expensive and difficult task and often results in models that do not accurately assess the uncertainty of the forecast. We propose an alternative re-formulation of the problem, in which the role of the reservoir model is reconsidered. Instead of using the model to match the historical production, and then forecasting, the model is used in combination with Monte Carlo sampling to establish a statistical relationship between the historical and forecast variables. The estimated relationship is then used in conjunction with the actual production data to produce a statistical forecast. This allows quantifying posterior uncertainty on the forecast variable without explicit inversion or history matching. The main rationale behind this is that the reservoir model is highly complex and even so, still remains a simplified representation of the actual subsurface. As statistical relationships can generally only be constructed in low dimensions, compression and dimension reduction of the reservoir models themselves would result in further oversimplification. Conversely, production data and forecast variables are time series data, which are simpler and much more applicable for dimension reduction techniques. We present a dimension reduction approach based on functional data analysis (FDA), and mixed principal component analysis (mixed PCA), followed by canonical correlation analysis (CCA) to maximize the linear correlation between the forecast and production variables. Using these transformed variables, it is then possible to apply linear Gaussian regression and estimate the statistical relationship between the forecast and historical variables. This relationship is used in combination with the actual observed historical data to estimate the posterior distribution of the forecast variable. Sampling from this posterior and reconstructing the corresponding forecast time series, allows assessing uncertainty on the forecast. This workflow will be demonstrated on a case based on a Libyan reservoir and compared with traditional history matching.

Keywords Reservoir · Forecasting · Functional data analysis · Uncertainty quantification

1 Introduction

Forecasting future reservoir performance from existing production data requires the integration of many disciplines. While a large variety of methods are available [13, 18], the practical application of these methods remains limited and often practice resorts to manual adjustment [31]. In reality, considerable complexity presents itself. First, reservoir models need to be constructed consisting of structural, lithological, and petrophysical components. These components may be interrelated to some degree. Several uncertainties exist, such as geological scenario uncertainty (discrete uncertainty); uncertainty on spatial distribution of

✉ Jef Caers
jcaers@stanford.edu

[1] Shell Oil, Houston, Texas, USA

[2] Energy Resources Engineering Department, Stanford University, Stanford, CA, USA

[3] Geological Sciences Department, Stanford University, Stanford, CA, USA

lithologies, porosity, and permeability (spatial uncertainty); and uncertainty in the structural framework in terms of number of faults (categorical uncertainty), fault hierarchy (scenario uncertainty), and fault throw (continuous parameter uncertainty). In addition, uncertainties may exist that affect fluid behavior, such as relative permeability, PVT properties or phase contacts, and initial and boundary conditions. Whether using a sampling approach to address the history-matching problem or an optimization approach, some forms of iteration must be done, in order to perturb all of these components to achieve multiple history-matched models that will hopefully represent realistic uncertainty. In this sense, sampling methods are preferred as they can incorporate prior uncertainty and hence, generate posterior solutions for uncertainty quantification, which are required for decision-making purposes. The amount of parameters involved however may become very large, including the issue of dealing with spatial uncertainty. For the latter, several dimension reduction techniques have been proposed [14, 20].

To date, there is not one single method that can address all these complexities. In other words, account properly for prior model uncertainty on all components of the model (fluid, rock, structure, petrophysics, and boundary) and sample models based on a Bayesian formulation of the posterior, matching all dynamic data, such as four-dimensional (4D) seismic, well test, production data, etc. Most publications address part of the problem, either by focusing on petrophysical properties only but not structure or by focusing on a fixed spatial model while perturbing simple engineering parameters. Generating usable and comprehensive software also remains elusive.

In this paper, we make a first small step towards reformulating and perhaps, rethinking and re-evaluating the practice of history matching in terms of real field applications involving full complexity, going beyond the usual simplified academic problems treated in many publications. Our approach builds on the previously published prediction focus analysis (PFA; see [11, 24, 25]). In this approach, it is advocated to *not* generate multiple models that match dynamic data, and *then only* run a forecast model on the history-matched models. Instead, the role of reservoir models is reconsidered; rather, models are generated in order to establish a direct statistical relationship between data variables and forecast/prediction variables, then using this estimated relationship and the actual production data to produce a statistical forecast. In this regard, we quantify posterior uncertainty on the forecast without history matching individual models. The rationale here is that reservoir modeling is complex and that reservoir models are extremely high dimensional and despite this, still remains a simplified representation of the actual subsurface geological and fluid complexity. Any sparse representation or dimension reduction method would further simplify an already simplified reality. On the other hand, production data and forecast variables are simple time-series on which statistical dimension reduction techniques as well multivariate modeling can be readily applied, without much loss of information. In this paper, we apply this idea on a real field case and show that roughly the same forecast can be obtained as with the traditional approach of generating multiple history-matched models. The structure of the paper is as follows. First, we review the prediction-focused analysis in broad terms as well as provide specific details. Next, we present a bootstrap method that allows evaluating the confidence in the method. We then apply the method to two versions of the same real field case, one involving structural uncertainty and one not. We provide appropriate limitations and discussion in concluding this paper and identify scenarios in which traditional inversion is still required.

2 Prediction-focused analysis

2.1 General overview

In prediction-focused analysis, one considers not just data variables and model variables but also the intended purpose in terms of prediction variables. Figure 1 illustrates the difference in approach between traditional model building by

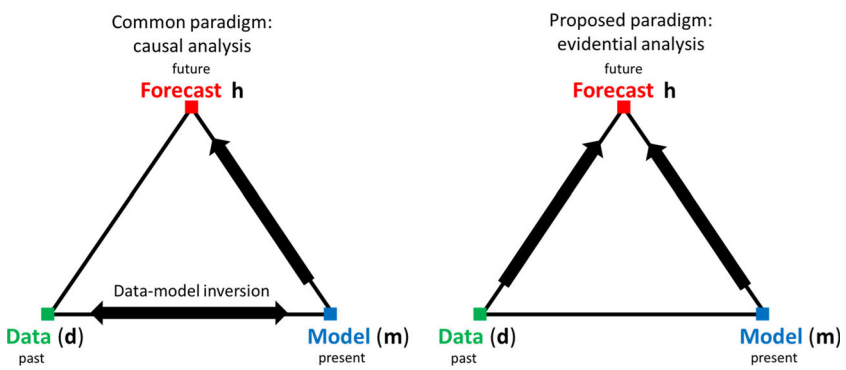

Fig. 1 Two views on addressing the forecasting issue. The traditional framework (*left*) applies causal analysis to match the models to the data, then use those matched models for forecasting. The proposed methodology (*right*) uses evidential analysis by which the model is used to construct a statistical relationship between the data and forecast

history matching and then forecasting and an approach that includes the forecast in an integrated fashion. In general terms, we represent time-varying historical data variables (such as watercut, oil rate, pressure, etc.) as vector **d**. The reservoir model is represented by **m**. The latter consists of the spatial model, the structural model, the fluid model, and all associated parameters, hence, is very high dimensional. The (future) forecast (such as cumulative oil, water cut, and volume) is represented by **h**. Clearly, the dimension of **m** is much larger than that of **h** and **d.**

$$\dim(\mathbf{m}) \gg \dim(\mathbf{h}),\ \dim(\mathbf{d}) \tag{1}$$

In addition, we observe actual field production data (or any other data), which we term $\mathbf{d}_{obs}$.

The idea of a prediction focus analysis is straightforward: using a stated prior on **m**, namely $f(\mathbf{m})$, one generates, by Monte Carlo (or other methods, such as quasi Monte Carlo), a set of N prior models: $\{\mathbf{m}^1\mathbf{m}^2\ldots\mathbf{m}^N\}$. These models are then evaluated through the data forward model:

$$\mathbf{d} = \mathcal{G}_d(\mathbf{m}) \tag{2}$$

and the forecast forward model:

$$\mathbf{h} = \mathcal{G}_h(\mathbf{m}) \tag{3}$$

These are forward functions are deterministic functions that generates the data and forecast variables for a given reservoir model **m**. These functions are determinsitric (assumed exact); observation error will be treated later. In practice, these functions are generally reservoir simulators that generate the expected historical and future production rates for a given reservoir model. Applying them to each prior model results in the pairs of data:

$$\left\{\{\mathbf{d}^1,\mathbf{h}^1\},\{\mathbf{d}^2,\mathbf{h}^2\},\ldots,\{\mathbf{d}^N,\mathbf{h}^N\}\right\} \tag{4}$$

Next, a statistical dimension reduction method is employed (e.g., MDS, PCA, kernel principal component analysis (KPCA), functional principal component analysis (FPCA), and canonical functional component analysis (CFCA), see later) to generate reduced dimension vectors, $\mathbf{d}^*$ and $\mathbf{h}^*$or in terms of the sample:

$$\begin{aligned}&\left\{\{\mathbf{d}_\mathbf{1}^*,\mathbf{h}_\mathbf{1}^*\},\{\mathbf{d}_\mathbf{2}^*,\mathbf{h}_\mathbf{2}^*\},\ldots,\{\mathbf{d}_N^*,\mathbf{h}_N^*\}\right\}\\ &\text{where } \dim(\mathbf{d}^*) << \dim(\mathbf{d});\ \dim(\mathbf{h}^*) << \dim(\mathbf{h})\end{aligned} \tag{5}$$

This joint sample is used to construct a multivariate distribution, assuming this is now possible because of the reduced dimensions (see next section on the specifics): $f\left(\mathbf{d}^*,\mathbf{h}^*\right)$. The observed data is reduced in dimension using the same dimension reduction method, obtaining $\mathbf{d}_{obs}^*$. The reduced observations are then used to condition the multivariate distribution as $f\left(\mathbf{d}^*,\mathbf{h}^*|\mathbf{d}_{obs}^*\right)$ which can be used to obtain:

$$f\left(\mathbf{h}^*|\mathbf{d}_{\text{obs}}^*\right) = \int_{\mathbf{d}} f\left(\mathbf{h}^*|\mathbf{d}_{\text{obs}}^*\right) f\left(\mathbf{d}^*\right)\ \mathrm{d}\mathbf{d}^* \tag{6}$$

After backtransformation, this results in the desired posterior uncertainty on the forecast $f(\mathbf{h}|\mathbf{d}_{\text{obs}})$. Note that the only requirement for the dimension reduction technique is for it to bijective namely having a uniqueness property:

$$\text{if } \mathbf{h}^* = r_h(\mathbf{h}); \mathbf{d}^* = r_d(\mathbf{d}) \text{ then } \mathbf{h} = r_h^{-1}\left(\mathbf{h}^*\right); \mathbf{d} = r_d^{-1}\left(\mathbf{d}^*\right) \tag{7}$$

with r_h and r_d the multivariate functions representing the respective mapping from higher to lower dimensions. The bijectivity allows moving from one lower to higher dimensional space and vice versa without encountering an ill-posed inverse problem. For example, multidimensional scaling and KPCA are not bijective as they require the solution of what is termed a pre-image problem (see [22]), while PCA, NLPCA, and FPCA are bijective operations (see Satija and Caers [24]).

