

Chapter 2

Assimilating data into scientific models: An optimal coupling perspective

Yuan Cheng and Sebastian Reich

Abstract Data assimilation is the task of combining mathematical models with observational data. From a mathematical perspective data assimilation leads to Bayesian inference problems which can be formulated in terms of Feynman-Kac formulae. In this paper we focus on the sequential nature of many data assimilation problems and their numerical implementation in form of Monte Carlo methods. We demonstrate how sequential data assimilation can be interpreted in terms of time-dependent dynamical systems or, more generally, time-dependent Markov processes, which is often referred to as the McKean approach to Feynman-Kac formulae. It is shown that the McKean approach has very natural links to coupling of probability measures and optimal transportation. This link allows one to propose novel sequential Monte Carlo methods/particle filters. In combination with localization these novel algorithms have the potential of beating the curse of dimensionality, which has prevented particle filters from being applied to spatially extended systems.

1 Introduction

This paper is concerned with Monte Carlo methods for approximating expectation values for sequences of probability density functions (PDFs) $\pi^n(z)$, $n \geq 0$, $z \in \mathcal{Z}$. We assume that these PDFs arise sequentially from a Markov process with given transition kernel $\pi(z|z')$ and are modified by weight functions $G^n(z) \geq 0$ at each iteration index $n \geq 1$. More precisely, the PDFs satisfy the recursion

Y. Cheng

Institut für Mathematik, Universität Potsdam, Am Neuen Palais 10, D-14469 Potsdam, Germany
e-mail: yuan.cheng@uni-potsdam.de

S. Reich (✉)

Institut für Mathematik, Universität Potsdam, Am Neuen Palais 10, D-14469 Potsdam, Germany

Department of Mathematics and Statistics, University of Reading, Whiteknights,
PO Box 220, Reading RG6 6AX, UK
e-mail: sreich@math.uni-potsdam.de

$$\pi^n(z) = \frac{1}{C} G^n(z) \int_{\mathcal{Z}} \pi(z|z') \pi^{n-1}(z') dz' \quad (2.1)$$

with the constant C chosen such that $\int_{\mathcal{Z}} \pi^n(z) dz = 1$.

A general mathematical framework for such problems is provided by the Feynman-Kac formalism as discussed in detail in del Moral (2004).¹ In order to apply Monte Carlo methods to (2.1) it is useful to reformulate (2.1) in terms of modified Markov processes with transition kernel $\pi^n(z|z')$, which satisfy the consistency condition

$$\pi^n(z) = \int_{\mathcal{Z}} \pi^n(z|z') \pi^{n-1}(z') dz'. \quad (2.2)$$

This reformulation has been called the McKean approach to Feynman-Kac models in del Moral (2004).² Once a particular McKean model is available, a Monte Carlo implementation reduces to sequences of particles $\{z_i^n\}_{i=1}^M$ being generated sequentially by

$$z_i^n \sim \pi^n(\cdot | z_i^{n-1}), \quad i = 1, \dots, M, \quad (2.3)$$

for $n = 0, 1, \dots, N$. In other words, z_i^n is the realization of a random variable with (conditional) PDF $\pi^n(z|z_i^{n-1})$. Such a Monte Carlo method constitutes a particular instance of the far more general class of sequential Monte Carlo methods (SMCMs) (Doucet et al., 2001).

While there are many applications that naturally give rise to Feynman-Kac formulae (del Moral, 2004), we will focus in this paper on Markov processes for which the underlying transition kernel $\pi(z|z')$ is determined by a deterministic dynamical system and that we wish to estimate its current state z^n from partial and noisy observations y_{obs}^n . The weight function $G^n(z)$ of a Feynman-Kac recursion (2.1) is in this case given by the likelihood of observing y_{obs}^n given z^n and we encounter a particular application of *Bayesian inference* (Jazwinski, 1970; Stuart, 2010). The precise mathematical setting and the Feynman-Kac formula for the associated data assimilation problem will be discussed in Section 2. Some of the standard Monte Carlo approaches to Feynman-Kac formulae will be summarized in Section 3.

¹The classic Feynman-Kac formulae provide a connection between stochastic processes and solutions to partial differential equations. Here we use a generalization which links discrete-time stochastic processes to sequences of marginal distributions and associated expectation values. In addition to sequential Bayesian inference, which primarily motivates this review article, applications of discrete-time Feynman-Kac formula of type (2.1) can, for example, be found in non-equilibrium molecular dynamics, where the weight functions G^n in (2.1) correspond to the incremental work exerted on a molecular system at time t_n . See Lelièvre et al. (2010) for more details.

²McKean (1966) pioneered the study of stochastic processes which are generated by stochastic differential equations for which the diffusion term depends on the time-evolving marginal distributions $\pi(z, t)$. Here we utilize a generalization of this idea to discrete-time Markov processes which allows for transition kernels $\pi^n(z|z')$ to depend on the marginal distributions $\pi^n(z)$.

It is important to note that the consistency condition (2.2) does not specify a McKean model uniquely. In other words, given a Feynman-Kac recursion (2.1) there are many options to define an associated McKean model $\pi^n(z|z')$. It has been suggested independently by Reich (2011); Reich and Cotter (2013) and Moselhy and Marzouk (2012) in the context of Bayesian inference that optimal transportation (Villani, 2003) can be used to couple the prior and posterior distributions. This idea generalizes to all Feynman-Kac formulae and leads to optimal in the sense of optimal transportation McKean models. This optimal transportation approach to McKean models will be developed in detail in Section 4 of this paper.

Optimal transportation problems lead to a nonlinear elliptic PDE, called the Monge-Ampere equation (Villani, 2003), which is very hard to tackle numerically in space dimensions larger than one. On the other hand, optimal transportation is an infinite-dimensional generalization (McCann, 1995; Villani, 2009) of the classic linear transport problem (Strang, 1986). This interpretation is very attractive in terms of Monte Carlo methods and gives rise to a novel SMCM of type (2.3), which we call the ensemble transform particle filter (ETPF) (Reich, 2013a). The ETPF is based on a linear transformation of the forecast particles

$$z_i^f \sim \pi(\cdot | z_i^{n-1}), \quad i = 1, \dots, M, \quad (2.4)$$

of type

$$z_j^n = \sum_{i=1}^M z_i^f s_{ij} \quad (2.5)$$

with the entries $s_{ij} \geq 0$ of the transform matrix $S \in \mathbb{R}^{M \times M}$ being determined by an appropriate linear transport problem. Even more remarkably, it turns out that SMCMs which resample in each iteration as well as the popular class of ensemble Kalman filters (EnKFs) (Evensen, 2006) also fit into the linear transform framework of (2.5). We will discuss particle/ensemble-based sequential data assimilation algorithms within the unifying framework of linear ensemble transform filters in Section 5. An extension of the ETPF to spatially extended dynamical systems using the concept of localization (Evensen, 2006) is proposed in Section 6. Section 7 provides numerical results for the Lorenz-63 (Lorenz, 1963) and the Lorenz-96 (Lorenz, 1996) models. The results for the 40 dimensional Lorenz-96 indicate that the ensemble transform particle filter with localization can beat the curse of dimensionality which has so far prevented SMCMs from being used for high-dimensional systems (Bengtsson et al., 2008). A brief historical account of data assimilation and filtering is given in Section 7.

We mention that, while the focus of this paper is on state estimation for deterministic dynamical systems, the results can easily be extended to stochastic models as well as combined state and parameter estimation problems. Furthermore, possible applications include all areas in which SMCMs have successfully been used. We refer, for example, to navigation, computer vision, and cognitive sciences (see, e.g., Doucet et al. (2001); Lee and Mumford (2003) and the references therein).

2 Data assimilation and Feynman-Kac formula

Consider a deterministic *dynamical system*³

$$z^{n+1} = \Psi(z^n) \quad (2.6)$$

with state variable $z \in \mathbb{R}^{N_z}$, iteration index $n \geq 0$, and given initial $z^0 \in \mathcal{Z} \subset \mathbb{R}^{N_z}$. We assume that Ψ is a diffeomorphism on \mathbb{R}^{N_z} and that $\Psi(\mathcal{Z}) \subseteq \mathcal{Z}$, which implies that the iterates z^n stay in \mathcal{Z} for all $n \geq 0$. Dynamical systems of the form (2.6) often arise as the time- Δt -flow maps of differential equations

$$\frac{dz}{dt} = f(z). \quad (2.7)$$

In many practical applications the initial state z^0 is not precisely known. We may then assume that our uncertainty about the correct initial state can, for example, be quantified in terms of ratios of frequencies of occurrence. More precisely, the *ratio of frequency of occurrence* (RFO) of two initial conditions $z_a^0 \in \mathcal{Z}$ and $z_b^0 \in \mathcal{Z}$ is defined as the RFOs for the two associated ε -neighborhoods $\mathcal{U}_\varepsilon(z_a^0)$ and $\mathcal{U}_\varepsilon(z_b^0)$, respectively, and taking the limit $\varepsilon \rightarrow 0$. Here $\mathcal{U}_\varepsilon(z)$ is defined as

$$\mathcal{U}_\varepsilon(z) = \{z' \in \mathbb{R}^{N_z} : \|z' - z\|^2 \leq \varepsilon\}$$

It is important to note that the volume of both neighborhoods is identical, *i.e.*, $V(\mathcal{U}_\varepsilon(z_a^0)) = V(\mathcal{U}_\varepsilon(z_b^0))$.

From a frequentist perspective, RFOs can be thought of as arising from repeated experiments with one and the same dynamical system (2.6) under varying initial conditions and upon counting how often $z^0 \in \mathcal{U}_\varepsilon(z_a^0)$ relative to $z^0 \in \mathcal{U}_\varepsilon(z_b^0)$. There will, of course, be many instances for which z^0 is neither in $\mathcal{U}_\varepsilon(z_a^0)$ nor $\mathcal{U}_\varepsilon(z_b^0)$. Alternatively, one can take a Bayesian perspective and think of RFOs as our subjective belief about a z_a^0 to actually arise as the initial condition in (2.6) relative to another initial condition z_b^0 . The later interpretation is also applicable in case only a single experiment with (2.6) is conducted.

Independent of such a statistical interpretation of the RFO, we assume the availability of a function $\tau(z) > 0$ such that the RFO can be expressed as

$$\text{RFO} = \frac{\tau(z_a^0)}{\tau(z_b^0)} \quad (2.8)$$

for all pairs of initial conditions from \mathcal{Z} .

³Even though this review article assumes deterministic evolution equations, the results presented here can easily be generalized to evolution equations with stochastic model errors.

