

Inverse Problems in a Bayesian Setting*

Hermann G. Matthies[†], Elmar Zander, Bojana V. Rosić,
Alexander Litvinenko, Oliver Pajonk

Institute of Scientific Computing
Technische Universität Braunschweig

20th October 2015

Abstract

In a Bayesian setting, inverse problems and uncertainty quantification (UQ) — the propagation of uncertainty through a computational (forward) model — are strongly connected. In the form of conditional expectation the Bayesian update becomes computationally attractive. We give a detailed account of this approach via conditional approximation, various approximations, and the construction of filters. Together with a functional or spectral approach for the forward UQ there is no need for time-consuming and slowly convergent Monte Carlo sampling. The developed sampling-free non-linear Bayesian update in form of a filter is derived from the variational problem associated with conditional expectation. This formulation in general calls for further discretisation to make the computation possible, and we choose a polynomial approximation. After giving details on the actual computation in the framework of functional or spectral approximations, we demonstrate the workings of the algorithm on a number of examples of increasing complexity. At last, we compare the linear and nonlinear Bayesian update in form of a filter on some examples.

Keywords: inverse identification, uncertainty quantification, Bayesian update, parameter identification, conditional expectation, filters, functional and spectral approximation

*Partly supported by the Deutsche Forschungsgemeinschaft (DFG) through SFB 880.

[†]corresponding author

Classification: —(MSC2010) 62F15, 65N21, 62P30, 60H15, 60H25,
74G75, 80A23, 74C05
—(PACS2010) 46.65.+g, 46.35.+z, 44.10.+i
—(ACM1998) G.1.8, G.3, J.2

1 Introduction

Inverse problems deal with the determination of parameters in computational models, by comparing the prediction of these models with either real measurements or observations, or other, presumably more accurate, computations. These parameters can typically not be observed or measured directly, only other quantities which are somehow connected to the one for which the information is sought. But it is typical that we can compute what the observed response should be, under the assumption that the unknown parameters have a certain value. And the difference between predicted or forecast response is obviously a measure for how well these parameters were identified.

There are different ways of attacking the problem of parameter identification theoretically and numerically. One way is to define some measure of discrepancy between predicted observation and the actual observation. Then one might use optimisation algorithms to make this measure of discrepancy as small as possible by changing the unknown parameters. Classical least squares approaches start from this point. The parameter values where a minimum is attained is then usually taken as the ‘best’ value and regarded as close to the ‘true’ value.

One of the problems is that for one the measure of discrepancy crops up pretty arbitrarily, and on the other hand the minimum is often not unique. This means that there are many parameter values which explain the observations in a ‘best’ way. To obtain a unique solution, some kind of ‘niceness’ of the optimal solution is required, or mathematically speaking, for the optimal solution some regularity is enforced, typically in competition with discrepancy measure to be minimised. This optimisation approach hence leads to regularisation procedures, a good overview of which is given by [5].

Here we take another tack, and base our approach on the Bayesian idea of updating the knowledge about something like the unknown parameters in a probabilistic fashion according to Bayes’s theorem. In order to apply this, the knowledge about the parameters has to be described in a *Bayesian* way through a probabilistic model [16], [41], [40]. As it turns out, such a probabilistic description of our previous knowledge can often be interpreted as a regularisation, thus tying these differing approaches together.

The Bayesian way is on one hand difficult to tackle, i.e. finding a com-

putational way of doing it; and on the other hand often becomes computationally very demanding. One way the Bayesian update may be achieved computationally is through sampling. On the other hand, we shall here use a functional approximation setting to address such stochastic problems. See [26] for a synopsis on our approach to such parametric problems.

It is well-known that such a Bayesian update is in fact closely related to *conditional expectation* [2], [11], and this will be the basis of the method presented. For these and other probabilistic notions see for example [30] and the references therein.

The functional approximation approach towards stochastic problems is explained e.g. in [24]. These approximations are in the simplest case known as Wiener’s so-called *homogeneous* or *polynomial chaos* expansion [43], which are polynomials in independent Gaussian RVs — the ‘chaos’ — and which can also be used numerically in a Galerkin procedure [10], [25], [24]. This approach has been generalised to other types of RVs [44]. It is a computational variant of *white noise analysis*, which means analysis in terms of independent RVs, hence the term ‘white noise’ [14], [15], [13], see also [25], [33], and [8] for here relevant results on stochastic regularity. Here we describe a computational extension of this approach to the inverse problem of Bayesian updating, see also [28], [35], [29], [34].

To be more specific, let us consider the following situation: we are investigating some physical system which is modelled by an evolution equation for its state:

$$\frac{d}{dt}u = A(q; u(t)) + \eta(q; t); \quad u(0) = u_a \text{ given.} \quad (1)$$

where $u(t) \in \mathcal{U}$ describes the state of the system at time $t \in [0, T]$ lying in a Hilbert space \mathcal{U} (for the sake of simplicity), A is a—possibly non-linear—operator modelling the physics of the system, and $\eta \in \mathcal{U}^*$ is some external influence (action / excitation / loading). Both A and ℓ may involve some *noise* — i.e. a random process — so that (1) is a stochastic evolution equation.

Assume that the model depends on some parameters $q \in \mathcal{Q}$, which are uncertain. These may actually include the initial conditions for the state, u_a . To have a concrete example of Eq. (1), consider the diffusion equation

$$\frac{\partial}{\partial t}u(x, t) - \operatorname{div}(\kappa(x)\nabla u(x, t)) = \eta(x, t), \quad x \in \mathcal{G}, \quad (2)$$

with appropriate boundary and initial conditions, where $\mathcal{G} \subset \mathbb{R}^n$ is a suitable domain. The diffusing quantity is $u(x, t)$ (heat, concentration) and the term $\eta(x, t)$ models sinks and sources. Similar examples will be used for the numerical experiments in Section 5 and Section 6. Here $\mathcal{U} = H_E^1(\mathcal{G})$, the subspace of the Sobolev space $H^1(\mathcal{G})$ satisfying the essential boundary

conditions, and we assume that the diffusion coefficient $\kappa(x)$ is uncertain. The parameters could be the positive diffusion coefficient field $\kappa(x)$, but for reasons to be explained fully later, we prefer to take $q(x) = \log(\kappa(x))$, and assume $q \in \mathcal{Q} = L_2(\mathcal{G})$.

The updating methods have to be well defined and stable in a continuous setting, as otherwise one can not guarantee numerical stability with respect to the PDE discretisation refinement, see [40] for a discussion of related questions. Due to this we describe the update before any possible discretisation in the simplest Hilbert space setting.

On the other hand, no harm will result for the basic understanding if the reader wants to view the occurring spaces as finite dimensional Euclidean spaces. Now assume that we observe a function of the state $Y(u(q), q)$, and from this observation we would like to identify the corresponding q . In the concrete example Eq. (2) this could be the value of $u(x_j, t)$ at some points $x_j \in \mathcal{G}$. This is called the *inverse* problem, and as the mapping $q \mapsto Y(q)$ is usually not invertible, the inverse problem is *ill-posed*. Embedding this problem of finding the best q in a larger class by modelling our knowledge about it with the help of probability theory, then in a Bayesian manner the task becomes to estimate conditional expectations, e.g. see [16], [41], [40], and the references therein. The problem now is *well-posed*, but at the price of ‘only’ obtaining probability distributions on the possible values of q , which now is modelled as a \mathcal{Q} -valued random variable (RV). On the other hand one naturally also obtains information about the remaining uncertainty. Predicting what the measurement $Y(q)$ should be from some assumed q is computing the *forward* problem. The *inverse* problem is then approached by comparing the forecast from the forward problem with the actual information.

Since the parameters of the model to be estimated are uncertain, all relevant information may be obtained via their stochastic description. In order to extract information from the posterior, most estimates take the form of expectations w.r.t. the posterior. These expectations — mathematically integrals, numerically to be evaluated by some quadrature rule — may be computed via asymptotic, deterministic, or sampling methods. In our review of current work we follow our recent publications [28], [35], [29], [34].

One often used technique is a Markov chain Monte Carlo (MCMC) method [21], [9], constructed such that the asymptotic distribution of the Markov chain is the Bayesian posterior distribution; for further information see [34] and the references therein.

These approaches require a large number of samples in order to obtain satisfactory results. Here the main idea here is to perform the Bayesian update directly on the polynomial chaos expansion (PCE) without any sampling [28], [35], [26], [29], [34]. This idea has appeared independently in [1] in a

simpler context, whereas in [37] it appears as a variant of the Kalman filter (e.g. [17]). A PCE for a push-forward of the posterior measure is constructed in [27].

From this short overview it may already have become apparent that the update may be seen abstractly in two different ways. Regarding the uncertain parameters

$$q : \Omega \rightarrow \mathcal{Q} \text{ as a RV on a probability space } (\Omega, \mathfrak{A}, \mathbb{P}) \quad (3)$$

where the set of elementary events is Ω , \mathfrak{A} a σ -algebra of events, and \mathbb{P} a probability measure, one set of methods performs the update by changing the probability measure \mathbb{P} and leaving the mapping $q(\omega)$ as it is, whereas the other set of methods leaves the probability measure unchanged and updates the function $q(\omega)$. In any case, the push forward measure $q_*\mathbb{P}$ on \mathcal{Q} defined by $q_*\mathbb{P}(\mathcal{R}) := \mathbb{P}(q^{-1}(\mathcal{R}))$ for a measurable subset $\mathcal{R} \subset \mathcal{Q}$ is changed from prior to posterior. For the sake of simplicity we assume here that \mathcal{Q} — the set containing possible realisations of q — is a Hilbert space. If the parameter q is a RV, then so is the state u of the system Eq. (1). In order to avoid a profusion of notation, unless there is a possibility of confusion, we will denote the random variables q, f, u which now take values in the respective spaces $\mathcal{Q}, \mathcal{U}^*$ and \mathcal{U} with the same symbol as the previously deterministic quantities in Eq. (1).

In our overview [34] on spectral methods in identification problems, we show that Bayesian identification methods [16], [41], [11], [40] are a good way to tackle the identification problem, especially when these latest developments in functional approximation methods are used. In the series of papers [35], [26], [29], [34], Bayesian updating has been used in a linearised form, strongly related to the Gauss-Markov theorem [20], in ways very similar to the well-known Kalman filter [17]. These similarities will be used to construct an abstract linear filter, which we term the **Gauss-Markov-Kalman** filter (GMKF). This turns out to be a linearised version of *conditional expectation*. Here we want to extend this to a non-linear form, and show some examples of linear (LBU) and non-linear (QBU) Bayesian updates.

The organisation of the remainder of the paper is as follows: in Section 2 we review the Bayesian update—classically defined via conditional probabilities—and recall the link between conditional probability measures and conditional expectation. In Section 3, first we point out in which way — through the conditional expectation — the posterior measure is characterised by Bayes’s theorem, and we point out different possibilities. Often, one does not only want a characterisation of the posterior measure, but actually an RV which has the posterior measure as push-forward or distribution measure.

Some of the well-known filtering algorithms start from this idea. Again by means of the conditional expectation, some possibilities of construction such as an RV are explored, leading to ‘filtering’ algorithms.

In most cases, the conditional expectation can not be computed exactly. We show how the abstract version of the conditional expectation is translated into the possibility of real computational procedures, and how this leads to various approximations, also in connection with the previously introduced filters.

We show how to approximate the conditional expectation up to any desired polynomial degree, not only the linearised version [20], [17] which was used in [28], [35], [26], [29], [34]. This representation in monomials is probably numerically not very advantageous, so we additionally show a version which uses general function bases for approximation.

The numerical realisation in terms of a functional or spectral approximations — here we use the well known Wiener-Hermite chaos — is shortly sketched in Section 4. In Section 5 we then show some computational examples with the *linear version (LBU)*, whereas in Section 6 we show how to compute with the non-linear or quadratic (QBU) version. Some concluding remarks are offered in Section 7.

2 Bayesian Updating

Here we shall describe the frame in which we want to treat the problem of Bayesian updating, namely a dynamical system with time-discrete observations and updates. After introducing the setting in Subsection 2.1, we recall Bayes’s theorem in Subsection 2.2 in the formulation of Laplace, as well as its formulation in the special case where densities exist, e.g. [2]. The next Subsection 2.3 treats the more general case and its connection with the notion of *conditional expectation*, as it was established by Kolmogorov, e.g. [2]. This notion will be the basis of our approach to characterise a RV which corresponds to the posterior measure.

2.1 Setting

In the setting of Eq. (1) consider the following problem: one makes observations y_n at times $0 < t_1 < \dots < t_n \dots \in [0, T]$, and from these one would like to infer what q (and possibly $u(q; t)$) is. In order to include a possible identification of the state $u(q; t_n)$, we shall define a new variable $x = (u, q)$, which we would thus like to identify:

Assume that $U : \mathcal{U} \times \mathcal{Q} \times [0, T] \ni (u_a, q, t) \mapsto u(q; t) \in \mathcal{U}$ is the flow or solution operator of Eq. (1), i.e. $u(q; t) = U(u_a, t_a, q, t)$, where u_a is the initial condition at time t_a . We then look at the operator which advances the variable $x = (u, q) \in \mathcal{X} = \mathcal{U} \times \mathcal{Q}$ from $x_n = (u_n, q)$ at time t_n to $x_{n+1} = (u_{n+1}, q)$ at t_{n+1} , where the Hilbert space \mathcal{X} carries the natural inner product implied from \mathcal{U} and \mathcal{Q} ,

$$x_n = (u_n, q) \mapsto x_{n+1} = (u_{n+1}, q) = (U(u_n, t_n, q, t_{n+1}), q) \in \mathcal{X},$$

or a bit more generally encoded in an operator \hat{f} :

$$\forall n \in \mathbb{N}_0 : \quad x_{n+1} = \hat{f}(x_n, w_n, n); \quad x_0 = x_a \in \mathcal{X} \text{ given.} \quad (4)$$

This is a discrete time step advance map, for example of the dynamical system Eq. (1), where a random ‘error’ term w_n is included, which may be used to model randomness in the dynamical system per se, or possible discretisation errors, or both, or similar things. Most dynamical — and also quasi-static and stationary systems, considering different loadings as a sequence in some pseudo-time — can be put in the form Eq. (4) when observed at discrete points in time. Obviously, for fixed model parameters like q in Eq. (1) the evolution is trivial and does not change anything, but the Eq. (4) allows to model everything in one formulation.

Often the dependence on the random term is assumed to be linear, so that one has

$$\forall n \in \mathbb{N}_0 : \quad x_{n+1} = f(x_n) + \varepsilon S_x(x_n) w_n; \quad x_0 = x_a \text{ given,} \quad (5)$$

where the scalar $\varepsilon \geq 0$ explicitly measures the size of the random term w_n , which is now assumed to be discrete white noise of unit variance and zero mean, and possible correlations are introduced via the linear operator $S_x(x_n)$.

But one can not observe the entity q or $u(q; t)$, i.e. $x = (q, u)$ directly — like in Plato’s cave allegory we can only see a ‘shadow’ — here denoted by a vector $y \in \mathcal{Y}$ — of it, formally given by a ‘measurement operator’

$$Y : \mathcal{X} = \mathcal{Q} \times \mathcal{U} \ni (q, u(t_n)) \mapsto y_{n+1} = Y(q; u(t_n)) \in \mathcal{Y}, \quad (6)$$

where for the sake of simplicity we assume \mathcal{Y} to be a Hilbert space.

Typically one considers also some observational ‘error’ εv_n , so that the observation may be expressed as

$$y_{n+1} = H(Y(q; u(t_n)), \varepsilon v_n) = \hat{h}(x_n, \varepsilon v_n), \quad (7)$$

where similarly as before v_n is a discrete white noise process, and the observer map H resp. \hat{h} combines the ‘true’ quantity $Y(q; u(t_n))$ to be measured with the error, to give the observation y_n .

Translating this into the notation of the discrete dynamical system Eq. (4), one writes

$$y_{n+1} = \hat{h}(x_n, \varepsilon v_n) \in \mathcal{Y}, \quad (8)$$

where again the operator \hat{h} is often assumed to be linear in the noise term, so that one has similarly to Eq. (5)

$$y_{n+1} = h(x_n) + \varepsilon S_y(x_n) w_n \in \mathcal{Y}. \quad (9)$$

The mappings Y in Eq. (6), H in Eq. (7), \hat{h} in Eq. (8), resp. h Eq. (9) are usually not invertible and hence the problem is called *ill-posed*. One way to address this is via regularisation (see e.g. [5]), but here we follow a different track. Modelling our lack of knowledge about q and $u(t_n)$ in a Bayesian way [41] by replacing them with a \mathcal{Q} - resp. \mathcal{U} -valued random variable (RV), the problem becomes well-posed [40]. But of course one is looking now at the problem of finding a probability distribution that best fits the data; and one also obtains a probability distribution, not just *one* pair $x_n = (q, u(t_n))$.