2.2 Specifics

Figure 2 provides an overview of the specific components that are presented in this section. Preferably, the multivariate distribution in Eq. 6 should be multivariate Gaussian. In such case, the conditional distribution can be obtained using Gaussian process regression, equivalent to simple kriging the forecast from the data [28]. A multivariate Gaussian can be reasonably assumed when the relationship between $\mathbf{d}^*$ and $\mathbf{h}^*$is linear:

$$\mathbf{d}^* = G \cdot \mathbf{h}^* \tag{8}$$

and when the marginal distributions are Gaussian. Gaussian marginal distributions can always be obtained by means of histogram transformations, but obtaining a linear relationship is more challenging. To that extent, Satija and Caers [24] recognize that output of flow simulators represent systematic, physical "signals," varying in time. For example, cumulative curves obtained from multiple reservoir models have similar functional behavior (they start at zero, break through, and then increase) and they are not purely stochastically varying time series. As a result, they propose to reduce dimension of such signals by means of a statistical dimension reduction method that capitalizes on such systematic behavior, namely FPCA. To that extent, consider the basis expansion of the forecast variables as follows:

$$\mathbf{h}(t) \cong \sum_{i=1}^{L} k_{\xi,i}\boldsymbol{\xi}_i(t) \tag{9}$$

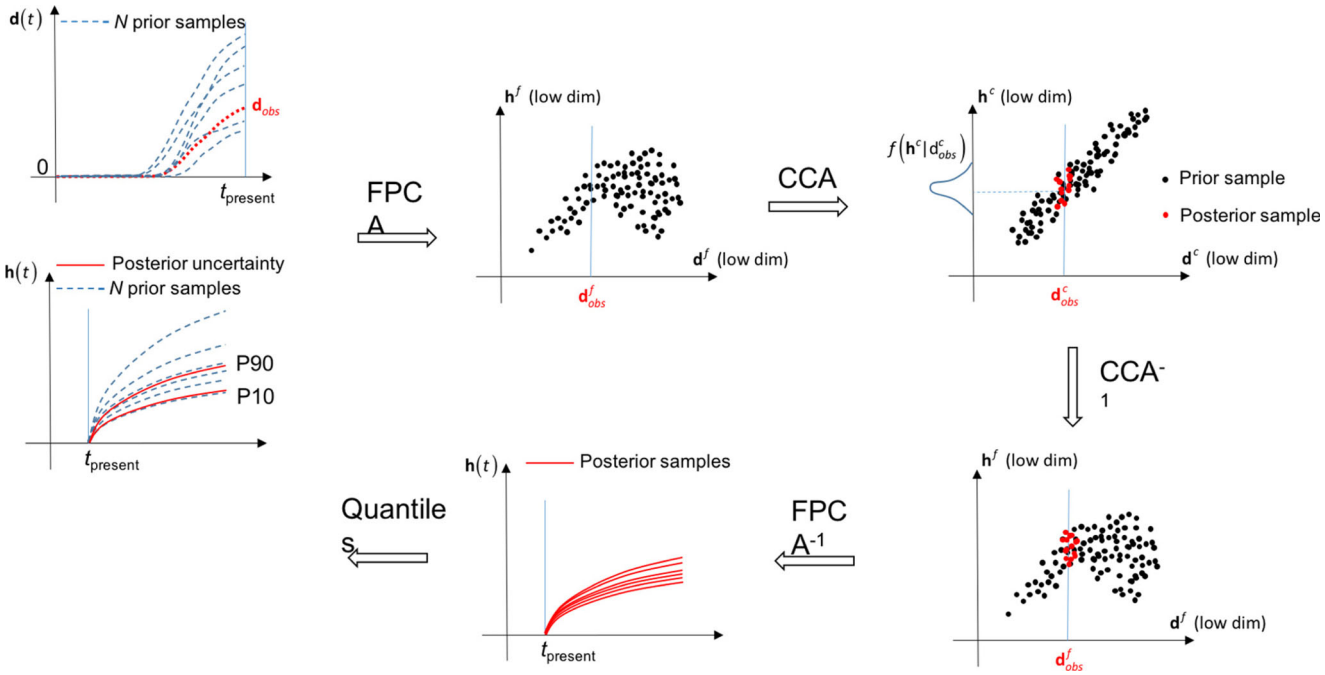

Fig. 2 General overview of the proposed methodology. The workflow uses Monte Carlo sampling and forward modeling to produce a set of prior historical and forecast-response curves. FPCA and CCA are used to reduce the dimension of the responses and maximize a linear correlation between the two variables. Performing Gaussian process regression and sampling from the resulting posterior distribution yields a set of updated forecasts conditioned to the observed historical data. Undoing the CCA/FPCA transformations produces the posterior forecasts as a set of time series responses from which updated quantiles can be computed

Using a spline basis has the advantage of computational ease of evaluation as well as establishing derivatives. The choice of the number of basis is a modeling choice and will need to be tuned for each case, usually using cross-validation [20]. We will use such spline basis throughout the paper. PCA on the coefficients of this linear combination characterizes the functional variations in the time series data and is referred to as FCA:

$$\mathbf{h}(t) \cong \sum_{i=1}^{K} h_i^f \boldsymbol{\phi}_{h,i}(t) \tag{10}$$

Hence, FPCA represents a time series as a linear combination of K orthonormal eigen-functions $\left\{\boldsymbol{\phi}_{h,1}(t), \boldsymbol{\phi}_{h,2}(t) \cdots \boldsymbol{\phi}_{h,K}(t)\right\}$ with coefficients $\mathbf{h}^f$. Note that FPCA, like PCA is bijective. A similar decomposition can be achieved for the data variables:

$$\mathbf{d}(t) \cong \sum_{i=1}^{K} d_i^f \boldsymbol{\phi}_{d,i}(t) \tag{11}$$

Applying FPCA to the set of N prior samples one obtains:

$$\left\{\{\boldsymbol{d}_1^f, \boldsymbol{h}_1^f\}, \{\boldsymbol{d}_2^f, \boldsymbol{h}_2^f\}, \ldots, \{\boldsymbol{d}_N^f, \boldsymbol{h}_N^f\}\right\} \tag{12}$$

Given the non-linear nature of the forecast and data response model, Eqs. 2 and 3, it is not necessarily guaranteed that one observes, after FPCA, a linear relationship between the pairwise components of $\mathbf{d}^f$ and $\mathbf{h}^f$. This is usually attributed to the presence of cross-correlations among the functional data variables. Therefore, a linearizing operation is performed by means of canonical correlation analysis (CCA), which is a more general form of partial least squares [29]. CCA relies on a linear transformation of both $\mathbf{d}^f$ and $\mathbf{h}^f$ such that the components in such transformation are maximally correlated, or in terms of notation:

$$\mathbf{d}^c = \mathbf{d}^f A^T \text{ and } \mathbf{h}^c = \mathbf{h}^f B^{\mathbf{T}} \tag{13}$$

where A and B are obtained as solution of:

$$\max_{\boldsymbol{A},\boldsymbol{B}} \frac{A \Sigma_{DH} B^T}{\sqrt{A \Sigma_{DD} A^T} \sqrt{B \Sigma_{HH} B^T}} \tag{14}$$

with

$$D = \begin{bmatrix} \mathbf{d}_1^f \\ \vdots \\ \mathbf{d}_N^f \end{bmatrix}^T, \mathbf{H} = \begin{bmatrix} \mathbf{h}_1^f \\ \vdots \\ \mathbf{h}_N^f \end{bmatrix}^T$$
$$\Sigma_{DH} = cov(D, H); \Sigma_{DD} = cov(D, D); \Sigma_{HH} = cov(H, H) \tag{15}$$

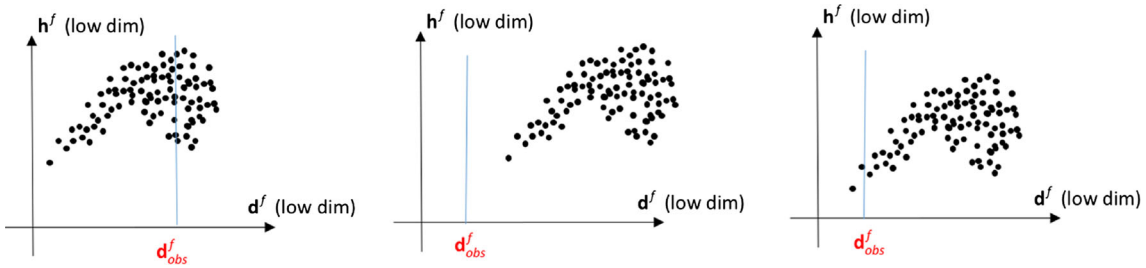

Fig. 3 Situations where CFCA could yield an improper estimate of the posterior. A non-linear relationship between the data and forecast (*left*), inconsistent prior (*middle*), and insufficient models that match the data in the prior (*right*)

thereby maximizing the correlations between pairwise components of $\mathbf{h}_i^c$ and $\mathbf{d}_i^c$ while constraining all the intercomponent correlations between $\mathbf{h}_i^c$ and $\mathbf{h}_{j\neq i}^c$, between $\mathbf{d}_i^c$ and $\mathbf{d}_{j\neq i}^c$, and between $\mathbf{h}_i^c$ and $\mathbf{d}_{j\neq i}^c$ to 0.

If a linear correlation is observed between $\mathbf{d}^c$ and $\mathbf{h}^c$, then a linear model G is regressed and a Gaussian likelihood model formulated as follows:

$$L\left(\mathbf{h}^c\right) = \exp\left(-\frac{1}{2}\left(G\mathbf{h}^c - \mathbf{d}_{obs}^c\right)^T C_{d^c}^{-1}\left(G\mathbf{h}^c - \mathbf{d}_{obs}^c\right)\right) \quad (16)$$

C_{d^c} refers to the data covariance matrix of the canonical components, which is estimated from the data error covariance in the original time domain using a Monte Carlo approach proposed in [12]. Since prior and likelihood are multivariate Gaussian, the posterior is also Gaussian and the posterior mean and covariance can be readily estimated using classical methods [28] as shown in Eqs. 17 and 18.

$$\tilde{h} = \bar{h}_{\text{prior}}^c + C_h G^T\left(GC_H G^T + C_{d^c} + C_T\right)^{-1}\left(d_{\text{obs}}^c - G\bar{h}_{\text{prior}}^c\right) \quad (17)$$

$$\tilde{C}_h = C_h - C_h G^T\left(GC_h G^T + C_{d^c} + C_T\right)^{-1} GC_h \quad (18)$$

C_T is the covariance of the error that arises due to the linear fitting in Eq. 8, which can be readily estimated empirically from the residuals in Eq. 8. Once the posterior mean and covariance is estimated, the posterior distribution is then sampled and the obtained samples are backtransformed into

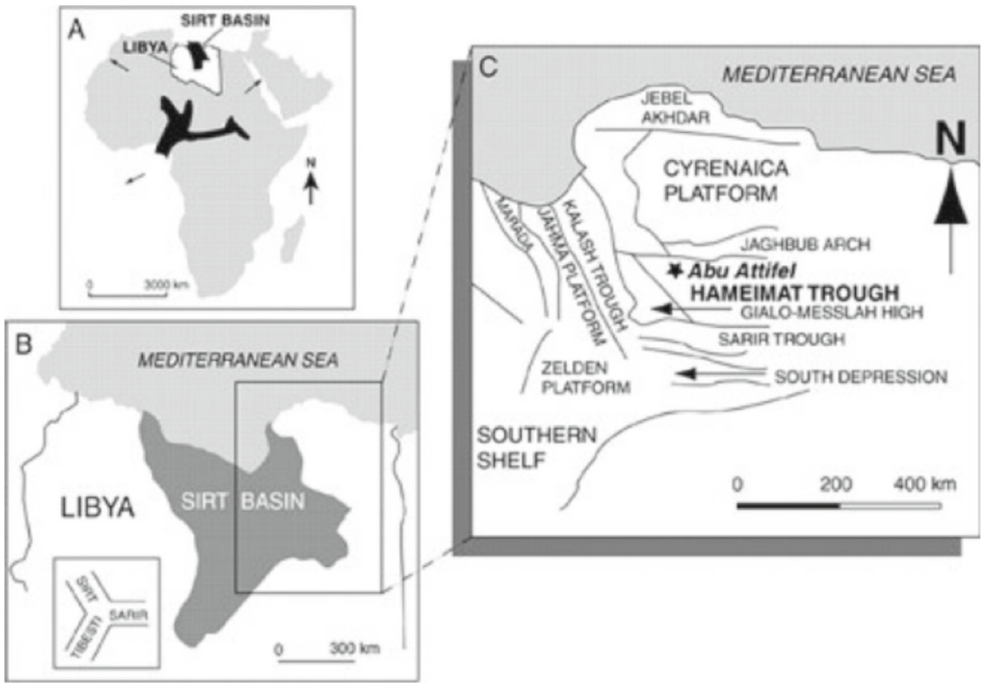

Fig. 4 Location of the Hameimat Trough in the Sirte Basin where the N-97 field is located. Image source: [2]

Table 1 Prior distribution of uncertain reservoir parameters used in generation of both cases of scenario 1

Parameter	OWC	μ_{oil}	T_{F1}	T_{F1}	T_{F1}	T_{F1}
Distribution	U[1061, 1076]	N(0.3, 0.2)	U [0.2, 0.8]	U [0.2, 0.8]	U [0.2, 0.8]	U [0.2, 0.8]
Parameter	K_{rw}	K_{ro}	S_{wir}	S_{or}	n_w	n_o
Distribution	N (0.3, 0.2)	N(0.7, 0.2)	N (0.2, 0.2)	N (0.2, 0.2)	N (2.5, 0.2)	N (2.0, 0.2)

The scoping set of 500 models were sampled from these distributions

actual signals, from which quantiles can be calculated (see Fig. 2).