Provided that

$$\int_{\mathcal{Z}} \tau(z) dz < \infty$$

we can introduce the probability density function (PDF)

$$\pi_{Z^0}(z) = \frac{\tau(z)}{\int_{\mathcal{Z}} \tau(z) dz}$$

and interpret initial conditions z^0 as realizations of a random variable $Z^0 : \Omega \rightarrow \mathbb{R}^{N_z}$ with PDF π_{Z^0} .⁴ We remark that most of our subsequent discussions carry through even if $\int_{\mathcal{Z}} \tau(z) dz$ is unbounded as long as the RFOs remain well-defined.

So far we have discussed RFOs for initial conditions. But one can also consider such ratios for any iteration index $n \geq 0$, *i.e.*, for solutions

$$z_a^n = \Psi^n(z_a^0)$$

and

$$z_b^n = \Psi^n(z_b^0).$$

Here Ψ^n denotes the n -fold application of Ψ . The RFO at iteration index n is now defined as the ratio of the frequencies of occurrence for the two associated ε -neighborhoods $\mathcal{U}_\varepsilon(z_a^n)$ and $\mathcal{U}_\varepsilon(z_b^n)$, respectively, in the limit $\varepsilon \rightarrow 0$. We pull this ratio back to $n = 0$ and find that

$$\text{RFO}(n) \approx \frac{\tau(\Psi^{-n}(z_a^n)) V_a}{\tau(\Psi^{-n}(z_b^n)) V_b}$$

for ε sufficiently small, where

$$V_{a/b} = V(\Psi^{-n}(\mathcal{U}_\varepsilon(z_{a/b}^n))) := \int_{\Psi^{-n}(\mathcal{U}_\varepsilon(z_{a/b}^n))} dz$$

denote the volumes of $\Psi^{-n}(\mathcal{U}_\varepsilon(z_{a/b}^n))$ and Ψ^{-n} refers to the inverse of Ψ^n . These two volumes can be approximated as

$$V_{a/b} \approx V(\mathcal{U}_\varepsilon(z_{a/b}^0)) \times |D\Psi^{-n}(z_{a/b}^n)|$$

⁴We have assumed the existence of an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The specific structure of this probability space does not play a role in the subsequent discussions.

for $\varepsilon > 0$ sufficiently small. Here $D\Psi^{-n}(z) \in \mathbb{R}^{N_z \times N_z}$ stands for the Jacobian matrix of partial derivatives of Ψ^{-n} at z and $|D\Psi^{-n}(z)|$ for its determinant. Hence, upon taking the limit $\varepsilon \rightarrow 0$, we obtain

$$\begin{aligned} \text{RFO}(n) &= \frac{\tau(\Psi^{-n}(z_a^n))|D\Psi^{-n}(z_a^n)|}{\tau(\Psi^{-n}(z_b^n))|D\Psi^{-n}(z_b^n)|} \\ &= \frac{\pi_{Z^0}(\Psi^{-n}(z_a^n))|D\Psi^{-n}(z_a^n)|}{\pi_{Z^0}(\Psi^{-n}(z_b^n))|D\Psi^{-n}(z_b^n)|}. \end{aligned}$$

Therefore we may interpret solutions z^n for fixed iteration index $n \geq 1$ as realizations of a random variable $Z^n : \Omega \rightarrow \mathbb{R}^{N_z}$ with PDF

$$\pi_{Z^n}(z) = \pi_{Z^0}(\Psi^{-n}(z))|D\Psi^{-n}(z)|. \quad (2.9)$$

These PDFs can also be defined recursively using

$$\pi_{Z^{n+1}}(z) = \int_{\mathcal{Z}} \delta(z - \Psi(z')) \pi_{Z^n}(z') dz'. \quad (2.10)$$

Here $\delta(\cdot)$ denotes the Dirac delta function, which satisfies

$$\int_{\mathbb{R}^{N_z}} f(z) \delta(z - \bar{z}) dz = f(\bar{z})$$

for all smooth functions $f : \mathbb{R}^{N_z} \rightarrow \mathbb{R}$.⁵ In other words, the dynamical system (2.6) induces a *Markov process*, which we can also write as

$$Z^{n+1} = \Psi(Z^n)$$

in terms of the random variables Z^n , $n \geq 0$.

The sequence of random variables $\{Z^n\}_{n=0}^N$ for fixed $N \geq 1$ gives rise to the finite-time stochastic process $Z^{0:N} : \Omega \rightarrow \mathcal{Z}^{N+1}$ with realizations

$$z^{0:N} := (z^0, z^1, \dots, z^N) = Z^{0:N}(\omega), \quad \omega \in \Omega,$$

that satisfy (2.6). The joint distribution of $Z^{0:N}$, denoted by $\pi_{Z^{0:N}}$, is formally⁶ given by

$$\pi_{Z^{0:N}}(z^0, \dots, z^N) = \pi_{Z^0}(z^0) \delta(z^1 - \Psi(z^0)) \cdots \delta(z^N - \Psi(z^{N-1})) \quad (2.11)$$

and (2.9) is the marginal of $\pi_{Z^{0:N}}$ in z^n , $n = 1, \dots, N$.

⁵The Dirac delta function $\delta(z - \bar{z})$ provides a convenient shorthand for the point measure $\mu_{\bar{z}}(dz)$.

⁶To be mathematically precise one should talk about the joint measure

$$\mu_{Z^{0:N}}(dz^0, \dots, dz^N) = \mu_{Z^0}(dz^0) \mu_{\Psi(z^0)}(dz^1) \cdots \mu_{\Psi(z^{N-1})}(dz^N)$$

with initial measure $\mu_{Z^0}(dz^0) = \pi_{Z^0}(z^0) dz^0$.

Let us now consider the situation where (2.6) serves as a model for an unknown physical process with realization

$$z_{\text{ref}}^{0:N} = (z_{\text{ref}}^0, z_{\text{ref}}^1, \dots, z_{\text{ref}}^N). \quad (2.12)$$

In the classic filtering/smoothing setting (Jazwinski, 1970; Bain and Crisan, 2009) one assumes that there exists an $\omega_{\text{ref}} \in \Omega$ such that

$$z_{\text{ref}}^{0:N} = Z^{0:N}(\omega_{\text{ref}}).$$

In practice such an assumption is highly unrealistic and the reference trajectory (2.12) may instead follow an iteration

$$z_{\text{ref}}^{n+1} = \Psi_{\text{ref}}(z_{\text{ref}}^n) \quad (2.13)$$

with unknown initial z_{ref}^0 and unknown Ψ_{ref} . Of course, it should hold that Ψ in (2.6) is close to Ψ_{ref} in an appropriate mathematical sense.

Independently of such assumptions, we assume that $z_{\text{ref}}^{0:N}$ is accessible to us through partial and noisy observations of the form

$$y_{\text{obs}}^n = h(z_{\text{ref}}^n) + \xi^n, \quad n = 1, \dots, N,$$

where $h : \mathcal{X} \rightarrow \mathbb{R}^{N_y}$ is called the *forward* or observation map and the ξ^n 's are realizations of independent and identically distributed Gaussian random variables with mean zero and covariance matrix $R \in \mathbb{R}^{N_y \times N_y}$. Estimating z_{ref}^n from y_{obs}^n constitutes a classic inverse problem (Tarantola, 2005).

The *ratio of fits to data* (RFD) of two realizations $z_a^{0:N}$ and $z_b^{0:N}$ from the stochastic process $Z^{0:N}$ is defined as

$$\text{RFD} = \frac{\prod_{n=1}^N e^{-\frac{1}{2} (h(z_a^n) - y_{\text{obs}}^n)^T R^{-1} (h(z_a^n) - y_{\text{obs}}^n)}}{\prod_{n=1}^N e^{-\frac{1}{2} (h(z_b^n) - y_{\text{obs}}^n)^T R^{-1} (h(z_b^n) - y_{\text{obs}}^n)}}.$$

Finally we define the *ratio of fits to model and data* (RFMD) of a $z_a^{0:N}$ versus a $z_b^{0:N}$ given the model and the observations as follows:

$$\begin{aligned} \text{RFMD} &= \text{RFD} \times \text{RFO}(0) \\ &= \frac{\prod_{n=1}^N e^{-\frac{1}{2} (h(z_a^n) - y_{\text{obs}}^n)^T R^{-1} (h(z_a^n) - y_{\text{obs}}^n)}}{\prod_{n=1}^N e^{-\frac{1}{2} (h(z_b^n) - y_{\text{obs}}^n)^T R^{-1} (h(z_b^n) - y_{\text{obs}}^n)}} \frac{\pi_{Z^0}(z_a^0)}{\pi_{Z^0}(z_b^0)} \\ &= \frac{\prod_{n=1}^N e^{-\frac{1}{2} (h(z_a^n) - y_{\text{obs}}^n)^T R^{-1} (h(z_a^n) - y_{\text{obs}}^n)}}{\prod_{n=1}^N e^{-\frac{1}{2} (h(z_b^n) - y_{\text{obs}}^n)^T R^{-1} (h(z_b^n) - y_{\text{obs}}^n)}} \frac{\pi_{Z^{0:N}}(z_a^{0:N})}{\pi_{Z^{0:N}}(z_b^{0:N})}. \end{aligned} \quad (2.14)$$

The simple product structure arises since the uncertainty in the initial conditions is assumed to be independent of the measurement errors and the last identity follows from the fact that our dynamical model is deterministic.

Again we may translate this combined ratio into a PDF

$$\pi_{Z^{0:N}}(z^{0:N} | y_{\text{obs}}^{1:N}) = \frac{1}{C} \prod_{n=1}^N e^{-\frac{1}{2}(h(z^n) - y_{\text{obs}}^n)^T R^{-1} (h(z^n) - y_{\text{obs}}^n)} \pi_{Z^0}(z^0), \quad (2.15)$$

where $C > 0$ is a normalization constant depending only on $y_{\text{obs}}^{1:N}$. This PDF gives the probability distribution in $z^{0:N}$ conditioned on the given set of observations

$$y_{\text{obs}}^{1:N} = (y_{\text{obs}}^1, \dots, y_{\text{obs}}^N).$$

The PDF (2.15) is, of course, also conditioned on (2.6) and the initial PDF π_{Z^0} . This dependence is not explicitly taken account of in order to avoid additional notational clutter.