We shall allow for \mathcal{X} to be an infinite-dimensional space, as well as for \mathcal{Y} ; although practically in any real situation only finitely many components are measured. But by allowing for the infinite-dimensional case, we can treat the case of partial differential equations — PDE models — like Eq. (1) directly and not just their discretisations, as its often done, and we only use arguments which are independent on the number of observations. In particular this prevents hidden dependencies on local compactness, the dimension of the model, or the number of measurements, and the possible break-down of computational procedures as these dimensions grow, as they will be designed for the infinite-dimensional case. The procedure practically performed in a real computation on a finite-dimensional model and a finite-dimensional observation may then be seen as an approximation of the infinite-dimensional case, and analysed as such.

Here we focus on the use of a Bayesian approach inspired by the ‘linear Bayesian’ approach of [11] in the framework of ‘white noise’ analysis [13], [14], [15], [22], [43], [44]. Please observe that although the unknown ‘truth’ x_n may be a deterministic quantity, the model for the observed quantity y_{n+1} involves randomness, and it therefore becomes a RV as well.

To complete the mathematical setup we assume that Ω is a measure space with σ -algebra \mathfrak{A} and with a probability measure \mathbb{P} , and that $x : \Omega \rightarrow \mathcal{X}$ and similarly q, u , and y are random variables (RVs). The corresponding *expectation* will be denoted by $\bar{x} = \mathbb{E}(x) = \int_{\Omega} x(\omega) \mathbb{P}(d\omega)$, giving the mean \bar{x} of the random variable, also denoted by $\langle x \rangle := \bar{x}$. The quantity $\tilde{x} := x - \bar{x}$ is the zero-mean or fluctuating part of the RV x .

The space of vector valued RVs, say $x : \Omega \rightarrow \mathcal{X}$, will for simplicity only be considered in the form $\mathcal{X} = \mathcal{X} \otimes \mathcal{S}$, where \mathcal{X} is a Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathcal{X}}$, \mathcal{S} is a Hilbert space of scalar RVs — here we shall simply take $\mathcal{S} = L_2(\Omega, \mathfrak{A}, \mathbb{P})$ — with inner product $\langle \cdot, \cdot \rangle_{\mathcal{S}}$, and the tensor product signifies the Hilbert space completion with the scalar product as usually defined for elementary tensors $x_1 \otimes s_1, x_2 \otimes s_2 \in \mathcal{X}$ with $x_1, x_2 \in \mathcal{X}$ and $s_1, s_2 \in \mathcal{S}$ by

$$\langle\langle x_1 \otimes s_1, x_2 \otimes s_2 \rangle\rangle_{\mathcal{X}} := \langle x_1, x_2 \rangle_{\mathcal{X}} \langle s_1, s_2 \rangle_{\mathcal{S}},$$

and extended to all of \mathcal{X} by linearity.

Obviously, we may also consider the expectation not only as a linear operator $\mathbb{E} : \mathcal{X} \rightarrow \mathcal{X}$, but, as \mathcal{X} is isomorphic to the subspace of constants $\mathcal{X}_c := \mathcal{X} \otimes \text{span}\{1\} \subset \mathcal{X}$, also as an orthogonal projection onto that subspace $\mathbb{E} = P_{\mathcal{X}_c}$, and we have the orthogonal decomposition

$$\mathcal{X} = \mathcal{X}_c \oplus \mathcal{X}_c^{\perp}, \text{ with } \mathcal{X}_c^{\perp} =: \mathcal{X}_0,$$

where \mathcal{X}_0 is the zero-mean subspace, so that

$$\forall x \in \mathcal{X} : \quad \bar{x} = \mathbb{E}(x) = P_{\mathcal{X}_c} x \in \mathcal{X}_c, \quad \tilde{x} = (I - P_{\mathcal{X}_c})x \in \mathcal{X}_0.$$

Later, the covariance operator between two Hilbert-space valued RVs will be needed. The covariance operator between two RVs x and y is denoted by

$$C_{xy} : \mathcal{Y} \ni v \mapsto \mathbb{E}(\tilde{x} \langle \tilde{y}, v \rangle_{\mathcal{Y}}) \in \mathcal{X} \cong \mathcal{X}_c.$$

For $x \in \mathcal{X} \otimes \mathcal{S}$ and $y \in \mathcal{Y} \otimes \mathcal{S}$ it is also often written as $C_{xy} = \mathbb{E}(\tilde{x} \otimes \tilde{y})$.

2.2 Recollection of Bayes's theorem

Bayes's theorem is commonly accepted as a consistent way to incorporate new knowledge into a probabilistic description [16], [41], and its present mathematical form is due to Laplace, so that a better denomination would be the *Bayes-Laplace* theorem.

The elementary textbook statement of the theorem is about conditional probabilities

$$\mathbb{P}(\mathcal{I}_x | \mathcal{M}_y) = \frac{\mathbb{P}(\mathcal{M}_y | \mathcal{I}_x)}{\mathbb{P}(\mathcal{M}_y)} \mathbb{P}(\mathcal{I}_x), \quad \mathbb{P}(\mathcal{M}_y) > 0, \quad (10)$$

where $\mathcal{I}_x \subseteq \mathcal{X}$ is some measurable subset of possible x 's, and the measurable subset $\mathcal{M}_z \subseteq \mathcal{Y}$ is the information provided by the measurement. Here the conditional probability $\mathbb{P}(\mathcal{I}_x | \mathcal{M}_y)$ is called the *posterior* probability, $\mathbb{P}(\mathcal{I}_x)$ is

called the *prior* probability, the conditional probability $\mathbb{P}(\mathcal{M}_y|\mathcal{I}_x)$ is called the *likelihood*, and $\mathbb{P}(\mathcal{M}_y)$ is called the evidence. The Eq. (10) is only valid when the set \mathcal{M}_y has non-vanishing probability measure, and becomes problematic when $\mathbb{P}(\mathcal{M}_y)$ approaches zero, cf. [16], [32]. This arises often when $\mathcal{M}_y = \{y_m\}$ is a one-point set representing a measured value $y_m \in \mathcal{Y}$, as such sets have typically vanishing probability measure. In fact the well-known *Borel-Kolmogorov paradox* has led to numerous controversies and shows the possible ambiguities [16]. Typically the posterior measure is singular w.r.t. the prior measure, precluding a formulation in densities. Kolmogorov's resolution of this situation shall be sketched later.

One well-known very special case where the formulation in densities is possible, which has particular requirements on the likelihood, is when \mathcal{X} —as here—is a metric space, and there is a background measure μ on $(\mathcal{X}, \mathfrak{B}_{\mathcal{X}})$ — $\mathfrak{B}_{\mathcal{X}}$ is the Borel- σ -algebra of \mathcal{X} — and similarly with ν and $(\mathcal{Y}, \mathfrak{B}_{\mathcal{Y}})$, and the RVs x and y have probability density functions (pdf) $\pi_x(x)$ w.r.t. μ and $\pi_y(y)$ w.r.t. ν resp., and a joint density $\pi_{xy}(x, y)$ w.r.t. $\mu \otimes \nu$. Then the theorem may be formulated as ([41] Ch. 1.5, [32], [16])

$$\pi_{(x|y)}(x|y) = \frac{\pi_{xy}(x, y)}{\pi_y(y)} = \frac{\pi_{(y|x)}(y|x)}{Z_y} \pi_x(x), \quad (11)$$

where naturally the marginal density $Z_y := \pi_y(y) = \int_{\mathcal{X}} \pi_{xy}(x, y) \mu(dx)$ (from German *Zustandssumme*) is a normalising factor such that the conditional density $\pi_{(x|y)}(\cdot|y)$ integrates to unity w.r.t. x . In this case the limiting case where $\mathbb{P}(\mathcal{M}_y)$ vanishes may be captured via the metric [32] [16]. The joint density

$$\pi_{xy}(x, y) = \pi_{(y|x)}(y|x) \pi_x(x)$$

may be factored into the likelihood function $\pi_{(y|x)}(y|x)$ and the prior density $\pi_x(x)$, like $\pi_y(y)$ a marginal density, $\pi_x(x) = \int_{\mathcal{Y}} \pi_{xy}(x, y) \nu(dy)$. These terms in the second equality in Eq. (11) are in direct correspondence with those in Eq. (10). Please observe that the model for the RV representing the error in Eq. (8) determines the likelihood functions $\mathbb{P}(\mathcal{M}_y|\mathcal{I}_x)$ resp. $\pi_{(y|x)}(y|x)$. To require the existence of the joint density is quite restrictive. As Eq. (8) shows, y is a function of x , and a joint density on $\mathcal{X} \times \mathcal{Y}$ will generally not be possible as $(x, y) \in \mathcal{X} \times \mathcal{Y}$ are most likely on a sub-manifold; but the situation of Eq. (9) is one possibility where a joint density may be established. The background densities are typically in finite dimensions the Lebesgue measure on \mathbb{R}^d , or more general Haar measures on locally compact Lie-groups [39]. Most computational approaches determine the pdfs [23], [40], [18].

However, to avoid the critical cases alluded to above, Kolmogorov already defined conditional probabilities via conditional expectation, e.g. see [2].

Given the conditional expectation operator $\mathbb{E}(\cdot|\mathcal{M}_y)$, the conditional probability is easily recovered as $\mathbb{P}(\mathcal{I}_x|\mathcal{M}_y) = \mathbb{E}(\chi_{\mathcal{I}_x}|\mathcal{M}_y)$, where $\chi_{\mathcal{I}_x}$ is the characteristic function of the subset \mathcal{I}_x . It may be shown that this extends the simpler formulation described by Eq. (10) or Eq. (11) and is the more fundamental notion, which we examine next. Its definition will lead directly to practical computational procedures.

2.3 Conditional expectation

The easiest point of departure for conditional expectation [2] in our setting is to define it not just for one piece of measurement \mathcal{M}_y —which may not even be possible unambiguously—but for sub- σ -algebras $\mathfrak{S} \subset \mathfrak{A}$ on Ω . A sub- σ -algebra \mathfrak{S} is a mathematical description of a reduced possibility of randomness — the smallest sub- σ -algebra $\{\emptyset, \Omega\}$ allows only the constants in \mathcal{X}_c — as it contains fewer events than the full algebra \mathfrak{A} . The connection with a measurement \mathcal{M}_y is to take $\mathfrak{S} := \sigma(y)$, the σ -algebra generated by the measurement $y = \hat{h}(x, \varepsilon v)$ from Eq. (8). These are all events which are consistent with possible observations of some value for y . This means that the observation of y allows only a certain ‘fineness’ of information to be obtained, and this is encoded in the sub- σ -algebra \mathfrak{S} .

2.3.1 Scalar random variables

For scalar RVs — functions $r(x)$ of x with finite variance, i.e. elements of $\mathcal{S} := L_2(\Omega, \mathfrak{A}, \mathbb{P})$ —the subspace corresponding to the sub- σ -algebra $\mathcal{S}_\infty := L_2(\Omega, \mathfrak{S}, \mathbb{P})$ is a closed subspace [2] of the full space \mathcal{S} . One example of such a scalar RV is the function

$$r(x) := \chi_{\mathcal{I}_x}(x) = \begin{cases} 1 & \text{if } x \in \mathcal{I}_x, \\ 0 & \text{otherwise,} \end{cases}$$

mentioned at the end of Subsection 2.2 used to define conditional probability of the subset $\mathcal{I}_x \subseteq \mathcal{X}$ once a conditional expectation operator is defined: $\mathbb{P}(\mathcal{I}_x|\mathfrak{S}) = \mathbb{E}(\chi_{\mathcal{I}_x}|\mathfrak{S})$.

Definition 1. For scalar functions of x — scalar RVs $r(x)$ — in \mathcal{S} , the conditional expectation $\mathbb{E}(\cdot|\mathfrak{S})$ is defined as the orthogonal projection onto the closed subspace \mathcal{S}_∞ , so that $\mathbb{E}(r(x)|\mathfrak{S}) \in \mathcal{S}_\infty$, e.g. see [2].

The question is now on how to characterise this subspace \mathcal{S}_∞ , in order to make it more accessible for possible numerical computations. In this regard, note that the *Doob-Dynkin* lemma [2] assures us that if a RV $s(x)$ — like

$\mathbb{E}(r(x)|\mathfrak{G})$ — is in the subspace \mathcal{S}_∞ , then $s(x) = \varphi(y)$ for some $\varphi \in L_0(\mathcal{Y})$, the space of measurable scalar functions on \mathcal{Y} . We state this key fact and the resulting new characterisation of the conditional expectation in

Proposition 2. *The subspace \mathcal{S}_∞ is given by*

$$\mathcal{S}_\infty = \overline{\text{span}}\{\varphi \mid \varphi(\hat{h}(x, \varepsilon v)); \varphi \in L_0(\mathcal{Y}) \text{ and } \varphi \in \mathcal{S}\}. \quad (12)$$

The conditional expectation of a scalar RV $r(x) \in \mathcal{S}$, being the orthogonal projection, minimises the distance to the original RV over the whole subspace:

$$\mathbb{E}(r(x)|\mathfrak{G}) := P_{\mathcal{S}_\infty}(r(x)) := \arg \min_{\tilde{r} \in \mathcal{S}_\infty} \|r(x) - \tilde{r}\|_{\mathcal{S}}, \quad (13)$$

where $P_{\mathcal{S}_\infty}$ is the orthogonal projector onto \mathcal{S}_∞ . The Eq. (12) and Eq. (13) imply the existence of a optimal map $\phi \in L_0(\mathcal{Y})$ such that

$$\mathbb{E}(r(x)|\mathfrak{G}) = P_{\mathcal{S}_\infty}(r(x)) = \phi(\hat{h}(x, \varepsilon v)). \quad (14)$$

In Eq. (13), one may equally well minimise the square of the distance, the loss-function

$$\beta_{r(x)}(\tilde{r}) = \frac{1}{2} \|r(x) - \tilde{r}\|_{\mathcal{S}}^2. \quad (15)$$

Taking the vanishing of the first variation / Gâteaux derivative of the loss-function Eq. (15) as a necessary condition for a minimum leads to a simple geometrical interpretation: the difference between the original scalar RV $r(x)$ and its projection has to be perpendicular to the subspace:

$$\forall \tilde{r} \in \mathcal{S}_\infty : \langle r(x) - \mathbb{E}(r(x)|\mathfrak{G}), \tilde{r} \rangle_{\mathcal{S}} = 0, \text{ i.e. } r(x) - \mathbb{E}(r(x)|\mathfrak{G}) \in \mathcal{S}_\infty^\perp. \quad (16)$$

Rephrasing Eq. (13) with account to Eq. (16) and Eq. (15) leads for the optimal map $\phi \in L_0(\mathcal{Y})$ to

$$\mathbb{E}(r(x)|\sigma(y)) = \phi(\hat{h}(x, \varepsilon v)) := \arg \min_{\varphi \in L_0(\mathcal{Y})} \beta_{r(x)}(\varphi(\hat{h}(x, \varepsilon v))), \quad (17)$$

and the orthogonality condition of Eq. (17) which corresponds to Eq. (16) leads to

$$\forall \varphi \in L_0(\mathcal{Y}) : \langle r(x) - \phi(\hat{h}(x, \varepsilon v)), \varphi(\hat{h}(x, \varepsilon v)) \rangle_{\mathcal{S}} = 0. \quad (18)$$

Proof. The Eq. (12) is a direct statement of the Doob-Dynkin lemma [2], and the Eq. (13) is equivalent to the definition of the conditional expectation being an orthogonal projection in $L_2(\Omega, \mathfrak{A}, \mathbb{P})$ — actually an elementary fact of Euclidean geometry.

The existence of the optimal map ϕ in Eq. (14) is a consequence of the minimisation of a continuous, coercive, and strictly convex function — the norm Eq. (13) — over the closed set \mathcal{S}_∞ in the complete space \mathcal{S} . The equivalence of minimising the norm Eq. (13) and Eq. (15) is elementary, which is re-stated in Eq. (17).

The two equivalent statements — the ‘Galerkin orthogonality’ conditions — Eq. (16) and Eq. (18) follow not only from requiring the Gâteaux derivative of Eq. (15) to vanish, but also express an elementary fact of Euclidean geometry. \square

The square of the distance $r(x) - \phi(y)$ may be interpreted as a difference in variance, tying conditional expectation with variance minimisation; see for example [30], [2], and the references therein for basic descriptions of conditional expectation. See also [20].