2.3 Assumptions

There exist a number of underlying assumptions for this methodology to work in actual practice. These pertain to the consistency of the prior, the effectiveness of the dimension reduction, and the existence of a linear relationship between data and forecast in the canonical space. We would like to emphasize these before presenting application cases.

2.3.1 Dimension reduction

The "compression" by means of FPCA must be significant. FPCA attempts to represent response variables from simulators by means of few principal components. The lower the amount of components, the easier the Gaussian process regression. This compression will become more difficult as more wells are involved; hence, the procedure is likely to apply at early development stages.

2.3.2 Consistency of the prior

The Bayesian formulation of the forecasting problem requires the specification of a subjective belief on a hypothesis (e.g., the reservoir model). Evidence is then gathered (data) and then the probability of the evidence under the hypothesis is evaluated (the likelihood). The importance of this subjective prior is well known, and some authors in the statistical world are working on methods of falsification [9, 10]. If the "data" falls outside "the prior" as stated, then the probability of the data under the prior hypothesis is very small; consequently, the posterior is also very small for the given hypothesis. It may be tempting to perform ad hoc modifications to the prior (such as multiplying permeability with some value around a well based on the production data) with the purpose of ensuring that the prior range encompasses the observed data. However, this has been refuted as "ad hoc" [5] and can lead to incorrect posteriors. Indeed, any ad hoc modification of the prior will only lead to posteriors to be again inconsistent with observation at a future date. Accordingly, the prior distribution should be selected

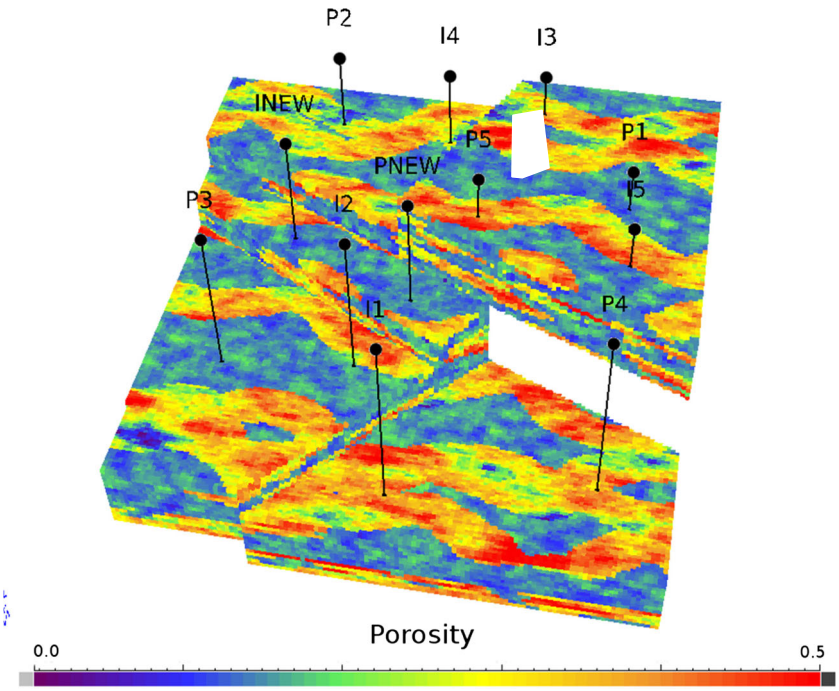

Fig. 5 Structural model and horizontal permeability used for scenario 1. The location of the existing producers are denoted by *P1 ... P5*, while the location of the new well to be drilled is denoted by *PNEW*. The uncertain reservoir parameters and their prior distributions are listed in Table 1

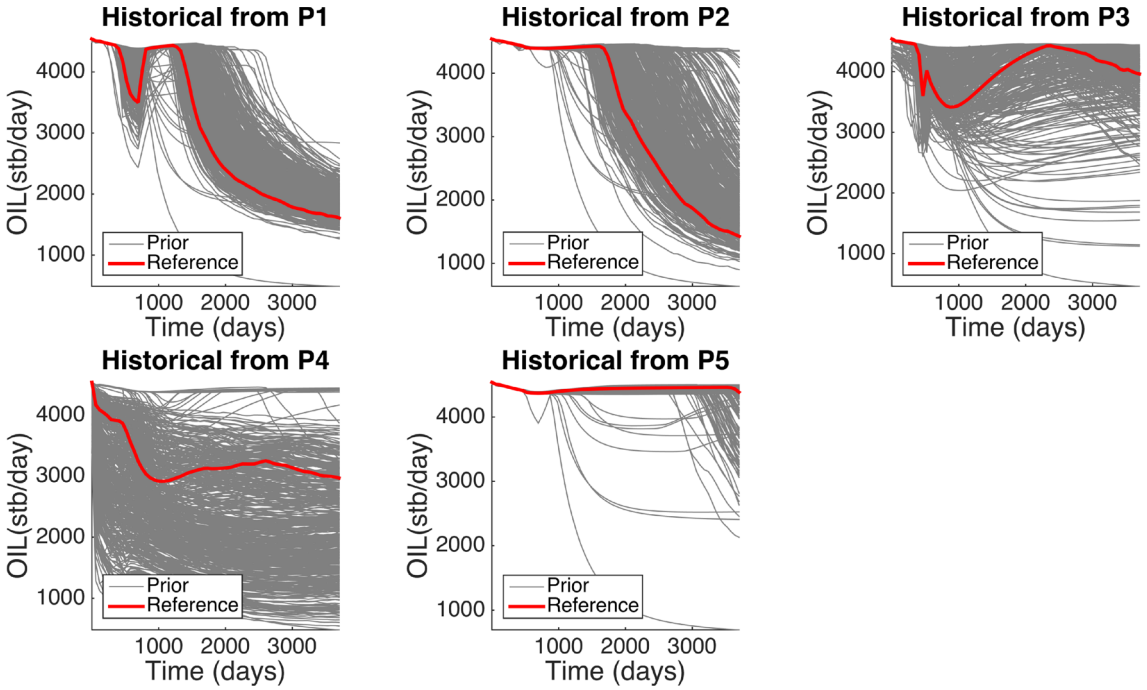

Fig. 6 Production data until day 3500 for each of the five existing producers for each prior model (*gray*). The production profiles were generated by forward simulating the prior models using a streamline simulator. The observed production is shown in *red*

to be wide enough and/or properly sampled to allow meaningful statistical analysis (see Fig. 3). A number of scenarios may still occur: (A) the data is not informative hence history matching is needed to predict **h** (the original intent of PFA [25] is to detect this); (B) the data may not be covered by the prior, which may be due to an inconsistent (e.g., too narrow) prior or the number of samples (N) from the prior is insufficient; and (C) the data is covered by the prior but an insignificant amount of samples is available to estimate the posterior of **h**. The latter may occur when the data lies in the extremes of the prior. These situations must be assessed prior to performing PFA.

2.3.3 Linearity in canonical space

The relationship between the canonical components of **d** and **h** can be modeled by means of a linear regression. Any deviation of this linear model will lead to increased model error, and hence, large wide uncertainty when this model error is included in the forecasting (as in [24]).

Any violations of these assumptions may lead to no reduction in posterior uncertainty or unreliable forecasts. These issues therefore beg for a quantitative assessment of the prediction power of the procedure, which is covered in the next section. In instances where the predictive power provided by PFA is low, traditional history matching will still be required.

2.4 Confidence vs uncertainty

An uncertainty statement for a forecast, can be as simple as stating a posterior PDF of that forecast. Since all reservoir modeling forecasts are obtained from limited amount of samples from that PDF, a question of confidence on the stated uncertainty is required. For example, if an uncertainty interval based on quantiles is provided (often in terms of P10–P90), then the confidence on those quantiles needs to be calculated. First, this is relevant for testing whether the posterior quantiles are different from the prior quantiles,

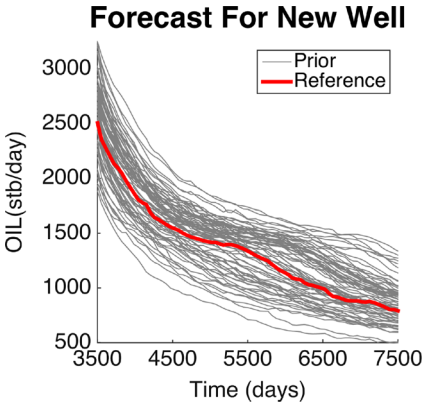

Fig. 7 The forecasted response for the new well from days 3500 to 7500, generated by forward modeling each of the prior models (*gray*). The true (in reality unknown) forecast is shown in *red*

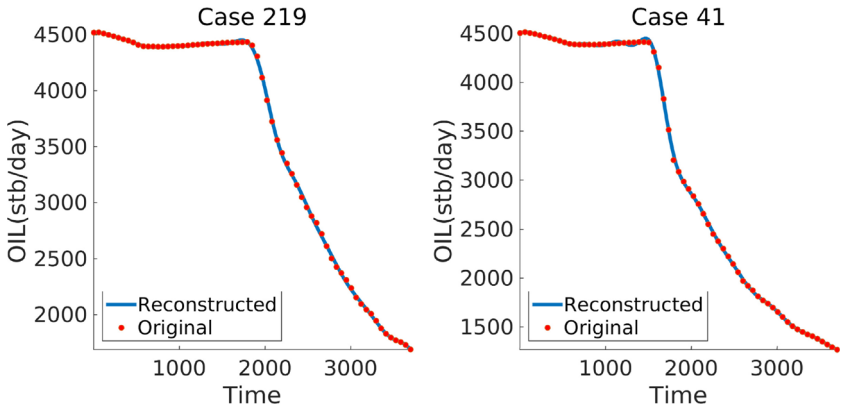

Fig. 8 Reconstructions of the original time series using only the first four components of FPCA along with the original production data from producer 1 for select prior models. The quality of the reconstruction serves as a metric for selecting appropriate basis splines

as this would indicate predictivity of the data towards that particular forecast. Note that uncertainty and confidence intervals need not be related. One may be very confident on a wide uncertainty and less confident on a narrow uncertainty. Since this depends on a number of factors (the prior uncertainty, the particular data and forecast, the flow model, etc.), a quantitative method is required to establish this relationship.

