

Analysis of the Ensemble and Polynomial Chaos Kalman Filters in Bayesian Inverse Problems*

Oliver G. Ernst[†], Björn Sprungk[†], and Hans-Jörg Starkloff[‡]

Abstract. We analyze the ensemble and polynomial chaos Kalman filters applied to nonlinear stationary Bayesian inverse problems. In a sequential data assimilation setting, such stationary problems arise in each step of either filter. We give a new interpretation of the approximations produced by these two popular filters in the Bayesian context and prove that, in the limit of large ensemble or high polynomial degree, both methods yield approximations which converge to a well-defined random variable termed the analysis random variable. We then show that this analysis variable is more closely related to a specific linear Bayes estimator than to the solution of the associated Bayesian inverse problem given by the posterior measure. This suggests limited or at least guarded use of these generalized Kalman filter methods for the purpose of uncertainty quantification.

Key words. inverse problems, Bayesian inverse problem, Bayes estimator, Kalman filter, ensemble Kalman filter, polynomial chaos, conditional distribution

AMS subject classifications. 60H35, 35R30, 35R60

DOI. 10.1137/140981319

1. Introduction. Due to increasing attention to uncertainty quantification (UQ) for complex systems, in particular as it relates to the study and solution of partial differential equations (PDEs) with random data, interest has also focused on inverse problems for random PDEs. In particular, the Bayesian approach to inverse problems has become popular in this context. From a UQ perspective the inverse problem is of tremendous interest since incorporating any available information into the probability law of an uncertain quantity will, in general, reduce uncertainty and lead to improved stochastic models.

We consider in this paper the fundamental task of inferring knowledge about an unknown element $u \in \mathcal{X}$ from a separable Hilbert space \mathcal{X} by observing finite-dimensional noisy data

$$(1) \quad z = G(u) + \varepsilon,$$

where $G : \mathcal{X} \rightarrow \mathbb{R}^d$ denotes the known (and deterministic) parameter-to-observation map and ε the observational noise. Adopting the Bayesian perspective, we assume a probability measure μ_0 on \mathcal{X} to be given describing our *prior* knowledge or belief about u which may be based, e.g., on physical reasoning, expert knowledge, or previously collected data. We wish to highlight the

*Received by the editors August 6, 2014; accepted for publication (in revised form) June 24, 2015; published electronically September 22, 2015. This work was supported by the DFG priority programme 1324. This paper is a substantially extended and revised version of the contribution [12].

<http://www.siam.org/journals/juq/3/98131.html>

[†]Fakultät für Mathematik, Technical University of Chemnitz, 09126 Chemnitz, Germany (oliver.ernst@mathematik.tu-chemnitz.de, bjoern.sprungk@mathematik.tu-chemnitz.de).

[‡]Fachgruppe Mathematik, University of Applied Sciences Zwickau, 08012 Zwickau, Germany (hans.joerg.starkloff@fh-zwickau.de).

distinction between the two main tasks associated with Bayesian inverse problems, namely *identification* and *inference*, where the latter may include the former. Identification refers to the task of determining an element $\hat{u} \in \mathcal{X}$ which best explains the observed data z in accordance with given a priori assumptions, yielding a best guess or best single approximation to the unknown u . By inference we mean the gain in knowledge by merging a prior probabilistic model μ_0 with new information $z \in \mathbb{R}^d$ to obtain an updated model μ^z which represents the new understanding or belief about u .

This incorporation of new information is realized mathematically by conditioning the prior probability measure μ_0 on the event $\{G(u) + \varepsilon = z\}$ and is thus rooted in Kolmogorov's fundamental concept of *conditional expectation* [35]. *Bayes' rule* provides an analytic expression for the resulting conditioned or posterior distribution in terms of the prior distribution and provides the main tool in Bayesian inference and Bayesian inverse problems (BIPs).

While BIPs enjoy a number of favorable theoretical properties compared with their deterministic counterparts, i.e., they are well-posed and their solution in the form of the a posteriori measure is, in a certain sense, explicitly characterized, they do pose significant computational challenges in that they entail calculations with highly correlated and complex distributions in high-dimensional spaces. The primary "workhorse" here is the Markov chain Monte Carlo (MCMC) method [18], whose continued improvement drives a highly active field of research. However, MCMC simulations can be quite costly, since the chain has to run long enough to give sufficiently accurate estimates, and each iteration typically requires one evaluation of the forward map G , e.g., one PDE solve. Thus, for online monitoring or control of complex dynamical systems such as those that arise in weather forecasting or oil reservoir management, MCMC methods are prohibitively expensive, and filtering methods like the Kalman filter (KF) or the ensemble Kalman filter (EnKF) are often applied to the associated state or parameter estimation problem. Moreover, in dynamical systems where observational data arrives sequentially in time, KF-type methods provide the significant advantage that their recursive structure is adapted to this sequential availability of data (see [41, section 5.4] for a nice discussion of this issue). So far, the KF [25] and its generalizations have mainly been used for state estimation, i.e., for identification rather than for Bayesian inference for quantifying uncertainty. In recent years, however, these methods have attracted the attention of the growing UQ community (cf. [23, 22, 26]) and are increasingly applied also to Bayesian inverse problems. The point of departure is typically the EnKF [13], an extension of the KF to nonlinear models, in which an ensemble of random elements in \mathcal{X} is propagated through the nonlinear state equation and, given observational data, updated by a procedure analogous to the KF applied individually to each ensemble member. The empirical distribution of the updated ensemble can be used for state estimation, and additional extracted statistical information is often used as an uncertainty indicator. As a further development, the authors of [5, 34, 37, 36, 39, 38] have combined the idea of the EnKF with the computationally attractive representation of random variables in a polynomial chaos expansion (PCE) with the goal of an efficient sampling-free method for solving Bayesian inverse problems. This method models the uncertain state as a random variable approximated by a PCE, for which the coefficients are updated with the arrival of each new set of observations. We will refer to this approach in the following as the polynomial chaos Kalman filter (PCKF). It was the study of the PCKF method which motivated this work, because, although its authors give a motivation

for deriving their algorithm, the random variable approximated by the PCKF is not clearly characterized, and it is sometimes implied that, in the limit of large polynomial degree, these random variables are distributed according to the posterior measure. Similarly, it is not always clearly stated which distribution is approximated by the empirical distribution generated by the EnKF in the limit of large sample size. While convergence results for EnKF have been established [20, 29, 27], the relation of the limit distribution to the posterior measure is seldom made explicit beyond occasional hints that the two, in general, will differ [2, 20, 17].

The present paper fills this gap and clarifies the stochastic model underlying the EnKF and PCKF in the special case of a single update. We show that both methods provide different types of approximations to the same random variable, which we term the *analysis variable* and denote by U^a . Specifically, the empirical distribution of the EnKF ensemble provides an approximation to the distribution of U^a , while the PCKF constructs an approximating PCE of U^a . By explicitly characterizing U^a , we can determine the precise quantities approximated by the EnKF and PCKF and how these relate to the solution of Bayesian inverse problems and Bayes estimators. In fact, both methods fail to approximate the posterior distribution, in general.

In addition, we prove convergence results for both methods in the limit of increasing “resolution,” i.e., for large ensemble size for the EnKF and large polynomial degree for the PCKF, respectively. In the case of EnKF, our work necessarily focuses on the large ensemble limit; this contrasts with the typical EnKF application, where computational constraints limit practitioners to small ensembles. To the best of our knowledge, a convergence analysis for PCKF is missing in the literature so far. Convergence results for EnKF are given in [20, 29, 27] in the context of sequential data assimilation for discrete-time dynamical systems with possibly nonlinear but (locally) Lipschitz-continuous drift and linear observation operators. Under suitable assumptions the almost sure weak convergence of the random empirical distribution provided by the EnKF to a limit distribution is shown in [20]. Moreover, the authors establish the rate of an L^p -type convergence for these empirical measures. In [29] L^p -convergence of the ensemble members to random variables distributed according to the posterior measure is proved in the case of linear dynamics. Furthermore, the authors of [27] extend the L^p -convergence result of [20] to somewhat more general forms of dynamics and state them in terms of a suitably chosen metric for random measures. In addition, they also propose new deterministic filtering algorithms with possibly higher efficiency than the EnKF. As the main issue discussed in this work is the suitability of the filtering methods EnKF and PCKF for approximating the solution of Bayesian inverse problems, the convergence results given in the following consider the specific setting of a single Bayesian update. This may be viewed as a filtering problem without dynamics and a single observation. In contrast with [29, 20, 27], we consider an abstract Hilbert space setting, and the nonlinear parameter-to-observation mapping G need not satisfy a Lipschitz condition.

The remainder of this paper is organized as follows: Section 2 briefly recalls the Bayesian approach to inverse problems as well as Bayes estimators. In section 3 we describe and analyze the PCKF and EnKF. In particular, we prove L^2 -convergence of the random variable constructed by the PCKF to the analysis variable U^a and an almost sure weak convergence of the empirical distributions provided by the EnKF to the distribution of U^a . A characterization of the analysis variable in light of Bayes estimators is further given in section 4,

where we show that its distribution, in general, differs from the desired posterior measure. Moreover, we illustrate the performance of the EnKF and PCKF and the difference between their approximations and the solution of the Bayesian inverse problem for a one-dimensional boundary value problem and a dynamical system in section 5. Section 6 provides a summary and conclusion.

2. Bayesian inverse problems and Bayes estimators. In this section we introduce the basic concepts of the Bayesian approach to inverse problems. Throughout, let $|\cdot|$ denote the Euclidean norm on \mathbb{R}^d , $\|\cdot\|$ the norm, and $\langle \cdot, \cdot \rangle$ the inner product in a general separable Hilbert space \mathcal{X} , and let \mathcal{Y} denote a second separable Hilbert space. By $\mathcal{L}(\mathcal{X}, \mathcal{Y})$ we denote the set of all bounded linear operators $A : \mathcal{X} \rightarrow \mathcal{Y}$.

In order to regularize the usually ill-posed least-squares formulation

$$u = \operatorname{argmin}_{v \in \mathcal{X}} |z - G(v)|^2$$

of the inverse problem (1) in the deterministic setting, one incorporates additional a priori information about the desired u into the identification problem by way of a regularization functional [10], $R : \mathcal{X} \rightarrow [0, \infty]$, and solves for

$$u_\alpha = \operatorname{argmin}_{v \in \mathcal{X}} |z - G(v)|^2 + \alpha R(v),$$

where $\alpha \in [0, \infty)$ serves as a regularization parameter to be chosen wisely [1]. A further possibility for regularization is to restrict u to a subset or subspace $\tilde{\mathcal{X}} \subset \mathcal{X}$, e.g., by using a stronger norm than $\|\cdot\|$ as the regularization functional. Broadly speaking, the Bayesian approach may be viewed as yet another way of modelling prior information on u and adding it to the inverse problem. In this case we express our prior belief about u through a probability distribution μ_0 on the Hilbert space \mathcal{X} , by which a quantitative preference of some solutions u over others may be given by assigning higher and lower probabilities. However, the goal in the Bayesian approach is not the *identification* of one specific $u \in \mathcal{X}$, but rather *inference* on u ; i.e., we would like to *learn from the data* in a statistical or probabilistic sense by *adjusting our prior belief* μ_0 about u in accordance with the newly available data z . The task of identification may also be achieved within the Bayesian framework through *Bayes estimates* and *Bayes estimators*, which are discussed in section 2.3.

In the Bayesian setting the deterministic model (1) becomes

$$(2) \quad Z = G(U) + \varepsilon,$$

where now ε , and hence Z , are \mathbb{R}^d -valued random variables. For the unknown random variable U with values in \mathcal{X} and prior probability distribution μ_0 , we seek the posterior probability distribution given the available observations $Z = z$. Before giving a precise definition of the posterior distribution we require some basic concepts from probability theory.

2.1. Probability measures and random variables. Let $(\Omega, \mathcal{F}, \mathbb{P})$ denote a probability space, and let $\mathcal{B}(\mathcal{X})$ denote the Borel σ -algebra of \mathcal{X} generated by the open sets in \mathcal{X} w.r.t. $\|\cdot\|$. A measurable mapping $X : (\Omega, \mathcal{F}) \rightarrow (\mathcal{X}, \mathcal{B}(\mathcal{X}))$ is called a random variable (RV), and the measure $\mathbb{P}_X := \mathbb{P} \circ X^{-1}$, i.e., $\mathbb{P}_X(A) = \mathbb{P}(X^{-1}(A))$ for all $A \in \mathcal{B}(\mathcal{X})$, defines the

distribution of X as the push-forward measure of \mathbb{P} under X . Conversely, given a probability measure μ on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$, we indicate by $X \sim \mu$ that $\mathbb{P}_X = \mu$. Further, let $\sigma(X) \subset \mathcal{F}$ denote the σ -algebra generated by X , i.e., $\sigma(X) = \{X^{-1}(A) : A \in \mathcal{B}(\mathcal{X})\}$.

The Bochner space of p -integrable \mathcal{X} -valued RVs, i.e., the space of (equivalence classes of) RVs $X : \Omega \rightarrow \mathcal{X}$ such that $\int_{\Omega} \|X(\omega)\|^p \mathbb{P}(d\omega) < \infty$, is denoted by $L^p(\Omega, \mathcal{F}, \mathbb{P}; \mathcal{X})$ or simply $L^p(\mathcal{X})$ when the context is clear.