2.3.2 Vector valued random variables

Now assume that $R(x)$ is a function of x which takes values in a vector space \mathcal{R} , i.e. a \mathcal{R} -valued RV, where \mathcal{R} is a Hilbert space. Two simple examples are given by the conditional mean where $R(x) := x \in \mathcal{X}$ with $\mathcal{R} = \mathcal{X}$, and by the conditional variance where one takes $R(x) := (x - \bar{x}) \otimes (x - \bar{x}) = (\tilde{x}) \otimes (\tilde{x})$, where $\mathcal{R} = \mathcal{L}(\mathcal{X})$. The Hilbert tensor product $\mathcal{R} = \mathcal{R} \otimes \mathcal{S}$ is again needed for such vector valued RVs, where a bit more formalism is required, as we later want to take linear combinations of RVs, but with linear operators as ‘coefficients’ [20], and this is most clearly expressed in a component-free fashion in terms of L -invariance, where we essentially follow [3], [4]:

Definition 3. *Let \mathcal{V} be a subspace of $\mathcal{R} = \mathcal{R} \otimes \mathcal{S}$. The subspace is called linearly closed, L -closed, or L -invariant, iff \mathcal{V} is closed, and $\forall v \in \mathcal{V}$ and $\forall L \in \mathcal{L}(\mathcal{R})$ it holds that $Lv \in \mathcal{V}$.*

In finite dimensional spaces one can just apply the notions for the scalar case in section 2.3.1 component by component, but this is not possible in the infinite dimensional case. Of course the vectorial description here collapses to the scalar case upon taking $\mathcal{R} = \mathbb{R}$. From [3] one has the following

Proposition 4. *It is obvious that the whole space $\mathcal{R} = \mathcal{R} \otimes \mathcal{S}$ is linearly closed, and that for a linearly closed subspace $\mathcal{V} \subseteq \mathcal{R}$ its orthogonal complement \mathcal{V}^\perp is also linearly closed. Clearly, for a closed subspace $\mathcal{S}_a \subseteq \mathcal{S}$, the tensor space $\mathcal{R} \otimes \mathcal{S}_a$ is linearly closed, and hence the space of constants $\mathcal{R}_c = \mathcal{R} \otimes \text{span}\{\chi_\Omega\} \cong \mathcal{R}$ is linearly closed, as well as its orthogonal complement $\mathcal{R}_0 = \mathcal{R}_c^\perp$, the subspace of zero-mean RVs.*

Let $v \in \mathcal{R}$ be a RV, and denote by

$$\mathcal{R}_v := \overline{\text{span}} v(\Omega), \quad \sigma(v) := \{v^{-1}(B) : B \in \mathfrak{B}_{\mathcal{R}}\} \quad (19)$$

the closure of the span of the image of v and the σ -algebra generated by v , where $\mathfrak{B}_{\mathcal{R}}$ is the Borel- σ -algebra of \mathcal{R} . Denote the closed subspace generated by $\sigma(v)$ by $\mathcal{S}_v := L_2(\Omega, \sigma(v), \mathbb{P}) \subseteq \mathcal{S}$. Let $\mathcal{R}_{Lv} := \overline{\text{span}} \{Lv : L \in \mathcal{L}(\mathcal{R})\} \subseteq \mathcal{R}$, the linearly closed subspace generated by v , and finally denote by $\mathcal{R}_v := \text{span}\{v\} \subseteq \mathcal{R}$, the one-dimensional ray and hence closed subspace generated by v . Obviously it holds that

$$v \in \mathcal{R}_v \subseteq \mathcal{R}_{Lv} \subseteq \mathcal{R} \otimes \mathcal{S}_v \subseteq \mathcal{R}, \quad \text{and } \bar{v} \in \mathcal{R}_v,$$

and $\mathcal{R} \otimes \mathcal{S}_v$ is linearly closed according to Proposition 4.

Definition 5. Let \mathcal{V} and \mathcal{W} be subspaces of \mathcal{R} , and $v, w \in \mathcal{R}$ two RVs.

- The two subspaces are weakly orthogonal or simply just orthogonal, denoted by
 $\mathcal{V} \perp \mathcal{W}$, iff $\forall v \in \mathcal{V}, \forall w \in \mathcal{W}$ it holds that $\langle\langle v, w \rangle\rangle_{\mathcal{R}} = 0$.
- A RV $v \in \mathcal{R}$ is weakly orthogonal or simply just orthogonal to the subspace \mathcal{W} , denoted by
 $v \perp \mathcal{W}$, iff $\mathcal{R}_v \perp \mathcal{W}$, i.e. $\forall w \in \mathcal{W}$ it holds that $\langle\langle v, w \rangle\rangle_{\mathcal{R}} = 0$.
- Two RVs $v, w \in \mathcal{R}$ are weakly orthogonal or as usual simply just orthogonal, denoted by
 $v \perp w$, iff $\langle\langle v, w \rangle\rangle_{\mathcal{R}} = 0$, i.e. $\mathcal{R}_v \perp \mathcal{R}_w$.
- The two subspaces \mathcal{V} and \mathcal{W} are strongly orthogonal or L-orthogonal, iff they are linearly closed—Definition 3—and it holds that $\langle\langle Lv, w \rangle\rangle_{\mathcal{R}} = 0, \forall v \in \mathcal{V}, \forall w \in \mathcal{W}$ and $\forall L \in \mathcal{L}(\mathcal{R})$. This is denoted by
 $\mathcal{V} \perp\!\!\!\perp \mathcal{W}$, and in other words $\mathcal{L}(\mathcal{R}) \ni C_{vw} = \mathbb{E}(v \otimes w) = 0$.
- The RV v is strongly orthogonal to a linearly closed subspace $\mathcal{W} \subseteq \mathcal{R}$, denoted by
 $v \perp\!\!\!\perp \mathcal{W}$, iff $\mathcal{R}_{Lv} \perp\!\!\!\perp \mathcal{W}$, i.e. $\forall w \in \mathcal{W}$ it holds that $C_{vw} = 0$.
- The two RVs v, w are strongly orthogonal or simply just uncorrelated, denoted by
 $v \perp\!\!\!\perp w$, iff $C_{vw} = 0$, i.e. $\mathcal{R}_{Lv} \perp\!\!\!\perp \mathcal{R}_{Lw}$.

- Let $\mathfrak{C}_1, \mathfrak{C}_2 \subseteq \mathfrak{A}$ be two sub- σ -algebras. They are independent, denoted by
 $\mathfrak{C}_1 \perp\!\!\!\perp \mathfrak{C}_2$, iff the closed subspaces of \mathcal{S} generated by them are orthogonal in \mathcal{S} :
 $L_2(\Omega, \mathfrak{C}_1, \mathbb{P}) \perp L_2(\Omega, \mathfrak{C}_2, \mathbb{P})$.
- The two subspaces \mathcal{V} and \mathcal{W} are stochastically independent, denoted by
 $\mathcal{V} \perp\!\!\!\perp \mathcal{W}$, iff the sub- σ -algebras generated are: $\sigma(\mathcal{V}) \perp\!\!\!\perp \sigma(\mathcal{W})$.
- The two RVs v, w are stochastically independent, denoted by
 $v \perp\!\!\!\perp w$, iff $\sigma(v) \perp\!\!\!\perp \sigma(w)$, i.e. $\mathcal{S}_v \perp \mathcal{S}_w$.

Proposition 6. Obviously $\mathcal{R}_c \perp \mathcal{R}_0$. It is equally obvious that for any two closed subspaces $\mathcal{S}_a, \mathcal{S}_b \subseteq \mathcal{S}$, the condition $\mathcal{S}_a \perp \mathcal{S}_b$ implies that the tensor product subspaces are strongly orthogonal:

$$\mathcal{R} \otimes \mathcal{S}_a \perp \mathcal{R} \otimes \mathcal{S}_b.$$

This implies that for a closed subspace $\mathcal{S}_s \subseteq \mathcal{S}$ the subspaces $\mathcal{R}_s = \mathcal{R} \otimes \mathcal{S}_s \subseteq \mathcal{R}$ and its orthogonal complement $\mathcal{R}_s^\perp = \mathcal{R} \otimes \mathcal{S}_s^\perp$ are linearly closed and strongly orthogonal.

We note from [3], [4] the following results which we collect in

Proposition 7. Let $v, w \in \mathcal{R}_0$ be two zero-mean RVs. Then

$$v \perp\!\!\!\perp w \Rightarrow v \perp w \Rightarrow v \perp w.$$

Strong orthogonality in general does not imply independence, and orthogonality does not imply strong orthogonality, unless \mathcal{R} is one-dimensional.

If $\mathcal{S} \subseteq \mathcal{R}$ is linearly closed, then

$$v \perp \mathcal{S} \Rightarrow v \perp\!\!\!\perp \mathcal{S}, \text{ i.e. } \mathcal{R}_v \perp \mathcal{S} \Rightarrow \mathcal{R}_v \perp\!\!\!\perp \mathcal{S} \Rightarrow \mathcal{R}_{Lv} \perp\!\!\!\perp \mathcal{S}.$$

From this we obtain the following:

Lemma 8. Set $\mathcal{R}_\infty := \mathcal{R} \otimes \mathcal{S}_\infty$ for the \mathcal{R} -valued RV $R(x)$ with finite variance on the sub- σ -algebra \mathfrak{S} , representing the new information.

Then \mathcal{R}_∞ is L -invariant or strongly closed, and for any zero mean RV $v \in \mathcal{R}$:

$$v \in \mathcal{R}_\infty^\perp \Leftrightarrow v \perp \mathcal{R}_\infty \Rightarrow v \perp\!\!\!\perp \mathcal{R}_\infty. \quad (20)$$

In addition, it holds — even if $v \in \mathcal{R}$ is not zero mean — that

$$v \in \mathcal{R}_\infty^\perp \Leftrightarrow v \perp \mathcal{R}_\infty \Rightarrow \forall w \in \mathcal{R}_\infty : \mathbb{E}(v \otimes w) = 0. \quad (21)$$

Proof. \mathcal{R}_∞ is of the type $\mathcal{R} \otimes \mathcal{S}_\infty$ where \mathcal{S}_∞ is a closed subspace, and \mathcal{R} is obviously closed. From the remarks above it follows that \mathcal{R}_∞ is L -invariant or linearly resp. strongly closed. The Eq. (20) is a direct consequence of Proposition 7.

To prove Eq. (21), take any $w \in \mathcal{R}_\infty$ and any $L \in \mathcal{L}(\mathcal{R})$, then

$$v \in \mathcal{R}_\infty^\perp \Rightarrow 0 = \langle\langle v, w \rangle\rangle_{\mathcal{R}} = \langle\langle v, Lw \rangle\rangle_{\mathcal{R}} = \mathbb{E}(\langle v, Lw \rangle_{\mathcal{R}}).$$

Now, for any $r_1, r_2 \in \mathcal{R}$, take the mapping $L : r_* \mapsto \langle r_2, r_* \rangle_{\mathcal{R}} r_1$, yielding

$$\begin{aligned} 0 &= \mathbb{E}(\langle v, Lw \rangle_{\mathcal{R}}) = \mathbb{E}(\langle v, \langle r_2, w \rangle_{\mathcal{R}} r_1 \rangle_{\mathcal{R}}) = \\ &\quad \mathbb{E}(\langle v, r_1 \rangle_{\mathcal{R}} \langle r_2, w \rangle_{\mathcal{R}}) = \langle r_1, \mathbb{E}(v \otimes w) r_2 \rangle_{\mathcal{R}} \Leftrightarrow \mathbb{E}(v \otimes w) \equiv 0. \end{aligned}$$

□

Extending the scalar case described in section 2.3.1, instead of

$$\mathcal{S} = L_2(\Omega, \mathbb{P}, \mathfrak{A}) = L_2(\Omega, \mathbb{P}, \mathfrak{A}; \mathbb{R}) \cong \mathbb{R} \otimes L_2(\Omega, \mathbb{P}, \mathfrak{A}) = \mathbb{R} \otimes \mathcal{S}$$

and its subspace generated by the measurement

$$\mathcal{S}_\infty = L_2(\Omega, \mathbb{P}, \mathfrak{S}) = L_2(\Omega, \mathbb{P}, \mathfrak{S}; \mathbb{R}) \cong \mathbb{R} \otimes L_2(\Omega, \mathbb{P}, \mathfrak{S}) = \mathbb{R} \otimes \mathcal{S}_\infty$$

one now considers the space Eq. (22) and its subspace Eq. (23)

$$L_2(\Omega, \mathbb{P}, \mathfrak{A}; \mathcal{R}) \cong \mathcal{R} \otimes L_2(\Omega, \mathbb{P}, \mathfrak{A}) = \mathcal{R} \otimes \mathcal{S} := \mathcal{R} \quad \text{and} \quad (22)$$

$$L_2(\Omega, \mathbb{P}, \mathfrak{S}; \mathcal{R}) \cong \mathcal{R} \otimes L_2(\Omega, \mathbb{P}, \mathfrak{S}) = \mathcal{R} \otimes \mathcal{S}_\infty := \mathcal{R}_\infty \subseteq \mathcal{R}. \quad (23)$$

The conditional expectation in the vector-valued case is defined completely analogous to the scalar case, see Definition 1:

Definition 9. For \mathcal{R} -valued functions of x — vectorial RVs $R(x)$ — in the Hilbert-space \mathcal{R} Eq. (22), the conditional expectation $\mathbb{E}(\cdot | \mathfrak{S}) : \mathcal{R} \rightarrow \mathcal{R}$ is defined as the orthogonal projection onto the closed subspace \mathcal{R}_∞ Eq. (23), denoted by $P_{\mathcal{R}_\infty}$, so that $\mathbb{E}(R(x) | \mathfrak{S}) = P_{\mathcal{R}_\infty}(R(x)) \in \mathcal{R}_\infty$, e.g. see [3], [2].

From this one may derive a characterisation of the conditional expectation similar to Proposition 2.

Theorem 10. The subspace \mathcal{R}_∞ is given by

$$\mathcal{R}_\infty = \{\varphi \mid \varphi(\hat{h}(x, \varepsilon v)) \in \mathcal{R}; \varphi \in L_0(\mathcal{Y}, \mathcal{R})\}. \quad (24)$$

The conditional expectation of a vector-valued RV $R(x) \in \mathcal{R}$, being the orthogonal projection, minimises the distance to the original RV over the whole subspace:

$$\mathbb{E}(R(x)|\mathfrak{G}) := P_{\mathcal{R}_\infty}(R(x)) := \arg \min_{\tilde{R} \in \mathcal{R}_\infty} \|R(x) - \tilde{R}\|_{\mathcal{R}}, \quad (25)$$

where $P_{\mathcal{R}_\infty}$ is the orthogonal projector onto \mathcal{R}_∞ . The Eq. (24) and Eq. (25) imply the existence of a optimal map $\Phi \in L_0(\mathcal{Y}, \mathcal{R})$ such that

$$\mathbb{E}(R(x)|\mathfrak{G}) = P_{\mathcal{R}_\infty}(R(x)) = \Phi(\hat{h}(x, \varepsilon v)). \quad (26)$$

In Eq. (25), one may equally well minimise the square of the distance, the loss-function

$$\beta_{R(x)}(\tilde{R}) = \frac{1}{2} \|R(x) - \tilde{R}\|_{\mathcal{R}}^2. \quad (27)$$

Taking the vanishing of the first variation / Gâteaux derivative of the loss-function Eq. (27) as a necessary condition for a minimum leads to a simple geometrical interpretation: the difference between the original vector-valued RV $R(x)$ and its projection has to be perpendicular to the subspace \mathcal{R}_∞ : $\forall \tilde{R} \in \mathcal{R}_\infty$:

$$\langle R(x) - \mathbb{E}(R(x)|\mathfrak{G}), \tilde{R} \rangle_{\mathcal{R}} = 0, \quad \text{i.e. } R(x) - \mathbb{E}(R(x)|\mathfrak{G}) \in \mathcal{R}_\infty^\perp. \quad (28)$$

Rephrasing Eq. (25) with account to Eq. (28) and Eq. (27) leads for the optimal map $\Phi \in L_0(\mathcal{Y}, \mathcal{R})$ to

$$\mathbb{E}(R(x)|\sigma(y)) = \Phi(\hat{h}(x, \varepsilon v)) := \arg \min_{\varphi \in L_0(\mathcal{Y}, \mathcal{R})} \beta_{R(x)}(\varphi(\hat{h}(x, \varepsilon v))), \quad (29)$$

and the orthogonality condition of Eq. (29) which corresponds to Eq. (28) leads to

$$\forall \varphi \in L_0(\mathcal{Y}, \mathcal{R}) : \langle R(x) - \Phi(\hat{h}(x, \varepsilon v)), \varphi(\hat{h}(x, \varepsilon v)) \rangle_{\mathcal{R}} = 0. \quad (30)$$

In addition, as \mathcal{R}_∞ is linearly closed, one obtains the useful statement

$$\forall \tilde{R} \in \mathcal{R}_\infty : \mathcal{L}(\mathcal{R}) \ni \mathbb{E}((R(x) - \mathbb{E}(R(x)|\mathfrak{G})) \otimes \tilde{R}) = 0. \quad (31)$$

or rephrased $\forall \varphi \in L_0(\mathcal{Y}, \mathcal{R})$:

$$\mathcal{L}(\mathcal{R}) \ni \mathbb{E}((R(x) - \Phi(\hat{h}(x, \varepsilon v))) \otimes \varphi(\hat{h}(x, \varepsilon v))) = 0. \quad (32)$$

Proof. The Eq. (24) is just a version of the Doob-Dynkin lemma again [2], this time for vector-valued functions. The Eq. (25), Eq. (26), Eq. (27), Eq. (28), Eq. (29), and Eq. (30) follow just as in the scalar case of Proposition 2.

As \mathcal{R}_∞ is linearly closed according to Lemma 8, the Eq. (20) causes Eq. (28) to imply Eq. (31), and Eq. (30) together with Eq. (21) from Lemma 8 to imply Eq. (32). \square

Already in [17] it was noted that the conditional expectation is the best estimate not only for the *loss function* ‘distance squared’, as in Eq. (15) and Eq. (27), but for a much larger class of loss functions under certain distributional constraints. However for the quadratic loss function this is valid without any restrictions.