In this section, we develop a hypothesis test, which tests whether the data informs the forecast (prediction variables **h**) using the above (Fig. 2) procedure. Based on the p value of this hypothesis test, we then plot confidence vs uncertainty for several combinations of field data, to investigate what combination of historical data should be used for that forecast. It is clear that the most optimal situation occurs when we have small posterior uncertainty in combination with high confidence. A low confidence or large posterior uncertainty suggests that conventional inversion techniques may be required.

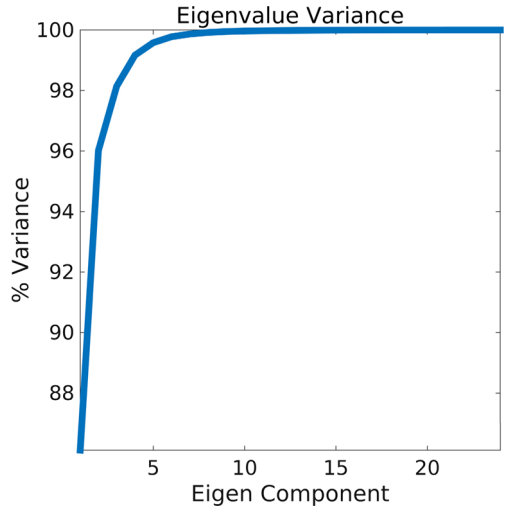

Fig. 9 Cumulative sum of the FPCA eigenvalues for the historical responses of producer 1. This indicates that over 99 % of the variability in P1's responses can be captured by just the first 5 eigencomponents of FPCA. This provides a metric for diagnosing the effectiveness of the dimension reduction

In general, the data $\mathbf{d}_{obs}$ is informative when there is a significant difference between the prior distribution $f(\mathbf{h})$ and the posterior $f(\mathbf{h}|\mathbf{d}_{obs})$. Since both these pdfs are multivariate and functional, it would be difficult to develop a simple measure. As a proxy, we propose using the first functional component h_1^f containing the maximum variability of H [16, 19, 21]. This would also be a necessary condition for $f(\mathbf{h})$ and $f(\mathbf{h}|\mathbf{d}_{obs})$ to be different. The problem of comparing pdfs is now simpler as only a comparison of univariate distributions is needed. We therefore define a theoretical difference between prior and posterior functional component as

$$\delta = \Delta\left(f\left(\boldsymbol{h}_1^f\right), f\left(\boldsymbol{h}_1^f|\boldsymbol{d}_{obs}\right)\right) \tag{19}$$

Δ can be any cdf- or pdf-based difference such as L1-norm difference [7], the Kolmogorov-Smirnov difference [15, 26], Jensen-Shannon or Kullback-Liebler divergence [17]. The prior distribution is estimated as $\hat{f}\left(\boldsymbol{h}_1^f\right)$ directly from the N scoping runs. The posterior cumulative distribution is estimated as $\hat{f}\left(\boldsymbol{h}_1^f|\boldsymbol{d}_{obs}\right)$ from the M posterior samples obtained from the Gaussian process regression and CCA backtransform. Thus, an empirical difference measure is estimated as

$$\hat{\delta} = \Delta\left(\hat{f}\left(\mathbf{h}_1^f\right), \hat{f}\left(\mathbf{h}_1^f|\mathbf{d}_{obs}\right)\right) \tag{20}$$

Fenwick [7] proposed a bootstrap-based hypothesis test to test for significant difference between two cdfs based on the L1 norm. For the purpose of hypothesis testing, the

Table 2 Percentage of variance captured by the first 4 eigenvalues of FPCA for each existing production well's historical data

Well	P1	P2	P3	P4	P5
Percentage of variance (%)	98.521	99.579	99.533	99.895	99.832

This suggests that the first 4 eigenvalues are sufficient to represent the majority of the variation in the prior models' historical response

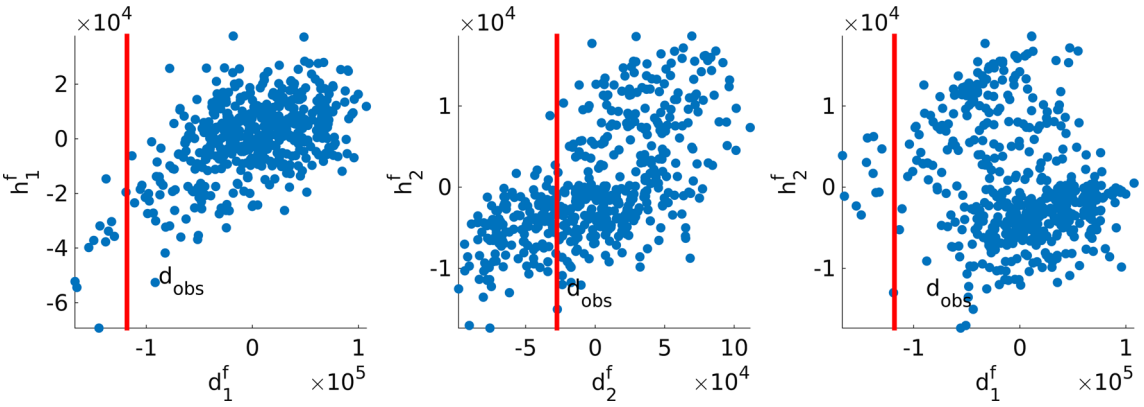

Fig. 10 Each of the prior models plotted in the reduced dimension (functional components) of data vs forecast. The correlation between the first components (*left*) shows a weak correlation between data and forecast $\rho = 0.2915$. Likewise, the cross-correlation between the first component of the data and the second component of the forecast shows a similarly weak correlation $\rho = -0.2577$. Performing regression and estimation using these components could result in inaccurate assessments of posterior forecast uncertainty

null hypothesis is that the prior and the posterior distributions of the first functional component are not different. For bootstrapping, B datasets are drawn from the existing N scoping runs without replacement, resulting in a B bootstrap estimate of the difference

$$\hat{\hat{\delta}}_b = \Delta_{CDF}\left(\hat{\hat{F}}_b\left(\mathbf{h}_1^f\right), \hat{\hat{F}}_b\left(\mathbf{h}_1^f|\mathbf{d}_{obs}\right)\right), \quad b=1,\ldots,\mathrm{B} \quad (21)$$

The number of times $\hat{\hat{\delta}}_b \leq \hat{\delta}$ is measure by how strong the predictivity is coinciding with the number of times the null hypothesis is rejected. The more the null hypothesis is rejected, the more confident we are on the informativeness of the data, hence a good measure is:

$$\omega = \frac{1}{B}\sum_{b=1}^{B} i\left(\hat{\hat{\delta}}_b \leq \hat{\delta}\right) \quad i = 1 \text{ if } \hat{\hat{\delta}}_b \leq \hat{\delta}, 0 \text{ else} \quad (22)$$

In the case study section, we illustrate the use of this measure for various combinations of historical production data.

3 Case study

3.1 Case description

The WintersHall Concession C97-I in the N-97 field is located in the Western Hameimat Trough of the Sirte Basin of north-central Libya (Fig. 4). The geological setting of the Sirte Basin is described in detail by Ahlbrandt [2]. The primary hydrocarbon source bed in the Sirte Basin has been identified by Ahlbrandt [2] as the Late Cretaceous Sirte Shale. The reservoir under consideration, the WintersHall Concession C97-I, is a fault-bounded horst block with the Upper Sarir Formation sandstone reservoir [3]. Complex interactions of the dextral slip movements within the rift system have led to the compartmentalization of the reservoir. Initial structural modeling attempts and interpretation from seismic data suggested the presence of up to four faults, each with unknown displacements and transmissibility, due to uncertainty in the interpretation of the seismic

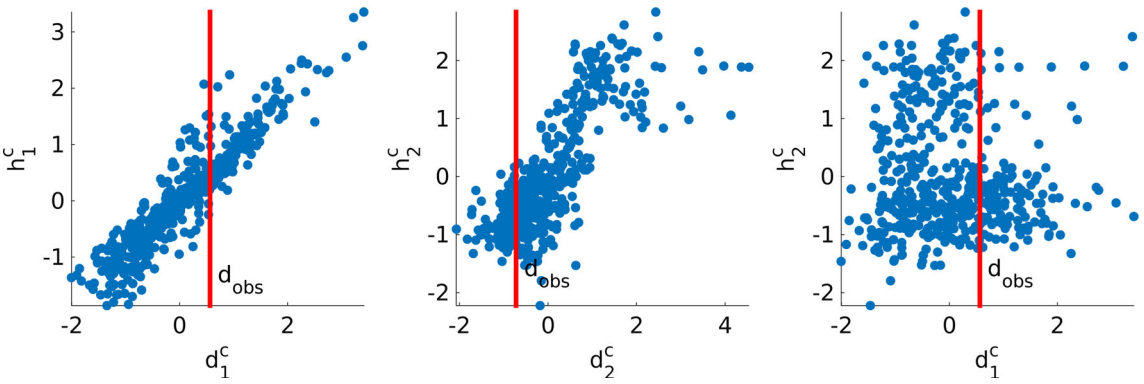

Fig. 11 The result of performing canonical correlation analysis on the functional components of data and forecast to transform the models into canonical space ($\mathbf{h}^c$ vs $\mathbf{d}^c$). The correlation between the first canonical components is much stronger than its functional counterpart ($\rho = 0.8941$). Furthermore, the cross-correlations between the first data canonical component and the second forecast canonical component is very low ($\rho = 0.0297$)

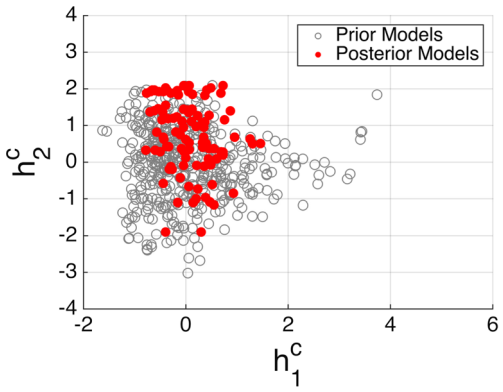

Fig. 12 Drawing from the posterior distribution yields posterior samples in canonical space. Only the first two dimensions are shown here, but each sample is actually a point in 4D space. The reduction in variance of the posterior samples in comparison with the prior models indicates that a reduction in forecast uncertainty is achieved

image. Based on this realistic settings and the reservoir model provided by the operator, we have generated a few virtual cases by which we hope to illustrate the validity of the direct forecasting approach.

3.2 Scenario 1: flow parameter uncertainty

In this first illustration, we consider the scenario where the structural geology and depositional environment are assumed to be well understood and the major sources of uncertainty reside in the flow parameters. The reservoir is thus composed of five distinct compartments resulting in four uncertain fault transmissibilities. A large aquifer is located in the lowest compartment, but the depth of this oil water contact also remains uncertain. Other uncertain reservoir parameters are relative permeabilities for the oil and water phases and oil viscosity. The relative permeabilites of the oil and water are modeled using the Corey Expressions [4], which requires three parameters for each phase (irreducible saturation, end point permeability, and Corey exponent). The prior distributions of these parameters are listed in Table 1. The structural model contains four faults as shown in Fig. 5 and is used for all realizations. Likewise, a three-facies training image containing sand channels is used with SNESIM [27] and Gaussian simulation to populate the grid with the appropriate facies and reservoir parameters (porosity/permeability) (see Fig. 5).