An element $m \in \mathcal{X}$ is called the *mean* of a RV X if for any $x \in \mathcal{X}$ there holds $\langle x, m \rangle = \mathbb{E}[\langle x, X \rangle]$. Here and in the following \mathbb{E} denotes the expectation operator w.r.t. \mathbb{P} . If $X \in L^1(\Omega, \mathcal{F}, \mathbb{P}; \mathcal{X})$, then its mean is given by the Bochner integral $m = \mathbb{E}[X] = \int_{\Omega} X(\omega) \mathbb{P}(d\omega)$. A bilinear form $C : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ is called the *covariance* $\text{Cov}(X, Y)$ of two RVs $X : \Omega \rightarrow \mathcal{X}$ and $Y : \Omega \rightarrow \mathcal{Y}$ if it satisfies $C(x, y) = \mathbb{E}[\langle x, X - \mathbb{E}[X] \rangle \langle y, Y - \mathbb{E}[Y] \rangle]$ for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, and we set $\text{Cov}(X) := \text{Cov}(X, X)$. We shall also employ the identity $\text{Cov}(X, Y) = \mathbb{E}[(X - \mathbb{E}[X]) \otimes (Y - \mathbb{E}[Y])]$ (see [41, section 6.2]) when convenient. The covariance $\text{Cov}(X, Y)$ can also be equivalently defined as an operator $\hat{C} : \mathcal{Y} \rightarrow \mathcal{X}$ such that $\langle x, \hat{C}y \rangle = C(x, y)$. We will mainly work with the latter definition in the following, but on occasion will also employ the tensor product form $\mathbb{E}[(X - \mathbb{E}[X]) \otimes (Y - \mathbb{E}[Y])]$. The definitions of mean and covariance extend to RVs with values in separable Banach spaces by considering the topological duals of \mathcal{X} and \mathcal{Y} , respectively.

We also require the notion of distance between probability measures, one of which is given by the *Hellinger metric* d_H : Given two probability measures μ_1 and μ_2 on the Hilbert space \mathcal{X} , it is defined as

$$d_H(\mu_1, \mu_2) := \left[\int_{\mathcal{X}} \left(\sqrt{\frac{d\mu_1}{d\nu}}(u) - \sqrt{\frac{d\mu_2}{d\nu}}(u) \right)^2 \nu(du) \right]^{1/2},$$

where ν is a dominating measure for μ_1 and μ_2 , e.g., $\nu = (\mu_1 + \mu_2)/2$. Note that the definition of the Hellinger metric is independent of the dominating measure. For relations of the Hellinger metric to other probability metrics such as total variation distance, we refer the reader to [19].

In the following, we will use uppercase Latin letters such as X, Y, Z, U to denote RVs on Hilbert spaces, and lowercase Latin letters such as x, y, z, u for elements in these Hilbert spaces or realizations of the associated RVs. Greek letters such as ε, η , and ξ will be used to denote RVs as well as their realizations, and μ and ν (with various subscripts) will denote measures on the Hilbert space \mathcal{X} and on \mathbb{R}^d .

2.2. Bayes' rule and the posterior measure. Bayesian inference consists in updating prior knowledge on the unknown quantity U , reflecting a gain in knowledge due to new observations. The distribution of the RV U , characterized by the probabilities $\mathbb{P}(U \in B)$ for $B \in \mathcal{B}(\mathcal{X})$, quantifies in stochastic terms our knowledge about the uncertainty associated with U . When new information becomes available, such as knowing that the event $Z = z$ has occurred, this is reflected in our quantitative description as the “conditional distribution of U given $\{Z = z\}$,” denoted $\mathbb{P}(U \in B | Z = z)$. Unfortunately, $\mathbb{P}(U \in B | Z = z)$ cannot be defined in an elementary fashion when $\mathbb{P}(Z = z) = 0$, in which case the conditional distribution is defined by an integral relation. The key concept here is that of *conditional expectation*: Given RVs $X \in L^1(\Omega, \mathcal{F}, \mathbb{P}; \mathcal{X})$ and $Y : \Omega \rightarrow \mathcal{Y}$, we define the conditional expectation $\mathbb{E}[X | Y]$ of X given

Y as any $\sigma(Y)$ -measurable mapping $\mathbb{E}[X|Y] : \Omega \rightarrow \mathcal{X}$ which satisfies

$$\int_A \mathbb{E}[X|Y] \mathbb{P}(d\omega) = \int_A X \mathbb{P}(d\omega) \quad \forall A \in \sigma(Y).$$

By the Doob–Dynkin lemma [24, Lemma 1.13] there exists a measurable function $\phi : \mathcal{Y} \rightarrow \mathcal{X}$ such that $\mathbb{E}[X|Y] = \phi(Y)$ \mathbb{P} -a.s. We note that this does not determine a unique function ϕ but rather an equivalence class of measurable functions, where $\phi_1 \sim \phi_2$ iff $\mathbb{P}(Y \in \{y \in \mathcal{Y} : \phi_1(y) \neq \phi_2(y)\}) = 0$. For a specific realization y of Y (and a specific ϕ), we define

$$\mathbb{E}[X|Y = y] := \phi(y) \in \mathcal{X}.$$

Setting $X = \mathbf{1}_{\{U \in B\}}$, one can then define for each fixed $B \in \mathcal{B}(\mathcal{X})$

$$(3) \quad \mathbb{P}(U \in B|Z = z) := \mathbb{E}[\mathbf{1}_{\{U \in B\}}|Z = z]$$

as an equivalence class of measurable functions $\mathbb{R}^d \rightarrow [0, 1]$. One would like to view this, conversely, as a family of probability measures with the realization z as a parameter, giving the posterior distribution of U resulting from having made the observation $Z = z$. Unfortunately, this construction need not, in general, yield a probability measure for each fixed value of z (cf. [35]). In case \mathcal{X} is a separable Hilbert space, a function

$$Q : \mathcal{B}(\mathcal{X}) \times \mathbb{R}^d \rightarrow \mathbb{R}$$

can be shown to exist (cf. [35]) such that the following hold:

- (a). For each $z \in \mathbb{R}^d$, $Q(\cdot, z)$ is a probability measure on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$.
- (b). For each $B \in \mathcal{B}(\mathcal{X})$ the function

$$\mathbb{R}^d \ni z \mapsto Q(B, z)$$

is a representative of the equivalence class (3); i.e., it is measurable and there holds

$$\mathbb{P}(U \in B, Z \in A) = \int_A Q(B, z) \mathbb{P}_Z(dz) \quad \forall A \in \mathcal{B}(\mathbb{R}^d).$$

Such a function Q , also denoted by $\mu_{U|Z}$, is called the *regular conditional distribution* of U given Z and is defined uniquely up to sets of z -values of \mathbb{P}_Z -measure zero. We have thus arrived at a consistent definition of the posterior probability $\mathbb{P}(U \in B|Z = z)$ as $\mu_{U|Z}(B, z)$.

It is helpful to maintain a clear distinction between *conditional* and *posterior* quantities: The former contain the—as yet unrealized—observation as a parameter, while in the latter the observation has been made. Specifically, $\mu_{U|Z}$ is the conditional measure of U conditioned on Z , whereas $\mu_{U|Z}(\cdot, z)$ denotes the posterior measure of U for the observation $Z = z$.

We now recall how Bayes' rule yields an explicit expression for the regular conditional distribution $\mu_{U|Z}$. To this end, we make the following assumptions for the model (2).

Assumption 2.1.

1. $U \sim \mu_0$, $\varepsilon \sim \nu_\varepsilon$, and $(U, \varepsilon) \sim \mu_0 \otimes \nu_\varepsilon$; in particular, U and ε are independent.
2. $\nu_\varepsilon = \rho(\varepsilon) d\varepsilon$, where $\rho(\varepsilon) = Ce^{-\ell(\varepsilon)}$ with $C > 0$ and $\ell : \mathbb{R}^d \rightarrow \mathbb{R}_0^+$ measurable and nonnegative. Here $d\varepsilon$ denotes Lebesgue measure on \mathbb{R}^d .
3. $G : \mathcal{X} \rightarrow \mathbb{R}^d$ is continuous.

By Assumption 2.1, the distribution ν_Z of Z in (2) is determined as $\nu_Z = C\gamma(z)dz$, where $C > 0$ and

$$\gamma(z) := \int_{\mathcal{X}} e^{-\ell(z-G(u))} \mu_0(du).$$

We note that $\gamma(z) > 0$ is well defined since $0 < |e^{-\ell(z-G(u))}| \leq 1$ and $\gamma \in L^1(\mathbb{R}^d)$ due to Fubini’s theorem [24, Theorem 1.27]. In particular, we have that $(U, Z) \sim \mu$ with $\mu(du, dz) = Ce^{-\ell(z-G(u))} \mu_0(du) \otimes dz$, where dz again denotes Lebesgue measure on \mathbb{R}^d . Further, we introduce the *potential*

$$\Phi(u; z) := \ell(z - G(u)),$$

for which we assume the following Lipschitz-like property.

Assumption 2.2. *The potential Φ is continuous in z in a mean-square sense w.r.t. μ_0 ; i.e., there exists a nondecreasing function $\psi : [0, \infty) \rightarrow [0, \infty)$ with $\lim_{s \rightarrow 0} \psi(s) = \psi(0) = 0$ such that*

$$\mathbb{E} [|\Phi(U; z_1) - \Phi(U; z_2)|^2] = \int_{\mathcal{X}} |\Phi(u; z_1) - \Phi(u; z_2)|^2 \mu_0(du) \leq \psi(|z_1 - z_2|).$$

For instance, there may exist a function $\theta \in L^2(\mathcal{X}, \mathcal{B}(\mathcal{X}), \mu_0; \mathbb{R})$ such that

$$|\Phi(u; z_1) - \Phi(u; z_2)| \leq \theta(u) \psi(|z_1 - z_2|).$$

Before stating the abstract version of Bayes’ rule in Theorem 2.3, we recall the finite-dimensional case $\mathcal{X} \simeq \mathbb{R}^n$, where it can be stated in terms of densities: Here $\mu_0(du) = \pi_0(u)du$, and *Bayes’ rule* takes the form

$$\pi^z(u) = \frac{1}{\gamma(z)} \exp(-\Phi(u; z)) \pi_0(u),$$

where $e^{-\Phi(u; z)} = e^{-\ell(z-G(u))}$ represents the *likelihood* of observing z when fixing u . The denominator $\gamma(z)$ can be interpreted as a *normalizing constant* to ensure $\int_{\mathcal{X}} \pi^z(u) du = 1$. We now show that, in the general setting, Bayes’ rule yields (a version of) the (regular) conditional measure $\mu_{U|Z}$ of U w.r.t. Z . The statement of Theorem 2.3 differs from related results in [41, Theorems 4.2 and 6.31] insofar as we explicitly characterize the posterior measure as a version of the regular conditional distribution, and as we also allow for a general prior μ_0 and log-likelihood ℓ .

Theorem 2.3. *Let Assumptions 2.1 and 2.2 be satisfied, and define for each $z \in \mathbb{R}^d$ a probability measure on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$ by*

$$(4) \quad \mu^z(du) := \frac{1}{\gamma(z)} \exp(-\Phi(u; z)) \mu_0(du).$$

Then the mapping $Q : \mathcal{B}(\mathcal{X}) \times \mathbb{R}^d \rightarrow [0, 1]$ given by

$$Q(B, z) := \mu^z(B) \quad \forall B \in \mathcal{B}(\mathcal{X})$$

is a regular conditional distribution of U given Z . We call μ^z the posterior measure (of U given $Z = z$). Moreover, μ^z depends continuously on z w.r.t. the Hellinger metric; i.e., for any $z_1, z_2 \in \mathbb{R}^d$ with $|z_1 - z_2| \leq r$, there holds

$$d_H(\mu^{z_1}, \mu^{z_2}) \leq C_r(z_1) \psi(|z_1 - z_2|),$$

where $C_r(z_1) = C(1 + \min\{\gamma(z') : |z_1 - z'| \leq r\})^{-1} < +\infty$.

Proof. Continuity with respect to the Hellinger metric is a slight generalization of [41, Theorem 4.2] and may be proved in the same way with obvious modifications. To show that Q is a regular conditional distribution, we verify the two properties (a) and (b). The first follows from the construction of μ^z . For the second property, note that measurability follows from continuity. The continuity of μ^z w.r.t. z in the Hellinger metric also implies that $\mu^z(B)$ depends continuously on z due to the relations between Hellinger metric and total variation distance (see [19]). Finally, we have for any $A \in \mathcal{B}(\mathbb{R}^d)$ and $B \in \mathcal{B}(\mathcal{X})$ that

$$\begin{aligned} \mathbb{P}(U \in B, Z \in A) &= \int_{A \times B} \mu(du, dz) = \int_A \int_B C e^{-\ell(z-G(u))} \mu_0(du) dz \\ &= \int_A C \gamma(z) Q(B, z) dz = \int_A Q(B, z) \mathbb{P}_Z(dz), \end{aligned}$$

which completes the proof. ■

Remark 2.4. Theorem 2.3 shows that the Lipschitz-like property of the potential stated in Assumption 2.2 carries over to the posterior for a general prior μ_0 and an additive noise ε with Lebesgue density proportional to $e^{-\ell(\varepsilon)}$. Roughly speaking, the negative log-likelihood ℓ and the posterior μ^z share the same local modulus of continuity.

By Theorem 2.3 the Bayesian inverse problem is well-posed under mild conditions. It is also possible to prove continuity of μ^z w.r.t. to the forward map G (see [41, section 4.4]), which is crucial when the forward map G is realized by numerical approximation.

To give meaning to the mean and covariance of $U \sim \mu_0$ and $Z = G(U) + \varepsilon$, we make the further assumption that all second moments exist.

Assumption 2.5. *There holds*

$$\int_{\mathcal{X}} (\|u\|^2 + |G(u)|^2) \mu_0(du) < +\infty \quad \text{and} \quad \int_{\mathbb{R}^d} |\varepsilon|^2 \nu_\varepsilon(d\varepsilon) < +\infty.$$

2.3. Bayes estimators. Although the posterior measure μ^z is, by definition, the solution to the Bayesian inverse problem, it is by no means easy to compute in practice. In special cases, such as when G is linear and μ_0 and ν_ε are Gaussian measures, or in the case of conjugate priors, closed-form expressions for μ^z are available. In general, however, μ^z can only be computed in an approximate sense. Moreover, when the dimension of \mathcal{X} is large or infinite, visualizing, exploring, and using μ^z for postprocessing are demanding tasks.