Requiring the derivative of the quadratic loss function in Eq. (15) and Eq. (27) to vanish may also be characterised by the *Lax-Milgram* lemma, as one is minimising a quadratic functional over the vector space \mathcal{R}_∞ , which is closed and hence a *Hilbert* space. For later reference, this result is recollected in

Theorem 11. *In the scalar case, there is a unique minimiser $\mathbb{E}(r(x)|\mathfrak{S}) = P_{\mathcal{S}_\infty}(r(x)) \in \mathcal{S}_\infty$ to the problem in Eq. (13), and it is characterised by the orthogonality condition Eq. (16)*

$$\forall \tilde{r} \in \mathcal{S}_\infty : \quad \langle r(x) - \mathbb{E}(r(x)|\mathfrak{S}), \tilde{r} \rangle_{\mathcal{S}} = 0. \quad (33)$$

The minimiser is unique as an element of \mathcal{S}_∞ , but the mapping $\phi \in L_0(\mathcal{Y})$ in Eq. (17) may not necessarily be. It also holds that

$$\|P_{\mathcal{S}_\infty}(r(x))\|_{\mathcal{S}}^2 = \|r(x)\|_{\mathcal{S}}^2 - \|r(x) - P_{\mathcal{S}_\infty}(r(x))\|_{\mathcal{S}}^2. \quad (34)$$

As in the scalar case, in the vector-valued case there is a unique minimiser $\mathbb{E}(R(x)|\mathfrak{S}) = P_{\mathcal{R}_\infty}(R(x)) \in \mathcal{R}_\infty$ to the problem in Eq. (25), which satisfies the orthogonality condition Eq. (28)

$$\forall \tilde{R} \in \mathcal{R}_\infty : \quad \langle R(x) - \mathbb{E}(R(x)|\mathfrak{S}), \tilde{R} \rangle_{\mathcal{R}} = 0, \quad (35)$$

which is equivalent to the the strong orthogonality condition Eq. (31)

$$\forall \tilde{R} \in \mathcal{R}_\infty : \quad \mathbb{E}(R(x) - \mathbb{E}(R(x)|\mathfrak{S}) \otimes \tilde{R}) = 0. \quad (36)$$

The minimiser is unique as an element of \mathcal{R}_∞ , but the mapping $\Phi \in L_0(\mathcal{Y}, \mathcal{R})$ in Eq. (29) may not necessarily be. It also holds that

$$\|P_{\mathcal{R}_\infty}(R(x))\|_{\mathcal{R}}^2 = \|R(x)\|_{\mathcal{R}}^2 - \|R(x) - P_{\mathcal{R}_\infty}(R(x))\|_{\mathcal{R}}^2. \quad (37)$$

Proof. It is all already contained in Proposition 2 resp. Theorem 10. Except for Eq. (36), this is just a re-phrasing of the *Lax-Milgram* lemma, as the bi-linear functional—in this case the inner product—is naturally coercive and continuous on the subspace \mathcal{R}_∞ , which is closed and hence a *Hilbert* space. The only novelty here are the Eq. (34) and Eq. (37) which follow from *Pythagoras’s* theorem. \square

3 Characterising the posterior

The information contained in the Bayesian update is encoded in the conditional expectation. And it only characterises the distribution of the posterior. A few different ways of characterising the distribution via the conditional expectation are sketched in Subsection 3.1. But in many situations, notably in the setting of Eq. (4) or Eq. (5), with the observations according to Eq. (8) or Eq. (9), we want to construct a new RV $z \in \mathcal{X}$ to serve as an approximation to the solution of Eq. (4) or Eq. (5). This then is a *filter*, and a few possibilities will be given in Subsection 3.2.

3.1 The posterior distribution measure

It was already mentioned at the beginning of section 2.3.1, that the scalar function $r_{\mathcal{I}_x}(x) = \chi_{\mathcal{I}_x}(x)$ may be used to characterise the conditional probability distribution of $x \in \mathcal{X}$. Indeed, if for a RV $R(x) \in \mathcal{R}$ one defines:

$$\forall \mathcal{E} \in \mathfrak{B}_{\mathcal{R}} : \mathbb{P}(\mathcal{E}|\mathfrak{S}) := \mathbb{E}(\chi_{\mathcal{E}}(R)|\mathfrak{S}), \quad (38)$$

one has completely characterised the posterior distribution, a version of which is under certain conditions—[2], [32], [16]—a measure on \mathcal{R} , the image space of the RV R .

One may also recall that the *characteristic function* in the sense of stochastics of a RV $R \in \mathcal{R}$, namely

$$\varphi_R : \mathcal{R}^* \ni r^* \mapsto \varphi_R(r^*) := \mathbb{E}(\exp(i \langle r^*, R \rangle_{\mathcal{R}})),$$

completely characterises the distribution of the RV R . As we assume that \mathcal{R} is a Hilbert space, we may identify \mathcal{R} with its dual space \mathcal{R}^* , and in this case take φ_R as defined on \mathcal{R} . If now a conditional expectation operator $\mathbb{E}(\cdot|\mathfrak{S})$ is given, it may be used to define the *conditional* characteristic function $\varphi_{R|\mathfrak{S}}$:

$$\forall r \in \mathcal{R} : \varphi_{R|\mathfrak{S}}(r) := \mathbb{E}(\exp(i \langle r, R \rangle_{\mathcal{R}})|\mathfrak{S}). \quad (39)$$

This again completely characterises the posterior distribution.

Another possible way, actually encompassing the previous two, is to look at all functions $\psi : \mathcal{R} \rightarrow \mathbb{R}$, and compute—when they are defined and finite—the quantities

$$\mu_{\psi} := \mathbb{E}(\psi(R)|\mathfrak{S}), \quad (40)$$

again completely characterising the posterior distribution. The two previous examples show that not *all* functions of R with finite conditional expectation are needed. The first example uses the set of functions

$$\{\psi \mid \psi(R) = \chi_{\mathcal{E}}(R), \mathcal{E} \in \mathfrak{B}_{\mathcal{R}}\},$$

whereas the second example uses the set

$$\{\psi \mid \psi(R) = \exp(i \langle r, R \rangle_{\mathcal{R}}), r \in \mathcal{R}\}.$$

3.2 A posterior random variable — filtering

In the context of a situation like in Eq. (4) resp. Eq. (5), which represents the *unknown* system and state vector x_n , and where one observes y_n according to Eq. (8) resp. Eq. (9), one wants to have an estimating or *tracking* model system, with a state estimate z_n for x_n which would in principle obey Eq. (4) resp. Eq. (5) with the noise w_n set to zero—as one only knows the structure of the system as given by the maps \hat{f} resp. f but not the initial condition x_0 nor the noise. The observations y_n can be used to correct the state estimate z_n , as will be shown shortly. The state estimate will be computed via Bayesian updating. But the Bayesian theory, as explained above, only characterises the posterior *distribution*; and there are many random variables which might have a given distribution. To obtain a RV z_n which can be used to predict the next state x_{n+1} through the estimate z_{n+1} one may use a filter based on Bayesian theory. The mean vehicle for this will be the notion of conditional expectation as described in the previous Section 2. As we will first consider only one update step, the time index n will be dropped for the sake of ease of notation: The true state is x , its *forecast* is x_f , and the forecast of the measurement is $y_f(x_f)$, whereas the observation is \hat{y} .

To recall, according to Definition 9, the Bayesian update is defined via the conditional expectation $\mathbb{E}(R(x)|\sigma(y(x)))$ through a measurement $y(x)$ —which will for the sake of simplicity be denoted just by $\mathbb{E}(R(x)|y)$ —of a \mathcal{R} -valued RV $R(x)$ is simply the orthogonal projection onto the subspace \mathcal{R}_∞ in Eq. (24),

$$\mathbb{E}(R(x)|y) = P_{\mathcal{R}_\infty}(R(x)) = \Phi_R(y(x)),$$

which is given by the optimal map Φ_R from Eq. (26), characterised by Eq. (32), where we have added an index R to signify that this is the optimal map for the conditional expectation of the RV $R \in \mathcal{R}$.

The linearly closed subspace \mathcal{R}_∞ induces a orthogonal decomposition

$$\mathcal{R} = \mathcal{R}_\infty \oplus \mathcal{R}_\infty^\perp,$$

where the orthogonal projection onto \mathcal{R}_∞^\perp is given by $I - P_{\mathcal{R}_\infty}$. Hence a RV in \mathcal{R} like $R(x)$ can be decomposed accordingly as

$$\begin{aligned} R(x) &= P_{\mathcal{R}_\infty}(R(x)) + (I - P_{\mathcal{R}_\infty})(R(x)) = \\ &\quad \Phi_R(y(x)) + (R(x) - \Phi_R(y(x))). \end{aligned} \quad (41)$$

This Eq. (41) is the starting point for the updating. A measurement \hat{y} will inform us about the component in \mathcal{R}_∞ , namely $\Phi_R(\hat{y})$, while we leave the component orthogonal to it unchanged: $R(x_f) - \Phi_R(y(x_f))$. Adding these two terms then gives an *updated* or *assimilated* RV $R_a \in \mathcal{R}$:

$$\begin{aligned} R_a &= \Phi_R(\hat{y}) + (R(x_f) - \Phi_R(y(x_f))) = \bar{R}_a^{\hat{y}} + \tilde{R}_a = \\ &R(x_f) + (\Phi_R(\hat{y}) - \Phi_R(y(x_f))) = R_f + R_\infty, \end{aligned} \quad (42)$$

where $R_f = R(x_f) \in \mathcal{R}$ is the *forecast* and $R_\infty = (\Phi_R(\hat{y}) - \Phi_R(y(x_f))) \in \mathcal{R}$ is the *innovation*. For $\bar{R}_a^{\hat{y}} = \Phi_R(\hat{y})$ and $\tilde{R}_a = R(x_f) - \Phi_R(y(x_f))$ one has the following result:

Proposition 12. *The assimilated RV R_a from Eq. (42) has the correct conditional expectation*

$$\mathbb{E}(R_a|y) = \Phi_R(\hat{y}) + \mathbb{E}((R(x_f) - \Phi_R(y(x_f)))|y) = \mathbb{E}(R(x_f)|\hat{y}) = \bar{R}_a^{\hat{y}}, \quad (43)$$

better would be posterior expectation—after the observation \hat{y} .

Proof. Observe that the conditional expectation of the second term \tilde{R}_a in Eq. (42) vanishes:

$$\begin{aligned} \mathbb{E}(\tilde{R}_a|y) &= \mathbb{E}((R(x) - \Phi_R(y(x)))|y) = \\ &P_{\mathcal{R}_\infty}(R(x) - P_{\mathcal{R}_\infty}(R(x))) = P_{\mathcal{R}_\infty}(R(x)) - P_{\mathcal{R}_\infty}(R(x)) = 0. \end{aligned}$$

This means that the conditional expectation of the second term in Eq. (43) is nought, whereas the remaining term $\Phi_R(\hat{y})$ is just $\mathbb{E}(R(x_f)|\hat{y})$. \square

From Eq. (42) one can now construct filters. As often the optimal map Φ_R is often not easy to compute, one may even want to replace it by an approximation, say g_R , so that the update equation is

$$\tilde{R}_a = R(x_f) + (g_R(\hat{y}) - g_R(y(x_f))). \quad (44)$$

Whichever way, either the Eq. (42) or Eq. (44), they are composed of the following elements, the prior knowledge, which gives the prediction or forecast $R_f = R(x_f)$ for the RV R , and the correction, innovation, or update

$$R_\infty = (\Phi_R(\hat{y}) - \Phi_R(y(x_f))) \approx (g_R(\hat{y}) - g_R(y(x_f))),$$

which is the update difference between the actual observation \hat{y} and the predicted or forecast observation $y(x_f)$.

3.2.1 Getting the mean right

The simplest function $R(x)$ to think of is the identity $R(x) := x$. This gives an update—a filter—for the RV x itself. The optimal map will be denoted by Φ_x in this case. From Eq. (42) one has:

$$x_a = x_f + (\Phi_x(\hat{y}) - \Phi_x(y(x_f))) = x_f + x_\infty, \quad (45)$$

and Proposition 12 ensures that the assimilated RV x_a has the correct conditional mean

$$\mathbb{E}(x_a|y) = \mathbb{E}(x_f|\hat{y}) = \Phi_x(\hat{y}) =: \bar{x}^{\hat{y}}. \quad (46)$$

The Eq. (45) is the basis for many filtering algorithms, and many variations on the basic prescription are possible. Often they will be such that the property according to Proposition 12, the correct conditional mean, is only approximately satisfied. This is due to the fact that for one the Eq. (45) is an equation for RVs, which in their entirety can not be easily handled, they are typically infinite dimensional objects and thus have to be discretised for numerical purposes.

It was also already pointed out that the optimal map Φ_x is not easy to compute, and thus approximations are used, $\Phi_x \approx g_x$, the simplest one being where $g_x = G_x \in \mathcal{L}(\mathcal{Y}, \mathcal{X})$ is taken as a linear map, leading to linear filters [3], [11]. The well-known Kalman filter (KF) [17] and its many variants and extensions — e.g. extended KF, Gauss-Markov-Kalman filter, square root KF, etc. — and simplifications — e.g. 3DVar, 4DVar, Kriging, Gaussian process emulation (GPE) — arise in this way (e.g. [30], [41], [6], [37], [1], [34], [35], [28], [29]).

As the conditional expectation of x_a in Eq. (45) is Eq. (46) $\mathbb{E}(x_a|\hat{y}) = \Phi_x(\hat{y}) = \bar{x}^{\hat{y}}$, the zero-mean part of x_a is $\tilde{x}_a = x_f - \Phi_x(y(x_f))$. The posterior variance of the RV x_a is thus

$$C_{x_a x_a|\hat{y}} = \mathbb{E}(\tilde{x}_a \otimes \tilde{x}_a|\hat{y}) = \mathbb{E}((x_f - \Phi_x(y(x_f))) \otimes (x_f - \Phi_x(y(x_f))))|\hat{y}), \quad (47)$$

and it has been noted many times that this does not depend on the observation \hat{y} . Still, one may note (e.g. [41])

Proposition 13. *Assume that x_f is a Gaussian RV, that the observation $y(x_f) = \hat{h}(x_f, v)$ — absorbing the scaling ε into v — is affine in x_f and in the uncorrelated Gaussian observational noise v , i.e. $v \perp x_f$ and $C_{vx_f} = 0$. Then the optimal map $\Phi_x = K_x \in \mathcal{L}(\mathcal{Y}, \mathcal{X})$ is linear, and the updated or assimilated RV x_a from Eq. (45) is also Gaussian, and has the correct posterior distribution, characterised by the mean Eq. (46), $\bar{x}^{\hat{y}}$, and the covariance Eq. (47). Setting $w = \hat{h}(x_f, v) := H(x_f) + v$ with $H \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$, one*

obtains from Eq. (47)

$$C_{x_a x_a | \hat{y}} = \mathbb{E}(\tilde{x}_a \otimes \tilde{x}_a | \hat{y}) = C_{x_f x_f} - K_x C_{w x_f} - C_{w x_f}^T K_x^T + K_x C_{w w} K_x^T = \\ (I - K_x H) C_{x_f x_f} (I - K_x H)^T + K_x C_{v v} K_x^T \quad (48)$$

for the covariance, and for the mean

$$\bar{x} | \hat{y} = \mathbb{E}(x_a | \hat{y}) = K_x \hat{y}. \quad (49)$$

Proof. As this is a well known result, we only show the connection of Eq. (48) with Eq. (47). Note that

$$\tilde{x}_a = x_f - \Phi_x(y(x_f)) = x_f - K_x(\hat{h}(x_f, v)) \\ = x_f - K_x(w) = (I - K_x H)x_f - K_x v.$$

This gives the Eq. (48), and the Eq. (49) follows directly from Eq. (46). \square

This means that in the purely linear Gaussian case described in Proposition 13 a RV with the correct posterior distribution is given simply by the process of projection.

In the context of the dynamical system Eq. (4) resp. Eq. (5), where the measurement is denoted by y_{n+1} , the update for the tracking equation is

$$z_{n+1} = \hat{f}(z_n, 0, n) + (\Phi_x(y_{n+1}) - \Phi_x(\hat{h}(\hat{f}(z_n, 0, n), 0))) \quad (50)$$

for the case Eq. (4), resp. for Eq. (5)

$$z_{n+1} = f(z_n) + (\Phi_x(y_{n+1}) - \Phi_x(h(f(z_n)))). \quad (51)$$

Again the assimilated or updated state estimate $x_a := z_{n+1}$ is composed of two components, the prediction or forecast $x_f := \hat{f}(z_n, 0, n)$ resp. $x_f := f(z_n)$, and the correction, innovation, or update

$$x_\infty := (\Phi_x(y_{n+1}) - \Phi_x(\hat{h}(\hat{f}(z_n, 0, n), 0)))$$

resp. $x_\infty := (\Phi_x(y_{n+1}) - \Phi_x(h(f(z_n))))$, which takes into account the difference resulting from the actual measurement $\hat{y} := y_{n+1}$ and the forecast measurement $\hat{h}(\hat{f}(z_n, 0, n), 0)$ resp. $h(f(z_n))$.