The formulation (2.15) is an instance of Bayesian inference on the one hand, and an instance of the Feynman-Kac formalism on the other. Within the Bayesian perspective, π_{Z^0} (or, equivalently, $\pi_{Z^{0:N}}$) represents the prior distribution,

$$\pi_{Y^{1:N}}(y^{1:N} | z^{0:N}) = \frac{1}{(2\pi)^{N_y N/2} |R|^{N/2}} \prod_{n=1}^N e^{-\frac{1}{2}(h(z^n) - y^n)^T R^{-1} (h(z^n) - y^n)}$$

the compounded likelihood function, and (2.15) the posterior PDF given an actually observed $y^{1:n} = y_{\text{obs}}^{1:n}$. The Feynman-Kac formalism is more general and includes a wide range of applications for which an underlying stochastic process is modified by weights $G^n(z^n) \geq 0$. These weights then replace the likelihood functions

$$\pi_Y(y_{\text{obs}}^n | z^n) = \frac{1}{(2\pi)^{N_y/2} |R|^{1/2}} e^{-\frac{1}{2}(h(z^n) - y_{\text{obs}}^n)^T R^{-1} (h(z^n) - y_{\text{obs}}^n)}$$

in (2.15). The functions $G^n : \mathbb{R}^{N_z} \rightarrow \mathbb{R}$ can depend on the iteration index, as in

$$G^n(z) := \pi_Y(y_{\text{obs}}^n | z)$$

or may be independent of the iteration index. See del Moral (2004) for further details on the Feynman-Kac formalism and Lelièvre et al. (2010) for a specific (non-Bayesian) application in the context of non-equilibrium molecular dynamics.

Formula (2.15) is hardly ever explicitly accessible and one needs to resort to numerical approximations whenever one wishes to either compute the expectation value

$$\mathbb{E}[f(Z^{0:N}) | y_{\text{obs}}^{1:N}] = \int_{\mathcal{Z}^{N+1}} f(z^{0:N}) \pi_{Z^{0:N}}(z^{0:N} | y_{\text{obs}}^{1:N}) dz^0 \cdots dz^N$$

of a given function $f : \mathcal{Z}^{N+1} \rightarrow \mathbb{R}$ or the

$$\text{RFMD} = \frac{\pi_{Z^{0:N}}(z_a^{0:N} | y_{\text{obs}}^{1:N})}{\pi_{Z^{0:N}}(z_b^{0:N} | y_{\text{obs}}^{1:N})}$$

for given trajectories $z_a^{0:N}$ and $z_b^{0:N}$. Basic Monte Carlo approximation methods will be discussed in Section 3. Alternatively, one may seek the *maximum a posteriori (MAP) estimator* z_{MAP}^0 , which is defined as the initial condition z^0 that maximizes (2.15) or, formulated alternatively,

$$z_{\text{MAP}}^0 = \arg \inf L(z^0), \quad L(z^0) := -\log \pi_{Z^{0:N}}(z^{0:N} | y_{\text{obs}}^{1:N})$$

(Kaipio and Somersalo, 2005; Lewis et al., 2006; Tarantola, 2005). The MAP estimator is closely related to *variational data assimilation techniques*, such as 3D-Var and 4D-Var, widely used in meteorology (Daley, 1993; Kalnay, 2002).

In many applications expectation values need to be computed for functions f which depend only on a single z^n . Those expectation values can be obtained by first integrating out all components in (2.15) except for z^n . We denote the resulting marginal PDF by $\pi_{Z^n}(z^n | y_{\text{obs}}^{1:N})$. The case $n = 0$ plays a particular role since

$$\text{RFMD} = \frac{\pi_{Z^0}(z_a^0 | y_{\text{obs}}^{1:N})}{\pi_{Z^0}(z_b^0 | y_{\text{obs}}^{1:N})}$$

and

$$\pi_{Z^n}(z^n | y_{\text{obs}}^{1:N}) = \pi_{Z^0}(\Psi^{-n}(z^n) | y_{\text{obs}}^{1:N}) |D\Psi^{-n}(z^n)|$$

for $n = 1, \dots, N$. These identities hold because our dynamical system (2.6) is deterministic and invertible. In Section 4 we will discuss recursive approaches for determining the marginal $\pi_{Z^n}(z^n | y_{\text{obs}}^{1:N})$ in terms of Markov processes. Computational techniques for implementing such recursions will be discussed in Section 5.

3 Monte Carlo methods in path space

In this section, we briefly summarize two popular Monte Carlo strategies for computing expectation values with respect to the complete conditional distribution $\pi_{Z^{0:N}}(\cdot | y_{\text{obs}}^{1:N})$. We start with the classic importance sampling Monte Carlo method.

3.1 Ensemble prediction and importance sampling

Ensemble prediction is a Monte Carlo method for assessing the marginal PDFs (2.9) for $n = 0, \dots, N$. One first generates z_i^0 , $i = 1, \dots, M$, independent samples from the initial PDF π_{Z^0} ; *i.e.*, samples are generated such that the probability of being in $\mathcal{U}_\varepsilon(z)$ is

$$\int_{\mathcal{U}_\varepsilon(z)} \pi_{Z^0}(z') dz' \approx V(\mathcal{U}_\varepsilon(z)) \times \pi_{Z^0}(z).$$

Furthermore, the expectation value of a function f with respect to Z^0 is approximated by the familiar empirical estimator

$$\bar{f}_M := \frac{1}{M} \sum_{i=1}^M f(z_i^0).$$

The initial ensemble $\{z_i^0\}_{i=1}^M$ is propagated independently under the dynamical system (2.6) for $n = 0, \dots, N-1$. This yields M trajectories

$$z_i^{0:N} = (z_i^0, z_i^1, \dots, z_i^N),$$

which provide independent samples from the $\pi_{Z^{0:N}}$ distribution. With each of these samples we associate the weight

$$w_i = \frac{1}{C} \prod_{n=1}^N e^{-\frac{1}{2} (h(z_i^n) - y_{\text{obs}}^n)^T R^{-1} (h(z_i^n) - y_{\text{obs}}^n)}$$

with the constant of proportionality chosen such that $\sum_{i=1}^M w_i = 1$.

The RFMD for any pair of samples $z_i^{1:N}$ and $z_j^{1:N}$ from $\pi_{Z^{0:N}}$ is now simply given by w_i/w_j and the expectation value of a function f with respect to $\pi_{Z^{0:N}}(z^{0:N} | y_{\text{obs}}^{1:N})$ can be approximated by the empirical estimator

$$\bar{f}_M = \sum_{i=1}^M w_i f(z_i^{0:N}).$$

This estimator is an instance of *importance sampling* since samples from a distribution different from the target distribution, here $\pi_{Z^{0:N}}$, are used to approximate the statistics of the target distribution, here $\pi_{Z^{0:N}}(\cdot | y_{\text{obs}}^{1:N})$. See Liu (2001) and Robert and Casella (2004) for further details.

3.2 Markov chain Monte Carlo (MCMC) methods

Importance sampling becomes inefficient whenever the *effective sample size*

$$M_{\text{eff}} = \frac{1}{\sum_{i=1}^M w_i^2} \in [1, M] \quad (2.16)$$

becomes much smaller than the sample size M . Under those circumstances it can be preferable to generate dependent samples $z_i^{0:N}$ using MCMC methods. MCMC methods rely on a proposal step and a Metropolis-Hastings acceptance criterion.

Note that only z_i^0 needs to be stored since the whole trajectory is then uniquely determined by $z_i^n = \Psi^n(z_i^0)$. Consider, for simplicity, the reversible proposal step

$$z_p^0 = z_i^0 + \xi,$$

where ξ is a realization of a random variable \mathcal{E} with PDF $\pi_{\mathcal{E}}$ satisfying $\pi_{\mathcal{E}}(\xi) = \pi_{\mathcal{E}}(-\xi)$ and z_i^0 denotes the last accepted sample. The associated trajectory $z_p^{0:N}$ is computed using (2.6). Next one evaluates the RFMD (2.14) for $z_a^{0:N} = z_p^{0:N}$ versus $z_b^{0:N} = z_i^{0:N}$, i.e.,

$$\alpha := \frac{\prod_{n=1}^N e^{-\frac{1}{2}(h(z_p^n) - y_{\text{obs}}^n)^T R^{-1}(h(z_p^n) - y_{\text{obs}}^n)} \pi_{Z^0}(z_p^0)}{\prod_{n=1}^N e^{-\frac{1}{2}(h(z_i^n) - y_{\text{obs}}^n)^T R^{-1}(h(z_i^n) - y_{\text{obs}}^n)} \pi_{Z^0}(z_i^0)}.$$

If $\alpha \geq 1$, then the proposal is accepted and the new sample for the initial condition is $z_{i+1}^0 = z_p^0$. Otherwise the proposal is accepted with probability α and rejected with probability $1 - \alpha$. In case of rejection one sets $z_{i+1}^0 = z_i^0$. Note that the accepted samples follow the $\pi_{Z^0}(z^0 | y_{\text{obs}}^{1:N})$ distribution and not the initial PDF $\pi_{Z^0}(z^0)$.

A potential problem with MCMC methods lies in low acceptance rates and/or highly dependent samples. In particular, if the distribution in $\pi_{\mathcal{E}}$ is too narrow, then exploration of phase space can be slow while a too wide distribution can potentially lead to low acceptance rates. Hence one should compare the effective sample size (2.16) from an importance sampling approach to the effective sample size of an MCMC implementation, which is inversely proportional to the integrated autocorrelation time of the samples. See Liu (2001) and Robert and Casella (2004) for more details.

We close this section by referring to Särkkä (2013) for further algorithms for filtering and smoothing.

4 McKean optimal transportation approach

We now restrict the discussion of the Feynman-Kac formalism to the marginal PDFs $\pi_{Z^n}(z^n | y_{\text{obs}}^{1:N})$. We have already seen in Section 2 that the marginal PDF with $n = N$ plays a particularly important role. We show that this marginal PDF can be recursively defined starting from the PDF π_{Z^0} for the initial condition z^0 of (2.6). For that reason we introduce the forecast and analysis PDF at iteration index n and denote them by $\pi_{Z^n}(z^n | y_{\text{obs}}^{1:n-1})$ and $\pi_{Z^n}(z^n | y_{\text{obs}}^{1:n})$, respectively. Those PDFs are defined recursively by

$$\pi_{Z^n}(z^n | y_{\text{obs}}^{1:n-1}) = \pi_{Z^{n-1}}(\Psi^{-1}(z^n) | y_{\text{obs}}^{1:n-1}) | D\Psi^{-1}(z^n) | \quad (2.17)$$

and

$$\pi_{Z^n}(z^n | y_{\text{obs}}^{1:n}) = \frac{\pi_Y(y_{\text{obs}}^n | z^n) \pi_{Z^n}(z^n | y_{\text{obs}}^{1:n-1})}{\int_{\mathcal{Z}} \pi_Y(y_{\text{obs}}^n | z) \pi_{Z^n}(z | y_{\text{obs}}^{1:n-1}) dz}. \quad (2.18)$$

Here (2.17) simply propagates the analysis from index $n - 1$ to the forecast at index n under the action of the dynamical system (2.6). Bayes' formula is then applied in (2.18) in order to transform the forecast into the analysis at index n by assimilating the observed y_{obs}^n .