More accessible quantities from Bayesian statistics [4], which are also closer in nature to the result of deterministic parameter identification procedures than the posterior measure, are *point estimates* for the unknown u . In the Bayesian setting a point estimate is a “best guess” \hat{u} of u based on posterior knowledge. Here “best” is determined by a *cost function* $c : \mathcal{X} \rightarrow \mathbb{R}_0^+$ satisfying $c(0) = 0$ and $c(u) \leq c(\lambda u)$ for any $u \in \mathcal{X}$ and $\lambda \geq 1$. Such a cost function describes

the loss or costs $c(u - \hat{u})$ incurred when \hat{u} is substituted for (the true) u for postprocessing or decision making. Also, more general forms of a cost function are possible; see, e.g., [3, 4].

For any realization $z \in \mathbb{R}^d$ of the observation RV Z we introduce the (*posterior*) *Bayes cost* of the estimate \hat{u} w.r.t. c as

$$B_c(\hat{u}; z) := \int_{\mathcal{X}} c(u - \hat{u}) \mu^z(du),$$

and define the *Bayes estimate* \hat{u} as a minimizer of this cost, i.e.,

$$\hat{u} := \operatorname{argmin}_{v \in \mathcal{X}} B_c(v; z),$$

assuming a unique minimizer exists. The *Bayes estimator* $\hat{\phi} : \mathbb{R}^d \rightarrow \mathcal{X}$ is then the mapping which assigns to an observation z the associated Bayes estimate \hat{u} , i.e.,

$$\hat{\phi} : z \mapsto \operatorname{argmin}_{v \in \mathcal{X}} B_c(v; z).$$

We assume measurability of $\hat{\phi}$ in the following and note that $\hat{\phi}$ is then also the minimizer of the *expected* or *prior Bayes cost*

$$B_c(\phi) := \mathbb{E}[B_c(\phi(Z); Z)] = \int_{\mathbb{R}^d} \int_{\mathcal{X}} c(u - \phi(z)) \mu^z(du) \nu_Z(dz) = \mathbb{E}[c(U - \phi(Z))];$$

i.e., for any other measurable $\phi : \mathbb{R}^d \rightarrow \mathcal{X}$ there holds

$$\mathbb{E}[c(U - \hat{\phi}(Z))] \leq \mathbb{E}[c(U - \phi(Z))].$$

Remark 2.6. Since $\hat{\phi} = \operatorname{argmin}_{\phi} \mathbb{E}[c(U - \phi(Z))]$, it is possible to determine the estimator $\hat{\phi}$, and hence also the estimate $\hat{u} = \hat{\phi}(z)$ for a given z , without actually computing the posterior measure μ^z , as the integration in $B_c(\hat{\phi})$ is carried out w.r.t. the prior measure. Therefore, Bayes estimators are typically easier to compute or approximate than μ^z .

We now recall two very common Bayes estimators: the *posterior mean estimator* and the *maximum a posteriori estimator*.

2.3.1. Posterior mean estimator. For the cost function $c(u) = \|u\|^2$ the posterior Bayes cost

$$B_c(\hat{u}; z) = \int_{\mathcal{X}} \|u - \hat{u}\|^2 \mu^z(du)$$

is minimized by the posterior mean $\hat{u} = u_{\text{CM}} := \int_{\mathcal{X}} u \mu^z(du)$, since for any Hilbert space-valued RV X its expectation $\mathbb{E}[X]$ is the minimizer of the functional $J_X(v) = \mathbb{E}[\|X - v\|^2]$, $v \in \mathcal{X}$. The corresponding Bayes estimator for $c(u) = \|u\|^2$ is then given by

$$\hat{\phi}_{\text{CM}}(z) := \int_{\mathcal{X}} u \mu^z(du).$$

In particular, $\hat{\phi}_{\text{CM}}(Z) = \mathbb{E}[U|Z]$ holds \mathbb{P} -a.s.

Remark 2.7. If \mathcal{X} is only a Banach space, then the expectation of an \mathcal{X} -valued RV X need not minimize the functional J_X ; i.e., we have in general

$$\mathbb{E}[X] \neq \operatorname{argmin}_{v \in \mathcal{X}} \mathbb{E}[\|X - v\|^2].$$

As a simple counterexample, consider $\mathcal{X} = \mathbb{R}^2$, $\|v\| = |v_1| + |v_2|$, and $X = (X_1, X_2)$ with independent RVs X_1, X_2 such that

$$\mathbb{P}(X_1 = -1) = p_1, \quad \mathbb{P}(X_1 = 1) = 1 - p_1 \quad \text{and} \quad \mathbb{P}(X_2 = -1) = p_2, \quad \mathbb{P}(X_2 = 1) = 1 - p_2.$$

Here $\mathbb{E}[X]$ minimizes $\mathbb{E}[\|X - v\|^2]$ iff $p_1 = p_2 = 0.5$. In fact, one can show that $\mathbb{E}[X] = \operatorname{argmin}_{v \in \mathcal{X}} \mathbb{E}[\|X - v\|^2]$ if X is distributed symmetrically w.r.t. its mean, i.e., if there holds $\mathbb{P}(X - \mathbb{E}[X] \in A) = \mathbb{P}(\mathbb{E}[X] - X \in A)$ for all $A \in \mathcal{B}(\mathcal{X})$.

2.3.2. Maximum a posteriori estimator. Another common estimator in Bayesian statistics is the *maximum a posteriori (MAP)* estimator $\hat{\phi}_{\text{MAP}}$. For finite-dimensional $\mathcal{X} \simeq \mathbb{R}^n$ and absolutely continuous prior μ_0 , i.e., $\mu_0(du) = \pi_0(u)du$, the MAP estimate is defined as

$$\hat{\phi}_{\text{MAP}}(z) = \operatorname{argmin}_{u \in \mathbb{R}^n} (\Phi(u; z) - \log \pi_0(u)),$$

provided that the minimum exists for all $z \in \mathbb{R}^d$. For the definition of the MAP estimate via a cost function and the Bayes cost, we refer the reader to the literature, e.g., [28, section 16.2] or the very recent work [6] for a novel approach; for MAP estimates in infinite dimensions, we refer the reader to [9].

There is an interesting link between the Bayes estimator $\hat{\phi}_{\text{MAP}}$ and the solution of the associated regularized least-squares problem: If $R : \mathbb{R}^n \rightarrow [0, \infty)$ is a regularizing functional which satisfies $\int_{\mathbb{R}^n} \exp(-\frac{\alpha}{\sigma^2} R(u)) du < +\infty$, then the solution $\hat{u}_\alpha = \operatorname{argmin}_u |z - G(u)|^2 + \alpha R(u)$ coincides with the MAP estimate $\hat{\phi}_{\text{MAP}}(z)$ for $\varepsilon \sim N(0, \sigma^2 I)$ and $\mu_0(du) \propto \exp(-\frac{\alpha}{\sigma^2} R(u)) du$.

3. Analysis of generalized Kalman filters. In this section we consider KFs and their application to the nonlinear Bayesian inverse problem (2). We begin with the classical KF for state estimation in linear dynamics and then consider two generalizations to the nonlinear setting which have been recently proposed for UQ in connection with inverse problems. We show that both methods can be understood as different discretizations of an updating scheme for a certain RV, and prove that both KF methods converge to this RV when the discretization is refined.

3.1. The Kalman filter. The KF [25] is a well-known method for sequential state estimation for incompletely observable, linear discrete-time dynamics of the form

$$(5) \quad U_n = A_n U_{n-1} + \eta_n, \quad Z_n = G_n U_n + \varepsilon_n, \quad n = 1, 2, \dots,$$

where $(U_n)_{n \in \mathbb{N}}$ denotes the unknown, unobservable state and $(Z_n)_{n \in \mathbb{N}}$ the observable process. The operators A_n and G_n are linear mappings in state space and from state to observation space, respectively, and the noise processes η_n, ε_n are usually assumed to have zero mean

with known covariances. In addition, the mean and covariance of U_0 need to be known, and the RVs $U_0, \eta_n, \varepsilon_n$ are taken to be mutually independent. Then, given observations $Z_1 = z_1, \dots, Z_n = z_n$, the KF yields recursive equations for the minimum variance estimates \hat{u}_n of U_n and their error covariances $\text{Cov}(U_n - \hat{u}_n)$; see, e.g., [8, 40] for an introduction and discussion.

Although the main advantage of the KF is its recursive structure, making it very efficient for state estimation in dynamical systems with sequentially arriving data, a detailed analysis of sequential methods is beyond the scope of this work. We focus instead on the application of the KF and its generalizations to time-independent systems of the form (2) and, in the linear case,

$$(6) \quad Z = GU + \varepsilon, \quad (U, \varepsilon) \sim \mu_0 \otimes \nu_\varepsilon.$$

We note that (6) can be seen as one step of the dynamical system (5) for $A_n \equiv I$, $\eta_n \equiv 0$, and $G_n = G$. Conversely, the state estimation problem for U_0, U_n , or (U_0, U_1, \dots, U_n) in (5) given $Z = (Z_1, \dots, Z_n) = (z_1, \dots, z_n) = z$ can be reformulated as (6).

If $\hat{u}_0 = \mathbb{E}[U]$ is taken as an initial estimate for the unknown U in (6) before observing $Z = z$, this results in the initial error covariance $\text{Cov}(U - \hat{u}_0) = \text{Cov}(U) =: C_0$. Given data $Z = z$, the KF provides a new estimate \hat{u}_1 and its error covariance $C_1 = \text{Cov}(U - \hat{u}_1)$ via the updates

$$(7) \quad \hat{u}_1 = \hat{u}_0 + K(z - G\hat{u}_0), \quad C_1 = C_0 - KGC_0,$$

where $K = C_0G^*(GC_0G^* + \Sigma)^{-1}$, $\Sigma = \text{Cov}(\varepsilon)$, is known as the *Kalman gain*. In fact, by assimilating the data $Z = z$, the KF produces an improved estimate, since its expected error is smaller than that of the initial estimate in the sense that $C_0 - C_1$ is positive definite.

If (U, Z) are jointly Gaussian RVs, i.e., $U \sim N(m_0, C_0)$ and $\varepsilon \sim N(0, \Sigma)$, the posterior measure μ^z of U given $Z = z$ also has a Gaussian distribution $\mu^z \sim N(m^z, C^z)$ with

$$m^z = m_0 + K(z - Gm_0), \quad C^z = C_0 - KGC_0;$$

see, e.g., [30]. Thus for G linear and U, ε independently Gaussian, the KF is seen to yield the solution of the Bayesian inverse problem by providing the posterior mean and covariance, which in this case also uniquely specify the Gaussian posterior measure μ^z . However, we emphasize that the KF does not directly approximate the posterior measure; it rather provides minimum variance estimates and their error covariances for linear problems (5). Without the assumption that μ_0 or ν_ε is Gaussian, the Kalman filter will not, in general, yield the first two posterior moments, nor is the posterior measure necessarily Gaussian.

In the following two subsections we consider generalizations of the KF to nonlinear problems (2). Historically, the first such method was the extended Kalman filter (EKF), which is based on local linearizations of the nonlinear map G , but which we will not consider here. We focus instead on the ensemble Kalman Filter (EnKF) introduced by Evensen [13] and the recently developed the polynomial chaos Kalman filter (PCKF) [5, 34, 36, 37, 38, 39].

3.2. The ensemble Kalman filter. Since its introduction in 1994, the EnKF has been investigated and evaluated in many publications [14, 7, 16, 15, 32]. However, the focus is usually on its application to state or parameter estimation rather than solving Bayesian inverse problems. Recently, interest in the EnKF for UQ in inverse problems has increased; see, e.g., [22, 23, 26].

If we consider the model (2) with $(U, \varepsilon) \sim \mu_0 \otimes \nu_\varepsilon$ and given observations $z \in \mathbb{R}^d$, the EnKF algorithm proceeds as follows:

1. **Initial ensemble:** Draw samples u_1, \dots, u_M of $U \sim \mu_0$.
2. **Forecast:** Draw samples $\varepsilon_1, \dots, \varepsilon_M$ of $\varepsilon \sim \nu_\varepsilon$ and set

$$z_j = G(u_j) + \varepsilon_j, \quad j = 1, \dots, M,$$

yielding samples z_1, \dots, z_M of $Z \sim \nu_Z$.

3. **Analysis:** Update the initial ensemble $\mathbf{u} = (u_1, \dots, u_M)$ member by member via

$$(8) \quad u_j^a = u_j + \tilde{K}(z - z_j), \quad j = 1, \dots, M,$$

where $\tilde{K} = \text{Cov}(\mathbf{u}, \mathbf{z})\text{Cov}(\mathbf{z})^{-1}$ and $\text{Cov}(\mathbf{u}, \mathbf{z})$ and $\text{Cov}(\mathbf{z})$ are the empirical covariances of the samples \mathbf{u} and $\mathbf{z} = (z_1, \dots, z_M)$, e.g.,

$$\text{Cov}(\mathbf{u}, \mathbf{z}) = \frac{1}{M-1} \sum_{j=1}^M (u_j - \bar{u}_M) \otimes (z_j - \bar{z}_M),$$

where $\bar{u}_M = \frac{1}{M}(u_1 + \dots + u_M)$ and $\bar{z}_M = \frac{1}{M}(z_1 + \dots + z_M)$. This yields an *analysis ensemble* $\mathbf{u}^a = (u_1^a, \dots, u_M^a)$ which in turn determines an *empirical analysis measure*

$$(9) \quad \tilde{\mu}_M^a = \frac{1}{M} \sum_{j=1}^M \delta_{u_j^a},$$

where $\delta_{u_j^a}$ denotes the Dirac measure at the point u_j^a . Moreover, the empirical mean of \mathbf{u}^a serves as an estimate for the unknown u , and the empirical covariance of \mathbf{u}^a as an indicator for the accuracy of the estimate.

For dynamical systems such as (5), the analysis ensemble \mathbf{u}^a would be propagated by the system dynamics and would then serve as the initial ensemble for the subsequent step n .