If the optimal map Φ_x is not easy to compute, one may want to replace it by an approximation as in Eq. (44), say g , so that for example the Eq. (51) would read

$$z_{n+1} = f(z_n) + (g(y_{n+1}) - g(h(f(z_n)))) = (f - g \circ h \circ f)(z_n) + g(y_{n+1}), \quad (52)$$

where one may hope to show that if the map $(f - g \circ h \circ f)$ is a contraction, the difference $x_n - z_n$ will decrease as $n \rightarrow \infty$ [19]. Many variations on this theme exist, especially in the case where both the observation map h and the update operator g are linear [38], [42].

3.2.2 Getting also the covariance right

An approximation to a RV which has the required posterior distribution was constructed in section 3.2.1, where at least the mean was correct. One may now go a step further and also get the correct posterior covariance. As a starting point take the *assimilated* RV x_a from Eq. (45) that has the correct conditional mean $\bar{x}^{|\hat{y}}$ from Eq. (46), but the covariance, from Eq. (47), is $C_{x_a x_a|\hat{y}} = \mathbb{E}(\tilde{x}_a \otimes \tilde{x}_a|\hat{y})$. To get the covariance and the mean right, we compute what the correct posterior covariance should be, by computing the optimal map for $R(x) := x \otimes x$. This gives for the posterior correlation

$$\hat{C}_p := \mathbb{E}(R(x_f)|\hat{y}) = \mathbb{E}((x_f \otimes x_f)|\hat{y}) = \Phi_{x \otimes x}(\hat{y}), \quad (53)$$

so that the posterior covariance is

$$C_p := \hat{C}_p - \bar{x}^{|\hat{y}} \otimes \bar{x}^{|\hat{y}} = \Phi_{x \otimes x}(\hat{y}) - \Phi_x(\hat{y}) \otimes \Phi_x(\hat{y}). \quad (54)$$

Proposition 14. *A new RV x_c with the correct posterior covariance Eq. (54) is built from $x_a = \bar{x}^{|\hat{y}} + \tilde{x}_a$ in Eq. (45) by taking*

$$x_c := \bar{x}^{|\hat{y}} + C_p^{1/2} C_{x_a x_a|\hat{y}}^{-1/2} \tilde{x}_a. \quad (55)$$

Proof. As $\mathbb{E}(x_c|\hat{y}) = \bar{x}^{|\hat{y}}$, one has

$$\begin{aligned} C_{x_c x_c} &= \mathbb{E}\left((C_p^{1/2} C_{x_a x_a|\hat{y}}^{-1/2} \tilde{x}_a) \otimes (C_p^{1/2} C_{x_a x_a|\hat{y}}^{-1/2} \tilde{x}_a)|\hat{y}\right) = \\ &C_p^{1/2} C_{x_a x_a|\hat{y}}^{-1/2} \mathbb{E}(\tilde{x}_a \otimes \tilde{x}_a|\hat{y}) C_{x_a x_a|\hat{y}}^{-1/2} C_p^{1/2} = \\ &C_p^{1/2} C_{x_a x_a|\hat{y}}^{-1/2} C_{x_a x_a|\hat{y}} C_{x_a x_a|\hat{y}}^{-1/2} C_p^{1/2} = C_p^{1/2} C_p^{1/2} = C_p, \end{aligned}$$

proving that the RV x_c in Eq. (55) has the correct posterior covariance. \square

Having achieved a RV which has the correct posterior mean and covariance, it is conceivable to continue in this fashion, building RVs which match the posterior better and better. A similar idea, but from a different starting point, is used in [27] and [31]. In the future, it is planned to combine these two approaches.

3.3 Approximations

In any actual inverse computations several kinds of approximations are usually necessary. Should one pursue an approach of sampling from the posterior distribution Eq. (11) in Subsection 2.2 for example, then typically a sampling and a quantisation or binning approximation is performed, often

together with some kernel-estimate of the density. All of these processes usually introduce approximation errors. Here we want to use methods based on the conditional expectation, which were detailed in Subsection 2.3.

Looking for example at Theorem 10 and Theorem 11 in section 2.3.2, one has to work in the usually infinite dimensional space $\mathcal{R} = \mathcal{R} \otimes \mathcal{S}$ from Eq. (22) and its subspace $\mathcal{R}_\infty = \mathcal{R} \otimes \mathcal{S}_\infty$ from Eq. (23), to minimise the functional in Eq. (27) to find the optimal map representing the conditional expectation for a desired function $R(x)$, Eq. (26) and Eq. (29), Φ in the space $L_0(\mathcal{Y}, \mathcal{R})$. Then one has to construct a RV whose distribution may be characterised by the conditional expectation, to represent the posterior measure. Approximations in this latter respect were discussed in Subsection 3.2. The space \mathcal{R}_∞ is computationally accessible via $L_0(\mathcal{Y}, \mathcal{R})$, which has to be approximated by some finite dimensional subspace. This will be discussed in this Subsection 3.3. Furthermore, the component spaces of $\mathcal{R} = \mathcal{R} \otimes \mathcal{S}$ are also typically infinite dimensional, and have in actual computations to be replaced by finite dimensional subspaces. This topic will be sketched in Section 4.

Computationally we will not be able to deal with the *whole* space \mathcal{R}_∞ , so we look at the effect of approximations. Assume that $L_0(\mathcal{Y}, \mathcal{R})$ in Eq. (29) or Eq. (30) is approximated by subspaces $L_{0,m} \subset L_0(\mathcal{Y}, \mathcal{R})$ with $\mathcal{L}(\mathcal{Y}, \mathcal{R}) \subseteq L_{0,m}$, where $m \in \mathbb{N}$ is a parameter describing the level of approximation and $L_{0,m} \subset L_{0,k}$ if $m < k$, such that the subspaces

$$\mathcal{R}_m = \{\varphi(y) \mid \varphi \in L_{0,m}; \varphi(\hat{h}(x, \varepsilon v)) \in \mathcal{R}\} \subseteq \mathcal{R}_\infty \quad (56)$$

are linearly closed and their union is dense

$$\overline{\bigcup_m \mathcal{R}_m} = \mathcal{R}_\infty, \quad (57)$$

a consistency condition.

To obtain results for the situation where the projection $P_{\mathcal{R}_\infty}$ is replaced by the orthogonal projection $P_{\mathcal{R}_m}$ of \mathcal{R} onto \mathcal{R}_m , all that is necessary is to reformulate the Theorem 10 and Theorem 11.

Theorem 15. *The orthogonal projection $P_{\mathcal{R}_m}$ of the RV $R(x) \in \mathcal{R}$ is characterised by:*

$$P_{\mathcal{R}_m}(R(x)) := \arg \min_{\tilde{R} \in \mathcal{R}_m} \frac{1}{2} \|R(x) - \tilde{R}\|_{\mathcal{R}}^2, \quad (58)$$

The Eq. (56) implies the existence of an optimal map $\Phi_m \in L_{0,m}(\mathcal{Y}, \mathcal{R})$ such that

$$P_{\mathcal{R}_m}(R(x)) = \Phi_m(\hat{h}(x, \varepsilon v)). \quad (59)$$

Taking the vanishing of the first variation / Gâteaux derivative of the loss-function as a necessary condition for a minimum leads to a simple geometrical interpretation: the difference between the original vector-valued RV $R(x)$ and its projection has to be perpendicular to the subspace \mathcal{R}_m : $\forall \tilde{R} \in \mathcal{R}_m$:

$$\langle R(x) - P_{\mathcal{R}_m}(R(x)), \tilde{R} \rangle_{\mathcal{R}} = 0, \quad \text{i.e. } R(x) - P_{\mathcal{R}_m}(R(x)) \in \mathcal{R}_m^\perp. \quad (60)$$

Rephrasing Eq. (58) with account to Eq. (60) leads for the optimal map $\Phi_m \in L_{0,n}(\mathcal{Y}, \mathcal{R})$ to

$$P_{\mathcal{R}_m}(R(x)) = \Phi_m(\hat{h}(x, \varepsilon v)) := \arg \min_{\varphi \in L_{0,m}(\mathcal{Y}, \mathcal{R})} \|R(x) - \varphi(\hat{h}(x, \varepsilon v))\|_{\mathcal{R}}^2, \quad (61)$$

and the orthogonality condition of Eq. (60) leads to

$$\forall \varphi \in L_{0,m}(\mathcal{Y}, \mathcal{R}) : \langle R(x) - \Phi_m(\hat{h}(x, \varepsilon v)), \varphi(\hat{h}(x, \varepsilon v)) \rangle_{\mathcal{R}} = 0. \quad (62)$$

In addition, as \mathcal{R}_m is linearly closed, one obtains the useful statement

$$\forall \tilde{R} \in \mathcal{R}_m : \mathcal{L}(\mathcal{R}) \ni \mathbb{E} \left((R(x) - P_{\mathcal{R}_m}(R(x))) \otimes \tilde{R} \right) = 0. \quad (63)$$

or rephrased $\forall \varphi \in L_{0,m}(\mathcal{Y}, \mathcal{R})$:

$$\mathcal{L}(\mathcal{R}) \ni \mathbb{E} \left((R(x) - \Phi_m(\hat{h}(x, \varepsilon v))) \otimes \varphi(\hat{h}(x, \varepsilon v)) \right) = 0. \quad (64)$$

There is a unique minimiser $P_{\mathcal{R}_m}(R(x)) \in \mathcal{R}_m$ to the problem in Eq. (58), which satisfies the orthogonality condition Eq. (60), which is equivalent to the the strong orthogonality condition Eq. (63). The minimiser is unique as an element of \mathcal{R}_m , but the mapping $\Phi_m \in L_{0,m}(\mathcal{Y}, \mathcal{R})$ in Eq. (61) may not necessarily be. It also holds that

$$\|P_{\mathcal{R}_m}(R(x))\|_{\mathcal{R}}^2 = \|R(x)\|_{\mathcal{R}}^2 - \|R(x) - P_{\mathcal{R}_m}(R(x))\|_{\mathcal{R}}^2. \quad (65)$$

Additionally, one has

$$\|P_{\mathcal{R}_m}(R(x))\|_{\mathcal{R}}^2 \leq \|P_{\mathcal{R}_\infty}(R(x))\|_{\mathcal{R}}^2. \quad (66)$$

Proof. It is all already contained in Theorem 10 and Theorem 11 when applied to \mathcal{R}_m . The stability condition Eq. (66) is due to the simple fact that $\mathcal{R}_m \subseteq \mathcal{R}_\infty$. \square

From the consistency condition, the stability Eq. (66) as shown in Theorem 15, and Céa's lemma, one immediately obtains:

Theorem 16. For all RVs $R(x) \in \mathcal{R}$, the sequence $R_m := P_{\mathcal{R}_m}(R(x))$ converges to $R_\infty := P_{\mathcal{R}_\infty}(R(x))$:

$$\lim_{m \rightarrow \infty} \|R_\infty - R_m\|_{\mathcal{R}}^2 = 0. \quad (67)$$

Proof. Well-posedness is a direct consequence of Theorem 11. As the $P_{\mathcal{R}_m}$ are orthogonal projections onto the subspaces \mathcal{R}_m , their norms are hence all equal to unity — a stability condition, as shown in Eq. (66). Application of Céa’s lemma then directly yields Eq. (67). \square

3.3.1 Approximation by polynomials

Here we choose the subspaces of polynomials up to degree m for the purpose of approximation, i.e.

$$\mathcal{R}_m := \overline{\text{span}} \{ \varphi \mid \varphi(\hat{h}(x, \varepsilon v)) \in \mathcal{R}, \varphi \in \mathcal{P}_m(\mathcal{Y}, \mathcal{X}) \},$$

where $\mathcal{P}_m(\mathcal{Y}, \mathcal{X}) \subset L_0(\mathcal{Y}, \mathcal{X})$ are the polynomials of degree at most m on \mathcal{Y} with values in \mathcal{X} . We may write $\psi_m \in \mathcal{P}_m(\mathcal{Y}, \mathcal{X})$ as

$$\psi_m(y) := {}^0H + {}^1Hy + \cdots + {}^kHy^{\vee k} + \cdots + {}^mHy^{\vee m}, \quad (68)$$

where ${}^kH \in \mathcal{L}_s^k(\mathcal{Y}, \mathcal{R})$ is symmetric and k -linear; and $y^{\vee k} := \overbrace{y \vee \cdots \vee y}^k := \text{Sym}(y^{\otimes k})$ is the symmetric tensor product of the y ’s taken k times with itself. Let us remark here that the form of Eq. (68), given in monomials, is numerically not a good form—except for very low m —and straightforward use in computations is not recommended. The relation Eq. (68) could be re-written in some orthogonal polynomials—or in fact any other system of multi-variate functions; this generalisation will be considered in section 3.3.3. For the sake of conceptual simplicity, we stay with Eq. (68) and then have that for any RV $R(x) \in \mathcal{R}$

$$\Phi_{R,m}(R(x)) := \psi_{R,m}(y) := {}^0H + \cdots + \cdots + {}^mHy^{\vee m} =: \Psi_{R,m}({}^0H, \dots, {}^mH) \quad (69)$$

the optimal map in Eq. (59) from Theorem 15 — where we have added an index R to indicate that it depends on the RV $R(x)$, but for simplicity omitted this index on the coefficient maps kH — is a function $\Psi_{R,m}$ of the coefficient maps kH . The stationarity or orthogonality condition Eq. (64) can then be written in terms of the kH . We need the following abbreviations for any $k, \ell \in \mathbb{N}_0$ and $p \in \mathcal{R}, v \in \mathcal{V}$:

$$\langle p \otimes v^{\vee k} \rangle := \mathbb{E} \left(p \otimes v^{\vee k} \right) = \int_{\Omega} p(\omega) \otimes v(\omega)^{\vee k} \mathbb{P}(\mathrm{d}\omega)$$

and

$${}^k H \langle y^{\vee(\ell+k)} \rangle := \langle y^{\vee \ell} \vee ({}^k H y^{\vee k}) \rangle = \mathbb{E} \left(y^{\vee \ell} \vee ({}^k H y^{\vee k}) \right).$$

We may then characterise the kH in the following way:

Theorem 17. *With $\Psi_{R,m}$ from Eq. (69), the stationarity condition Eq. (64) becomes, by the chain rule, for any $m \in \mathbb{N}_0$*

$$\forall \ell = 0, \dots, m : \quad \sum_{k=0}^m {}^k H \langle y^{\vee(\ell+k)} \rangle = \langle R(x) \otimes y^{\vee \ell} \rangle. \quad (70)$$

The Hankel operator matrix $(\langle y^{\vee(\ell+k)} \rangle)_{\ell,k}$ in the linear equations Eq. (70) is symmetric and positive semi-definite, hence the system Eq. (70) has a solution, unique in case the operator matrix is actually definite.

Proof. The relation Eq. (70) is the result of straightforward application of the chain rule to the Eq. (64).

The symmetry of the operator matrix is obvious—the $\langle y^{\vee k} \rangle$ are the coefficients—and positive semi-definiteness follows easily from the fact that it is the gradient of the functional in Eq. (61), which is convex. \square

Observe that the operator matrix is independent of the RV $R(x)$ for which the ncomputation is performed. Only the right hand side is influenced by $R(x)$.

The system of operator equations Eq. (70) may be written in more detailed form as:

$$\begin{array}{lll} \ell = 0 : & {}^0H & \cdots + {}^k H \langle y^{\vee k} \rangle \quad \cdots + {}^m H \langle y^{\vee m} \rangle = \langle R(x) \rangle, \\ \ell = 1 : & {}^0H \langle y \rangle & \cdots + {}^k H \langle y^{\vee(1+k)} \rangle \quad \cdots + {}^m H \langle y^{\vee(1+m)} \rangle = \langle R(x) \otimes y \rangle, \\ \vdots & & \vdots \\ \ell = m : & {}^0H \langle y^{\vee m} \rangle \cdots + {}^k H \langle y^{\vee(k+m)} \rangle \cdots + {}^m H \langle y^{\vee 2m} \rangle = \langle R(x) \otimes y^{\vee m} \rangle. \end{array}$$

Using ‘*symbolic index*’ notation a la Penrose — the reader may just think of indices in a finite dimensional space with orthonormal basis — the system Eq. (70) can be given yet another form: denote in symbolic index notation $R(x) = (R^i)$, $y = (y^j)$, and ${}^kH = ({}^kH_{j_1 \dots j_k})$, then Eq. (70) becomes, with the use of the Einstein convention of summation (a tensor contraction) over repeated indices, and with the symmetry explicitly indicated:

$$\begin{aligned} \forall \ell = 0, \dots, m; j_1 \leq \dots \leq j_\ell \leq \dots \leq j_{\ell+k} \leq \dots \leq j_{\ell+m} : \\ \langle y^{j_1} \dots y^{j_\ell} \rangle ({}^0 H^i) + \dots + \langle y^{j_1} \dots y^{j_{\ell+1}} \dots y^{j_{\ell+k}} \rangle ({}^k H^i_{j_{\ell+1} \dots j_{\ell+k}}) + \\ \dots + \langle y^{j_1} \dots y^{j_{\ell+1}} \dots y^{j_{\ell+m}} \rangle ({}^m H^i_{j_{\ell+1} \dots j_{\ell+m}}) = \langle R^i y^{j_1} \dots y^{j_\ell} \rangle. \quad (71) \end{aligned}$$

We see in this representation that the matrix does *not* depend on ι —it is identically *block diagonal* after appropriate reordering, which makes the solution of Eq. (70) or Eq. (71) much easier.