We consider the situation where five producers and three injectors have already been drilled at the locations depicted in Fig. 5. The field has been in production for 3500 days, and production data is available for all five wells. A decision needs to be made regarding the economic feasibility drilling of a 6th producer in the smallest reservoir compartment (denoted PNEW) and an additional injector (denoted INEW). Specifically, this decision will be made based on the forecasted performance of this new well over the next 4000 days. Therefore, we will seek to estimate the P10–P50–P90 forecasts of PNEW based on the first 10 years of production data from the existing five producers.

3.2.1 Generating prior realizations of model, data, and forecast variables

A prior set of models is required to establish a statistical relationship between the data and forecast. In this case, a set of 500 prior reservoir models is generated by applying Monte Carlo to the prior distributions in Table 1. The number of models was selected to ensure that the prior distributions were sufficiently sampled. The prior models were forward modelled using a streamline simulator (3DSL) over all 7500 days to encompass both the 3500 days of production data, as well as the 4000 days of forecast required to

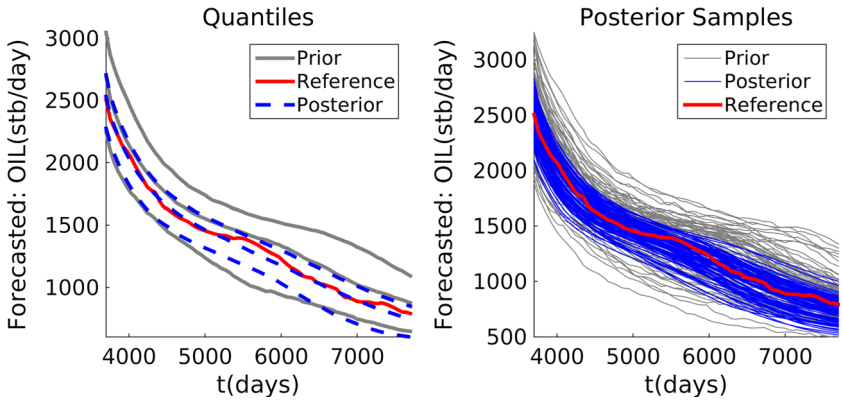

Fig. 13 By undoing the CCA and FPCA, the posterior forecast samples are transformed from canonical space back into the time domain. The posterior P10, P50, and P90 forecasts are shown along with the prior P10, P50, and P90 curves (*left*). The sampled forecasts are shown with the original forecasts from the prior models as well as the reference (*right*). This illustrates that CFCA does indeed provide a reduction in the forecast uncertainty

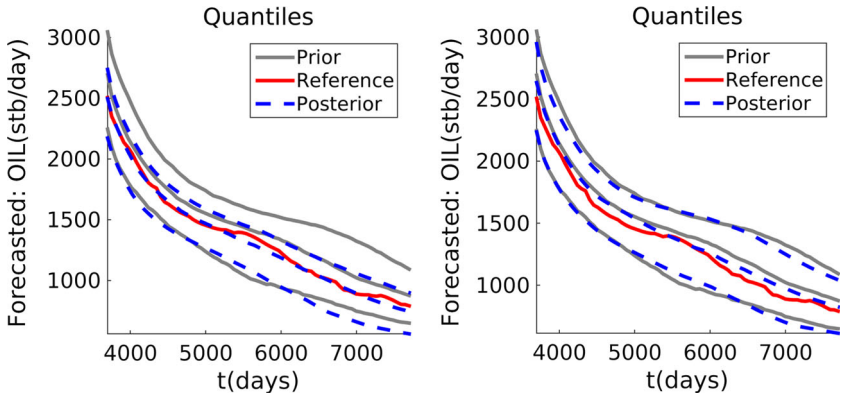

Fig. 14 Prior and posterior P10, P50, and P90 forecasts when the measurement error $\mathbf{C}_D$ is 100 $(\text{stb/day})^2$ (*left*) and 300 $(\text{stb/day})^2$. As the measurement error is increased, the posterior forecasts exhibits smaller reduction in uncertainty compared with the prior

make the decision regarding the new well. For illustration, one of the generated models is used as a reference case, and its production data used as the "observed" production data. The prior distribution of reservoir performance for each of the existing wells as well as the forecast for the new well is shown in Figs. 6 and 7.

3.2.2 Dimension reduction

The first requirement in establishing a relationship between the data and forecast is low dimensionality of both variables. In this case, the data and forecast are time series responses, that while technically are infinite dimensional, have been discretized into vectors of lengths 150 and 170, respectively. Directly establishing a statistical correlation between variables of such dimension is still infeasible, and thus a projection into a lower dimension using FPCA is first performed. The critical step of FPCA is the selection of an appropriate basis functions. In this particular instance, 6th-order splines with 20 knots as basis functions are found to provide reasonable approximations for both the data and forecast variables. The effectiveness of the projection is verified by performing FPCA, and then reconstructed the original time series using the basis functions and harmonic scores seen in Fig. 8.The choice of splines is then iteratively adjusted to minimize the average RMS between the original and reconstructed curves. It is also important to ensure that oscillations caused by Runge's phenomena [8] do not occur over any of the models, when high-order splines are used. The cumulative sum of the FPCA eigenvalues is used to ascertain the effectiveness of this compression. From Fig. 9, one observes that the compression is indeed significant, as 98.52 % of the variability in Producer 1's response is captured by the first 4 eigenvalues. Table 2 shows the percentage of the variability represented by the first 4 eigenvalues for each of the existing producers $\left(\mathbf{d}^f_{\text{P1}}, \mathbf{d}^f_{\text{P2}} \ldots \mathbf{d}^f_{\text{P5}}\right)$ in addition to the well to be drilled $\left(\mathbf{h}^{\text{f}}_{\text{PNEW}}\right)$.

Since the production data is composed of multiple wells, the relationship between each of the existing wells and the new well must be quantified. However, redundancy between the data from each production well exists, as they are obtained from the same underlying reservoir model. For instance, a shallow oil water contact in the aquifer would cause both P3 and P4 to experience early water breakthrough. Consequently, a second dimension reduction is applied to the production data in the form of a mixed PCA on the matrix $\mathbf{d}^f_{\text{Producers}}$ obtained by concatenating $(\mathbf{d}^f_{\text{P1}}, \mathbf{d}^f_{\text{P2}} \ldots \mathbf{d}^f_{\text{P5}})$. In our example, 97.42 % of the variability of the wells is captured by the first eight components of the mixed PCA. This effectively reduces the 20-dimensional $\mathbf{d}^f_{\text{Producers}}$ into an 8-dimensional $\mathbf{d}^f$, thus each prior model is represented by a single point in 12-dimensional space (eight components from historical $\mathbf{d}^f$ and four components from the forecast $\mathbf{h}^f$). This is illustrated in two dimensions corresponding to the two largest eigenvalues in Fig. 10.

3.2.3 Canonical correlation of d^f –hf and regression

Dimension reduction from an infinite dimensional time series into a 12-dimensional space enables the possibility of performing regression between data and forecast. However, this is contingent on a sufficiently strong linear correlation between the variables or else the resulting estimates from regression will be inconclusive. In this example, the correlation between first functional components of data and

Table 3 Difference between P10 and P90 averaged over the forecast period with varying levels of measurement error.

	Prior	$\sigma^2_\epsilon = 0$ stb/day	$\sigma^2_\epsilon = 50$ stb/day	$\sigma^2_\epsilon = 100$ stb/day	$\sigma^2_\epsilon = 300$ stb/day
Average P10–P90 stb/day	565.102	271.23	283.71	344.65	527.95

As the measurement error is increased to extreme values, the posterior forecast returns to that of the prior

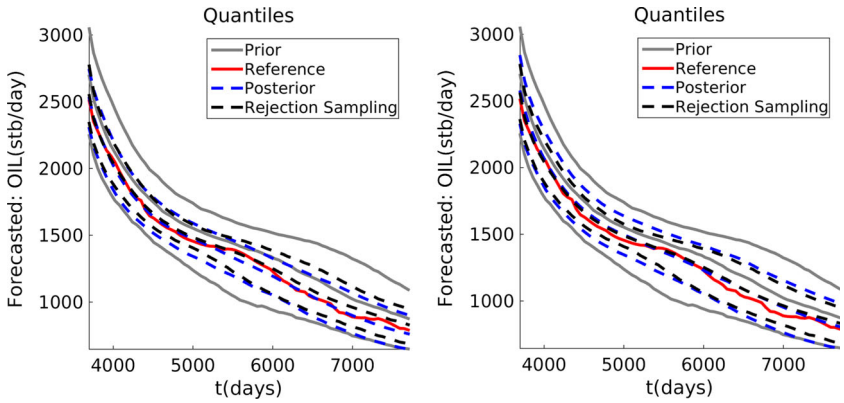

Fig. 15 P10, P50, and P90 forecasts from direct forecasting and rejection sampling. Two error covariances were illustrated: 50 (stb/day)^2 (*left*) and 150 (stb/day)^2 (*right*). The number of models remaining after rejection was 690 and 987 in each case. A similar reduction in uncertainty was achieved by direct forecasting using substantially less computational effort

forecast is rather poor ($\rho = 0.2915$). This is due to the presence of cross-correlations between the first and second components ($\rho = -0.2577$) (Fig. 10). To fully capitalize on the full multivariate correlation between all component of $\mathbf{h}$ and all components of $\mathbf{d}$, a CCA is performed to transform the models into a canonical space ($\mathbf{d}^c$ and $\mathbf{h}^c$). This subsequently increases the correlation between the first components from 0.2915 in the functional domain to 0.8941 in canonical space, similarly the cross-correlation between the first and second components is reduced to 0.0297 (Fig. 11). Now that a linear correlation has been established in low dimensions, the corresponding linear relationship G is obtained via Eq. 8. The application of linear Gaussian regression (Eq. 16) to estimate the posterior on the forecast components requires that $\mathbf{h}^c$ must be transformed using a normal score transform first to obtain $\mathbf{h}^c_{\text{Gauss}}$. Gaussian regression thus produces a multivariate normal posterior $f(\mathbf{h}^c_{\text{Gauss}}|\mathbf{d}_{\text{obs}})$ that is easily sampled to produce forecast components conditioned to $\mathbf{d}^c_{\text{obs}}$(Fig. 12). To obtain the corresponding forecasts as time series, the normal score transform and canonical transform must be backtransformed. The resulting posterior samples now in functional space, are used in conjunction with the original eigenfunctions from FPCA to reconstruct h(t). One hundred posterior forecasts were sampled and converted into time series as shown in Fig. 13 along with the P10–P50–P90 curves of the posterior forecasts. The posterior quantiles exhibit a significant reduction in uncertainty in comparison with the prior.

3.2.4 Accounting for measurement error

In the previous example, the observed data was assumed to be error-free; however, this is often not the case in real applications. This error is accounted for in the direct forecasting workflow by the $\mathbf{C}^c_D$ term in Eqs. 17 and 18. However, as the measurement error can only be estimated in the original time domain (modeled as a zero mean Gaussian with diagonal covariance matrix $\mathbf{C}_D$), the procedure described in Section 2.2 and Hermans et al. [12] must be applied to obtain $\mathbf{C}^c_D$ It is evident from Eq. 18, that increasing magnitudes of measurement error will result in larger posterior uncertainty That is to say, the less reliable the observed data, the less informative it is of the forecast As an illustration varying levels of measurement noise (100 and 300 (stb/day)^2) were assumed, and posterior distributions for each were computed (Fig. 14; Table 3). This also shows that increasing the measurement noise to extreme values causes the posterior uncertainty to approaches the prior uncertainty. This is in accordance with Eq. 18, as any new information, regardless of how reliable, cannot increase uncertainty.