3.3. The polynomial chaos Kalman filter. In [5, 34, 36, 37, 38, 39] the authors propose a sampling-free Kalman filtering scheme for nonlinear systems. Rather than updating samples of the unknown, this is carried out for the coefficient vector of a polynomial chaos expansion (PCE) of the unknown. This necessitates the construction of a PCE distributed according to the prior measure μ_0 : We assume there exist countably many independent real-valued RVs $\xi = (\xi_m)_{m \in \mathbb{N}}$ and *chaos coefficients* $u_\alpha \in \mathcal{X}$, $\varepsilon_\alpha \in \mathbb{R}^d$ for each multi-index

$$\alpha \in \mathbb{J} := \{\alpha \in \mathbb{N}_0^{\mathbb{N}} : \alpha_j \neq 0 \text{ for only finitely many } j\}$$

such that

$$\sum_{\alpha \in \mathbb{J}} \|u_\alpha\|^2 < +\infty \quad \text{and} \quad \sum_{\alpha \in \mathbb{J}} |\varepsilon_\alpha|^2 < +\infty$$

and

$$\left(\sum_{\alpha \in \mathbb{J}} u_\alpha P_\alpha(\boldsymbol{\xi}), \sum_{\alpha \in \mathbb{J}} \varepsilon_\alpha P_\alpha(\boldsymbol{\xi}) \right) \sim \mu_0 \otimes \nu_\varepsilon.$$

Here, $P_\alpha(\boldsymbol{\xi}) = \prod_{m \geq 1} P_{\alpha_m}^{(m)}(\xi_m)$ denotes the product of univariate orthogonal polynomials $P_{\alpha_m}^{(m)}$, where we require $\{P_\alpha^{(m)}\}_{\alpha \in \mathbb{N}}$ to be a complete orthonormal system (CONS) in $L^2(\mathbb{R}, \mathcal{B}(\mathbb{R}), \mathbb{P}_{\xi_m}; \mathbb{R})$. We note that the completeness of orthogonal polynomials will depend in general on properties of the measure \mathbb{P}_{ξ_m} ; see [11] for a complete characterization.

We then define $U := \sum_{\alpha \in \mathbb{J}} u_\alpha P_\alpha(\boldsymbol{\xi})$ and $\varepsilon := \sum_{\alpha \in \mathbb{J}} \varepsilon_\alpha P_\alpha(\boldsymbol{\xi})$, given the chaos coefficients $(u_\alpha)_{\alpha \in \mathbb{J}}$ and $(\varepsilon_\alpha)_{\alpha \in \mathbb{J}}$. However, for numerical simulations we have to truncate the PCE and, therefore, introduce the projection

$$P_J U := \sum_{\alpha \in J} u_\alpha P_\alpha(\boldsymbol{\xi}), \quad J \subset \mathbb{J}.$$

To simplify notation we further define for $J \subseteq \mathbb{J}$ the following two RVs:

$$U_J := P_J U \quad \text{and} \quad Z_J := P_J(G(U_J) + \varepsilon).$$

Due to the nonlinearity of G there holds in general $P_J G(U) \neq G(P_J U) \neq P_J G(U_J)$, and, hence, $Z_J \neq P_J Z$! In particular, we will consider finite subsets J , and for convergence studies we usually assume a monotone and exhaustive sequence of such finite subsets $(J_n)_{n \in \mathbb{N}}$, i.e., $J_m \subset J_n$ for $m \leq n$ and $J_n \uparrow \mathbb{J}$, e.g., the sequence

$$J_n := \left\{ \alpha \in \mathbb{J} : \alpha_j = 0 \ \forall j > n, \sum_{j=1}^{\infty} |\alpha_j| \leq n \right\}.$$

We note that for $n \rightarrow \infty$ the error $\|U - U_{J_n}\|_{L^2(\mathcal{X})}$ will tend to zero since $J_n \uparrow \mathbb{J}$. However, the L^2 -convergence is in general not preserved under continuous mappings (unlike convergence in the almost sure sense, in probability and in distribution). Thus, although there holds $\|U - U_{J_n}\|_{L^2(\mathcal{X})} \rightarrow 0$ and, of course, $\|G(U) - P_{J_n} G(U)\|_{L^2(\mathbb{R}^d)} \rightarrow 0$, the continuity of G does not imply $\|G(U) - P_{J_n} G(U_{J_n})\|_{L^2(\mathbb{R}^d)}^2 \rightarrow 0$ in general. If we assume that for a $\delta > 0$ there exists a $C < +\infty$ such that

$$(10) \quad \mathbb{E} \left[|G(U_{J_n})|^{2+\delta} \right] \leq C \quad \forall n \in \mathbb{N},$$

then the desired convergence of $\|Z - Z_{J_n}\|_{L^2(\mathbb{R}^d)} \rightarrow 0$ follows; see the proof of Theorem 3.2 for details.

For the same problem considered for the EnKF, the PCKF algorithm now reads as follows:

1. **Initialization:** Choose a finite subset $J \subset \mathbb{J}$ and compute the chaos coefficients $(u_\alpha)_{\alpha \in J}$ of $U \sim \mu_0$.

2. **Forecast:** Compute the chaos coefficients $(g_{J,\alpha})_{\alpha \in J}$ of $G(U_J)$ and set

$$z_{J,\alpha} := g_{J,\alpha} + \varepsilon_\alpha \quad \forall \alpha \in J,$$

where $(\varepsilon_\alpha)_{\alpha \in J}$ are the chaos coefficients of ε . By linearity $z_{J,\alpha}$ are the chaos coefficients of Z_J .

3. **Analysis:** Update the initial chaos coefficients by

$$(11) \quad u_{J,\alpha}^a := u_\alpha + K_J (\delta_{\alpha \mathbf{0}} z - z_{J,\alpha}) \quad \forall \alpha \in J,$$

where $\delta_{\alpha \mathbf{0}}$ is the Kronecker symbol for multi-indices, $(\delta_{\alpha \mathbf{0}} z)_{\alpha \in J} = (z, 0, \dots, 0)$ are the chaos coefficients of the observed data $z \in \mathbb{R}^d$, and $K_J := \text{Cov}(U_J, Z_J) \text{Cov}(Z_J)^{-1}$. The action of the covariances as linear operators can be described in the case of $\text{Cov}(U_J, Z_J) : \mathbb{R}^d \rightarrow \mathcal{X}$ by

$$\text{Cov}(U_J, Z_J)x = \sum_{\alpha \in J} \sum_{\beta \in J} z_{J,\beta}^\top x u_\alpha, \quad x \in \mathbb{R}^d.$$

Thus, the result of one step of the PCKF algorithm is an *analysis chaos coefficient vector* $(u_\alpha^a)_{\alpha \in J}$, which in turn determines a RV

$$U_J^a := \sum_{\alpha \in J} u_{J,\alpha}^a P_\alpha(\boldsymbol{\xi}).$$

Remark 3.1. An expansion in polynomials $P_\alpha(\boldsymbol{\xi})$ is not crucial for the application of the PCKF. In principle, any countable CONS $(\Psi_\alpha)_{\alpha \in \mathbb{N}}$ of the space $L^2(\mathbb{R}^N, \mathcal{B}(\mathbb{R}^N), \mathbb{P}_\boldsymbol{\xi}; \mathbb{R})$ such that $(\sum_\alpha u_\alpha \Psi_\alpha(\boldsymbol{\xi}), \sum_\alpha \varepsilon_\alpha \Psi_\alpha(\boldsymbol{\xi})) \sim \mu_0 \otimes \nu_\varepsilon$ would be suitable.

3.4. The analysis variable. Both EnKF and PCKF perform discretized versions of an update for RVs, namely,

$$(12) \quad U^a = U + K(z - Z), \quad K = \text{Cov}(U, Z) \text{Cov}(Z)^{-1},$$

where $Z := G(U) + \varepsilon$ and $(U, \varepsilon) \sim \mu_0 \otimes \nu_\varepsilon$, providing samples \mathbf{u}^a or chaos coefficients u_α^a of U^a , respectively. However, the output of both methods is corrupted by the approximation of the Kalman gain operator K by the empirical covariances and the operator K_J , respectively. That both methods do indeed converge to U^a in some sense for increasing sample size M or increasing chaos coefficient subset J_n is stated in the next two theorems.

Theorem 3.2. *Consider the model (2), and let Assumptions 2.1, 2.2, and 2.5 be satisfied. If $(J_n)_{n \in \mathbb{N}}$ is a monotone and exhaustive sequence of finite subsets of \mathbb{J} with $\mathbf{0} \in J_1$ such that (10) holds, then $\|Z - Z_{J_n}\|_{L^2(\mathbb{R}^d)} \rightarrow 0$ for $n \rightarrow \infty$. Moreover, if*

$$U_{J_n}^a = \sum_{\alpha \in J_n} u_{J_n,\alpha}^a P_\alpha(\boldsymbol{\xi})$$

denotes the RV generated by the PCKF in the analysis step for the subset $J = J_n$, we have

$$(13) \quad \|U^a - U_{J_n}^a\|_{L^2(\mathcal{X})} \in \mathcal{O} \left(\|U - U_{J_n}\|_{L^2(\mathcal{X})} + \|Z - Z_{J_n}\|_{L^2(\mathbb{R}^d)} \right),$$

which means in particular that $U_{J_n}^a \rightarrow U^a$ in $L^2(\mathcal{X})$ as $n \rightarrow \infty$.

Proof. In the following we use $\|\cdot\|_{L^2}$ as shorthand for $\|\cdot\|_{L^2(\mathcal{X})}$ and $\|\cdot\|_{L^2(\mathbb{R}^d)}$, respectively. Since $(J_n)_{n \in \mathbb{N}}$ is exhaustive, we have $U_{J_n} \rightarrow U$ in $L^2(\mathcal{X})$, and hence $U_{J_n} \xrightarrow{\mathbb{P}} U$, where $\xrightarrow{\mathbb{P}}$ denotes convergence in probability. Since G is continuous, it follows by the continuous mapping theorem [24, Lemma 3.3] that also $G(U_{J_n}) \xrightarrow{\mathbb{P}} G(U)$. Now the boundedness assumption (10) implies the uniform integrability of the RVs $|G(U_{J_n})|^2$, $n \in \mathbb{N}$ (see [24, p. 44]), and by [24, Proposition 3.12] we then obtain $G(U_{J_n}) \rightarrow G(U)$ in $L^2(\mathcal{X})$. Thus,

$$\|Z - Z_{J_n}\|_{L^2} \leq \underbrace{\|Z - P_{J_n} Z\|_{L^2}}_{\rightarrow 0} + \underbrace{\|P_{J_n}(Z - G(U_{J_n}) - \varepsilon)\|_{L^2}}_{\leq \|G(U) - G(U_{J_n})\|_{L^2} \rightarrow 0} \rightarrow 0.$$

Now consider J as an arbitrary subset. Since $U^a = U + K(z - Z)$ and $U_J^a = U_J + K_J(z - Z_J)$, we have

$$\|U^a - U_J^a\|_{L^2} \leq \|U - U_J\|_{L^2} + \|K - K_J\| \|z - Z_J\|_{L^2} + \|K\| \|Z - Z_J\|_{L^2},$$

where the norm for K and $K - K_J$ is the usual operator norm for linear mappings from $\mathbb{R}^d \rightarrow \mathcal{X}$. It is clear that we can estimate

$$\|z - Z_J\|_{L^2} \leq |z| + \|Z\|_{L^2},$$

because $\|Z_J\|_{L^2} \leq \|Z\|_{L^2}$. Considering $\|K - K_J\|$, we can further split this error into

$$\begin{aligned} \|K - K_J\| &\leq \|\text{Cov}(U, Z) - \text{Cov}(U_J, Z_J)\| \|\text{Cov}^{-1}(Z)\| \\ &\quad + \|\text{Cov}(U_J, Z_J)\| \|\text{Cov}^{-1}(Z) - \text{Cov}^{-1}(Z_J)\|. \end{aligned}$$

Next, we recall that the covariance operator $\text{Cov}(X, Y)$ depends continuously on X and Y ; in particular, for zero-mean Hilbert space-valued RVs $X_1, X_2 \in L^2(\mathcal{X})$ and $Y_1, Y_2 \in L^2(\mathcal{Y})$, we have

$$\begin{aligned} \|\text{Cov}(X_1, Y_1) - \text{Cov}(X_2, Y_2)\| &= \|\mathbb{E}[X_1 \otimes Y_1] - \mathbb{E}[X_2 \otimes Y_2]\| \\ &\leq \mathbb{E}[\|(X_1 - X_2) \otimes Y_1\| + \|X_2 \otimes (Y_1 - Y_2)\|] \\ &= \mathbb{E}[\|X_1 - X_2\| \|Y_1\|] + \mathbb{E}[\|X_2\| \|Y_1 - Y_2\|] \\ &\leq (\|Y_1\|_{L^2} + \|X_2\|_{L^2}) (\|X_1 - X_2\|_{L^2} + \|Y_1 - Y_2\|_{L^2}), \end{aligned}$$

where we have used Jensen's and the triangle inequalities in the second line and the Cauchy-Schwartz inequality in the last line. Since $\text{Cov}(X, Y) = \text{Cov}(X - \mathbb{E}[X], Y - \mathbb{E}[Y])$ and $\|X - \mathbb{E}[X]\|_{L^2} \leq \|X\|_{L^2}$, the above estimate holds also for nonzero-mean RVs. Thus, we get

$$\|\text{Cov}(U, Z) - \text{Cov}(U_J, Z_J)\| \leq (\|U\|_{L^2} + \|Z\|_{L^2}) (\|U - U_J\|_{L^2} + \|Z - Z_J\|_{L^2})$$

and

$$\|\text{Cov}(Z) - \text{Cov}(Z_J)\| \leq 4\|Z\|_{L^2} \|Z - Z_J\|_{L^2},$$

due to $\|Z_J\|_{L^2} \leq \|Z\|_{L^2}$. Now consider again the assumed monotone and exhaustive sequence $(J_n)_{n \in \mathbb{N}}$, and recall that, by taking a sufficiently large n , the errors $\|U - U_{J_n}\|_{L^2}$ and $\|Z -$