Some special cases are: for $m = 0$ —*constant* functions. One does not use any information from the measurement — and from Eq. (70) or Eq. (71) one has

$$\Phi_{R,0}(R(x)) = \psi_{R,0}(y) = {}^0H = \langle R \rangle = \mathbb{E}(R) = \bar{R}.$$

Without any information, the conditional expectation is equal to the unconditional expectation. The update corresponding to Eq. (42) — actually Eq. (44) as we are approximating the map Φ_R by $g_R = \Phi_{R,0}$ — then becomes $R_a \approx R_{a,0} = R(x_f) = R_f$, as $R_\infty = 0$ in this case; the assimilated quantity stays equal to the forecast. This was to be expected, is not of much practical use, but is a consistency check.

The case $m = 1$ in Eq. (70) or Eq. (71) is more interesting, allowing up to *linear* terms:

$$\begin{aligned} {}^0H + {}^1H \langle y \rangle &= \langle R(x) \rangle = \bar{R} \\ {}^0H \langle y \rangle + {}^1H \langle y \vee y \rangle &= \langle R(x) \otimes y \rangle. \end{aligned}$$

Remembering that $C_{Ry} = \langle R(x) \otimes y \rangle - \langle R \rangle \otimes \langle y \rangle$ and analogous for C_{yy} , one obtains by tensor multiplication with $\langle R(x) \rangle$ and symbolic Gaussian elimination

$$\begin{aligned} {}^0H &= \langle R \rangle - {}^1H \langle y \rangle = \bar{R} - {}^1H \bar{y} \\ {}^1H (\langle y \vee y \rangle - \langle y \rangle \vee \langle y \rangle) &= {}^1H C_{yy} = \langle R(x) \otimes y \rangle - \langle R \rangle \otimes \langle y \rangle = C_{Ry}. \end{aligned}$$

This gives

$${}^1H = C_{Ry} C_{yy}^{-1} =: K \quad (72)$$

$${}^0H = \bar{R} - K \bar{y}. \quad (73)$$

where K in Eq. (72) is the well-known *Kalman* gain operator [17], so that finally

$$\Phi_{R,1}(R(x)) = \psi_{R,1}(y) = {}^0H + {}^1H y = \bar{R} + C_{Ry} C_{yy}^{-1} (y - \bar{y}) = \bar{R} + K(y - \bar{y}). \quad (74)$$

The update corresponding to Eq. (42) — again actually Eq. (44) as we are approximating the map Φ_R by $g_R = \Phi_{R,1}$ — then becomes

$$\begin{aligned} R_a \approx R_{a,1} &= R(x_f) + \left((\bar{R} + K(\hat{y} - \bar{y})) - (\bar{R} + K(y(x_f) - \bar{y})) \right) = \\ &= R_f + K(\hat{y} - y(x_f)) = R_f + R_{\infty,1}. \end{aligned} \quad (75)$$

This may be called a *linear* Bayesian update (LBU), and is similar to the ‘Bayes linear’ approach [11]. It is important to see Eq. (75) as a symbolic expression, especially the inverse C_{yy}^{-1} indicated in Eq. (74) should not really be computed, especially when C_{yy} is ill-conditioned or close to singular. The inverse can in that case be replaced by the *pseudo-inverse*, or rather the computation of K , which is in linear algebra terms a *least-squares* approximation, should be done with orthogonal transformations and not by elimination. We will not dwell on these well-known matters here. It is also obvious that the constant term in Eq. (74) — or even Eq. (69) for that matter — is of no consequence for the update filter, as it cancels out.

The case $m = 2$ can still be solved symbolically, the system to be solved is from Eq. (70) or Eq. (71):

$$\begin{aligned} {}^0H &+ {}^1H \langle y \rangle + {}^2H \langle y^{\vee 2} \rangle = \langle R \rangle \\ {}^0H \langle y \rangle &+ {}^1H \langle y^{\vee 2} \rangle + {}^2H \langle y^{\vee 3} \rangle = \langle R \otimes y \rangle \\ {}^0H \langle y^{\vee 2} \rangle &+ {}^1H \langle y^{\vee 3} \rangle + {}^2H \langle y^{\vee 4} \rangle = \langle R \otimes y^{\vee 2} \rangle. \end{aligned}$$

After some symbolic elimination steps one obtains

$$\begin{aligned} {}^0H + {}^1H \langle y \rangle + {}^2H \langle y^{\vee 2} \rangle &= \bar{R} \\ 0 + {}^1H + {}^2H \mathbf{F} &= K \\ 0 + 0 + {}^2H \mathbf{G} &= E, \end{aligned}$$

with the Kalman gain operator $K \in (\mathcal{R} \otimes \mathcal{Y})^*$ from Eq. (72), the third order tensors $\mathbf{F} \in (\mathcal{Y}^{\otimes 3})^*$ given in Eq. (76), and $E \in (\mathcal{R} \otimes \mathcal{Y}^{\otimes 2})^*$ given in Eq. (77), and the fourth order tensor $\mathbf{G} \in (\mathcal{Y}^{\otimes 4})^*$ given in Eq. (78):

$$\mathbf{F} = \left(\langle y^{\vee 3} \rangle - \langle y^{\vee 2} \rangle \vee \langle y \rangle \right) C_{yy}^{-1}, \quad (76)$$

$$E = \langle R \otimes y^{\vee 2} \rangle - \bar{R} \otimes \langle y^{\vee 2} \rangle - K \left(\langle y^{\vee 3} \rangle - \langle y \rangle \vee \langle y^{\vee 2} \rangle \right) \quad (77)$$

$$\mathbf{G} = \left(\langle y^{\vee 4} \rangle - \langle y^{\vee 2} \rangle^{\vee 2} \right) - \mathbf{F} \cdot \left(\langle y^{\vee 3} \rangle - \langle y \rangle \vee \langle y^{\vee 2} \rangle \right), \quad (78)$$

where the single central dot ‘ \cdot ’ denotes as usual a contraction over the appropriate indices, and a colon ‘ $:$ ’ a double contraction. From this one easily obtains the solution

$${}^2H = E : \mathbf{G}^{-1} \quad (79)$$

$${}^1H = K - {}^2H \mathbf{F} \quad (80)$$

$${}^0H = \bar{R} - (K - {}^1H) \bar{y} - {}^2H \langle y^{\vee 2} \rangle = \bar{R} - {}^2H (\mathbf{F} \cdot \bar{y} + \langle y^{\vee 2} \rangle). \quad (81)$$

The update corresponding to Eq. (42) — again actually Eq. (44) as we are approximating the map Φ_R now by $g_R = \Phi_{R,2}$ — then becomes

$$\begin{aligned} R_a \approx R_{a,2} &= R(x_f) + \left(({}^2H\hat{y}^{\vee 2} + {}^1H\hat{y}) - ({}^2Hy(x_f)^{\vee 2} + {}^1Hy(x_f)) \right) = \\ &= R_f + \left(E : \mathbf{G}^{-1} : (\hat{y}^{\vee 2} - y(x_f)^{\vee 2}) + (K - E : \mathbf{G}^{-1} : \mathbf{F})(\hat{y} - y(x_f)) \right) \\ &= R_f + R_{\infty,2}. \end{aligned} \quad (82)$$

This may be called a *quadratic* Bayesian update (QBU), and it is clearly an extension of Eq. (75).

3.3.2 The Gauss-Markov-Kalman filter

The $m = 1$ version of Theorem 17 is well-known for the special case $R(x) := x$, and we rephrase this generalisation of the well-known *Gauss-Markov* theorem from [20] Chapter 4.6, Theorem 3:

Proposition 18. *The update $x_{a,1}$, minimising $\|x_f - \cdot\|_{\mathcal{X}}^2$ over all elements generated by affine mappings (the up to $m = 1$ case of Theorem 17) of the measurement \hat{y} with predicted measurement $y(x_f)$ is given*

$$x_{a,1} = x_f + K(\hat{y} - y(x_f)), \quad (83)$$

where the operator K is the Kalman gain from Eq. (72) and Eq. (75).

The Eq. (83) is reminiscent — actually an extension — not only of the well-known *Gauss-Markov* theorem [20], but also of the *Kalman* filter [17] [30], so that we propose to call Eq. (83) the **Gauss-Markov-Kalman** (GMK) filter (GMKF).

We point out that $x_{a,1}$, x_f , and $y(x_f)$ are RVs, i.e. Eq. (83) is an equation in $\mathcal{X} = \mathcal{X} \otimes \mathcal{S}$ between RVs, whereas the traditional Kalman filter is an equation in \mathcal{X} . If the mean is taken in Eq. (83), one obtains the familiar Kalman filter formula [17] for the update of the mean, and one may show [28] that Eq. (83) also contains the Kalman update for the covariance by computing Eq. (47) for this case, which gives the familiar result of Kalman, i.e. the Kalman filter is a low-order part of Eq. (83).

The computational strategy for a typical filter is now to replace and approximate the—only abstractly given—computation of x_a Eq. (45) by the practically possible calculation of $x_{a,m}$ as in Eq. (69). This means that we approximate x_a by $x_{a,m}$ by using $\mathcal{X}_m \subseteq \mathcal{X}_\infty$, and rely on Theorem 16. This corresponds to some loss of information from the measurement as one uses a smaller subspace for the projection, but yields a manageable computation. If the assumptions of Theorem 16 are satisfied, then one can expect for m large enough that the terms in Eq. (69) converge to zero, thus providing an error indicator on when a sufficient accuracy has been reached.

3.3.3 Approximation by general functions

The derivation in section 3.3.1 was for the special case where polynomials are used to find a subspace $L_{0,m}(\mathcal{Y}, \mathcal{X})$ for the approximation. It had the advantage of showing the connection to the ‘Bayes linear’ approach [11], to the Gauss-Markov theorem [20], and to the *Kalman* filter [17] [30], giving in Eq. (83) of Proposition 18 the *Gauss-Markov-Kalman* filter (GMKF).

But for a more general approach not limited to polynomials, we proceed similarly as in Eq. (56), but now concretely assume a set of linearly independent functions, not necessarily orthonormal,

$$\mathcal{B} := \{\psi_\alpha \mid \alpha \in \mathcal{A}, \psi_\alpha \in L_0(\mathcal{Y}); \psi_\alpha(\hat{h}(x, \varepsilon v)) \in \mathcal{S}\} \subseteq \mathcal{S}_\infty \quad (84)$$

where \mathcal{A} is some countable index set. Assume now that

$$\mathcal{S}_\infty = \overline{\text{span}} \mathcal{B},$$

i.e. \mathcal{B} is a Hilbert basis of \mathcal{S}_∞ , again a consistency condition.

Denote by \mathcal{A}_k a finite part of \mathcal{A} of cardinality k , such that $\mathcal{A}_k \subset \mathcal{A}_\ell$ for $k < \ell$ and $\bigcup_k \mathcal{A}_k = \mathcal{A}$, and set

$$\mathcal{R}_k := \mathcal{R} \otimes \mathcal{S}_k \subseteq \mathcal{R}_\infty, \quad (85)$$

where the finite dimensional and hence closed subspaces \mathcal{S}_k are given by

$$\mathcal{S}_k := \text{span}\{\psi_\alpha \mid \alpha \in \mathcal{A}_k, \psi_\alpha \in \mathcal{B}\} \subseteq \mathcal{S}. \quad (86)$$

Observe that the spaces \mathcal{R}_k from Eq. (85) are linearly closed according to Proposition 4.

Theorem 15 and Theorem 16 apply in this case. For a RV $R(x) \in \mathcal{R}$ we make the following ‘ansatz’ for the optimal map $\Phi_{R,k}$ such that $P_{\mathcal{R}_k}(R(x)) = \Phi_{R,k}(\hat{h}(x, \varepsilon v))$:

$$\Phi_{R,k}(y) = \sum_{\alpha \in \mathcal{A}_k} v_\alpha \psi_\alpha(y), \quad (87)$$

with as yet unknown coefficients $v_\alpha \in \mathcal{R}$. This is a normal *Galerkin*-ansatz, and Eq. (64) from Theorem 15 can be used to determine these coefficients.

Take $\mathcal{Z}_k := \mathbb{R}^{\mathcal{A}_k}$ with canonical basis $\{e_\alpha \mid \alpha \in \mathcal{A}_k\}$, and let

$$\mathbf{G}_k := (\langle \psi_\alpha(y(x)), \psi_\beta(y(x)) \rangle_{\mathcal{S}})_{\alpha, \beta \in \mathcal{A}_k} \in \mathcal{L}(\mathcal{Z}_k)$$

be the symmetric positive definite Gram matrix of the basis of \mathcal{S}_k ; also set

$$\begin{aligned} \mathbf{v} &:= \sum_{\alpha \in \mathcal{A}_k} e_\alpha \otimes v_\alpha \in \mathcal{Z}_k \otimes \mathcal{R}, \\ \mathbf{r} &:= \sum_{\alpha \in \mathcal{A}_k} e_\alpha \otimes \mathbb{E}(\psi_\alpha(y(x))R(x)) \in \mathcal{Z}_k \otimes \mathcal{R}. \end{aligned}$$

Theorem 19. For any $k \in \mathbb{N}$, the coefficients $\{v_\alpha\}_{\alpha \in \mathcal{A}_k}$ of the optimal map $\Phi_{R,k}$ in Eq. (87) are given by the unique solution of the Galerkin equation

$$(\mathbf{G}_k \otimes I_{\mathcal{R}})\mathbf{v} = \mathbf{r}. \quad (88)$$

It has the formal solution

$$\mathbf{v} = (\mathbf{G}_k \otimes I_{\mathcal{R}})^{-1} \mathbf{r} = (\mathbf{G}_k^{-1} \otimes I_{\mathcal{R}}) \mathbf{r} \in \mathcal{Z}_k \otimes \mathcal{R}.$$

Proof. The Galerkin Eq. (88) is a simple consequence of Eq. (64) from Theorem 15. As the Gram matrix \mathbf{G}_k and the identity $I_{\mathcal{R}}$ on \mathcal{R} are positive definite, so is the tensor operator $(\mathbf{G}_k \otimes I_{\mathcal{R}})$, with inverse $(\mathbf{G}_k^{-1} \otimes I_{\mathcal{R}})$. \square

As in Eq. (71), the block structure of the equations is clearly visible. Hence, to solve Eq. (88), one only has to deal with the ‘small’ matrix \mathbf{G}_n .

The update corresponding to Eq. (42) — again actually Eq. (44) as we are approximating the map Φ_R now by a new map $g_R = \Phi_{R,k}$ — then becomes

$$R_a \approx R_{a,k} = R(x_f) + (\Phi_{R,k}(\hat{y}) - \Phi_{R,k}(y(x_f))) = R_f + R_{\infty,k}. \quad (89)$$

This may be called a ‘general Bayesian update’. Applying Eq. (89) now again to the special case $R(x) := x$, one obtains a possibly nonlinear filter based on the basis \mathcal{B} :

$$x_a \approx x_{a,k} = x_f + (\Phi_{x,k}(\hat{y}) - \Phi_{x,k}(y(x_f))) = x_f + x_{\infty,k}. \quad (90)$$

In case the $\mathcal{Y}^* \subseteq \text{span}\{\psi_\alpha\}_{\alpha \in \mathcal{A}_k}$, i.e. the basis generates all the linear functions on \mathcal{Y} , this is a true extension of the Kalman filter.

4 Numerical realisation

In the instances where we want to employ the theory detailed in the previous Section 2 and Section 3, the spaces \mathcal{U} and \mathcal{Q} and hence \mathcal{X} are usually infinite dimensional, as is the space $\mathcal{S} = L_2(\Omega)$. For an actual computation they all have to be discretised or approximated by finite dimensional subspaces.

In our examples we will chose finite element discretisations for \mathcal{U} , \mathcal{Q} , and hence \mathcal{X} , and corresponding subspaces. Hence let $\mathcal{X}_M := \text{span}\{\varrho_m : m = 1, \dots, M\} \subset \mathcal{X}$ be an M -dimensional subspace with basis $\{\varrho_m\}_{m=1}^M$. An element of \mathcal{X}_M will be represented by the vector $\mathbf{x} = [x^1, \dots, x^M]^T \in \mathbb{R}^M$ such that $\sum_{m=1}^M x^m \varrho_m \in \mathcal{X}_M$. To avoid a profusion of notations, the corresponding random vector in $\mathbb{R}^M \otimes \mathcal{S}$ — a mapping $\Omega \rightarrow \mathbb{R}^M \cong \mathcal{X}_M$ — will also be denoted by \mathbf{x} , as the meaning will be clear from the context.

The norm $\|\mathbf{x}\|_M$ one has to take on \mathbb{R}^M results from the inner product $\langle \mathbf{x}_1 | \mathbf{x}_2 \rangle_M := \mathbf{x}_1^T \mathbf{Q} \mathbf{x}_2$ with $\mathbf{Q} = (\langle \varrho_m | \varrho_n \rangle_{\mathcal{X}})$, the Gram matrix of the basis. We will later choose an orthonormal basis, so that $\mathbf{Q} = \mathbf{I}$ is the identity matrix. Similarly, on $\mathcal{X}_M = \mathbb{R}^M \otimes \mathcal{S}$ the inner product is $\langle \langle \mathbf{x}_1 | \mathbf{x}_2 \rangle \rangle_{\mathcal{X}_M} := \mathbb{E}(\langle \mathbf{x}_1 | \mathbf{x}_2 \rangle_M)$.