Theorem 1. *Consider the sequence of forecast and analysis PDFs defined by the recursion (2.17)–(2.18) for $n = 1, \dots, N$ with $\pi_{z^0}(z^0)$ given. Then the analysis PDF at $n = N$ is equal to the Feynman-Kac PDF (2.15) marginalized down to the single variable z^N .*

Proof. We prove the theorem by induction in N . We first verify the claim for $N = 1$. Indeed, by definition,

$$\begin{aligned}\pi_{z^1}(z^1|y_{\text{obs}}^1) &= \frac{1}{C_1} \int_{\mathcal{Z}} \pi_Y(y_{\text{obs}}^1|z^1) \delta(z^1 - \Psi(z^0)) \pi_{z^0}(z^0) dz^0 \\ &= \frac{1}{C_1} \pi_Y(y_{\text{obs}}^1|z^1) \pi_{z^0}(\Psi^{[-1]}(z^1)) |D\Psi^{-1}(z^1)| \\ &= \frac{1}{C_1} \pi_Y(y_{\text{obs}}^1|z^1) \pi_{z^1}(z^1)\end{aligned}$$

and $\pi_{z^1}(z^1)$ is the forecast PDF at index $n = 1$. Here C_1 denotes the constant of proportionality which only depends on y_{obs}^1 .

The induction step from N to $N + 1$ follows from the following line of arguments. We know by the induction assumption that the marginal PDF at $N + 1$ can be written as

$$\begin{aligned}\pi_{z^{N+1}}(z^{N+1}|y_{\text{obs}}^{1:N+1}) &= \int_{\mathcal{Z}} \pi_{z^{N:N+1}}(z^{N:N+1}|y_{\text{obs}}^{1:N+1}) dz^N \\ &= \frac{1}{C_{N+1}} \pi_Y(y_{\text{obs}}^{N+1}|z^{N+1}) \int_{\mathcal{Z}} \delta(z^{N+1} - \Psi(z^N)) \pi_{z^N}(z^N|y_{\text{obs}}^{1:N}) dz^N \\ &= \frac{1}{C_{N+1}} \pi_Y(y_{\text{obs}}^{N+1}|z^{N+1}) \pi_{z^{N+1}}(z^{N+1}|y_{\text{obs}}^{1:N})\end{aligned}$$

in agreement with (2.18) for $n = N + 1$. Here C_{N+1} denotes the constant of proportionality that depends only on y_{obs}^{N+1} and we have made use of the fact that the forecast PDF at index $n = N + 1$ is, according to (2.17), defined by

$$\pi_{z^{N+1}}(z^{N+1}|y_{\text{obs}}^{1:N}) = \pi_{z^N}(\Psi^{-1}(z^{N+1})|y_{\text{obs}}^{1:N}) |D\Psi^{-1}(z^{N+1})|.$$

□

While the forecast step (2.17) is in the form of a Markov process with transition kernel

$$\pi_{\text{model}}(z|z') = \delta(z - \Psi(z')),$$

this does not hold for the analysis step (2.18). The *McKean approach* to (2.17)–(2.18) is based on introducing Markov transition kernels $\pi_{\text{data}}^n(z|z')$, $n = 1, \dots, N$, for the analysis step (2.18). In other words, the transition kernel π_{data}^n at iteration index n has to satisfy the consistency condition

$$\pi_{Z^n}(z|y_{\text{obs}}^{1:n}) = \int_{\mathcal{Z}} \pi_{\text{data}}^n(z|z') \pi_{Z^n}(z'|y_{\text{obs}}^{1:n-1}) dz'. \quad (2.19)$$

These transition kernels are not unique. The following kernel

$$\pi_{\text{data}}^n(z|z') = \epsilon \pi_Y(y_{\text{obs}}^n|z') \delta(z - z') + (1 - \epsilon \pi_Y(y_{\text{obs}}^n|z')) \pi_{Z^n}(z|y_{\text{obs}}^{1:n}) \quad (2.20)$$

has, for example, been considered in del Moral (2004). Here $\epsilon \geq 0$ has to be chosen such that

$$1 - \epsilon \pi_Y(y_{\text{obs}}^n|z) \geq 0$$

for all $z \in \mathbb{R}^{N_z}$. Indeed, we find that

$$\begin{aligned} \int_{\mathcal{Z}} \pi_{\text{data}}^n(z|z') \pi_{Z^n}(z'|y_{\text{obs}}^{1:n-1}) dz' &= \pi_{Z^n}(z|y_{\text{obs}}^{1:n}) + \epsilon \pi_Y(y_{\text{obs}}^n|z) \pi_{Z^n}(z|y_{\text{obs}}^{1:n-1}) - \\ &\quad \epsilon \pi_{Z^n}(z|y_{\text{obs}}^{1:n}) \pi_Y(y_{\text{obs}}^n|y_{\text{obs}}^{1:n-1}) \\ &= \pi_{Z^n}(z|y_{\text{obs}}^{1:n}). \end{aligned}$$

The intuitive interpretation of this transition kernel is that one stays at z' with probability $p = \epsilon \pi_Y(y_{\text{obs}}^n|z')$ while with probability $(1 - p)$ one samples from the analysis PDF $\pi_{Z^n}(z|y_{\text{obs}}^{1:n})$.

Let us define the combined McKean-Markov transition kernel

$$\begin{aligned} \pi^n(z^n|z^{n-1}) &:= \int_{\mathcal{Z}} \pi_{\text{data}}^n(z^n|z) \pi_{\text{model}}(z|z^{n-1}) dz \\ &= \int_{\mathcal{Z}} \pi_{\text{data}}^n(z^n|z) \delta(z - \Psi(z^{n-1})) dz \\ &= \pi_{\text{data}}^n(z^n|\Psi(z^{n-1})) \end{aligned} \quad (2.21)$$

for the propagation of the analysis PDF from iteration index $n-1$ to n . The combined McKean-Markov transition kernels π^n , $n = 1, \dots, N$, define a finite-time stochastic process $\hat{Z}^{0:N} = \{\hat{Z}^n\}_{n=0}^N$ with $\hat{Z}^0 = Z^0$. The marginal PDFs satisfy

$$\begin{aligned} \pi_{\hat{Z}^n}(z^n|y_{\text{obs}}^{1:n}) &= \int_{\mathcal{Z}} \pi^n(z^n|z^{n-1}) \pi_{\hat{Z}^{n-1}}(z^{n-1}|y_{\text{obs}}^{1:n-1}) dz^{n-1} \\ &= \pi_{Z^n}(z^n|y_{\text{obs}}^{1:n}) \end{aligned}$$

Corollary 1. *The final time marginal distribution $\pi_{Z^N}(z^N|y_{\text{obs}}^{1:N})$ of the Feynman-Kac formulation (2.15) is identical to the final time marginal distribution $\pi_{\hat{Z}^N}(z^N|y_{\text{obs}}^{1:N})$ of the finite-time stochastic process $\hat{Z}^{0:N}$ induced by the McKean-Markov transition kernels $\pi^n(z^n|z^{n-1})$.*

We summarize our discussion on the McKean approach in terms of analysis and forecast random variables, which constitute the basic building blocks for most current sequential data assimilation methods.

Definition 1. Given a dynamic iteration (2.6) with PDF π_{Z^0} for the initial conditions and observations y_{obs}^n , $n = 1, \dots, N$, the McKean approach leads to a recursive definition of forecast $Z^{n,f}$ and analysis $Z^{n,a}$ random variables. The iteration is started by declaring Z^0 the analysis $Z^{0,a}$ at $n = 0$. The forecast at iteration index $n > 0$ is defined by propagating the analysis at $n - 1$ forward under the dynamic model (2.6), i.e.,

$$Z^{n,f} = \Psi(Z^{n-1,a}). \quad (2.22)$$

The analysis $Z^{n,a}$ at iteration index n is defined by applying the transition kernel $\pi_{\text{data}}^n(z|z')$ to $Z^{n,f}$. In particular, if $z^{n,f} = Z^{n,f}(\omega)$ is a realized forecast at iteration index n , then the analysis is distributed according to

$$Z^{n,a}|z^{n,f} \sim \pi_{\text{data}}^n(z|z^{n,f}). \quad (2.23)$$

If the marginal PDFs of $Z^{n,f}$ and $Z^{n,a}$ are denoted by $\pi_{Z^{n,f}}(z)$ and $\pi_{Z^{n,a}}(z)$, respectively, then the transition kernel π_{data}^n has to satisfy the compatibility condition (2.19), i.e.,

$$\int_{\mathcal{Z}} \pi_{\text{data}}^n(z|z') \pi_{Z^{n,f}}(z') dz' = \pi_{Z^{n,a}}(z) \quad (2.24)$$

with

$$\pi_{Z^{n,a}}(z) = \frac{\pi_Y(y_{\text{obs}}^n|z) \pi_{Z^{n,f}}(z)}{\int_{\mathcal{Z}} \pi_Y(y_{\text{obs}}^n|z) \pi_{Z^{n,f}}(z) dz}.$$

The data related transition step (2.23) introduces randomness into the analysis of a given forecast value. This appears counterintuitive and, indeed, the main purpose of the rest of this section is to demonstrate that the transition kernel π_{data}^n can be chosen such that

$$Z^{n,a} = \nabla_z \phi^n(Z^{n,f}), \quad (2.25)$$

where $\phi^n : \mathbb{R}^{N_z} \rightarrow \mathbb{R}$ is an appropriate convex potential. In other words, the data-driven McKean update step can be reduced to a (deterministic) map and the stochastic process $\hat{Z}^{0:N}$ is induced by the deterministic recursion (or dynamical system)

$$\hat{Z}^n = \nabla_z \phi^n(\Psi(\hat{Z}^{n-1}))$$

with $\hat{Z}^0 = Z^0$.

The compatibility condition (2.24) with $\pi_{\text{data}}^n(z|z') = \delta(z - \nabla_z \phi^n(z'))$ reduces to

$$\pi_{Z^{n,a}}(\nabla_z \phi^n(z)) |D \nabla_z \phi^n(z)| = \pi_{Z^{n,f}}(z), \quad (2.26)$$

which constitutes a highly nonlinear elliptic PDE for the potential ϕ^n . In the remainder of this section we discuss under which conditions this PDE has a solution. This discussion will also guide us towards a numerical approximation technique that circumvents the need for directly solving (2.26) either analytically or numerically.