$Z_{J_n}\|_{L^2}$ can be made arbitrarily small. Thus, also $\|\text{Cov}(Z) - \text{Cov}(Z_{J_n})\|$ will tend to zero as $n \rightarrow \infty$. We now apply the continuity of the matrix inverse to estimate $\|\text{Cov}^{-1}(Z) - \text{Cov}^{-1}(Z_{J_n})\|$. Recall that $\text{Cov}(Z), \text{Cov}(Z_{J_n}) \in \mathbb{R}^{d \times d}$. Let n be sufficiently large such that

$$\|\text{Cov}(Z) - \text{Cov}(Z_{J_n})\| < \frac{1}{2\|\text{Cov}^{-1}(Z)\|};$$

then there holds

$$\|\text{Cov}^{-1}(Z) - \text{Cov}^{-1}(Z_{J_n})\| \leq 2\|\text{Cov}^{-1}(Z)\|^2 \|\text{Cov}(Z) - \text{Cov}(Z_{J_n})\|$$

(see [21, sect. 5.8]). Summing up all previous estimates, we obtain

$$\|K - K_{J_n}\| \leq C_1(\|U - U_{J_n}\|_{L^2} + \|Z - Z_{J_n}\|_{L^2}) + C_2\|Z - Z_{J_n}\|_{L^2},$$

with $C_1 = \|\text{Cov}^{-1}(Z)\|(\|U\|_{L^2} + \|Z\|_{L^2})$ and $C_2 = 8\|U\|_{L^2} \|\text{Cov}^{-1}(Z)\|^2 \|Z\|_{L^2}^2$, where we have used

$$\|\text{Cov}(U_J, Z_J)\| \leq \|U_J\|_{L^2} \|Z_J\|_{L^2} \leq \|U\|_{L^2} \|Z\|_{L^2}$$

to obtain C_2 . Finally, we arrive at

$$\begin{aligned} \|U^a - U_{J_n}^a\|_{L^2} &\leq \|U - U_{J_n}\|_{L^2} + (|z| + \|Z\|_{L^2})\|K - K_{J_n}\| + \|K\| \|Z - Z_{J_n}\|_{L^2} \\ &\leq C(\|U - U_{J_n}\|_{L^2} + \|Z - Z_{J_n}\|_{L^2}), \end{aligned}$$

with $C = 1 + \|K\| + |z| + \|Z\|_{L^2} + C_1 + C_2$, and the assertion follows. \blacksquare

Remark 3.3. Since for many applications evaluating the forward map G corresponds to solving a differential or integral equation, an additional error arises due to numerical approximations G_h of G . This error affects the filters by sampling or computing chaos coefficients of $Z_h = G_h(U) + \varepsilon$ instead of Z . We neglect this error in our analysis since it is beyond the scope of this work. However, if G is the solution operator for differential equations, we expect that (10) could be verified in many cases, such as for elliptic boundary value problems with U a random diffusion coefficient or source term.

The first large ensemble convergence analysis for the EnKF was carried out in [20]. There the authors considered filtering for finite-dimensional locally Lipschitz dynamical systems. They proved that the empirical mean of locally Lipschitz-continuous functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$ applied to the analysis ensemble converges \mathbb{P} -a.s. to the expectation of f w.r.t to a limit distribution which in our setting relates to the distribution of the above introduced analysis variable U^a . We will derive a similar result for the considered Bayesian inverse problem setting; i.e., we extend it to Hilbert spaces and to one update given data obtained by general nonlinear continuous observation operators. Although the proof employs ideas similar to that of [20, Theorem 5.1], we present it for completion.

Theorem 3.4. *Given the model (2) under Assumptions 2.1, 2.2, and 2.5, let (u_1^a, \dots, u_M^a) denote the analysis ensemble resulting from the EnKF, and let $\tilde{\mu}_M^a$ denote the associated empirical measure (9). Further, let μ^a denote the push-forward measure of the analysis variable U^a . Then, for any $f : \mathcal{X} \rightarrow \mathcal{Y}$ which satisfies*

$$\|f(u) - f(v)\|_{\mathcal{Y}} \leq C(1 + \|u\|_{\mathcal{X}} + \|v\|_{\mathcal{X}}) \|u - v\|_{\mathcal{X}} \quad \forall u, v \in \mathcal{X},$$

where \mathcal{Y} is an arbitrary separable Hilbert space, there holds

$$\lim_{M \rightarrow \infty} \int_{\mathcal{X}} f(u) \tilde{\mu}_M^a(du) = \int_{\mathcal{X}} f(u) \mu^a(du) \quad \mathbb{P}\text{-a.s.}$$

This implies in particular for $\bar{u}_M^a = \frac{1}{M} \sum_{j=1}^M u_j^a$ that

$$\lim_{M \rightarrow \infty} \bar{u}_M^a = \mathbb{E}[U^a] \quad \text{and} \quad \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{j=1}^M (u_j^a - \bar{u}_M^a) \otimes (u_j^a - \bar{u}_M^a) = \text{Cov}(U^a) \quad \mathbb{P}\text{-a.s.}$$

Proof. We denote by U_i and ε_i , $i \in \mathbb{N}$, independent and identically distributed (i.i.d.) RVs such that $(U_i, \varepsilon_i) \sim \mu_0 \otimes \nu_\varepsilon$. Further, we define

$$U_i^a := U_i + K(z - Z_i), \quad K = \text{Cov}(U_1, Z_1)\text{Cov}^{-1}(Z_1),$$

where $Z_i := G(U_i) + \varepsilon_i$, and

$$X_{M,i}^a := U_i + K_M(z - Z_i), \quad K_M = \text{Cov}(\mathbf{U}_M, \mathbf{Z}_M)\text{Cov}^{-1}(\mathbf{Z}_M),$$

where $\text{Cov}(\mathbf{U}_M, \mathbf{Z}_M)$ and $\text{Cov}(\mathbf{Z}_M)$ are empirical covariances, e.g.,

$$\text{Cov}(\mathbf{U}_M, \mathbf{Z}_M) = \frac{1}{M-1} \sum_{i=1}^M (U_i - \bar{\mathbf{U}}_M) \otimes (Z_i - \bar{\mathbf{Z}}_M)$$

with $\bar{\mathbf{U}}_M = \frac{1}{M}(U_1 + \dots + U_M)$ and $\bar{\mathbf{Z}}_M = \frac{1}{M}(Z_1 + \dots + Z_M)$. Note that (X_1^a, \dots, X_M^a) represents the random analysis ensemble of the EnKF and that $U_i^a \sim \mu^a$ i.i.d. For any function $f : \mathcal{X} \rightarrow \mathcal{Y}$ which fulfills the assumptions stated in the theorem, we have

$$\frac{1}{M} \sum_{i=1}^M f(X_{M,i}^a) = \frac{1}{M} \sum_{i=1}^M f(X_{M,i}^a) - f(U_i^a) + \frac{1}{M} \sum_{i=1}^M f(U_i^a).$$

Due to the strong law of large numbers (SLLN) [33] there holds

$$\lim_{M \rightarrow \infty} \frac{1}{M} \sum_{i=1}^M f(U_i^a) = \int_{\mathcal{X}} f(u) \mu^a(du) \quad \mathbb{P}\text{-a.s.}$$

Hence, we need only ensure that

$$\begin{aligned} \left\| \frac{1}{M} \sum_{i=1}^M f(X_{M,i}^a) - f(U_i^a) \right\| &\leq \frac{1}{M} \sum_{i=1}^M C(1 + \|U_i^a\| + \|X_{M,i}^a\|) \|X_{M,i}^a - U_i^a\| \\ &\leq \left(\frac{C}{M} \sum_{i=1}^M (1 + \|U_i^a\| + \|X_{M,i}^a\|)^2 \right)^{1/2} \left(\frac{C}{M} \sum_{i=1}^M \|X_{M,i}^a - U_i^a\|^2 \right)^{1/2} \end{aligned}$$

converges \mathbb{P} -a.s. to 0 as $M \rightarrow \infty$ to prove the first statement. We estimate

$$\|X_{M,i}^a - U_i^a\| \leq \|K - K_M\| \|z - Z_i\| \quad \forall i \in \mathbb{N},$$

where we can further split

$$K - K_M = (\text{Cov}(U, Z) - \text{Cov}(\mathbf{U}_M, \mathbf{Z}_M))\text{Cov}^{-1}(Z) \\ + \text{Cov}(\mathbf{U}_M, \mathbf{Z}_M) (\text{Cov}^{-1}(Z) - \text{Cov}^{-1}(\mathbf{Z}_M)).$$

Next, we recall that the empirical covariance converges \mathbb{P} -a.s. to the true covariance, which follows easily (see [31, Satz 3.14] for the scalar case) by writing

$$\text{Cov}(\mathbf{U}_M, \mathbf{Z}_M) = \frac{1}{M-1} \sum_{i=1}^M (U_i - \mathbb{E}[U]) \otimes (Z_i - \mathbb{E}[Z]) - \frac{M}{M-1} (\bar{U}_M - \mathbb{E}[U]) \otimes (\bar{Z}_M - \mathbb{E}[Z]).$$

Then by the SLLN we get

$$\frac{1}{M-1} \sum_{i=1}^M (U_i - \mathbb{E}[U]) \otimes (Z_i - \mathbb{E}[Z]) \xrightarrow{M \rightarrow \infty} \mathbb{E}[(U - \mathbb{E}[U]) \otimes (Z - \mathbb{E}[Z])],$$

and $\frac{M}{M-1}(\bar{U}_M - \mathbb{E}[U]) \otimes (\bar{Z}_M - \mathbb{E}[Z]) \xrightarrow{M \rightarrow \infty} 0$ \mathbb{P} -a.s. Thus, we have

$$\text{Cov}(U, Z) - \text{Cov}(\mathbf{U}_M, \mathbf{Z}_M) \xrightarrow{M \rightarrow \infty} 0 \quad \text{and} \quad \text{Cov}(Z) - \text{Cov}(\mathbf{Z}_M) \xrightarrow{M \rightarrow \infty} 0 \quad \mathbb{P}\text{-a.s.}$$

Since the matrix inverse is a continuous mapping, it also follows that

$$\text{Cov}^{-1}(Z) - \text{Cov}^{-1}(\mathbf{Z}_M) \xrightarrow{M \rightarrow \infty} 0 \quad \mathbb{P}\text{-a.s.},$$

and hence $K - K_M \rightarrow 0$ as $M \rightarrow \infty$ \mathbb{P} -a.s. We thus have for $p \in [1, 2]$ \mathbb{P} -a.s.

$$\lim_{M \rightarrow \infty} X_{M,i}^a = U_i^a \quad \forall i \in \mathbb{N} \quad \text{and} \quad \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{i=1}^M \|X_{M,i}^a - U_i^a\|^p = 0,$$

since by the SLLN $\frac{1}{M} \sum_{i=1}^M \|z - Z_i\|^p$ will tend to $\mathbb{E}[\|z - Z\|^p]$ \mathbb{P} -a.s. Moreover, there holds

$$(1 + \|U_i^a\| + \|X_{M,i}^a\|)^2 \leq (1 + 2\|U_i^a\| + \|X_{M,i}^a - U_i^a\|)^2 \leq 2(1 + 2\|U_i^a\|)^2 + \|X_{M,i}^a - U_i^a\|^2,$$

which yields, again by the SLLN and the above arguments,

$$\frac{1}{M} \sum_{i=1}^M (1 + \|U_i^a\| + \|X_{M,i}^a\|)^2 \leq \frac{1}{M} \sum_{i=1}^M (1 + 2\|U_i^a\|)^2 + \frac{1}{M} \sum_{i=1}^M \|X_{M,i}^a - U_i^a\|^2 \rightarrow \mathbb{E}[(1 + 2\|U^a\|)^2]$$

as $M \rightarrow \infty$ \mathbb{P} -a.s. We thus finally obtain

$$\left\| \frac{1}{M} \sum_{i=1}^M f(X_{M,i}^a) - f(U_i^a) \right\| \xrightarrow{M \rightarrow \infty} 0 \quad \mathbb{P}\text{-a.s.},$$

proving the first statement of the theorem. The remaining two then follow immediately. ■

4. Bayesian interpretation of generalized Kalman filters. In the previous section we have characterized the limit of the EnKF and PCKF approximations for increasing sample size or polynomial degree, respectively. We now investigate how this limit, the analysis variable U^a , may be understood in the context of Bayesian inverse problems. By analyzing the properties of this RV we are able to characterize the approximations provided by the two Kalman filtering methods. In particular, we show that these do not, in general, solve the nonlinear Bayesian inverse problem, nor can they be justified even as approximations to its solution. They are, rather, related to a linear approximation of the Bayes estimator $\hat{\phi}_{\text{CM}}$ and its estimation error.

4.1. The linear conditional mean. The quantity known in classical statistics as the best linear unbiased estimator corresponds in the Bayesian setting to the *linear posterior mean estimator* $\hat{\phi}_{\text{LCM}}$ defined as

$$(14) \quad \hat{\phi}_{\text{LCM}} = \underset{\phi \in \mathcal{P}_1(\mathbb{R}^d; \mathcal{X})}{\operatorname{argmin}} \mathbb{E} [\|U - \phi(Z)\|^2],$$

where $\mathcal{P}_1(\mathbb{R}^d; \mathcal{X}) = \{\phi : \phi(z) = b + Az \text{ with } b \in \mathcal{X}, A \in \mathcal{L}(\mathbb{R}^d, \mathcal{X})\}$ denotes the set of all linear mappings from \mathbb{R}^d to \mathcal{X} . Moreover, we refer to the RV $\hat{\phi}_{\text{LCM}}(Z)$ as the *linear conditional mean*. Recall that the conditional mean $\hat{\phi}_{\text{CM}}(Z) = \mathbb{E}[U|Z]$ is the best approximation of U in $L^2(\Omega, \sigma(Z), \mathbb{P}; \mathcal{X})$ w.r.t. the $L^2(\mathcal{X})$ -norm. Thus $\hat{\phi}_{\text{LCM}}(Z)$ can be seen as the best approximation of U in the subspace $\mathcal{P}_1(Z; \mathcal{X}) \subset L^2(\Omega, \sigma(Z), \mathbb{P}; \mathcal{X})$, where $\mathcal{P}_1(Z; \mathcal{X})$ is short for $\mathcal{P}_1(\mathbb{R}^d; \mathcal{X}) \circ Z = \{\phi(Z), \phi \in \mathcal{P}_1(\mathbb{R}^d; \mathcal{X})\}$.