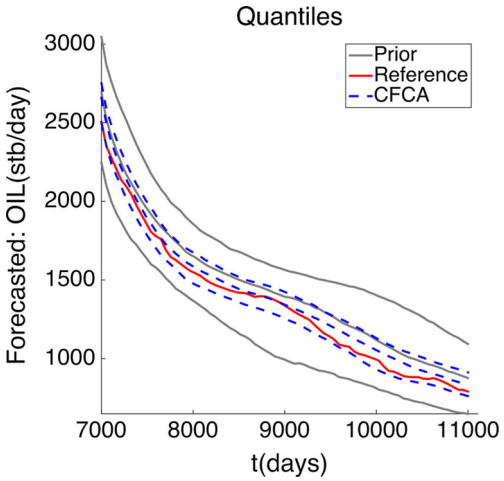

Fig. 16 Posterior forecast quantiles for the situation where the new well is to be after 7000 days instead of 3600 days. The inclusion of additional informative data reduces the forecast uncertainty

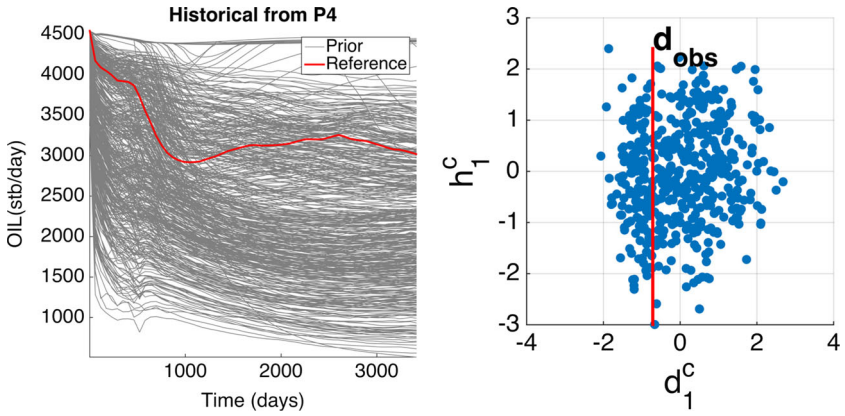

Fig. 17 Prior production data for P4 along with the true observed (*left*) and the subsequent models in canonical space using only P4 as the data (*right*). In this instance, the data does not inform the forecast due to the compartment in which P4 is located having very little communication with the compartment of the new well. This manifests as poor correlation in the canonical space. This means the posterior will provide little reduction in forecast uncertainty in comparison with the prior

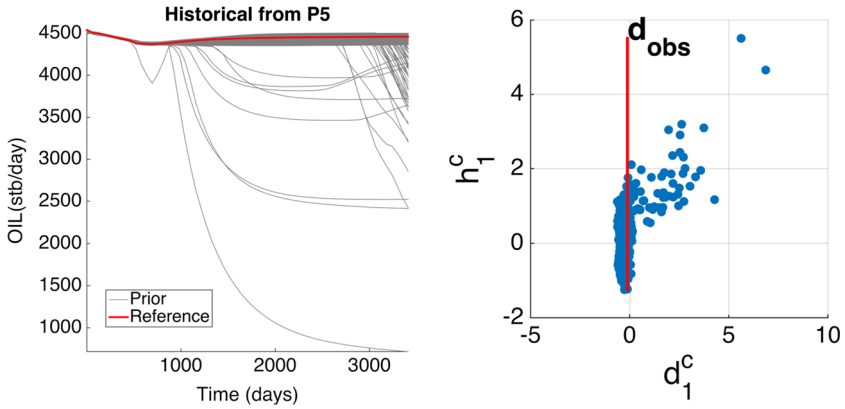

Fig. 18 Prior production data for P5 along with the true observed (*left*) and the subsequent models in canonical space using only P5 as the data (*right*). Water breakthrough has yet to occur in P5 for most of the prior models, and as a result, the majority of the models are producing at the production limits. This means that for a given production profile in P5, the forecasts for the new well could vary widely. This manifests in the canonical plot as a vertical cluster of models around the observed data

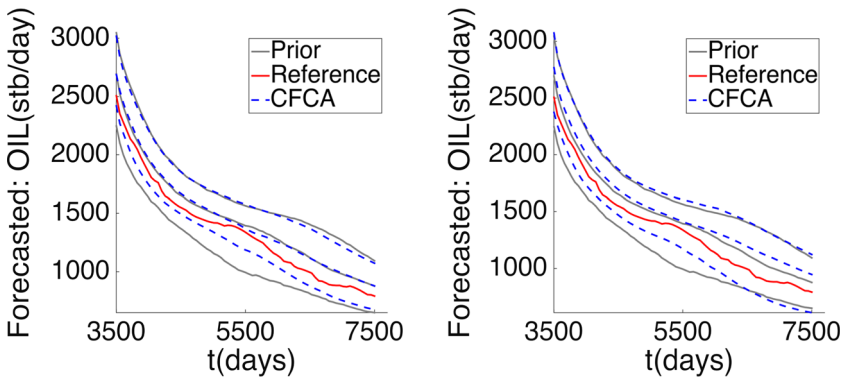

Fig. 19 Due to the lack of informativty between P4 (*left*) and P5 (*right*), the posterior uncertainty does not show marked reduction from the prior uncertainty

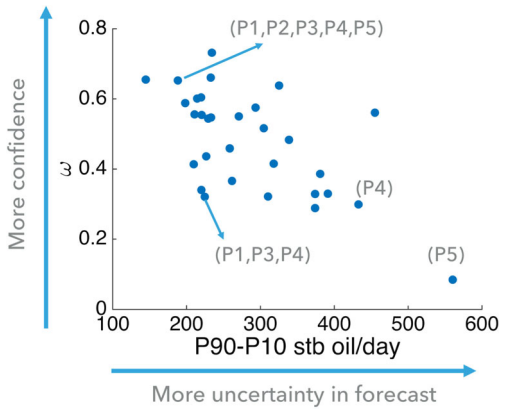

Fig. 20 By using a combinatorial of historical data from a subset of producers, and running the bootstrap test for statistical significance, we obtain not only an updated forecast uncertainty but also a confidence in this uncertainty. Each point represents some combination of wells, its posterior forecast uncertainty (P10–P90 averaged over 4000 days of forecast)

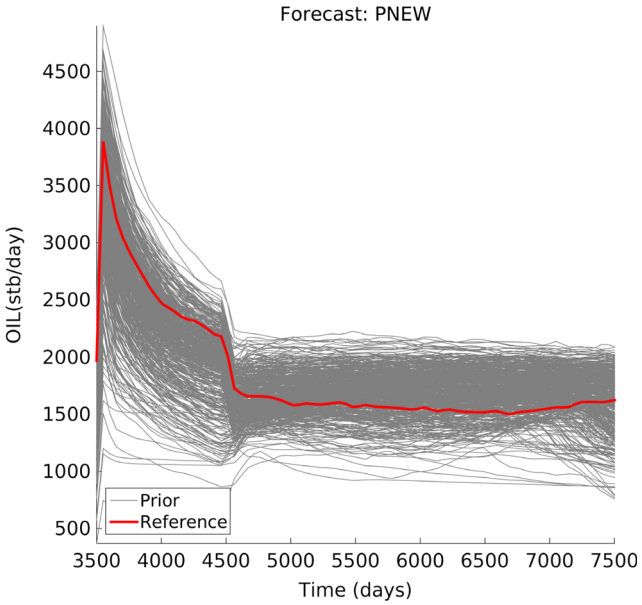

Fig. 22 The forecasted response for the new well from days 3500 to 7500, generated by forward modeling each of the prior models (*gray*). The true (in reality unknown) forecast is shown in *red*. An injector has been added in the same compartment as this producer

3.2.5 Comparison with rejection sampling

To compare the reduction in uncertainty provided by direct forecasting with an idealized case, where multiple history-matched models are available, 15,000 realizations were sampled from the prior distribution and forward simulated. Rejection sampling is performed on the resulting responses to identify the realizations that match the production data using the same methodology as described in Satija and

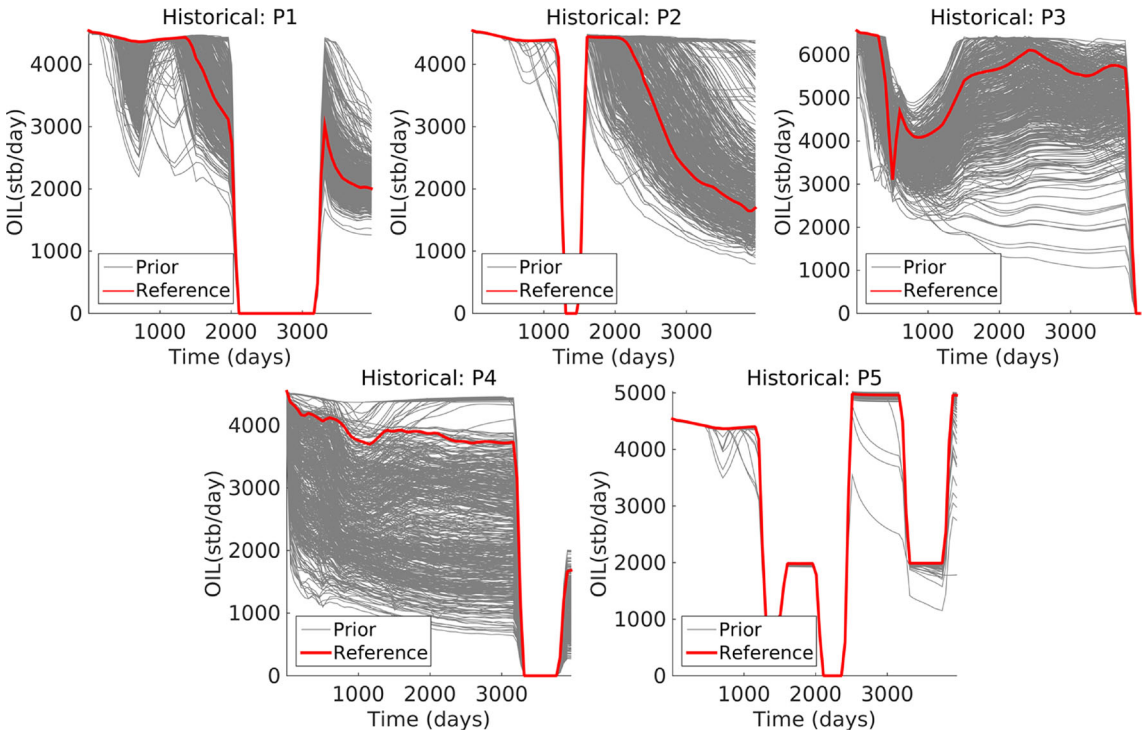

Fig. 21 Production data until day 3500 for each of the five existing producers for each prior model (*gray*). The production profiles were generated by forward simulating the prior models using a streamline simulator and variable well schedules. The observed production is also shown (*red*)