Consider the forecast PDF $\pi_{Z^{n,f}}$ and the analysis PDF $\pi_{Z^{n,a}}$ at iteration index n . For simplicity of notion we drop the iteration index and simply write π_{Z^f} and π_{Z^a} , respectively.

Definition 2. A *coupling* of π_{Z^f} and π_{Z^a} consists of a pair $Z^{f:a} = (Z^f, Z^a)$ of random variables such that $Z^f \sim \pi_{Z^f}$, $Z^a \sim \pi_{Z^a}$, and $Z^{f:a} \sim \pi_{Z^{f:a}}$. The joint PDF $\pi_{Z^{f:a}}(z^f, z^a)$ on the product space $\mathbb{R}^{N_z} \times \mathbb{R}^{N_z}$ is called the *transference plan* for this coupling. The set of all transference plans is denoted by $\Pi(\pi_{Z^f}, \pi_{Z^a})$.⁷

Clearly, couplings always exist since one can use the trivial product coupling

$$\pi_{Z^{f:a}}(z^f, z^a) = \pi_{Z^f}(z^f) \pi_{Z^a}(z^a),$$

in which case the associated random variables Z^f and Z^a are independent and each random variable follows its given marginal distribution. Once a coupling has been found, a McKean transition kernel is determined by

$$\pi_{\text{data}}(z|z') = \frac{\pi_{Z^{f:a}}(z', z)}{\int_{\mathcal{Z}^f} \pi_{Z^{f:a}}(z', z'') dz''}.$$

Reversely, any transition kernel $\pi_{\text{data}}(z|z')$, such as (2.20), also induces a coupling.

A diffeomorphism $T : \mathcal{Z}^f \rightarrow \mathcal{Z}^a$ is called a *transport map* if the induced random variable $Z^a = T(Z^f)$ satisfies

$$\int_{\mathcal{Z}^a} f(z^a) \pi_{Z^a}(z^a) dz^a = \int_{\mathcal{Z}^f} f(T(z^f)) \pi_{Z^f}(z^f) dz^f$$

for all suitable functions $f : \mathcal{Z}^a \rightarrow \mathbb{R}$. The associated coupling

$$\pi_{Z^{f:a}}(z^f, z^a) = \delta(z^a - T(z^f)) \pi_{Z^f}(z^f)$$

is called a *deterministic coupling*. Indeed, one finds that

$$\int_{\mathcal{Z}^a} \pi_{Z^{f:a}}(z^f, z^a) dz^a = \pi_{Z^f}(z^f)$$

⁷Couplings should be properly defined in terms of probability measures. A coupling between two measures $\mu_{Z^f}(dz^f)$ and $\mu_{Z^a}(dz^a)$ consists of a pair of random variables with joint measure $\mu_{Z^{f:a}}(dz^f, dz^a)$ such that $\mu_{Z^f}(dz^f) = \int_{\mathcal{Z}^a} \mu_{Z^{f:a}}(dz^f, dz^a)$ and $\mu_{Z^a}(dz^a) = \int_{\mathcal{Z}^f} \mu_{Z^{f:a}}(dz^f, dz^a)$, respectively.

and

$$\pi_{Z^a}(z^a) = \int_{\mathcal{Z}^f} \pi_{Z^f:a}(z^f, z^a) dz^f = \pi_{Z^f}(T^{-1}(z^a)) |DT^{-1}(z^a)|,$$

respectively.

When it comes to actually choosing a particular coupling from the set $\Pi(\pi_{Z^f}, \pi_{Z^a})$ of all admissible ones, it appears preferable to pick the one that maximizes the covariance or correlation between Z^f and Z^a . But maximizing their covariance for given marginals has an important geometric interpretation. Consider, for simplicity, univariate random variables Z^f and Z^a , then

$$\begin{aligned} \mathbb{E}[|Z^f - Z^a|^2] &= \mathbb{E}[|Z^f|^2] + \mathbb{E}[|Z^a|^2] - 2\mathbb{E}[Z^f Z^a] \\ &= \mathbb{E}[|Z^f|^2] + \mathbb{E}[|Z^a|^2] - 2\mathbb{E}[(Z^f - \bar{z}^f)(Z^a - \bar{z}^a)] - 2\bar{z}^f \bar{z}^a \\ &= \mathbb{E}[|Z^f|^2] + \mathbb{E}[|Z^a|^2] - 2\bar{z}^f \bar{z}^a - 2\text{cov}(Z^f, Z^a), \end{aligned}$$

where $\bar{z}^f/a = \mathbb{E}[Z^f/a]$. Hence finding a joint PDF $\pi_{Z^f:a} \in \Pi(\pi_{Z^f}, \pi_{Z^a})$ that minimizes the expectation of $|z^f - z^a|^2$ simultaneously maximizes the covariance between Z^f and Z^a . This geometric interpretation leads to the celebrated Monge-Kantorovich problem.

Definition 3. Let $\Pi(\pi_{Z^f}, \pi_{Z^a})$ denote the set of all possible couplings between π_{Z^f} and π_{Z^a} . A transference plan $\pi_{Z^f:a}^* \in \Pi(\pi_{Z^f}, \pi_{Z^a})$ is called the solution to the *Monge-Kantorovich problem* with cost function $c(z^f, z^a) = \|z^f - z^a\|^2$ if

$$\pi_{Z^f:a}^* = \arg \inf_{\pi_{Z^f:a} \in \Pi(\pi_{Z^f}, \pi_{Z^a})} \mathbb{E}[\|Z^f - Z^a\|^2]. \quad (2.27)$$

The associated functional $W(\pi_{Z^f}, \pi_{Z^a})$, defined by

$$W(\pi_{Z^f}, \pi_{Z^a})^2 = \mathbb{E}[\|Z^f - Z^a\|^2] \quad (2.28)$$

is called the L^2 -Wasserstein distance between π_{Z^f} and π_{Z^a} .

Example 1. Let us consider the discrete set

$$\mathbb{Z} = \{z_1, z_2, \dots, z_M\}, \quad z_i \in \mathbb{R}, \quad (2.29)$$

and two probability vectors $\mathbb{P}(z_i) = 1/M$ and $\mathbb{P}(z_i) = w_i$, respectively, on \mathbb{Z} with $w_i \geq 0$, $i = 1, \dots, M$, and $\sum_i w_i = 1$. Any coupling between these two probability vectors is characterized by a matrix $T \in \mathbb{R}^{M \times M}$ such that its entries $t_{ij} = (T)_{ij}$ satisfy $t_{ij} \geq 0$ and

$$\sum_{i=1}^M t_{ij} = 1/M, \quad \sum_{j=1}^M t_{ij} = w_i. \quad (2.30)$$

These matrices characterize the set of all couplings Π in the definition of the Monge-Kantorovich problem. Given a coupling T and the mean values

$$\bar{z}^f = \frac{1}{M} \sum_{i=1}^M z_i, \quad \bar{z}^a = \sum_{i=1}^M w_i z_i,$$

the covariance between the associated discrete random variables $Z^f : \Omega \rightarrow \mathbb{Z}$ and $Z^a : \Omega \rightarrow \mathbb{Z}$ is defined by

$$\text{cov}(Z^f, Z^a) = \sum_{i,j=1}^M (z_i - \bar{z}^a) t_{ij} (z_j - \bar{z}^f). \quad (2.31)$$

The particular coupling defined by $t_{ij} = w_i/M$ leads to zero correlation between Z^f and Z^a . On the other hand, maximizing the correlation leads to a *linear transport problem* in the M^2 unknowns $\{t_{ij}\}$. More precisely, the unknowns t_{ij} have to satisfy the inequality constraints $t_{ij} \geq 0$, the equality constraints (2.30), and should minimize

$$J(T) = \sum_{i,j=1}^M t_{ij} |z_i - z_j|^2,$$

which is equivalent to maximizing (2.31). See Strang (1986) and Nocedal and Wright (2006) for an introduction to linear transport problems and, more generally, to linear programming.

We now return to continuous random variables and the desired coupling between forecast and analysis PDFs. The following theorem is an adaptation of a more general result on optimal couplings from Villani (2003).

Theorem 2. *If the forecast PDF π_{Z^f} has bounded second-order moments, then the optimal transference plan that solves the Monge-Kantorovich problem gives rise to a deterministic coupling with transport map*

$$Z^a = \nabla_z \phi(Z^f),$$

where $\phi : \mathbb{R}^{N_z} \rightarrow \mathbb{R}$ is a convex potential.

Below we sketch the basic line of arguments that lead to this theorem. We first introduce the *support* of a coupling $\pi_{Z^f:a}$ on $\mathbb{R}^{N_z} \times \mathbb{R}^{N_z}$ as the smallest closed set on which $\pi_{Z^f:a}$ is concentrated, i.e.,

$$\text{supp}(\pi_{Z^f:a}) := \bigcap \{S \subset \mathbb{R}^{N_z} \times \mathbb{R}^{N_z} : S \text{ closed and } \mu_{Z^f:a}(\mathbb{R}^{N_z} \times \mathbb{R}^{N_z} \setminus S) = 0\}$$

with the measure of $\mathbb{R}^{N_z} \times \mathbb{R}^{N_z} \setminus S$ defined by

$$\mu_{Z^f:a}(\mathbb{R}^{N_z} \times \mathbb{R}^{N_z} \setminus S) = \int_{\mathbb{R}^{N_z} \times \mathbb{R}^{N_z} \setminus S} \pi_{Z^f:a}(z^f, z^a) \, dz^f \, dz^a.$$

The support of $\pi_{Z^f:a}$ is called *cyclically monotone* if for every set of points $(z_i^f, z_i^a) \in \text{supp}(\pi_{Z^f:a}) \subset \mathbb{R}^{N_z} \times \mathbb{R}^{N_z}$, $i = 1, \dots, I$, and any permutation σ of $\{1, \dots, I\}$ one has

$$\sum_{i=1}^I \|z_i^f - z_i^a\|^2 \leq \sum_{i=1}^I \|z_i^f - z_{\sigma(i)}^a\|^2. \quad (2.32)$$

Note that (2.32) is equivalent to

$$\sum_{i=1}^I (z_i^f)^T (z_{\sigma(i)}^a - z_i^a) \leq 0.$$

It can be shown that any coupling whose support is not cyclically monotone can be modified into another coupling with lower transport cost. Hence it follows that a solution $\pi_{Z^f:a}^*$ of the Monge-Kantorovich problem (2.27) has cyclically monotone support.