Lemma 4.1. *The linear conditional mean as defined in (14) is given by*

$$\hat{\phi}_{\text{LCM}}(z) = \mathbb{E}[U] + \operatorname{Cov}(U, Z)\operatorname{Cov}(Z)^{-1}(z - \mathbb{E}[Z]).$$

Proof. The assertion follows by verifying that

$$\hat{\phi}_{\text{LCM}}(Z) = \mathbb{E}[U] + K(Z - \mathbb{E}[Z]), \quad K = \operatorname{Cov}(U, Z)\operatorname{Cov}(Z)^{-1}$$

coincides with the orthogonal projection of U to $\mathcal{P}_1(Z; \mathcal{X})$. To do so, we will show that $U - \hat{\phi}_{\text{LCM}}(Z)$ is orthogonal to $\mathcal{P}_1(Z; \mathcal{X})$ w.r.t. the inner product in $L^2(\mathcal{X})$.

Let $b \in \mathcal{X}$ and $A \in \mathcal{L}(\mathbb{R}^d, \mathcal{X})$ be arbitrary. Then there holds

$$\begin{aligned} \mathbb{E} [\langle U - \hat{\phi}_{\text{LCM}}(Z), b + AZ \rangle] &= \underbrace{\mathbb{E} [\langle U - \mathbb{E}[U], b \rangle]}_{=0} + \mathbb{E} [\langle U - \mathbb{E}[U], AZ \rangle] \\ &\quad - \mathbb{E} [\langle K(Z - \mathbb{E}[Z]), AZ \rangle] - \underbrace{\mathbb{E} [\langle K(Z - \mathbb{E}[Z]), b \rangle]}_{=0} \\ &= \mathbb{E} [\langle U - \mathbb{E}[U], A(Z - \mathbb{E}[Z]) \rangle] - \mathbb{E} [\langle K(Z - \mathbb{E}[Z]), A(Z - \mathbb{E}[Z]) \rangle] \\ &= \operatorname{Cov}(U, Z)A^* - K\operatorname{Cov}(Z)A^* = 0, \end{aligned}$$

since

$$\mathbb{E} [\langle U - \mathbb{E}[U], A\mathbb{E}[Z] \rangle] = \mathbb{E} [\langle K(Z - \mathbb{E}[Z]), A\mathbb{E}[Z] \rangle] = 0$$

and $\operatorname{Cov}(AX, BY) = A\operatorname{Cov}(X, Y)B^*$ for Hilbert space-valued RVs X, Y and bounded, linear operators A, B . ■

We note that Lemma 4.1 fails to hold in the case when \mathcal{X} is only a separable Banach space, since then the expectation $\mathbb{E}[U]$ and covariance $\operatorname{Cov}(U, Z)$ no longer minimize $\mathbb{E}[\|U - b\|^2]$, $b \in \mathcal{X}$, and $\mathbb{E}[\|U - AZ\|^2]$, $A \in \mathcal{L}(\mathbb{R}^d, \mathcal{X})$, respectively; see also Remark 2.7.

4.2. Interpretation of the analysis variable. Lemma 4.1 immediately yields a characterization of the analysis variable U^a defined in (12).

Theorem 4.2. *Let Assumptions 2.1, 2.2, and 2.5 be satisfied for the model (2). Then for any $z \in \mathbb{R}^d$ the analysis variable $U^a = U + K(z - Z)$, $K = \text{Cov}(U, Z)\text{Cov}(Z)^{-1}$, coincides with*

$$U^a = \hat{\phi}_{\text{LCM}}(z) + (U - \hat{\phi}_{\text{LCM}}(Z)).$$

In particular, there hold

$$\mathbb{E}[U^a] = \hat{\phi}_{\text{LCM}}(z) \quad \text{and} \quad \text{Cov}(U^a) = \text{Cov}(U) - K\text{Cov}(Z, U).$$

We summarize the consequences of Theorem 4.2 as follows:

- The analysis variable U^a , to which the EnKF and the PCKF provide approximations, is the sum of a Bayes estimate $\hat{\phi}_{\text{LCM}}(z)$ and the prior error $U - \hat{\phi}_{\text{LCM}}(Z)$ of the corresponding Bayes estimator $\hat{\phi}_{\text{LCM}}$.
- The mean of the EnKF analysis ensemble or PCKF analysis vector provides approximations to the linear posterior mean estimate. How far the latter deviates from the true posterior mean depends on the model and observation z .
- The covariance approximated by the empirical covariance of the EnKF analysis ensemble, as well as that of the PCKF analysis vector, is independent of the actual observational data $z \in \mathbb{R}^d$. It therefore constitutes a prior rather than a posterior measure of uncertainty.
- In particular, the randomness in U^a is entirely determined by the prior measures μ_0 and ν_ε . Only the location, i.e., the mean, of U^a is influenced by the observation data z ; the randomness of U^a is independent of z and determined only by the projection error $U - \hat{\phi}_{\text{LCM}}(Z)$ w.r.t. the prior measures.
- In view of the last two items, the analysis variable U^a , and therefore the EnKF analysis ensemble or the result of the PCKF, are in general not distributed according to the posterior measure μ^z . Moreover, the difference between μ^z and the distribution of U^a depends on the data z and can become quite large for nonlinear problems, see Example 4.4.

Remark 4.3. In particular, the second and third items above explain the observations made in [26] that “... (i) with appropriate parameter choices, approximate filters can perform well in reproducing the mean of the desired probability distribution, (ii) they do not perform as well in reproducing the covariance ...”

We illustrate the conceptual difference between the distribution of the analysis variable U^a and the posterior measure μ^z with a simple yet striking example (cf. also [2, sect. 7], [20, Example 2.2], or [17, sect. 5]).

Example 4.4. We consider $U \sim N(0, 1)$, $\varepsilon \sim N(0, \sigma^2)$, and $G(u) = u^2$. Given data $z \in \mathbb{R}$, the posterior measure, obtained from Bayes’ rule for the densities, is

$$\mu^z(du) = C \exp\left(-\frac{\sigma^2 u^2 + (z - u^2)^2}{2\sigma^2}\right) du.$$

Due to the symmetry of μ^z we have $\hat{u}_{\text{CM}} = \int_{\mathcal{X}} u \mu^z(du) = 0$ for any $z \in \mathbb{R}^d$. Thus, $\mathbb{E}[U|Z] \equiv 0$,

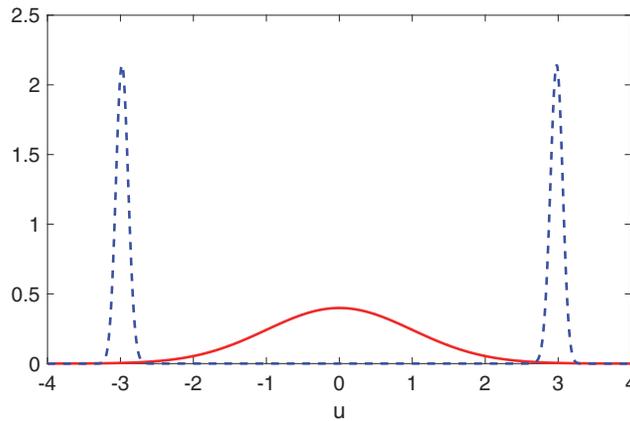


Figure 1. Density of the posterior μ^z (dashed blue line) and the probability density of the analysis variable U^a (solid red line) for $z = 9$ and $\sigma = 0.5$.

and $\hat{\phi}_{\text{LCM}} \equiv \hat{\phi}_{\text{CM}}$. In particular, we have $K = 0$ due to

$$\text{Cov}(U, Z) = \text{Cov}(U, U^2) = \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} u(u^2 - 1)e^{-u^2/2} du = 0,$$

which in turn yields $U^a = U \sim N(0, 1)$. Thus, the analysis variable is distributed according to the prior measure. This is not surprising as, by definition, its mean is the best linear approximation to the posterior mean according to μ^z , and its fluctuation is simply the prior estimation error $U - \hat{\phi}_{\text{LCM}}(Z) = U - 0 = U$. This illustrates that U^a is suited for approximating the posterior mean, but is not appropriate for the purpose of UQ for the nonlinear inverse problem. As displayed in Figure 1, the distribution of U^a can be markedly different from the true posterior distribution.

5. Numerical examples. To illustrate the application of the EnKF and PCKF to simple Bayesian inverse problems, we consider in the following a one-dimensional elliptic boundary value problem and a time-dependent RLC circuit model.

5.1. One-dimensional elliptic boundary value problem. Let $D = [0, 1]$ and

$$(15) \quad -\frac{d}{dx} \left(\exp(u_1) \frac{d}{dx} p(x) \right) = f(x), \quad p(0) = p_0, \quad p(1) = u_2,$$

be given, where $u = (u_1, u_2)$ are unknown scalar parameters. The solution of (15) is

$$(16) \quad p(x) = p_0 + (u_2 - p_0)x + \exp(-u_1) (S_x(F) - S_1(F)x),$$

where $S_x(g) := \int_0^x g(y) dy$ and $F(x) = S_x(f) = \int_0^x f(y) dy$. For simplicity we choose $f \equiv 1$, $p_0 = 0$ in the following and assume noisy measurements have been made of p at $x_1 = 0.25$ and $x_2 = 0.75$ with values $z = (27.5, 79.7)$. We seek to infer u based on this data and on a priori information modelled by $(u_1, u_2) \sim N(0, 1) \otimes \text{Uni}(90, 110)$, where $\text{Uni}(a, b)$ denotes the uniform distribution on the interval $[a, b]$. Thus the forward map here is $G(u) = (p(x_1), p(x_2))$,

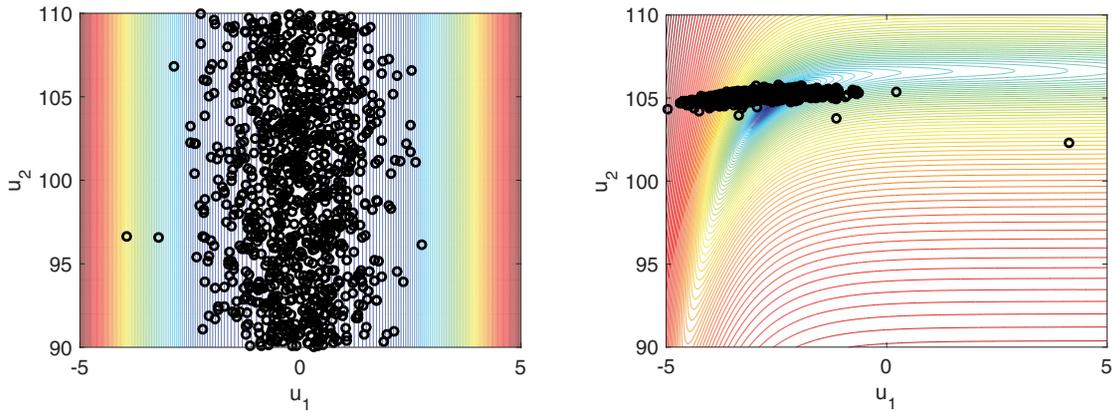


Figure 2. Left: Contour plot of the negative logarithm of the prior density and the locations of 1000 ensemble members of the initial EnKF ensemble. Right: Contour plot of the logarithm of the negative logarithm of the posterior density and the locations of the updated 1000 ensemble members in the analysis EnKF ensemble.

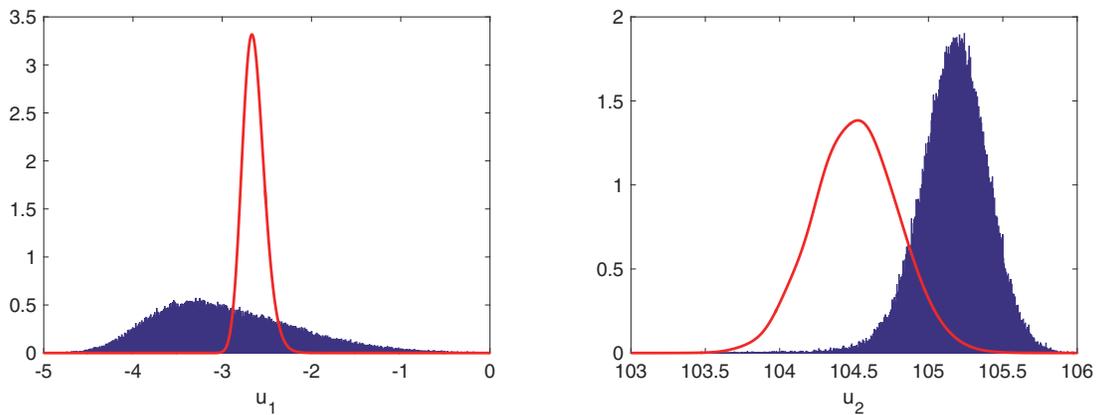


Figure 3. Posterior marginals and relative frequencies in the analysis ensemble for u_1 (left) and u_2 (right).

where p is given in (16) with $f \equiv 1$ and $p_0 = 0$. As the model for the measurement noise we take $\varepsilon \sim N(0, 0.01 I_2)$.

In Figure 2 we show the level curves of the prior and posterior densities as well as 1000 ensemble members of the initial and analysis ensemble obtained by the EnKF. A total ensemble size of $M = 10^5$ was chosen in order to reduce the sampling error to a negligible level. It can be seen, however, that the analysis EnKF ensemble does not follow the posterior distribution, although its mean $(-2.92, 105.14)$ is quite close to the true posterior mean $(-2.65, 104.5)$ (computed by quadrature). To illustrate the difference between the distribution of the analysis ensemble/variable and the true posterior distribution, we present the marginal posterior distributions of u_1 and u_2 in Figure 3. For the posterior the marginals were evaluated by quadrature, whereas for the analysis ensemble we show a relative frequency plot.