The space of possible measurements \mathcal{Y} can usually be taken to be finite dimensional, otherwise we take similarly as before a R -dimensional subspace \mathcal{Y}_R , whose elements are similarly represented by a vector of coefficients $\mathbf{y} \in \mathbb{R}^R$. For the discretised version of the RV $y(x_f) = y(\hat{h}(x_f, \varepsilon v))$ we will often use the shorthand $\mathbf{y}_f := \mathbf{y}(x_f) = \mathbf{y}(\hat{h}(\mathbf{x}_f, \varepsilon \mathbf{v}))$.

As some of the most efficient ways of doing the update are linear filters based on the general idea of orthogonal decomposition — Eq. (42) in Subsection 3.2 — applied to the mean — Eq. (45) in section 3.2.1 — but in the modified form Eq. (44) where g is a linear map, and especially the optimal linear map of the Gauss-Markov-Kalman (GMK) filter Eq. (83), we start from Proposition 18 in section 3.3.2. For other approximations the finite dimensional discretisation would be largely analogous.

On \mathbb{R}^M , representing \mathcal{X}_M , the Kalman gain operator in Proposition 18 in Eq. (83) becomes a matrix $\mathbf{K} \in \mathbb{R}^{M \times R}$. Then the update corresponding to Eq. (83) is

$$\mathbf{x}_a = \mathbf{x}_f + \mathbf{K}(\hat{\mathbf{y}} - \mathbf{y}(x_f)), \text{ with } \mathbf{K} = \mathbf{C}_{xy} \mathbf{C}_{yy}^{-1}. \quad (91)$$

Here the covariances are $\mathbf{C}_{xy} := \mathbb{E}(\tilde{\mathbf{x}}_f \tilde{\mathbf{y}}(x_f))$, and similarly for \mathbf{C}_{yy} . Often the measurement error v in the measurement model $\tilde{h}(x_f, \varepsilon v) = h(x_f) + \varepsilon S_y(x_f)v$ is independent of \mathbf{x} — actually *uncorrelated* would be sufficient, i.e. $\mathbf{C}_{xv} = \mathbf{0}$ — hence, assuming that S_y does not depend on x , $\mathbf{C}_{xx} = \mathbf{C}_{hh} + \varepsilon^2 \mathbf{S}_y \mathbf{C}_{vv} \mathbf{S}_y^T$ and $\mathbf{C}_{xy} = \mathbf{C}_{xh}$, where $h = h(x_f)$.

It is important to emphasise that the theory presented in the forgoing Section 2 and Section 3 is independent of any discretisation of the underlying spaces. But one usually can still not numerically compute with objects like $\mathbf{x} \in \mathcal{X}_M = \mathbb{R}^M \otimes \mathcal{S}$, as $\mathcal{S} = L_2(\Omega)$ is normally an infinite dimensional space, and has to be discretised. One well-known possibility are samples, i.e. the RV $\mathbf{x}(\omega)$ is represented by its value at certain points ω_z , and the points usually come from some quadrature rule. The well-known Monte Carlo (MC) method uses random samples, the quasi-Monte Carlo (QMC) method uses low discrepancy samples, and other rules like sparse grids (Smolyak rule) are possible. Using MC samples in the context of the linear update Eq. (83) is known as the *Ensemble Kalman Filter* (EnKF), see [34] for a general overview in this context, and [6], [7] for a thorough description and analysis. This method is conceptually fairly simple and is currently a favourite for problems

where the computation of the predicted measurement $\mathbf{y}(\mathbf{x}_f(\omega_z))$ is difficult or expensive. It needs far fewer samples for meaningful results than MCMC, but on the other hand it uses the linear approximation inherent in Eq. (91).

Here we want to use so-called *functional* or *spectral* approximations, so similarly as for \mathcal{X}_M , we pick a finite set of linearly independent vectors in \mathcal{S} . As $\mathcal{S} = L_2(\Omega)$, these abstract vectors are in fact RVs with finite variance. Here we will use the best known example, namely *Wiener's polynomial chaos* expansion (PCE) as basis [43], [10], [14], [15], [22], [24], this allows us to use Eq. (91) without sampling, see [34], [28], [35], [26], [29], and also [37], [1].

The PCE is an expansion in multivariate *Hermite polynomials* [10], [14], [15], [22], [24]; we denote by $H_\alpha(\boldsymbol{\theta}) = \prod_{k \in \mathbb{N}} h_{\alpha_k}(\theta_k) \in \mathcal{S}$ the multivariate polynomial in standard and independent Gaussian RVs $\boldsymbol{\theta}(\omega) = (\theta_1(\omega), \dots, \theta_k(\omega), \dots)_{k \in \mathbb{N}}$, where h_j is the usual uni-variate Hermite polynomial, and $\alpha = (\alpha_1, \dots, \alpha_k, \dots)_{k \in \mathbb{N}} \in \mathcal{N} := \mathbb{N}_0^{(\mathbb{N})}$ is a multi-index of generally infinite length but with only finitely many entries non-zero. As $h_0 \equiv 1$, the infinite product is effectively finite and always well-defined.

The *Cameron-Martin* theorem assures us [14], [22], [15] that the set of these polynomials is dense in $\mathcal{S} = L_2(\Omega)$, and in fact $\{H_\alpha / \sqrt{(\alpha)!}\}_{\alpha \in \mathcal{N}}$ is a complete orthonormal system (CONS), where $\alpha! := \prod_{k \in \mathbb{N}} (\alpha_k!)$ is the product of the individual factorials, also well-defined as except for finitely many k one has $\alpha_k! = 0! = 1$. So one may write $\mathbf{x}(\omega) = \sum_{\alpha \in \mathcal{N}} \mathbf{x}^\alpha H_\alpha(\boldsymbol{\theta}(\omega))$ with $\mathbf{x}^\alpha \in \mathbb{R}^M$, and similarly for \mathbf{y} and all other RVs. In this way the RVs are expressed as functions of other, known RVs $\boldsymbol{\theta}$ —hence the name *functional* approximation—and not through samples.

The space \mathcal{S} may now be discretised by taking a finite subset $\mathcal{J} \subset \mathcal{N}$ of size $J = |\mathcal{J}|$, and setting $\mathcal{S}_J = \text{span}\{H_\alpha : \alpha \in \mathcal{J}\} \subset \mathcal{S}$. The orthogonal projection P_J onto \mathcal{S}_J is then simply

$$P_J : \mathcal{X}_M \otimes \mathcal{S} \ni \sum_{\alpha \in \mathcal{N}} \mathbf{x}^\alpha H_\alpha \mapsto \sum_{\alpha \in \mathcal{J}} \mathbf{x}^\alpha H_\alpha \in \mathcal{X}_M \otimes \mathcal{S}_J. \quad (92)$$

Taking Eq. (91), one may rewrite it as

$$\mathbf{x}_a = \mathbf{x}_f + \mathbf{K}(\hat{\mathbf{y}} - \mathbf{y}_f) = \quad (93)$$

$$\sum_{\alpha \in \mathcal{N}} \mathbf{x}_a^\alpha H_\alpha(\boldsymbol{\theta}) = \sum_{\alpha \in \mathcal{N}} \left(\mathbf{x}_f^\alpha + \mathbf{K}(\hat{\mathbf{y}}^\alpha - \mathbf{y}_f^\alpha) \right) H_\alpha(\boldsymbol{\theta}). \quad (94)$$

Observe, that as the measurement or observation $\hat{\mathbf{y}}$ is a constant, one has in Eq. (94) that only $\hat{\mathbf{y}}^0 = \hat{\mathbf{y}}$, all other coefficients $\hat{\mathbf{y}}^\alpha = \mathbf{0}$ for $\alpha \neq 0$.

Projecting both sides of Eq. (94) onto $\mathcal{X}_M \otimes \mathcal{S}_J$ is very simple and results in

$$\sum_{\alpha \in \mathcal{J}} \mathbf{q}_a^\alpha H_\alpha = \sum_{\alpha \in \mathcal{J}} \left(\mathbf{q}_f^\alpha + \mathbf{K}(\mathbf{z}^\alpha - \mathbf{y}_f^\alpha) \right) H_\alpha. \quad (95)$$

Obviously the projection P_J commutes with the Kalman operator K and hence with its finite dimensional analogue \mathbf{K} . One may actually concisely write Eq. (95) as

$$P_J \mathbf{x}_a = P_J \mathbf{x}_f + P_J \mathbf{K}(\hat{\mathbf{y}} - \mathbf{y}_f) = P_J \mathbf{x}_f + \mathbf{K}(P_J \hat{\mathbf{y}} - P_J \mathbf{y}_f). \quad (96)$$

Elements of the discretised space $\mathcal{X}_{M,J} = \mathcal{X}_M \otimes \mathcal{S}_J \subset \mathcal{X}$ thus may be written fully expanded as $\sum_{m=1}^M \sum_{\alpha \in \mathcal{J}} x^{\alpha,m} \varrho_m H_{\alpha}$. The tensor representation is $\mathbf{x} := \sum_{\alpha \in \mathcal{J}} \mathbf{x}^{\alpha} \otimes \mathbf{e}^{\alpha}$, where the $\{\mathbf{e}^{\alpha}\}$ are the canonical basis in \mathbb{R}^J , and may be used to express Eq. (95) or Eq. (96) succinctly as

$$\mathbf{x}_a = \mathbf{x}_f + \mathbf{K}(\hat{\mathbf{y}} - \mathbf{y}_f), \quad (97)$$

again an equation between the tensor representations of some RVs, where $\mathbf{K} = \mathbf{K} \otimes \mathbf{I}$, with \mathbf{K} from Eq. (91). Hence the update equation is naturally in a tensorised form. This is how the update can finally be computed in the PCE representation without any sampling [34], [28], [35], [26]. Analogous statements hold for the forms of the update Eq. (68) with higher order terms $n > 1$, and do not have to be repeated here. Let us remark that these updates go very seamlessly with very efficient methods for sparse or low-rank approximation of tensors, c.f. the monograph [12] and the literature therein. These methods are PCE-forms of the Bayesian update, and in particular the Eq. (97), because of its formal affinity to the Kalman filter (KF), may be called the polynomial chaos expansion based Kalman filter (PCEKF).

It remains to say how to compute the terms ${}^k H$ in the update equation Eq. (68)—or rather the terms in the defining Eq. (70) in Theorem 17—in this approach. Given the PCEs of the RVs, this is actually quite simple as any moment can be computed directly from the PCE [24], [28], [35]. A typical term $\langle y^{\vee k} \rangle = \langle \text{Sym}(y^{\otimes k}) \rangle = \text{Sym}(\langle y^{\otimes k} \rangle)$ in the operator matrix Eq. (70), where $\mathbf{y} = \sum_{\alpha} \mathbf{y}^{\alpha} H_{\alpha}(\boldsymbol{\theta})$, may be computed through

$$\begin{aligned} \langle \mathbf{y}^{\otimes k} \rangle &= \mathbb{E} \left(\bigotimes_{i=1}^k \sum_{\alpha_i} (\mathbf{y}^{\alpha_i} H_{\alpha_i}) \right) = \\ &= \mathbb{E} \left(\sum_{\alpha_1, \dots, \alpha_k} \bigotimes_{i=1}^k \mathbf{y}^{\alpha_i} \prod_{i=1}^k H_{\alpha_i} \right) = \sum_{\alpha_1, \dots, \alpha_k} \bigotimes_{i=1}^k \mathbf{y}^{\alpha_i} \mathbb{E} \left(\prod_{i=1}^k H_{\alpha_i} \right) \end{aligned} \quad (98)$$

As here the H_{α} are *polynomials*, the last expectation in Eq. (98) is finally over products of powers of pairwise independent normalised Gaussian variables, which actually may be done analytically [14], [22], [15]. But some simplifications come from remembering that $\mathbf{y}^0 = \mathbb{E}(\mathbf{y}) = \bar{\mathbf{y}}$, $H_0 \equiv 1$, the orthogonality relation $\langle H_{\alpha} | H_{\beta} \rangle = \delta_{\alpha, \beta} \alpha!$, and that the Hermite polynomials

are an *algebra*. Hence $H_\alpha H_\beta = \sum_\gamma c_{\alpha,\beta}^\gamma H_\gamma$, where the *structure* coefficients $c_{\alpha,\beta}^\gamma$ are known analytically [22], [24], [28], [35].

Similarly, for a RV $R = R(x)$, for a typical right-hand-side term $\langle R(x) \otimes y^{\otimes k} \rangle = \langle R \otimes \text{Sym}(y^{\otimes k}) \rangle$ in Eq. (70) with $\mathbf{R} = \sum_\beta \mathbf{R}^\beta H_\beta(\boldsymbol{\theta})$ one has

$$\langle R \otimes \text{Sym}(y^{\otimes k}) \rangle = \sum_{\beta, \alpha_1, \dots, \alpha_k} \mathbf{R} \otimes \text{Sym} \left(\bigotimes_{i=1}^k \mathbf{y}^{\alpha_i} \right) \mathbb{E} \left(H_\beta \prod_{i=1}^k H_{\alpha_i} \right). \quad (99)$$

As these relations may seem a bit involved — they are actually just an intricate combination of *known* terms — we show here how simple they become for the case of the covariance needed in the linear update formula Eq. (83) or rather Eq. (91):

$$\mathbf{C}_{yy} = \sum_{\alpha \in \mathcal{N}, \alpha \neq 0} (\alpha!) \mathbf{y}^\alpha \otimes \mathbf{y}^\alpha \approx \sum_{\alpha \in \mathcal{J}, \alpha \neq 0} (\alpha!) \mathbf{y}^\alpha \otimes \mathbf{y}^\alpha, \quad (100)$$

$$\mathbf{C}_{xy} = \sum_{\alpha \in \mathcal{N}, \alpha \neq 0} (\alpha!) \mathbf{x}^\alpha \otimes \mathbf{y}^\alpha \approx \sum_{\alpha \in \mathcal{J}, \alpha \neq 0} (\alpha!) \mathbf{x}^\alpha \otimes \mathbf{y}^\alpha. \quad (101)$$

Looking for example at Eq. (91) and our setup as explained in Section 1, we see that the coefficients of $\mathbf{y}(\mathbf{x}_f) = \sum_\alpha \mathbf{y}_f^\alpha H_\alpha$ have to be computed from those of $\mathbf{x}_f = \sum_\beta \mathbf{x}_f^\beta H_\beta$. This propagation of uncertainty through the system is known as *uncertainty quantification* (UQ), e.g. [24] and the references therein. For the sake of brevity, we will not touch further on this subject, which nevertheless is the bedrock on which the whole computational procedure is built.

We next concentrate in Section 5 on examples of updating with ψ_m for the case $m = 1$ in Eq. (68), whereas in Section 6 an example for the case $m = 2$ in Eq. (68) will be shown.

5 The linear Bayesian update

All the examples in this Section 5 have been computed with the case $m = 1$ of up to linear terms in Eq. (68), i.e. this is the LBU with PCEKF. As the traditional Kalman filter is highly geared towards Gaussian distributions [17], and also its Monte Carlo variant EnKF which was mentioned in Section 4 tilts towards Gaussianity, we start with a case—already described in [28]—where the the quantity to be identified has a strongly non-Gaussian distribution, shown in black—the ‘truth’—in Fig. 1. The operator describing the system is the identity—we compute the quantity directly, but there is a Gaussian measurement error. The ‘truth’ was represented as a 12th degree PCE. We




Figure 1: pdfs for linear Bayesian update (LBU), from [28]

use the methods as described in Section 4, and here in particular the Eq. (91) and Eq. (97), the PCEKF.

The update is repeated several times (here ten times) with new measurements—see Fig. 1. The task is here to identify the distribution labelled as ‘truth’ with ten updates of N samples (where $N = 10, 100, 1000$ was used), and we start with a very broad Gaussian prior (in blue). Here we see the ability of the polynomial based LBU, the PCEKF, to identify highly non-Gaussian distributions, the posterior is shown in red and the pdf estimated from the samples in green; for further details see [28].

The next example is also from [28], where the system is the well-known Lorenz-84 chaotic model, a system of three nonlinear ordinary differential equations operating in the chaotic regime. This is truly an example along the description of Eq. (5) and Eq. (9) in Subsection 2.1. Remember that this was originally a model to describe the evolution of some amplitudes of a spherical harmonic expansion of variables describing world climate. As the original scaling of the variables has been kept, the time axis in Fig. 2 is in *days*. Every ten days a noisy measurement is performed and the state description is updated. In between the state description evolves according to the chaotic dynamic of the system. One may observe from Fig. 2 how the uncertainty—the width of the distribution as given by the quantile lines—shrinks every time a measurement is performed, and then increases again due to the chaotic and hence noisy dynamics. Of course, we did not really measure world climate, but rather simulated the ‘truth’ as well, i.e. a *virtual*




Figure 2: Time evolution of Lorenz-84 state and uncertainty with the LBU, from [28]

experiment, like the others to follow. More details may be found in [28] and the references therein.