A fundamental theorem (*Rockafellar's theorem* (Villani, 2003)) of convex analysis now states that cyclically monotone sets $S \subset \mathbb{R}^{N_z} \times \mathbb{R}^{N_z}$ are contained in the *subdifferential* of a convex function $\phi : \mathbb{R}^{N_z} \rightarrow \mathbb{R}$. Here the subdifferential $\partial\phi$ of a convex function ϕ at a point $z \in \mathbb{R}^{N_z}$ is defined as the compact, non-empty and convex set of all $m \in \mathbb{R}^{N_z}$ such that

$$\phi(z') \geq \phi(z) + m(z' - z)$$

for all $z' \in \mathbb{R}^{N_z}$. We write $m \in \partial\phi(z)$. An optimal transport map is obtained whenever the convex potential ϕ for a given optimal coupling $\pi_{Z^f:a}^*$ is sufficiently regular in which case the subdifferential $\partial\phi(z)$ reduces to the classic gradient $\nabla_z\phi$ and $z^a = \nabla_z\phi(z^f)$. This regularity is ensured by the assumptions of the above theorem. See McCann (1995) and Villani (2003) for more details.

We summarize the *McKean optimal transportation approach* in the following definition.

Definition 4. Given a dynamic iteration (2.6) with PDF π_{Z^0} for the initial conditions and observations y_{obs}^n , $n = 1, \dots, N$, the forecast $Z^{n,f}$ at iteration index $n > 0$ is defined by (2.22) and the analysis $Z^{n,a}$ by (2.25). The convex potential ϕ^n is the solution to the Monge-Kantorovich optimal transportation problem for coupling $\pi_{Z^{f,n}}$ and $\pi_{Z^{a,n}}$. The iteration is started at $n = 0$ with $Z^{0,a} = Z^0$.

The application of optimal transportation to Bayesian inference and data assimilation has first been discussed by Reich (2011), Moselhy and Marzouk (2012), and Reich and Cotter (2013).

In the following section we discuss data assimilation algorithms from a McKean optimal transportation perspective.

5 Linear ensemble transform methods

In this section, we discuss SMCs, EnKFs, and the recently proposed (Reich, 2013a) ETPF from a coupling perspective. All three data assimilation methods have in common that they are based on an ensemble z_i^n , $i = 1, \dots, M$, of model states. In the absence of observations the M ensemble members propagate independently under the model dynamics (2.6), *i.e.*, an analysis ensemble at time-level $n - 1$ gives rise to a forecast ensemble at time-level n via

$$z_i^{nf} = \Psi(z_i^{n-1,a}), \quad i = 1, \dots, M.$$

The three methods differ in the way the forecast ensemble $\{z_i^{nf}\}_{i=1}^M$ is transformed into an analysis ensemble $\{z_i^{n,a}\}_{i=1}^M$ under an observation y_{obs}^n . However, all three methods share a common mathematical structure which we outline next. We drop the iteration index in order to simplify the notation.

Definition 5. The class of *linear ensemble transform filters* (LETs) is defined by

$$z_j^a = \sum_{i=1}^M z_i^f s_{ij}, \quad (2.33)$$

where the coefficients s_{ij} are the M^2 entries of a matrix $S \in \mathbb{R}^{M \times M}$.

The concept of LETs is well established for EnKF formulations (Tippett et al., 2003; Evensen, 2006). It will be shown below that SMCs and the ETPF also belong to the class of LETs. In other words, these three methods differ only in the definition of the corresponding transform matrix S .

5.1 Sequential Monte Carlo methods (SMCMs)

A central building block of an SMCM is the *proposal density* $\pi_{\text{prop}}(z|z')$, which produces a *proposal ensemble* $\{z_i^p\}_{i=1}^M$ from the last analysis ensemble. Here we assume, for simplicity, that the proposal density is given by the model dynamics itself, *i.e.*,

$$\pi_{\text{prop}}(z|z') = \delta(z - \Psi(z')),$$

and then

$$z_i^p = z_i^f, \quad i = 1, \dots, M.$$

One associates with the proposal/forecast ensemble two discrete measures on

$$\mathbb{Z} = \{z_1^f, z_2^f, \dots, z_M^f\}, \quad (2.34)$$

namely the uniform measure $\mathbb{P}(z_i^f) = 1/M$ and the non-uniform measure

$$\mathbb{P}(z_i^f) = w_i,$$

defined by the importance weights

$$w_i = \frac{\pi_Y(y_{\text{obs}}|z_i^f)}{\sum_{j=1}^M \pi_Y(y_{\text{obs}}|z_j^f)}. \quad (2.35)$$

The *sequential importance resampling* (SIR) filter (Gordon et al., 1993) resamples from the weighted forecast ensemble in order to produce a new equally weighted analysis ensemble $\{z_i^a\}$. Here we only consider SIR filter implementations with resampling performed after each data assimilation cycle.

An in-depth discussion of the SIR filter and more general SMCs can be found, for example, in Doucet et al. (2001); Doucet and Johansen (2011). Here we focus on the coupling of discrete measures perspective of a resampling step. We first note that any resampling strategy effectively leads to a coupling of the uniform and the non-uniform measure on (2.34). As previously discussed, a coupling is defined by a matrix $T \in \mathbb{R}^{M \times M}$ such that $t_{ij} \geq 0$, and

$$\sum_{i=1}^M t_{ij} = 1/M, \quad \sum_{j=1}^M t_{ij} = w_i. \quad (2.36)$$

The resampling strategy (2.20) leads to

$$t_{ij} = \frac{1}{M} (\epsilon w_j \delta_{ij} + (1 - \epsilon w_j) w_i)$$

with $\epsilon \geq 0$ chosen such that $\epsilon w_j \leq 1$ for all $j = 1, \dots, M$. Monomial resampling corresponds to the special case $\epsilon = 0$, i.e. $t_{ij} = w_i/M$. The associated transformation matrix S in (2.33) is the realization of a random matrix with entries $s_{ij} \in \{0, 1\}$ such that each column of S contains exactly one entry $s_{ij} = 1$. Given a coupling T , the probability for the s_{ij} entry to be the one selected in the j th column is

$$\mathbb{P}(s_{ij} = 1) = M t_{ij}$$

and $M t_{ij} = w_i$ in case of monomial resampling. Any such resampling procedure based on a coupling matrix T leads to a consistent coupling for the underlying forecast and analysis probability measures as $M \rightarrow \infty$, which, however, is non-optimal in the sense of the Monge-Kantorovich problem (2.27). We refer to Bain and Crisan (2009) for resampling strategies which satisfy alternative optimality conditions.

We emphasize that the transformation matrix S of a SIR particle filter analysis step satisfies

$$\sum_{i=1}^M s_{ij} = 1 \quad (2.37)$$

and

$$s_{ij} \in [0, 1]. \quad (2.38)$$

In other words, each realization of the resampling step yields a Markov chain S . Furthermore, the weights $\hat{w}_i = M^{-1} \sum_{j=1}^M s_{ij}$ satisfy $\mathbb{E}[\hat{W}_i] = w_i$ and the analysis ensemble defined by $z_j^a = z_i^f$ if $s_{ij} = 1, j = 1, \dots, M$, is contained in the convex hull of the forecast ensemble (2.34).

A forecast ensemble $\{z_i^f\}_{i=1}^M$ leads to the following estimator

$$\bar{z}^f = \frac{1}{M} \sum_{i=1}^M z_i^f$$

for the mean and

$$P_{zz}^f = \frac{1}{M-1} \sum_{i=1}^M (z_i^f - \bar{z}^f)(z_i^f - \bar{z}^f)^T$$

for the covariance matrix. In order to increase the *robustness* of a SIR particle filter one often augments the resampling step by the *particle rejuvenation step* (Pham, 2001)

$$z_j^a = z_i^f + \xi_j, \quad (2.39)$$

where the ξ_j 's are realizations of M independent and identically distributed Gaussian random variables $N(0, h^2 P_{zz}^f)$ and $s_{ij} = 1$. Here $h > 0$ is the *rejuvenation parameter* which determines the magnitude of the stochastic perturbations. Rejuvenation helps to avoid the creation of identical analysis ensemble members which would remain identical under the deterministic model dynamics (2.6) for all times. Furthermore, rejuvenation can be used as a heuristic tool in order to compensate for model errors as encoded, for example, in the difference between (2.6) and (2.13).

In this paper we only discuss SMCMs which are based on the proposal step (2.4). Alternative proposal steps are possible and recent work on alternative implementations of SMCMs include van Leeuwen (2010), Chorin et al. (2010), Morzfeld et al. (2012), Morzfeld and Chorin (2012), van Leeuwen and Ades (2012), Reich (2013b).

5.2 Ensemble Kalman filter (EnKF)

The historically first version of the EnKF uses perturbed observations in order to transform a forecast ensemble into an analysis ensemble. The key requirement of any EnKF is that the transformation step is consistent with the classic Kalman update step in case the forecast and analysis PDFs are Gaussian. The, so-called, *EnKF with perturbed observations* is explicitly given by the simple formula

$$z_j^a = z_j^f - K(y_j^f + \xi_j - y_{\text{obs}}), \quad j = 1, \dots, M,$$

where $y_j^f = h(z_j^f)$, the ξ_j 's are realizations of independent and identically distributed Gaussian random variables with mean zero and covariance matrix R , and K denotes the Kalman gain matrix, which in case of the EnKF is determined by the forecast ensemble as follows:

$$K = P_{zy}^f (P_{yy}^f + R)^{-1}$$

with empirical covariance matrices

$$P_{zy}^f = \frac{1}{M-1} \sum_{i=1}^M (z_i^f - \bar{z}^f)(y_i^f - \bar{y}^f)^T$$

and

$$P_{yy}^f = \frac{1}{M-1} \sum_{i=1}^M (y_i^f - \bar{y}^f)(y_i^f - \bar{y}^f)^T,$$

respectively. Here the ensemble mean in observation space is defined by

$$\bar{y}^f = \frac{1}{M} \sum_{i=1}^M y_i^f.$$

In order to shorten subsequent notations, we introduce the $N_y \times M$ matrix of ensemble deviations

$$A_y^f = [y_1^f - \bar{y}^f, y_2^f - \bar{y}^f, \dots, y_M^f - \bar{y}^f]$$

in observation space and the $N_z \times M$ matrix of ensemble deviations

$$A_z^f = [z_1^f - \bar{z}^f, z_2^f - \bar{z}^f, \dots, z_M^f - \bar{z}^f]$$

in state space, respectively. In terms of these ensemble deviation matrices, it holds that

$$P_{zz}^f = \frac{1}{M-1} A_z^f (A_z^f)^T \quad \text{and} \quad P_{zy}^f = \frac{1}{M-1} A_z^f (A_y^f)^T,$$

respectively.