We remark that slightly changing the observational data to $\tilde{z} = (23.8, 71.3)$ moves the analysis ensemble as well as the distribution of the analysis RV much closer to the true

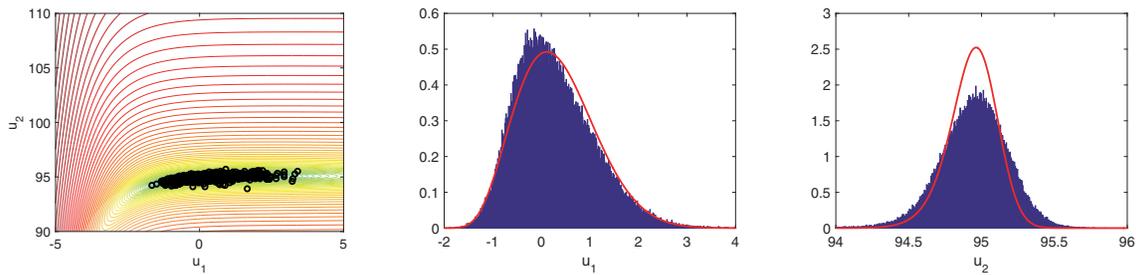


Figure 4. Left: Contours of the logarithm of the negative log posterior density and locations of 1000 members of the analysis EnKF ensemble. Middle, right: Posterior marginals and relative frequencies in the analysis ensemble for u_1 (middle) and u_2 (right).

posterior, as shown in Figure 4. Moreover, for these measurement values the mean of the analysis ensemble (0.33, 94.94) provides a better fit to the true posterior mean (0.33, 94.94).

To reaffirm the fact that only the mean of the analysis variable U^a depends on the actual data, in Figure 5 we show density estimates for the marginals of u_1 and u_2 of U^a obtained from the observational data $z = (27.5, 79.7)$ (blue lines) and $\tilde{z} = (23.8, 71.3)$ (green lines), respectively. The density estimates were obtained by normal kernel density estimation (in this case the MATLAB `ksdensity` routine) based on the resulting analysis ensembles $(\mathbf{u}_1^a, \mathbf{u}_2^a)$ and $(\tilde{\mathbf{u}}_1^a, \tilde{\mathbf{u}}_2^a)$ for the data sets z and \tilde{z} , respectively. We observe that the marginal distributions of the centered ensembles coincide, in agreement with Theorem 4.2.

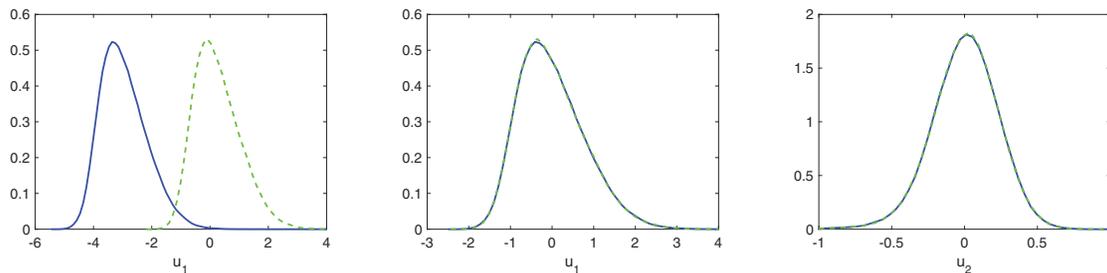


Figure 5. Left: Kernel density estimates for \mathbf{u}_1^a (solid blue line) and $\tilde{\mathbf{u}}_1^a$ (dashed green line). Middle, right: Kernel density estimates for $\mathbf{u}_i^a - \mathbb{E}[\mathbf{u}_i^a]$ (solid blue) and $\tilde{\mathbf{u}}_i^a - \mathbb{E}[\tilde{\mathbf{u}}_i^a]$ (dashed green), $i = 1, 2$.

In addition, whenever the prior and thus also the posterior support for u_2 are bounded—as in this example—the EnKF may generate members in the analysis ensemble which are outside this support. This is a further consequence of Theorem 4.2: Since the analysis ensemble of the EnKF follows the distribution of the analysis variable rather than the true posterior distribution, ensemble members lying outside the posterior support can always occur whenever the support of the analysis variable is not a subset of the support of the posterior.

Finally, we emphasize that whether or not the distribution of the analysis variable is a good fit to the true posterior distribution depends entirely on the observed data—which can be neither controlled nor known a priori.

The calculations for applying the PCKF to this simple example problem can be carried out analytically: We require four basic independent RVs $\xi_1 \sim N(0, 1)$, $\xi_2 \sim \text{Uni}(0, 1)$, $\xi_3 \sim N(0, 1)$,

and $\xi_4 \sim N(0, 1)$ to define PCEs which yield RVs distributed according to the prior and error distributions:

$$U := (\xi_1, 90 + 20\xi_2)^\top \sim \mu_0, \quad \varepsilon := (0.1\xi_3, 0.1\xi_4)^\top \sim \nu_\varepsilon.$$

Moreover, due to (16), $G(U)$ is also available in closed form as

$$G(U) = \begin{pmatrix} c_{11}(90 + 20\xi_2) + c_{12} \sum_{n=0}^{\infty} (-1)^n \frac{\sqrt{\varepsilon}}{\sqrt{n!}} H_n(\xi_1) \\ c_{21}(90 + 20\xi_2) + c_{22} \sum_{n=0}^{\infty} (-1)^n \frac{\sqrt{\varepsilon}}{\sqrt{n!}} H_n(\xi_1) \end{pmatrix},$$

where H_n denotes the n th normalized Hermite polynomial and $c_{11}, c_{12}, c_{21}, c_{22}$ can be deduced from inserting $x = 0.25$ and $x = 0.75$ into (16). Here we have used the Hermite expansion of $\exp(-\xi)$; see also [42, Example 2.2.7]. Thus, the chaos coefficient vectors of U and $G(U) + \varepsilon$ w.r.t. the polynomials

$$P_\alpha(\boldsymbol{\xi}) = H_{\alpha_1}(\xi_1) L_{\alpha_2}(\xi_2) H_{\alpha_3}(\xi_3) H_{\alpha_4}(\xi_4), \quad \boldsymbol{\alpha} \in \mathbb{N}_0^4,$$

can be obtained explicitly, where H_α and L_α denote the normalized Hermite and Legendre polynomials of degree α , respectively. In particular, the nonvanishing chaos coefficients involve only the basis polynomials

$$P_0(\boldsymbol{\xi}) \equiv 1, \quad P_1(\boldsymbol{\xi}) = L_1(\xi_2), \quad P_2(\boldsymbol{\xi}) = H_1(\xi_3), \quad P_3(\boldsymbol{\xi}) = H_1(\xi_4),$$

and $P_\alpha(\boldsymbol{\xi}) = H_{\alpha-3}(\xi_1)$ for $\alpha \geq 4$. Arranging the two-dimensional chaos coefficients of U and $G(U)$ as the column vectors of the matrices $[U], [G(U) + \varepsilon] \in \mathbb{R}^{2 \times \mathbb{N}_0}$ and denoting by $[\widetilde{U}]$ the matrix $[u_1, u_2, \dots] \in \mathbb{R}^{2 \times \mathbb{N}}$, we obtain

$$K = [\widetilde{U}][\widetilde{G(U)}]^\top \left([\widetilde{G(U)}][\widetilde{G(U)}]^\top + 0.01I_2 \right)^{-1}.$$

Thus, the only numerical error incurred in applying the PCKF in this example is the truncation of the PCE. We have carried out this calculation using a truncated PCE of length $J = 4 + 50$ according to the reduced basis above. In particular, we evaluated the approximation K_J to K by using the truncated vector $[P_J G(U)]$ in the formula above and then performed the update of the chaos coefficients according to (11). Subsequently, $M = 10^5$ samples of the resulting RV U_j^q were drawn, but, since the empirical distributions were essentially indistinguishable from those obtained by the EnKF described previously, they are omitted here.

Remark 5.1. Although a detailed complexity analysis of these methods is beyond the scope of this work, we mention that the EnKF calls for M evaluations of the forward map $G(u_j)$, $j = 1, \dots, M$, whereas the PCKF requires computing the chaos coefficients of $G(U)$ by, e.g., the stochastic Galerkin method. Thus the former yields, in general, many small systems to solve, whereas the latter typically requires the solution of a large coupled system. Moreover, we emphasize the computational savings by applying KF-type methods compared to a “full Bayesian update,” i.e., sampling from the posterior measure by MCMC methods. In particular, each MCMC run may require calculating many hundreds of thousands of forward maps $G(u)$, e.g., for each iteration u_j of the Markov chain as in the case of Metropolis–Hastings MCMC. Hence, if one is interested in only the posterior mean as a Bayes estimate, then EnKF and PCKF provide substantially less expensive alternatives to MCMC for its approximation by the linear posterior mean.

5.2. Dynamical system: RLC circuit. We apply the EnKF to sequential data assimilation in a simple dynamical system: a damped LC circuit or RLC circuit. Denoting the initial voltage by U_0 , the resistance by R , the inductance by L , and the capacitance by C , and assuming $R < 2\sqrt{LC}$, the voltage and current in the circuit can be modelled as

$$(17a) \quad U(t) = U_0 e^{\delta t} \left(\cos(w_e t) + \frac{\delta}{w_e} \sin(w_e t) \right),$$

$$(17b) \quad I(t) = -\frac{U_0}{w_e L} e^{\delta t} \sin(w_e t),$$

where $\delta = R/(2L)$, $w_e = \sqrt{w_0^2 - \delta^2}$, and $w_0 = 1/\sqrt{LC}$. The data assimilation setting is now as follows. We observe the state of the system (17) at four time points $t_n = 5n$, $n = 1, \dots, 4$, where all observations $z \in \mathbb{R}^8$ are corrupted by measurement noise $\varepsilon \sim N(0, \text{diag}(\sigma_1^2, \dots, \sigma_8^2))$. Here we have chosen $\sigma_{2n-1}^2 = 0.1|U(t_n)|$ and $\sigma_{2n}^2 = 0.1|I(t_n)|$ for $n = 1, \dots, 4$. We want to infer U_0 and L based on these observations; i.e, the unknown is $u = (U_0, L)$, and we take as the prior $(U_0, L) \sim N(0.5, 0.25) \otimes \text{Uni}(1, 5)$. Given observations $z \in \mathbb{R}^8$, we compare two assimilation strategies for applying the EnKF.

- *Simultaneous:* We apply the EnKF to the inverse problem

$$z = G(u) + \varepsilon,$$

where G maps (U_0, L) to the states $(U(t_1), I(t_1), \dots, U(t_4), I(t_4)) \in \mathbb{R}^8$. Thus, we perform one EnKF update using all the available data at once, resulting in one EnKF analysis ensemble.

- *Sequential:* We apply the EnKF to the inverse problem

$$z_n = G_n(u) + \varepsilon_n, \quad n = 1, \dots, 4,$$

where G_n maps (U_0, L) to the state $(U(t_n), I(t_n)) \in \mathbb{R}^2$. In particular, we will perform four EnKF updates using at each update only the corrupted data $z_n = (U(t_n) + \varepsilon_{2n-1}, I(t_n) + \varepsilon_{2n})$. This yields, for each update, one EnKF analysis ensemble which, in turn, serves as the initial ensemble for the next update.

Again we use two different data sets z, \tilde{z} ,¹ obtained by two realizations of ε given the solution of (17) for $U_0 = 0.75, R = 0.5, L = 1.5, C = 0.5$.

The resulting posteriors and EnKF analysis ensembles for the simultaneous and sequential update are presented in Figure 6. We observe again that, for different data sets, the EnKF results in an ensemble which follows a distribution which is, in one case, quite close to, and in the other, quite far away from, the true posterior distribution. Interestingly, the difference between the two updating schemes does not seem to be too large. This also holds true for the means of the EnKF analysis ensembles when compared to the true posterior means in Table 1 for both data sets z and \tilde{z} .

Finally, we are again interested in the marginals of the posterior and the associated histograms of the EnKF analysis ensembles which give us a rough impression of the difference

¹ $z = (0.505, 0.237, 0.014, 0.096, 0.036, 0.011, -0.002, -0.003), \tilde{z} = (0.265, 0.066, 0.058, 0.002, 0.021, 0.012, 0.007, -0.01)$.

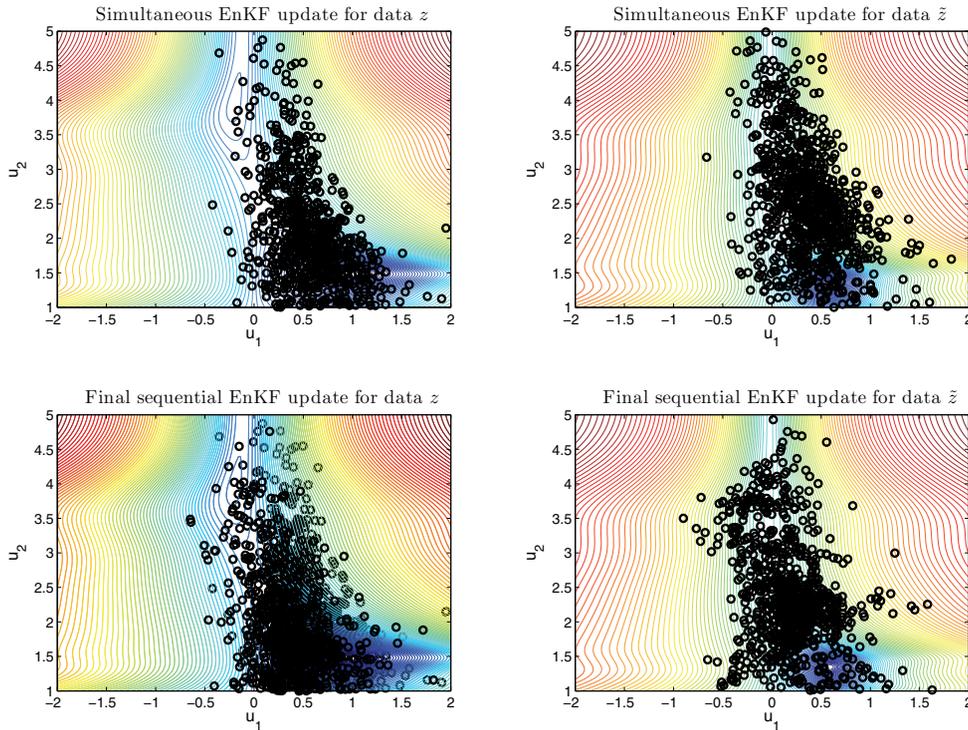


Figure 6. Contours of the logarithm of the negative log posterior density and locations of 1000 members of the analysis EnKF ensembles resulting from simultaneous and sequential updating for the two different sets z and \tilde{z} of observation data.