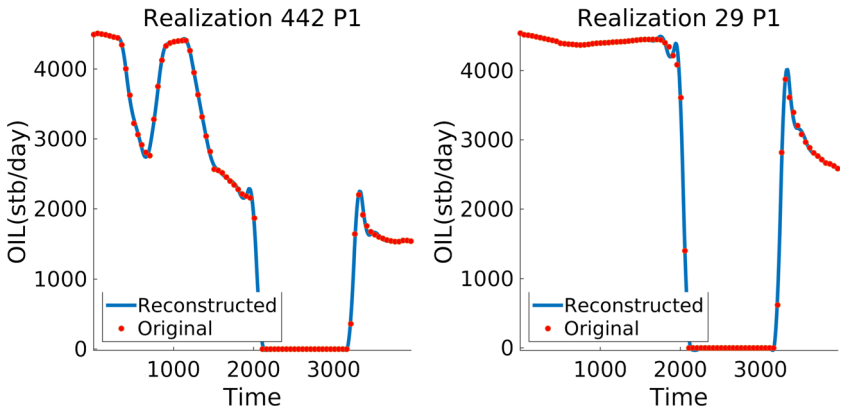

Fig. 23 Reconstructions of the time original time series using only the first five components of FPCA along with the original production data from producer 1 for select prior models. A larger number of knots is required in this scenario due to the additional complexities in the rate

Caers [24]. The likelihood is defined for the data observed in each realization in the rejection sampling set using Eq. 23.

$$l\left(\mathbf{d}_i|\mathbf{d}_{\text{obs}}\right) \propto \exp\left(\mathbf{d}_i - \mathbf{d}_{\text{obs}}\right)^T \mathbf{C}_D^{-1}\left(\mathbf{d}_i - \mathbf{d}_{\text{obs}}\right) \tag{23}$$

$\mathbf{C}_D$ is a diagonal matrix, as the noise is modeled as a zero mean Gaussian. Rejection sampling was performed for two $\mathbf{C}_D$ values, ($\sigma^2 = 50$ and 150 (stb/day)^2) which resulted in 690 and 987 models, respectively. The P10–P50–P90 curves for the corresponding forecasts of these matched models were computed and shown in Fig. 15, along with the computed P10–P50–P90 curves from direct forecasting. In both instances, direct forecasting has comparable uncertainty reduction with rejection sampling, while utilizing only a fraction of the computational cost.

3.2.6 Updating forecasts with additional data

Assimilation of additional observed data remains a challenge in convention history-matching problems. Techniques based on the Ensemble Kalman Filter [1], have been proposed, but handling large systems with non-linearities and non-Gaussian geostatistical priors remains a topic of research. Due to the complexity of the reservoir models, the history-matched models are often inconsistent with actual future observations. This requires another round of costly history matching and model updating in order to obtain an updated forecast. Conversely, in PFA, the prior models just need to be forward modeled for a longer simulation time to account for the new data. Consider the case where the new well is to be drilled after 7000 days rather than 3500. To obtain an updated forecast, Direct Forecasting is repeated, but with the original prior models forward simulated to 11,000 days (7000-day production and 4000-day forecast). The resulting quantiles are shown in Fig. 16 and exhibit a reduction in uncertainty compared with the previous 3500-day case. This should be expected as the incorporation of additional informative data should provide a reduction in the forecast uncertainty.

3.2.7 Confidence vs uncertainty

To understand if data is informative on the forecast, consider the scenario in which we only use a single production well's data to forecast the new well. In Fig. 17, only data from producer 4 is used to make the forecast. As seen in Fig. 5, producer 4 is separated from the compartment in which the new well is located by two faults and potentially has a low level of communication. Performing CFCA shows that only a 0.0341 correlation exists between $\mathbf{h}^c$ and $\mathbf{d}^c$, which is

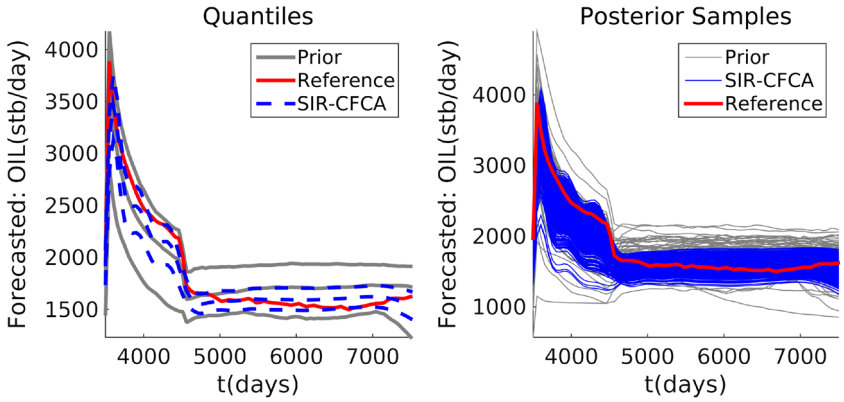

Fig. 24 The posterior P10, P50, and P90 forecasts are shown along with the prior P10, P50, and P90 curves (*left*). The sampled forecasts are shown with the original forecasts from the prior models as well as the reference (*right*). No measurement error was assumed here

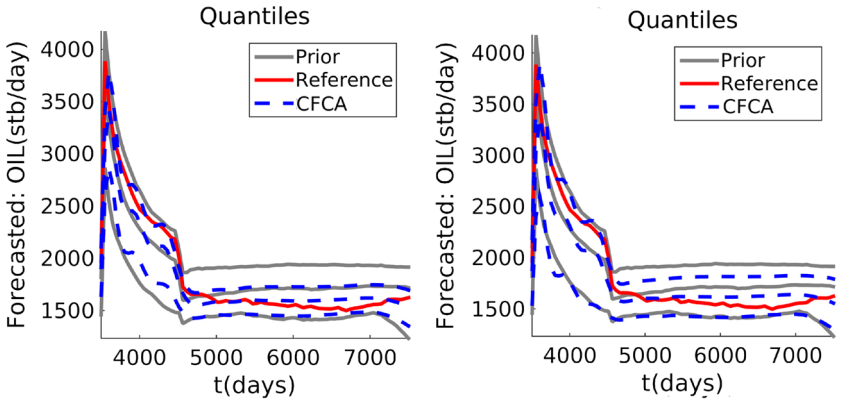

Fig. 25 Prior and posterior P10, P50, and P90 forecasts when the measurement error $\mathbf{C}_D$ is 100 stb/day (*left*) and 300 stb/day

insufficient to reduce forecast uncertainty (Fig. 19). This represents a case where the data is not related to the forecast. Conversely, considering only producer 5 as the data, and examining the production profiles from the prior models Fig. 18, one observes that most of the prior models have not experienced water breakthrough after 3500 days. Consequently, they are all producing oil at the production limits. This manifests in the canonical domain as a cluster of models that yield widely different forecasts. It follows that the quantiles obtained from CFCA using just P5, do not show a marked decrease in uncertainty when compared with the prior (Fig. 19).

To quantify the level of confidence on the updated forecast quantiles, the bootstrap procedure outlined in Section 2.3 is applied. To assess both the reduction in uncertainty as well as the confidence we have in each updated forecast, a combinatorial selection of wells was used to update the forecast (e.g., P1, P1–P2, etc.). The bootstrap was performed 500 times, and the average forecast uncertainty (mean difference between P10 and P90 over the duration of the forecast) is plotted against the computed ω of Eq. 22 (Fig. 20). This shows that in general lower confidence correlates with larger posterior forecast uncertainty, and it can be seen that using just P4 and P5 results in smaller uncertainty reduction and poor confidence. Conversely, using all five wells produces both higher confidence and lower uncertainty. It should also be noted that some combinations of wells such as (P1, P3, and P4) do produce reduced uncertainty, but with low confidence, and this must be taken into account when decisions are made.

3.3 Scenario 2: flow parameter uncertainty with variable well schedules

To illustrate a scenario where the data exhibits discontinuities and complexities commonly seen in real field cases, both the injector and producer schedules from the previous case were modified to include shut-ins and pressure changes (Fig. 21). Since the historical well schedule is already known when constructing the prior models, it is incorporated directly into the forward model ($\mathbf{g}_d$ in Eq. 2). An additional injector was also added to the same compartment as the producer to be forecasted, to induce additional variability in the forecasts (Fig. 22). Once again, 500 prior models were sampled according to the prior parameter distributions in Table 2.

3.3.1 Dimension reduction of discontinuous time series data

Due to the shut-ins and additional complexities in the historical production data, the 6th-order 20-knot B-splines used as basis functions in the previous scenario are insufficient to capture the variability among the prior models. However, this is resolved by increasing the number of knots to 30, as evident in the reconstructions shown in Fig. 23. Despite this increase in number of knots, the FPCA and mixed PCA

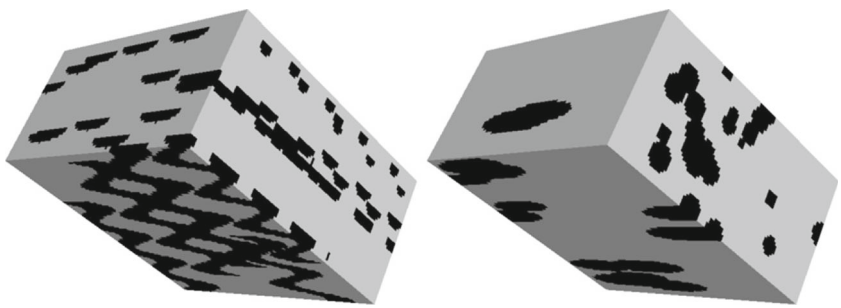

Fig. 26 The two different depositional scenarios depicted by training images containing channels (*left*) and ellipsoids (*right*)

procedures outlined in Section 2.2 are still able to identify redundancies in the data. For each producer, FPCA identified that the first 5 eigenvalues were required to represent 99 % of the variability (in contrast to 4 in the previous scenario), while the mixed PCA indicated that 12 total scores (vs 8 previously), are required to represent the entire set of historical production rates.

3.3.2 Posterior forecasts

Applying CCA on the scores obtained by FPCA/mixed PCA resulted in correlations of 0.8821 in the first canonical components of the data and forecast. Consequently, the reduction in uncertainty in the posterior forecasts (Fig. 24) is comparable with that of the previous scenario. Likewise, increasing the measurement error to 100 and 300 stb/day causes the posterior uncertainty to revert towards the prior uncertainty (Fig. 25). This demonstrates that if the compression provided by FPCA is significant, and CCA is successful in maximizing linear correlations, direct forecasting is able to provide reductions in uncertainty regardless of the nature of the historical and forecast data variables.