It can be verified by explicit calculations that the EnKF with perturbed observations fits into the class of LETFs with

$$z_j^a = \sum_{i=1}^M z_i^f \left(\delta_{ij} - \frac{1}{M-1} (y_i^f - \bar{y}_i^f)^T (P_{yy}^f + R)^{-1} (y_j^f + \xi_j - y_{\text{obs}}) \right)$$

and, therefore,

$$s_{ij} = \delta_{ij} - \frac{1}{M-1} (y_i^f - \bar{y}_i^f)^T (P_{yy}^f + R)^{-1} (y_j^f + \xi_j - y_{\text{obs}}).$$

Here we have used that

$$\frac{1}{M-1} \sum_{i=1}^M (z_i^f - \bar{z}^f) (y_i^f - \bar{y}^f)^T = \frac{1}{M-1} \sum_{i=1}^M z_i^f (y_i^f - \bar{y}^f)^T.$$

We note that the transform matrix S is the realization of a random matrix. The class of *ensemble square root filters* (ESRF) leads instead to deterministic transformation matrices S . More precisely, an ESRF uses separate transformation steps for the ensemble mean \bar{z}^f and the ensemble deviations $z_i^f - \bar{z}^f$. The mean is simply updated according to the classic Kalman formula, *i.e.*

$$\bar{z}^a = \bar{z}^f - K(\bar{y}^f - y_{\text{obs}}) \quad (2.40)$$

with the Kalman gain matrix defined as before.

Upon introducing the analysis matrix of ensemble deviations $A_z^a \in \mathbb{R}^{N_z \times M}$, one obtains

$$\begin{aligned} P_{zz}^a &= \frac{1}{M-1} A_z^a (A_z^a)^T \\ &= P_{zz}^f - K(P_{zy}^f)^T = \frac{1}{M-1} A_z^f Q (A_z^f)^T \end{aligned} \quad (2.41)$$

with the $M \times M$ matrix Q defined by

$$Q = I - \frac{1}{M-1} (A_y^f)^T (P_{yy}^f + R)^{-1} A_y^f.$$

Let us denote the matrix square root⁸ of Q by D and its entries by d_{ij} .

⁸The matrix square root of a symmetric positive semi-definite matrix Q is the unique symmetric matrix D which satisfies $DD = Q$.

We note that $\sum_{i=1}^M d_{ij} = 1$ and it follows that

$$\begin{aligned}
 z_j^a &= \bar{z}^f - K(\bar{y}^f - y_{\text{obs}}) + \sum_{i=1}^M (z_i^f - \bar{z}^f) d_{ij} \\
 &= \sum_{i=1}^M z_i^f \left(\frac{1}{M-1} (y_i^f - \bar{y}^f)^T (P_{yy}^f + R)^{-1} (y_{\text{obs}} - \bar{y}^f) + d_{ij} \right) \\
 &= \sum_{i=1}^M z_i^f \left(\frac{1}{M-1} (y_i^f - \bar{y}^f)^T (P_{yy}^f + R)^{-1} (y_{\text{obs}} - \bar{y}^f) + d_{ij} \right). \tag{2.42}
 \end{aligned}$$

The appropriate entries for the transformation matrix S of an ESRF can now be read off of (2.42). See Tippett et al. (2003); Wang et al. (2004); Livings et al. (2008); Ott et al. (2004); Nerger et al. (2012) and Evensen (2006) for further details on other ESRF implementations such as the ensemble adjustment Kalman filter. We mention in particular that an application of the Sherman-Morrison-Woodbury formula (Golub and van Loan, 1996) leads to the equivalent square root formula

$$D = \left\{ I + \frac{1}{M-1} (A_y^f)^T R^{-1} A_y^f \right\}^{-1/2}, \tag{2.43}$$

which avoids the need for inverting the $N_y \times N_y$ matrix $P_{yy}^f + R$, which is desirable whenever $N_y \gg M$. Furthermore, using the equivalent Kalman gain matrix representation

$$K = P_{zy}^a R^{-1},$$

the Kalman update formula (2.40) for the mean becomes

$$\begin{aligned}
 \bar{z}^a &= \bar{z}^f - P_{zy}^a R^{-1} (\bar{y}^f - y_{\text{obs}}) \\
 &= \bar{z}^f + \frac{1}{M-1} A_z^f Q (A_y^f)^T R^{-1} (y_{\text{obs}} - \bar{y}^f).
 \end{aligned}$$

This reformulation gives rise to

$$s_{ij} = \frac{1}{M-1} q_{ij} (y_j^f - \bar{y}^f) R^{-1} (y_{\text{obs}} - \bar{y}^f) + d_{ij}, \tag{2.44}$$

which forms the basis of the *local ensemble transform Kalman filter* (LETKF) (Ott et al., 2004; Hunt et al., 2007) to be discussed in more detail in Section 6.

We mention that the EnKF with perturbed observations or an ESRF implementation leads to transformation matrices S which satisfy (2.37) but the entries s_{ij} can take positive as well as negative values. This can be problematic in case the

state variable z should be non-negative. Then it is possible that a forecast ensemble $z_i^f \geq 0$, $i = 1, \dots, M$, is transformed into an analysis z_i^a , which contains negative entries. See Janjić et al. (2014) for modifications to EnKF type algorithms in order to preserve positivity.

One can discuss the various EnKF formulations from an optimal transportation perspective. Here the coupling is between two Gaussian distributions; the forecast PDF $N(\bar{z}^f, P_{zz}^f)$ and analysis PDF $N(\bar{z}^a, P_{zz}^a)$, respectively, with the analysis mean given by (2.40) and the analysis covariance matrix by (2.41). We know that the optimal coupling must be of the form

$$z^a = \nabla_z \phi(z^f)$$

and, in case of Gaussian PDFs, the convex potential $\phi : \mathbb{R}^{N_z} \rightarrow \mathbb{R}$ is furthermore bilinear, *i.e.*,

$$\phi(z) = b^T z + \frac{1}{2} z^T A z$$

with the vector b and the matrix A appropriately defined. The choice

$$z^a = b + A z^f = \bar{z}^a + A(z^f - \bar{z}^f)$$

leads to

$$b = \bar{z}^a - A \bar{z}^f$$

for the vector $b \in \mathbb{R}^{N_z}$. The matrix $A \in \mathbb{R}^{N_z \times N_z}$ then needs to satisfy

$$P_{zz}^a = A P_{zz}^f A^T.$$

The optimal, in the sense of Monge-Kantorovich with cost function $c(z^f, z^a) = \|z^f - z^a\|^2$, matrix A is given by

$$A = (P_{zz}^a)^{1/2} [(P_{zz}^a)^{1/2} P_{zz}^f (P_{zz}^a)^{1/2}]^{-1/2} (P_{zz}^a)^{1/2}.$$

See Olkin and Pukelsheim (1982). An efficient implementation of this optimal coupling in the context of ESRFs has been discussed in Reich and Cotter (2013). The essential idea is to replace the matrix square root of P_{zz}^a by the analysis matrix of ensemble deviations $A_z^a = A_z^f D$ scaled by $1/\sqrt{M-1}$.

Note that different cost functions $c(z^f, z^a)$ lead to different solutions to the associated Monge-Kantorovich problem (2.27). Of particular interest is the weighted inner product

$$c(z^f, z^a) = ((z^f - z^a)^T B^{-1} (z^f - z^a))^2$$

for an appropriate positive definite matrix $B \in \mathbb{R}^{N_z \times N_z}$ (Reich and Cotter, 2013).

As for SMCs particle rejuvenation can be applied to the analysis from an EnKF or ESRF. However, the more popular method for increasing the robustness of EnKFs is to apply *multiplicative ensemble inflation*

$$z_i^f \rightarrow \bar{z}^f + \alpha(z_i^f - \bar{z}^f), \quad \alpha \geq 1, \quad (2.45)$$

to the forecast ensemble prior to the application of an EnKF or ESRF. The parameter α is called the inflation factor. An adaptive strategy for determining the factor α has, for example, been proposed by Anderson (2007); Miyoshi (2011). The inflation factor α can formally be related to the rejuvenation parameter h in (2.39) through

$$\alpha = \sqrt{1 + h^2}.$$

This relation becomes exact as $M \rightarrow \infty$.

We mention that the *rank histogram filter* of Anderson (2010), which uses a nonlinear filter in observation space and linear regression from observation onto state space, also fits into the framework of the LETFs. See Reich and Cotter (2015) for more details. The *nonlinear ensemble adjustment filter* of Lei and Bickel (2011), on the other hand, falls outside the class of LETFs.

5.3 Ensemble transform particle filter (ETPF)

We now return to the SIR filter described in Section 5.1. Recall that a SIR filter relies on importance resampling which we have interpreted as a coupling between the uniform measure on (2.34) and the measure defined by (2.35). Any coupling is characterized by a matrix T such that its entries are non-negative and (2.36) hold.

Definition 6. The ETPF is based on choosing the T which minimizes

$$J(T) = \sum_{i,j=1}^M t_{ij} \|z_i^f - z_j^f\|^2 \quad (2.46)$$

subject to (2.36) and $t_{ij} \geq 0$. Let us denote the minimizer by T^* . Then the transform matrix S of an ETPF is defined by

$$S = MT^*,$$

which satisfies (2.37) and (2.38).

Let us give a geometric interpretation of the ETPF transformation step. Since T^* from Definition 6 provides an optimal coupling, Rockafellar's theorem implies the existence of a convex potential $\phi_M : \mathbb{R}^{N_z} \rightarrow \mathbb{R}$ such that

$$z_i^f \in \partial \phi_M(z_j^f) \quad \text{for all } i \in \mathcal{I}_j := \{i' \in \{1, \dots, M\} : t_{i'j}^* > 0\},$$

$j = 1, \dots, M$. In fact, ϕ_M can be chosen to be piecewise affine and a constructive formula can be found in Villani (2003). The ETPF transformation step

$$z_j^a = M \sum_{i=1}^M z_i^f t_{ij}^* = \sum_{i=1}^M z_i^f s_{ij} \quad (2.47)$$

corresponds to a particular selection from the linear space $\partial\phi_M(z_j^f)$, $j = 1, \dots, M$; namely the expectation value of the discrete random variable

$$Z_j^a : \Omega \rightarrow \{z_1^f, z_2^f, \dots, z_M^f\}$$

with probabilities $\mathbb{P}(z_i^f) = s_{ij}$, $i = 1, \dots, M$. Hence it holds that

$$\bar{z}^a := \frac{1}{M} \sum_{j=1}^M z_j^a = \sum_{i=1}^M w_i z_i^f.$$

See Reich (2013a) for more details, where it has also been shown that the potentials ϕ_M converge to the solution of the underlying continuous Monge-Kantorovich problem as the ensemble size M approaches infinity.