Table 1

Means of the EnKF analysis ensembles and corresponding true posterior means for data z (left) and \tilde{z} (right).

Update	EnKF mean for data z	Posterior mean for data z	EnKF mean for data \tilde{z}	Posterior mean for data \tilde{z}
1	(0.42, 1.56)	(0.42, 2.42)	(0.27, 2.25)	(0.35, 2.61)
2	(0.44, 1.53)	(0.39, 2.36)	(0.20, 2.20)	(0.32, 2.56)
3	(0.43, 1.59)	(0.38, 2.34)	(0.19, 2.26)	(0.31, 2.52)
4	(0.43, 1.59)	(0.38, 2.32)	(0.19, 2.24)	(0.30, 2.50)
Simultaneous	(0.58, 1.84)	(0.38, 2.32)	(0.38, 2.40)	(0.30, 2.50)

between the distribution of the analysis variable and the true posterior. In Figure 7 we compare both marginals for the fourth update and the simultaneous analysis ensemble for both data sets. The distribution of the simultaneous EnKF analysis ensemble should not depend on the data, whereas the distribution of the final EnKF analysis ensemble for the sequential updating clearly does in this example. This is certainly caused by the nonlinearity of the forward map G : In the sequential updating the former analysis variable U_n^a serves as the initial one for the current update step $n + 1$; therefore, the difference in the mean of the former analysis

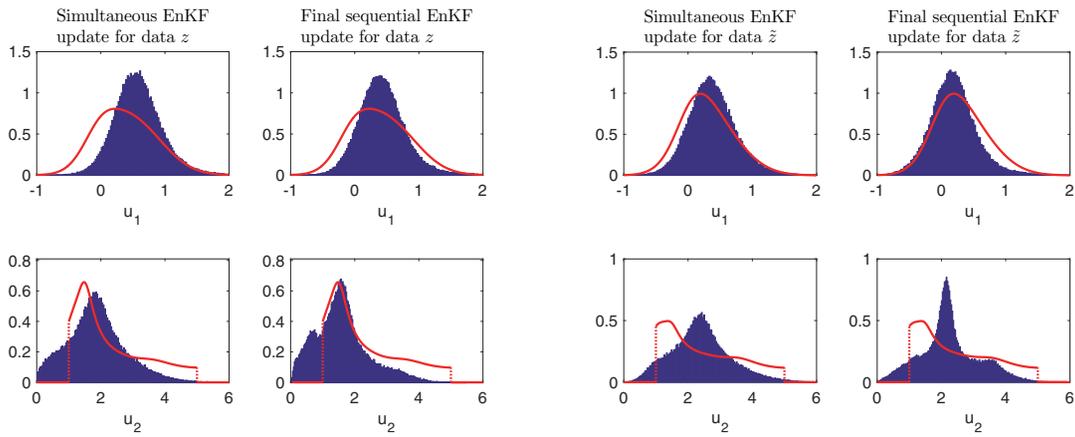


Figure 7. Posterior marginals and relative frequencies of the final EnKF analysis ensembles in u_1, u_2 for z (left group) and \tilde{z} (right group).

variables U_n^a, \tilde{U}_n^a for different data sets z, \tilde{z} might yield different forecast RVs $G(U_n^a), G(\tilde{U}_n^a)$ due to the nonlinearity of G , which yields different next analysis variables $U_{n+1}^a, \tilde{U}_{n+1}^a$.

6. Conclusions. We have given a detailed analysis of two recently popularized generalized Kalman filtering methods, the EnKF and PCKF, applied to nonlinear (stationary) Bayesian inverse problems. We recalled the Bayesian approach to inverse problems and its solution, the posterior measure, in a Hilbert space setting, for which we slightly generalized existing results concerning the well-posedness of Bayesian inverse problems. To characterize Kalman filtering methods in the Bayesian framework we related these to Bayes estimators and highlighted the distinction between the two objectives of inference and identification in Bayesian inversion, which are realized by the posterior measure and Bayes estimators, respectively. We then proved the convergence of the approximations provided by the EnKF and PCKF to what we have dubbed the analysis random variable in the limit of large ensemble and large polynomial degree, respectively, reaffirming the fact that both methods are merely different numerical discretizations of the same updating scheme for RVs. Moreover, the relation of both KF methods to a specific Bayes estimator, the linear posterior mean estimator, followed from the fact that the analysis variable is simply the associated Bayes estimate plus the random a priori error of the Bayes estimator. In particular, the analysis variable does not, in general, follow the posterior distribution. The conclusion of this work is that the EnKF and PCKF are methods suited for identification which provide additional statistical information in the form of the random a priori estimation error, but they are not appropriate methods for rigorous inference in the sense of the conditional or posterior measure. Several carefully chosen numerical examples were given to illustrate these basic differences.

REFERENCES

[1] S. ANZENGRUBER, B. HOFMANN, AND P. MATHÉ, *Regularization properties of the sequential discrepancy principle for Tikhonov regularization in Banach spaces*, Appl. Anal., 93 (2014), pp. 1382–1400.

- [2] A. APTE, M. HAIRER, A. M. STUART, AND J. VOSS, *Sampling the posterior: An approach to non-Gaussian data assimilation*, Phys. D, 230 (2007), pp. 50–64.
- [3] J. O. BERGER, *Statistical Decision Theory and Bayesian Analysis*, 2nd ed., Springer, New York, 1985.
- [4] J. M. BERNARDO, *Bayesian statistics*, in Probability and Statistics, R. Viertl, ed., Encyclopedia of Life Support Systems (EOLSS), UNESCO, Oxford, UK, 2003.
- [5] E. D. BLANCHARD, A. SANDU, AND C. SANDU, *A polynomial chaos-based Kalman filter approach for parameter estimation of mechanical systems*, J. Dyn. Syst. Meas. Control, 132 (2010), 061404.
- [6] M. BURGER AND F. LUCKA, *Maximum a posteriori estimates in linear inverse problems with log-concave priors are proper Bayes estimators*, Inverse Problems, 30 (2014), 114004.
- [7] G. BURGERS, P. J. VAN LEEUWEN, AND G. EVENSEN, *Analysis scheme in the ensemble Kalman filter*, Mon. Weather Rev., 126 (1998), pp. 1719–1724.
- [8] D. E. CATLIN, *Estimation, Control, and the Discrete Kalman Filter*, Springer, New York, 1989.
- [9] M. DASHTI, K. J. H. LAW, A. M. STUART, AND J. VOSS, *MAP estimators and their consistency in Bayesian nonparametric inverse problems*, Inverse Problems, 29 (2013), 095017.
- [10] H. W. ENGL, M. HANKE, AND A. NEUBAUER, *Regularization of Inverse Problems*, Kluwer Academic Publishers, Dordrecht, The Netherlands, 2000.
- [11] O. G. ERNST, A. MUGLER, H.-J. STARKLOFF, AND E. ULLMANN, *On the convergence of generalized polynomial chaos expansions*, ESAIM Math. Model. Numer. Anal., 46 (2012), pp. 317–339.
- [12] O. G. ERNST, B. SPRUNGK, AND H.-J. STARKLOFF, *Bayesian inverse problems and Kalman filters*, in Extraction of Quantifiable Information from Complex Systems, S. Dahlke et al., eds., Lect. Notes Comput. Sci. Eng. 102, Springer International Publishing, Cham, Switzerland, 2014, pp. 133–159.
- [13] G. EVENSEN, *Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics*, J. Geophys. Res., 99 (1994), pp. 10143–10162.
- [14] G. EVENSEN, *The ensemble Kalman filter: Theoretical formulation and practical implementation*, Ocean Dynam., 53 (2003), pp. 343–367.
- [15] G. EVENSEN, *Data Assimilation: The Ensemble Kalman Filter*, 2nd ed., Springer, New York, 2009.
- [16] G. EVENSEN, *The ensemble Kalman filter for combined state and parameter estimation: Monte Carlo techniques for data assimilation in large systems*, IEEE Control Syst. Mag., 29 (2009), pp. 83–104.
- [17] G. EVENSEN AND P. J. VAN LEEUWEN, *An ensemble Kalman smoother for nonlinear dynamics*, Mon. Weather Rev., 128 (2000), pp. 1852–1867.
- [18] C. J. GEYER, *Introduction to Markov Chain Monte Carlo*, in Handbook of Markov Chain Monte Carlo, Chapman Hall/CRC Handb. Mod. Stat. Methods, S. Brooks, A. Gelman, G. J. Jones, and X.-L. Meng, eds., CRC Press, Boca Raton, FL, 2011, pp. 3–48.
- [19] A. L. GIBBS AND F. E. SU, *On choosing and bounding probability metrics*, Internat. Statist. Rev., 70 (2001), pp. 419–435.
- [20] F. LE GLAND, V. MONBET, AND V.-D. TRAN, *Large sample asymptotics for the ensemble Kalman filter*, in The Oxford Handbook of Nonlinear Filtering, D. Crisan and B. Rozovskii, eds., Oxford University Press, Oxford, UK, 2011, pp. 598–631.
- [21] R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, Cambridge University Press, Cambridge, UK, 1990.
- [22] M. A. IGLESIAS, K. J. H. LAW, AND A. M. STUART, *Ensemble Kalman methods for inverse problems*, Inverse Problems, 29 (2013), 045001.
- [23] M. A. IGLESIAS, K. J. H. LAW, AND A. M. STUART, *Evaluation of Gaussian approximations for data assimilation in reservoir models*, Comput. Geosci., 17 (2013), pp. 851–885.
- [24] O. KALLENBERG, *Foundations of Modern Probability*, Springer, New York, 1997.
- [25] R. E. KALMAN, *A new approach to linear filtering and prediction problems*, Trans. AMSE Ser. D J. Basic Engrg., 82 (1960), pp. 35–45.
- [26] K. J. H. LAW AND A. M. STUART, *Evaluating data assimilation algorithms*, Mon. Weather Rev., 140 (2012), pp. 3757–3782.
- [27] K. J. H. LAW, H. TEMBINE, AND R. TEMPONE, *Deterministic Mean-field Ensemble Kalman Filtering*, preprint, [arXiv:1409.0628v3 \[math.PR\]](https://arxiv.org/abs/1409.0628v3), 2015.
- [28] J. M. LEWIS, S. LAKSHMIVARAHAN, AND S. DHALL, *Dynamic Data Assimilation—A Least Squares Approach*, Cambridge University Press, Cambridge, UK, 2006.
- [29] J. MANDEL, L. COBB, AND J. D. BEEZLEY, *On the convergence of the ensemble Kalman filter*, Appl. Math., 56 (2011), pp. 533–541.

- [30] A. MANDELBAUM, *Linear estimators and measurable linear transformations on a Hilbert space*, Z. Wahrsch. Verw. Gebiete, 65 (1984), pp. 385–397.
- [31] T. MÜLLER-GRONBACH, E. NOVAK, AND K. RITTER, *Monte Carlo Algorithmen*, Springer, Berlin, 2012 (in German).
- [32] I. MYRSETH AND H. OMRE, *The ensemble Kalman filter and related filters*, in Large-Scale Inverse Problems and Quantification of Uncertainty, Wiley Ser. Comput. Stat., L. Biegler, ed., Wiley, Chichester, UK, 2011, pp. 217–246.
- [33] W. J. PADGETT AND R. L. TAYLOR, *Laws of large numbers for normed linear spaces and certain Fréchet spaces*, Springer-Verlag, Berlin, New York, 1973.
- [34] O. PAJONK, B. V. ROSIĆ, A. LITVINENKO, AND H. G. MATTHIES, *A deterministic filter for non-Gaussian Bayesian estimation—Applications to dynamical system estimation with noisy measurements*, Phys. D, 241 (2012), pp. 775–788.
- [35] M. M. RAO, *Conditional Measures and Applications*, Chapman & Hall/CRC, Boca Raton, FL, 2010.
- [36] B. V. ROSIĆ, A. KUČEROVÁ, J. SÝKORA, O. PAJONK, A. LITVINENKO, AND H. G. MATTHIES, *Parameter identification in a probabilistic setting*, Engrg. Structures, 60 (2013), pp. 179–196.
- [37] B. V. ROSIĆ, A. LITVINENKO, O. PAJONK, AND H. G. MATTHIES, *Sampling-free linear Bayesian update of polynomial chaos representations*, J. Comput. Phys., 231 (2012), pp. 5761–5787.
- [38] G. SAAD AND R. GHANEM, *Characterization of reservoir simulation models using a polynomial chaos-based ensemble Kalman filter*, Water Resour. Res., 45 (2009), W04417.
- [39] G. SAAD, R. GHANEM, AND S. MASRI, *Robust system identification of strongly non-linear dynamics using a polynomial chaos-based sequential data assimilation technique*, in Collection of Technical Papers—48th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, Vol. 6, Springer, New York, 2007.
- [40] D. SIMON, *Optimal State Estimation: Kalman, H_∞ , and Nonlinear Approaches*, Wiley, Hoboken, NJ, 2006.
- [41] A. M. STUART, *Inverse problems: A Bayesian perspective*, Acta Numer., 19 (2010), pp. 451–559.
- [42] E. ULLMANN, *Solution Strategies for Stochastic Finite Element Discretizations*, Ph.D. thesis, TU Bergakademie Freiberg, Freiberg, Germany, 2008.