3.4 Scenario 3: with structural and depositional uncertainty

In the third scenario, both structural and spatial uncertainties are considered. A combination of the structural uncertainty modeling approaches suggested [6, 30] is applied. As

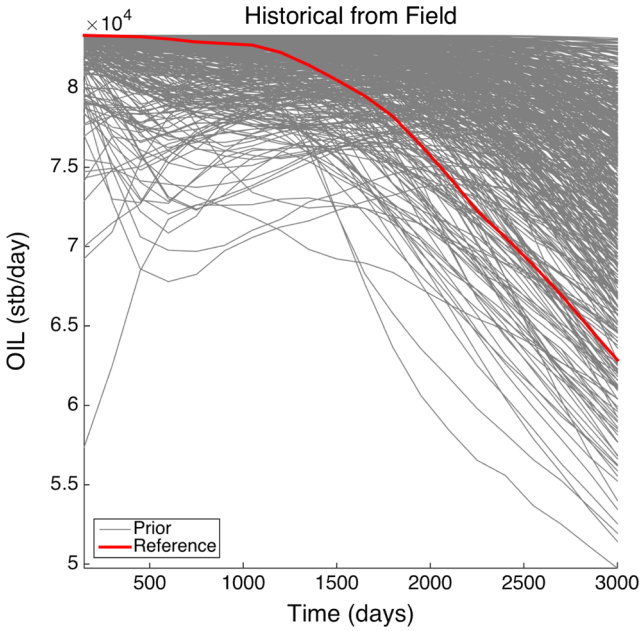

Fig. 27 Field-level production data until day 3000 for each prior model (*gray*). The production profiles were generated by forward simulating the prior models using a streamline simulator. The observed production is shown in *red*

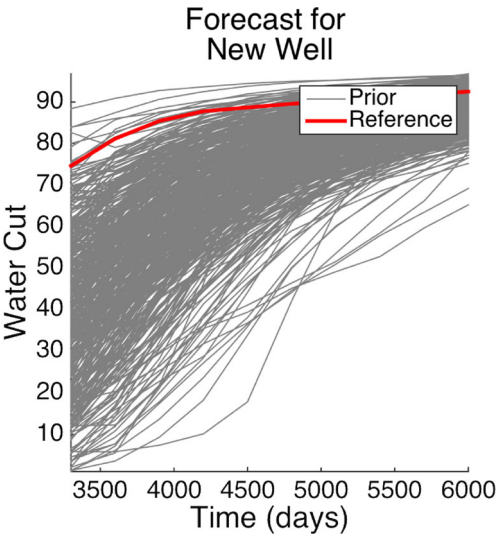

Fig. 28 The forecasted response for the new well from days 3000 to 6000, generated by forward modeling each of the prior models (*gray*). The true (in reality unknown) forecast is shown in *red*

described in Bellman et al. [3], there could be between one and four faults in the reservoir at the given locations. Therefore, the number of fault parameter (numFaults) is modeled as an integer random number. The throws or displacements of each fault (Throw_i) is uncertain and must abide by rules such as geologically older faults must have higher displacements than newer faults. This indicates a joint distribution between the Throw_i and numFaults parameters:

$$\text{Throw}_{i+1} \leq \text{Throw}_i \forall i \leq \text{numFaults}$$

$$\text{Throw}_i = 0 \forall i > \text{numFaults}$$

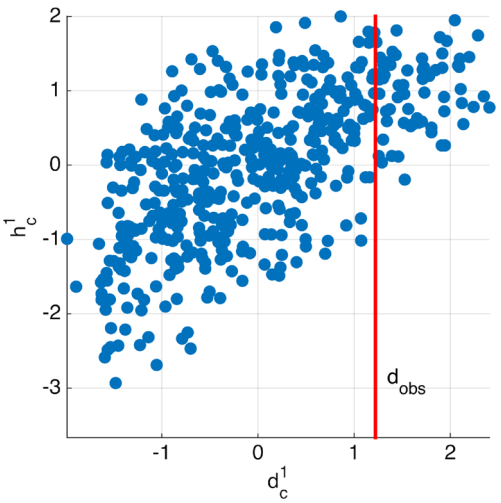

Fig. 29 Scatter plot of the first and second canonical components of the data and forecast variables. The observed data projected into canonical space is indicated by the red line. The correlation in this instance is moderate ($\rho = 0.65$)

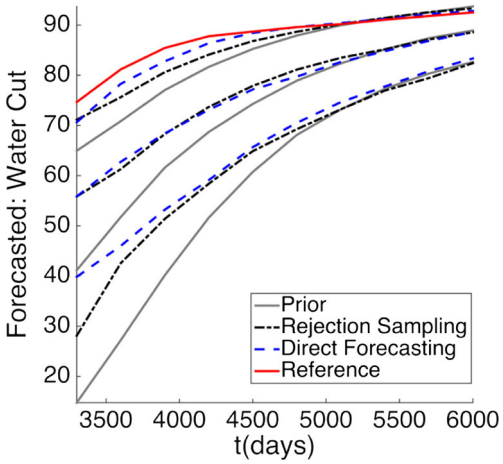

Fig. 30 Comparison of posterior quartiles obtained from direct forecasting and rejection sampling with prior quartiles of the predicted water cut in the prediction

where Throw_1 is the displacement in the oldest fault. The stochastic model-generation algorithm detailed in Satija [23] is used to sample a value for each flow displacement parameter under these constraints. To account for the spatial uncertainty, two depositional scenarios are modeled using two training images shown in Fig. 26. One TI contains sinusoidal channels, while the other contains ellipsoidal lobes. Finally, two separate marginal distributions of net to gross are considered (50 and 70 %) and used as an input to SNESIM along with the TI to generate the resulting realizations. In the scenario, one producer and one injector have already been in place and producing for 3000 days, a decision is required regarding drilling a new producer at a given location in the highest fault block in the reservoir. Such a decision would be based on the new well-forecasted performance over the next 3000 days, such as if the operator considers the well to be unprofitable if the water cut rises above a certain threshold. To perform PFA, 500 models are sampled from the prior and simulated using 3DSL. The resulting responses are shown in Figs. 27 and 28. The same workflow as the first scenario was applied, and resulting h^c vs d^c plot is illustrated in Fig. 29. It should be noted that the correlation in this example is not as linear as the previous cases without structural uncertainty (0.61). As before, rejection sampling consisting of 4700 models, resulting in 221 history-matched models, was used and the posterior bounds in both cases show similar results as in Fig. 30.

4 Discussions and conclusions

Forecasting problems in the oil and gas industry are generally formulated as iterative data inversion or history-matching problems that are computationally expensive and difficult. However, as the goal of forecasting is to inform a decision, rather than obtaining the matched models, we present a reformulation of the forecasting problem where the role of the reservoir model is reconsidered. Instead of attempting to use it to match the data, the model is instead used to establish a statistical relationship between the historical data and forecast. This estimated relationship is thus used to obtain a statistical forecast based on actual observed production data. To establish such a relationship, a strong linear correlation is required in a low dimension to render the statistical procedure feasible. In this paper, we use FPCA followed by CCA to achieve these requirements. The application of this methodology to three scenarios based on the WintersHall Concession C97-I reservoir in Libyan demonstrated that CFCA is able to provide updated estimates of the forecast uncertainty that is comparable with rejection sampling, but at a fraction of the computational cost. The confidence in these updated forecasts can be gauged through a bootstrap test of statistical significance that we have also presented in this paper. This level of confidence in the updated uncertainty should be taken into consideration when making decisions based off the forecast, as well as identifying scenarios where traditional inversion (history matching) is still required. While PFA was presented in this paper using time series as the responses, an extension of this research could be the application to other types of responses common in the Earth Sciences such as saturation maps or time-lapse data. The reduction in required computational time and complexity of this methodology in comparison with history matching could have considerable impact on how forecasting is done in the Earth Sciences. An additional appealing characteristic of this non-iterative procedure is that (1) it can be perfectly parallelizes making it extremely computationally efficient and (2) it can be done "offline" outside of reservoir modeling and flow simulation software thereby making it completely general from a software implementation point of view.

References

1. Aanonsen, S.I.: The ensemble Kalman filter in reservoir engineering—a review. SPE J. **14**(3), 393–412 (2009)
2. Ahlbrandt, T.S.: The Sirte Basin Province of Libya: Sirte-Zelten total petroleum system. US Geological Survey. CO: US Department of the Interior, Denver (2001)
3. Bellman, L.H., Kouwe, W., Yielding, G.: Fault seal analysis in the N-97 Oil Field in Wintershall Concession C97-I, Libya. 2nd EAGE International Conference on Fault and Top Seals—from pore to basin scale 2009, pp. 101–104. EAGE, Montpellier (2009)
4. Brooks, R.H., Corey, A.T.: Hydraulic properties of porous media and their relation to drainage design. Trans. ASAE **7**(1), 26–28 (1964)

5. Chalmers, A.F.: What is this thing called science? St Lucia, Hackett Publishing (2013)
6. England, W.A.: The effects of faulting on production form a shallow marine reservoir—a study of the relative importance of fault parameters. SPE Annual Technical Conference and Exhibition, pp. 279–292. Society of Petroleum Engineers, New Orleans (1988)
7. Fenwick, D.S.: Quantifying asymmetric parameter interactions in sensitivity analysis: application to reservoir modeling. Math. Geosci. **46**(4), 493–511 (2014)
8. Fornberg, B., Zuev, J.: The Runge phenomenon and spatially variable shape parameters in RBF interpolation. Computers & Mathematics with Applications **54**(3), 379–398 (2007)
9. Gelman, A., Meng, X.-L., Stern, H.: Posterior predictive assessment of model fitness via realized discrepancies. Stat. Sin., 733–760 (1996)
10. Gelman, A., Shalizi, C.R.: Philosophy and the practice of Bayesian statistics. Br. J. Math. Stat. Psychol. **66**(1), 8–38 (2013)
11. Hermans, T., Klepikova, M., Caers, J.: Assessing heat tracing experiment data sets for direct forecast of temperature evolution in subsurface models: an example of well and geophysical monitoring data. Geophysical Research Abstracts, p. 5105. European Geosciences Union, Vienna (2016)
12. Hermans, T., Erasmus, O., Caers, J.: Direct prediction of spatially and temporally varying physical properties from time-lapse electrical resistance data. Water Resources Research (in press) (2017)
13. Hohl, D.S.: Field experiences with history matching an offshore turbiditic reservoir using inverse modeling. SPE Annual Technical Conference and Exhibition, p. 9. Society of Petroleum Engineers, San Antonio (2006)
14. Jolliffe, I.: Principal component analysis Hoboken. John Wiley & Sons, Ltd (2002)
15. Justel, A.D.: A multivariate Kolmogorov-Smirnov test of goodness of fit. Statistics & Probability Letters **35**(3), 251–259 (1997)
16. Krzanowski, W.: Principles of multivariate analysis. OUP Oxford, Oxford (2000)
17. Kullback, S., Leibler, R.: On information and sufficiency. Ann. Math. Stat. **22**(1), 79–86 (1951)
18. Oliver, D.S., Chen, Y.: Recent progress on reservoir history matching: a review. Comput. Geosci. **15**(1), 185–221 (2011)
19. Ramsay, J.O.: Applied functional data analysis: methods and case studies. Springer, New York (2002)
20. Ramsay, J.O.: Functional data analysis. John Wiley & Sons, Inc, Hoboken (2006)
21. Ramsay, J.O.: Some tools for functional data analysis. J. R. Stat. Soc., 539–572 (1991)
22. Sarma, P.L.: Kernel principal component analysis for efficient, differentiable parameterization of multipoint geostatistics. Math. Geosci. **40**(1), 3–32 (2008)
23. Satija, A.: Reservoir forecasting based on statistical functional analysis of data and prediction variables 162. Stanford, Stanford University (2015)
24. Satija, A., Caers, J.: Direct forecasting of subsurface flow response from non-linear dynamic data by linear least-squares in canonical functional principal component space. Adv. Water Resour. **77**, 69–81 (2015)
25. Scheidt, C.R.: Prediction-focused subsurface modeling: investigating the need for accuracy in flow-based inverse modeling. Math. Geosci., 1–19 (2014)
26. Smirnov, N.: Table for estimating the goodness of fit of empirical distributions. Ann. Math. Stat., 279–281 (1948)
27. Strebelle, S.: Conditional simulation of complex geological structures using multiple-point statistics. Math. Geol. **34**(1), 1–24 (2002)
28. Tarantola, A.: Inverse problem theory and methods for model parameter estimation. SIAM, Philadelphia (2005)
29. Thompson, B.: Canonical correlation analysis: uses and interpretation. Sage, London (1984)
30. Thore, P., Shtuka, A., Lecour, M., Ait-Ettajer, T., Cognot, R.: Structural uncertainties: Determination, management, and applications. Geophysics **67**, 840–852 (2002)
31. Yang, C.N.: Reservoir model uncertainty quantification through computer-assisted history matching. SPE Annual Technical Conference and Exhibition, p. 12. Society of Petroleum Engineers, Anaheim (2007)