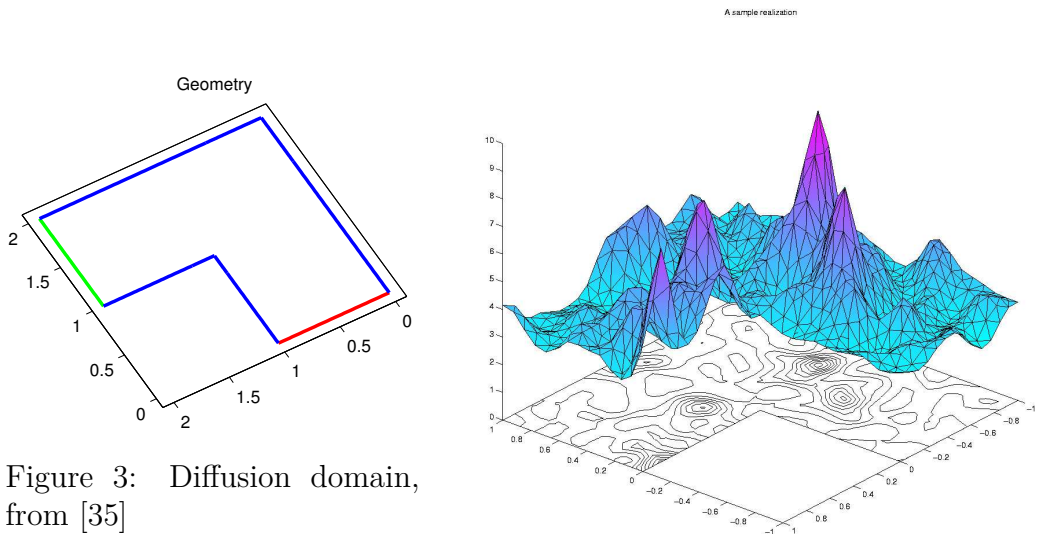


Figure 3: Diffusion domain, from [35]

Figure 4: Conductivity field, from [35]

From [35] we take the example shown in Fig. 3, a linear stationary diffusion equation on an L-shaped plane domain as alluded to in Section 1. The diffusion coefficient κ in Eq. (2) is to be identified. As argued in [34], it is

better to work with $q = \log \kappa$ as the diffusion coefficient has to be positive, but the results are shown in terms of κ .

One possible realisation of the diffusion coefficient is shown in Fig. 4. More realistically, one should assume that κ is a symmetric positive definite tensor field, unless one knows that the diffusion is *isotropic*. Also in this case one should do the updating on the logarithm. For the sake of simplicity we stay with the scalar case, as there is no principal novelty in the non-isotropic case. The virtual experiments use different right-hand-sides f in Eq. (2), and the measurement is the observation of the solution u averaged over little patches.




Figure 5: Convergence of identification, from [35]




Figure 6: Prior and posterior, from [35]

In Fig. 5 one may observe the decrease of the error with successive updates, but due to measurement error and insufficient information from just a few patches, the curves level off, leaving some residual uncertainty. The pdfs of the diffusion coefficient at some point in the domain before and after the updating is shown in Fig. 6, the ‘true’ value at that point was $\kappa = 2$. Further details can be found in [35].

6 The nonlinear Bayesian update

In this Section we want to show a computation with the case $m = 2$ of up to quadratic terms in ψ_m in Eq. (68). We go back to the example of the chaotic Lorentz-84 [28] model already shown in Section 5, from Eq. (5) and Eq. (9) in Subsection 2.1. This kind of experiment has several advantages but at the same time also challenges for identification procedures: it has only a three-dimensional state space, these are the uncertain ‘parameters’, i.e. $\mathbf{x} = (x_1, x_2, x_3) = (x, y, z) \in \mathcal{X} = \mathbb{R}^3$, the corresponding operator A resp. f in the abstract Eq. (1) resp. Eq. (5) is sufficiently nonlinear to make the problem difficult, and adding to this we operate the equation in its chaotic

regime, so that new uncertainty from the numerical computation is added between measurements.




Figure 7: Linear measurement: Comparison posterior for LBU ($m = 1$) and QBU ($m = 2$) after one update

As a first set of experiments we take the measurement operator to be linear in \mathbf{x} , i.e. we can observe the *whole* state directly. At the moment we consider updates after each day—whereas in Section 5 the updates were performed every 10 days. The update is done once with the linear Bayesian update (LBU), and again with a *quadratic* nonlinear BU (QBU) with $m = 2$. The results for the posterior pdfs are given in Fig. 7, where the linear update is dotted in blue, and the full red line is the quadratic QBU; there is hardly any difference between the two, most probably indicating that the LBU is already very accurate.

As the differences between LBU and QBU were small — we take this as an indication that the LBU is not too inaccurate an approximation to the conditional expectation — we change the experiment and take a nonlinear measurement function, which is now cubic: $h(\mathbf{x}) = (x^3, y^3, z^3)$. We now observe larger differences between LBU and QBU.




Figure 8: Cubic measurement: Comparison posterior for LBU ($m = 1$) and QBU ($m = 2$) after one update

These differences in posterior pdfs after one update may be gleaned from Fig. 8, and they are indeed larger than in the linear case Fig. 7, due to the strongly nonlinear measurement operator, showing that the QBU may provide much more accurate tracking of the state, especially for non-linear observation operators.




Figure 9: Deformations, from [34], [36]

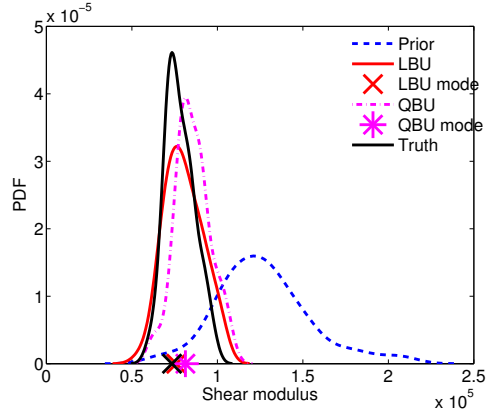


Figure 10: LBU and QBU for the shear modulus

As a last example we take a strongly nonlinear and also non-smooth situation, namely elasto-plasticity with linear hardening and large deformations and a *Kirchhoff-St. Venant* elastic material law [34], [36]. This example is known as *Cook's membrane*, and is shown in Fig. 9 with the undeformed mesh (initial), the deformed one obtained by computing with average values of the elasticity and plasticity material constants (deterministic), and finally the average result from a stochastic forward calculation of the probabilistic model (stochastic), which is described by a variational inequality [36].

The shear modulus G , a random field and not a deterministic value in this case, has to be identified, which is made more difficult by the non-smooth non-linearity. In Fig. 10 one may see the ‘true’ distribution at one point in the domain in an unbroken black line, with the mode — the maximum of the pdf — marked by a black cross on the abscissa, whereas the prior is shown in a dotted blue line. The pdf of the LBU is shown in an unbroken red line, with its mode marked by a red cross, and the pdf of the QBU is shown in a broken purple line with its mode marked by an asterisk. Again we see a difference between the LBU and the QBU. But here a curious thing happens; the mode of the LBU-posterior is actually closer to the mode of the ‘truth’ than the mode of the QBU-posterior. This means that somehow the QBU takes the prior more into account than the LBU, which is a kind of overshooting which has been observed at other occasions. On the other

hand the pdf of the QBU is narrower — has less uncertainty — than the pdf of the LBU.

7 Conclusion

The connection between inverse problems and uncertainty quantification was shown. An abstract model of a system was introduced, together with a measurement operator, which provides a possibility to predict — in a probabilistic sense — a measurement. The framework chosen is that of Bayesian analysis, where uncertain quantities are modelled as random variables. New information leads to an update of the probabilistic description via Bayes’s rule.

After elaborating on the — often not well-known — connection between conditional probabilities as in Bayes’s rule and conditional expectation, we set out to compute and — necessarily — approximate the conditional expectation. As a polynomial approximation was chosen, there is the choice up to which degree one should go. The case with up to linear terms — the linear Bayesian update (LBU) — is best known and intimately connected with the well-known Kalman filter. We call this update the Gauss-Markov-Kalman filter. In addition, we show how to compute approximations of higher order, in particular the quadratic Bayesian update (QBU).

There are several possibilities on how one may choose a numerical realisation of these theoretical concepts, and we decided on functional or spectral approximations. It turns out that this approach goes very well with recent very efficient approximation methods building on separated or so-called low-rank tensor approximations.

Starting with the linear Bayesian update, a series of examples of increasing complexity is shown. The method works well in all cases. Some examples are then chosen to show the nonlinear or rather quadratic Bayesian update, where we go up to quadratic terms. A series of experiments is chosen with different measurement operators, which have quite a marked influence on whether the linear and quadratic update are close to each other.

References

- [1] E. D. Blanchard, A. Sandu, and C. Sandu, *A polynomial chaos-based Kalman filter approach for parameter estimation of mechanical systems*, Journal of Dynamic Systems, Measurement, and Control **132** (2010), no. 6, 061404, doi:10.1115/1.4002481.

- [2] A. Bobrowski, *Functional analysis for probability and stochastic processes*, Cambridge University Press, Cambridge, 2005.
- [3] D. Bosq, *Linear processes in function spaces. theory and applications.*, Lecture Notes in Statistics, vol. 149, Springer, Berlin, 2000.
- [4] ———, *General linear processes in Hilbert spaces and prediction*, Journal of Statistical Planning and Inference **137** (2007), 879–894, doi:10.1016/j.jspi.2006.06.014.
- [5] H. W. Engl, M. Hanke, and A. Neubauer, *Regularization of inverse problems*, Kluwer, Dordrecht, 2000.
- [6] G. Evensen, *Data assimilation — the ensemble Kalman filter*, Springer, Berlin, 2009.
- [7] ———, *The ensemble Kalman filter for combined state and parameter estimation*, IEEE Control Systems Magazine **29** (2009), 82–104, doi:10.1109/MCS.2009.932223.
- [8] J. Galvis and M. Sarkis, *Regularity results for the ordinary product stochastic pressure equation*, SIAM Journal on Mathematical Analysis **44** (2012), 2637–2665, doi:10.1137/110826904.
- [9] D. Gamerman and H. F. Lopes, *Markov Chain Monte Carlo: Stochastic simulation for Bayesian inference*, Chapman & Hall, Boca Raton, FL, 2006.
- [10] R. Ghanem and P. D. Spanos, *Stochastic finite elements—a spectral approach*, Springer, Berlin, 1991.
- [11] M. Goldstein and D. Wooff, *Bayes linear statistics—theory and methods*, Wiley Series in Probability and Statistics, John Wiley & Sons, Chichester, 2007.
- [12] W. Hackbusch, *Tensor spaces and numerical tensor calculus*, Springer, Berlin, 2012.
- [13] T. Hida, H. H. Kuo, J. Potthoff, and L. Streit, *White noise—an infinite dimensional calculus*, Kluwer, Dordrecht, 1999.
- [14] H. Holden, B. Øksendal, J. Ubøe, and T.-S. Zhang, *Stochastic partial differential equations*, Birkhäuser, Basel, 1996.

- [15] S. Janson, *Gaussian Hilbert spaces*, Cambridge Tracts in Mathematics, 129, Cambridge University Press, Cambridge, 1997.
- [16] E. T. Jaynes, *Probability theory, the logic of science*, Cambridge University Press, Cambridge, 2003.
- [17] R. E. Kálmán, *A new approach to linear filtering and prediction problems*, Transactions of the ASME—J. of Basic Engineering (Series D) **82** (1960), 35–45.
- [18] A. Kučerová and H. G. Matthies, *Uncertainty updating in the description of heterogeneous materials*, Technische Mechanik **30** (2010), no. 1–3, 211–226.
- [19] K. H. J. Law, A. Litvinenko, and H. G. Matthies, *Nonlinear evolution, observation, and update*, 2015.
- [20] D. G. Luenberger, *Optimization by vector space methods*, John Wiley & Sons, Chichester, 1969.
- [21] N. Madras, *Lectures on Monte Carlo methods*, American Mathematical Society, Providence, RI, 2002.
- [22] P. Malliavin, *Stochastic analysis*, Springer, Berlin, 1997.
- [23] Y. M. Marzouk, H. N. Najm, and L. A. Rahn, *Stochastic spectral methods for efficient Bayesian solution of inverse problems*, Journal of Computational Physics **224** (2007), no. 2, 560–586, doi:10.1016/j.jcp.2006.10.010.
- [24] H. G. Matthies, *Uncertainty quantification with stochastic finite elements*, Encyclopaedia of Computational Mechanics (E. Stein, R. de Borst, and T. J. R. Hughes, eds.), John Wiley & Sons, Chichester, 2007, doi:10.1002/0470091355.ecm071.
- [25] H. G. Matthies and A. Keese, *Galerkin methods for linear and nonlinear elliptic stochastic partial differential equations*, Computer Methods in Applied Mechanics and Engineering **194** (2005), no. 12-16, 1295–1331. MR MR2121216 (2005j:65146)
- [26] H. G. Matthies, A. Litvinenko, O. Pajonk, B. V. Rosić, and E. Zander, *Parametric and uncertainty computations with tensor product representations*, Uncertainty Quantification in Scientific Computing (Berlin) (A. Dienstfrey and R. Boisvert, eds.), IFIP Advances in Information

and Communication Technology, vol. 377, Springer, 2012, pp. 139–150, doi:10.1007/978-3-642-32677-6.

- [27] T. A. Moselhy and Y. M. Marzouk, *Bayesian inference with optimal maps*, Journal of Computational Physics **231** (2012), 7815–7850, doi:10.1016/j.jcp.2012.07.022.
- [28] 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*, Physica D **241** (2012), 775–788, doi:10.1016/j.physd.2012.01.001.
- [29] O. Pajonk, B. V. Rosić, and H. G. Matthies, *Sampling-free linear Bayesian updating of model state and parameters using a square root approach*, Computers and Geosciences **55** (2013), 70–83, doi:10.1016/j.cageo.2012.05.017.
- [30] A. Papoulis, *Probability, random variables, and stochastic processes*, third ed., McGraw-Hill Series in Electrical Engineering, McGraw-Hill, New York, 1991.
- [31] M. Parno, T. Moselhy, and Y. Marzouk, *A multiscale strategy for Bayesian inference using transport maps*, arXiv:1507.07024v1 [stat.CO], 2015, Available from: <http://arxiv.org/abs/1507.07024>.
- [32] M. M. Rao, *Conditional measures and applications*, CRC Press, Boca Raton, FL, 2005.
- [33] L. Roman and M. Sarkis, *Stochastic Galerkin method for elliptic SPDEs: A white noise approach*, Discrete Cont. Dyn. Syst. Ser. B **6** (2006), 941–955.
- [34] B. V. Rosić, A. Kučerová, J. Sýkora, O. Pajonk, A. Litvinenko, and H. G. Matthies, *Parameter identification in a probabilistic setting*, Engineering Structures **50** (2013), 179–196, doi:10.1016/j.engstruct.2012.12.029.
- [35] B. V. Rosić, A. Litvinenko, O. Pajonk, and H. G. Matthies, *Sampling-free linear Bayesian update of polynomial chaos representations*, Journal of Computational Physics **231** (2012), 5761–5787, doi:10.1016/j.jcp.2012.04.044.
- [36] B. V. Rosić and H. G. Matthies, *Identification of properties of stochastic elastoplastic systems*, Computational Methods in Stochastic Dynamics

(Berlin) (M. Papadrakakis, G. Stefanou, and V. Papadopoulos, eds.), Computational Methods in Applied Sciences, vol. 26, Springer, 2013, pp. 237–253, doi:10.1007/978-94-007-5134-7_14.

- [37] G. Saad and R. Ghanem, *Characterization of reservoir simulation models using a polynomial chaos-based ensemble Kalman filter*, Water Resources Research **45** (2009), W04417, doi:10.1029/2008WR007148.
- [38] D. Sanz-Alonso and A. M. Stuart, *Long-time asymptotics of the filtering distribution for partially observed chaotic dynamical systems*, arXiv:1411.6510v1 [math.DS], 2014, Available from: <http://arxiv.org/abs/1411.6510>.
- [39] I. E. Segal and R. A. Kunze, *Integrals and operators*, Springer, Berlin, 1978.
- [40] A. M. Stuart, *Inverse problems: A Bayesian perspective*, Acta Numerica **19** (2010), 451–559, doi:10.1017/S0962492910000061.
- [41] A. Tarantola, *Inverse problem theory and methods for model parameter estimation*, SIAM, Philadelphia, PA, 2004.
- [42] T.-J. Tarn and Y. Rasis, *Observers for nonlinear stochastic systems*, IEEE Transactions on Automatic Control **21** (1976), 441–448.
- [43] N. Wiener, *The homogeneous chaos*, American Journal of Mathematics **60** (1938), no. 4, 897–936.
- [44] D. Xiu and G. E. Karniadakis, *The Wiener-Askey polynomial chaos for stochastic differential equations*, SIAM Journal of Scientific Computing **24** (2002), 619–644.