It should be noted that standard algorithms for finding the minimizer of (2.46) suffer from a $\mathcal{O}(M^3 \log M)$ computational complexity. This complexity has been reduced to $\mathcal{O}(M^2 \log M)$ by Pele and Werman (2009). There are also fast iterative methods for finding approximate minimizers of (2.46) using the *Sinkhorn distance* (Cuturi, 2013).

The particle rejuvenation step (2.39) for SMCs can be extended to the ETPF as follows:

$$z_j^a = \sum_{i=1}^M z_i^f s_{ij} + \xi_j, \quad j = 1, \dots, M. \quad (2.48)$$

As before the ξ_j 's are realizations of M independent Gaussian random variables with mean zero and appropriate covariance matrices P_j^a . We use $P_j^a = h^2 P_{zz}^f$ with rejuvenation parameter $h > 0$ for the numerical experiments conducted in this paper. Another possibility would be to locally estimate P_j^a from the coupling matrix T^* , i.e.,

$$P_j^a = \sum_{i=1}^M s_{ij} (z_i^f - \bar{z}_j^a)(z_i^f - \bar{z}_j^a)^T$$

with mean $\bar{z}_j^a = \sum_{i=1}^M s_{ij} z_i^f$.

5.4 Quasi-Monte Carlo (QMC) convergence

The expected rate of convergence for standard Monte Carlo methods is $M^{-1/2}$ where M denotes the ensemble size. QMC methods have an upper bound of $\log(M)^d M^{-1}$ where d stands for the dimension (Caflish, 1988). For the purpose of this paper, $d = N_z$. Unlike Monte Carlo methods QMC methods also depend on the dimension of the space which implies a better performance for small N_z or/and large M . However, in practice QMC methods perform considerably better than the theoretical bound for the convergence rate and outperform Monte Carlo methods even for small ensemble sizes and in very high dimensional models. The latter may be explained by the concept of *effective dimension* introduced by Caflish et al. (1997).

The following simulation investigates the convergence rate of the estimators for the first and second moment of the posterior distribution after applying a single analysis step of a SIR particle filter and an ETPF. The prior is chosen to be a uniform distribution on the unit square and the sum of both components is observed with additive noise drawn from a centered Gaussian distribution with variance equal to two. Reference values for the posterior moments are generated using Monte Carlo importance sampling with sample size $M = 2^{26}$. QMC samples of different sizes are drawn from the prior distribution and a single residual resampling step is compared to a single transformation step using an optimal coupling T^* . Figure 1 shows the root mean square errors (RMSEs) of the different posterior estimates with respect to their reference values. We find that the transform method preserves the optimal M^{-1} convergence rate of the prior QMC samples while resampling reduces the convergence rate to the $M^{-1/2}$.

We mention that replacing the deterministic transformation step in (2.47) by drawing ensemble member j from the prior ensemble according to the weights given by the j -th column of S leads to a stochastic version of the ETPF. This variant, despite being stochastic like the importance resampling step, results again in a QMC convergence rate.

6 Spatially extended dynamical systems and localization

Let us start this section with a simple thought experiment on the curse of dimensionality. Consider a state space of dimension $N_z = 100$ and a prior Gaussian distribution with mean zero and covariance matrix $P^f = I$. The reference solution is $z_{\text{ref}} = 0$ and we observe every component of the state vector subject to independent measurement errors with mean zero and variance $R = 0.16$. If one applies a single importance resampling step to this problem with ensemble size $M = 10$, one finds that the effective sample size collapses to $M_{\text{eff}} \approx 1$ and the resulting analysis ensemble is unable to recover the reference solution. However, one also quickly realizes that the stated problem can be decomposed into N_z independent data assimilation problems in each component of the state vector alone. If importance

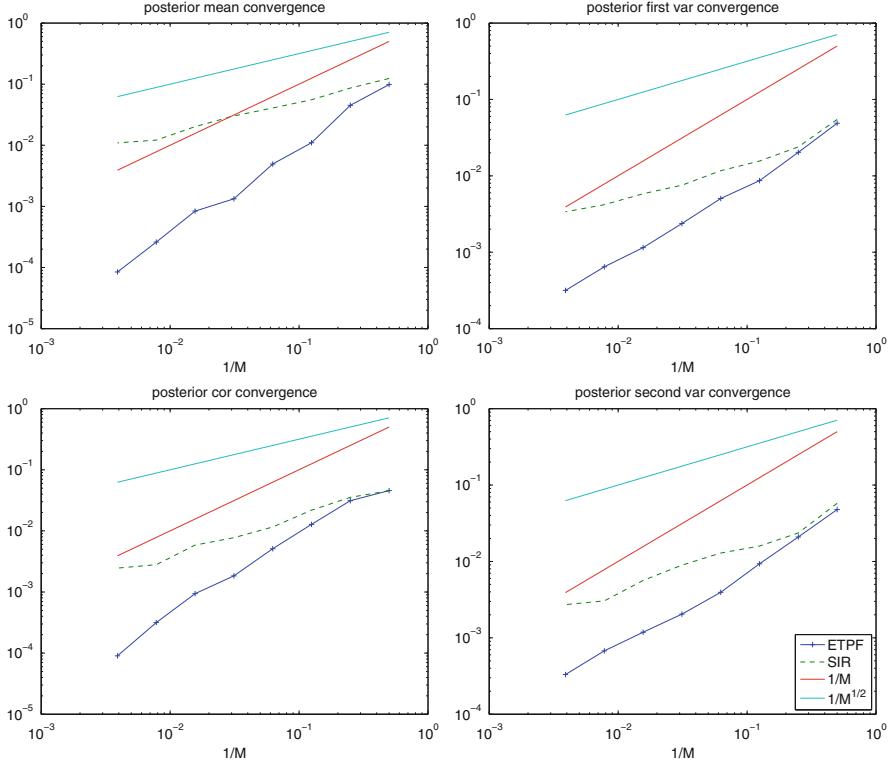


Fig. 1 RMSEs of estimates for the posterior mean, variances (var), and correlation (cor) using importance resampling (SIR) and optimal transformations (ETPF) plotted on a log-log scale as a function of ensemble sizes M .

resampling is now performed in each component of the state vector independently, then the effective sample size for each of the N_z analysis problems remains close to $M = 10$ and the reference solution can be recovered from the given set of observations. This is the idea of *localization*. Note that localization has increased the total sample size to $M \times N_z = 1000$ for this problem!

We now formally extend LETFs to spatially extended systems which may be viewed as an infinite-dimensional dynamical system (Robinson, 2001) and formulate an appropriate localization strategy. Consider the linear advection equation

$$u_t + u_x = 0$$

as a simple example for such a scenario. If $u_0(x)$ denotes the solution at time $t = 0$, then

$$u(x, t) = u_0(x + t)$$

is the solution of the linear advection equation for all $t \geq 0$. Given a time-increment $\Delta t > 0$, the associated dynamical system maps a function $u(x)$ into $u(x + \Delta t)$. A finite-dimensional dynamical system is obtained by discretizing in space with mesh-size $\Delta x > 0$. For example, the Box scheme (Morton and Mayers, 2005) leads to

$$\frac{u_j^{k+1} + u_{j+1}^{k+1} - u_j^k - u_{j+1}^k}{2\Delta t} + \frac{u_{j+1}^{k+1} + u_{j+1}^k - u_j^{k+1} - u_j^k}{2\Delta x} = 0$$

and, for J spatial grid points, the state vector at $t_k = k\Delta t$ becomes

$$z^k = (u_1^k, u_2^k, \dots, u_J^k)^T \in \mathbb{R}^J.$$

We may take the formal limit $J \rightarrow \infty$ and $\Delta x \rightarrow 0$ in order to return to functions $z^k(x)$. The dynamical system (2.6) is then defined as the map that propagates such functions (or their finite-difference approximations) from one observation instance to the next in accordance with the specified numerical method. Here we assume that observations are taken in intervals of $\Delta t_{\text{obs}} = N_{\text{out}}\Delta t$ with $N_{\text{out}} \geq 1$ a fixed integer. The index $n \geq 1$ in (2.6) is the counter for those observation instances.

In other words, forecast or analysis ensemble members, $z^{f/a}(x)$, now become functions of $x \in \mathbb{R}$, belong to some appropriate function space \mathcal{H} , and the dynamical system (2.6) is formally replaced by a map or evolution equation on \mathcal{H} (Robinson, 2001). For simplicity of exposition we assume periodic boundary conditions, *i.e.*, $z(x) = z(x + L)$ for some appropriate $L > 0$.

The curse of dimensionality (Bengtsson et al., 2008) implies that, generally speaking, none of the LETFs discussed so far is suitable for data assimilation of spatially extended systems. In order to overcome this situation, we now discuss the concept of *localization* as first introduced in Houtekamer and Mitchell (2001, 2005) for EnKFs. While we will focus on a particular localization, called *R-localization*, suggested by Hunt et al. (2007), our methodology can be extended to *B-localization* as proposed by Hamill et al. (2001).

In the context of the LETFs R-localization amounts to modifying (2.33) to

$$z_j^a(x) = \sum_{i=1}^M z_i^f(x) s_{ij}(x),$$

where the associated transform matrices $S(x) \in \mathbb{R}^M \times \mathbb{R}^M$ depend now on the spatial location $x \in [0, L]$. It is crucial that the transformation matrices $S(x)$ are sufficiently smooth in x in order to produce analysis ensembles with sufficient regularity for the evolution problem under consideration and, in particular $z_j^a \in \mathcal{H}$. In case of an SMC with importance resampling, the resulting $S(x)$ would, in general, not even be continuous for almost all $x \in [0, L]$. Hence we only discuss localization for the ESRF and the ETPF.

Let us, for simplicity, assume that the forward operator $h : \mathcal{H} \rightarrow \mathbb{R}^{N_y}$ for the observations y_{obs} is defined by

$$h_k(z) = z(x_k), \quad k = 1, \dots, N_y.$$

Here the $x_k \in [0, L)$ denote the spatial location at which the observation is taken. The measurement errors are Gaussian with mean zero and covariance matrix $R \in \mathbb{R}^{N_y \times N_y}$. We assume for simplicity that R is diagonal.

In the sequel we assume that $z(x)$ has been extended to $x \in \mathbb{R}$ by periodic extension from $x \in [0, L)$ and introduce the time-averaged and normalized *spatial correlation function*

$$C(x, s) := \frac{\sum_{n=0}^N z^n(x+s) z^n(x)}{\sum_{n=0}^N (z^n(x))^2} \quad (2.49)$$

for $x \in [0, L)$ and $s \in [-L/2, L/2)$. Here we have assumed that the underlying solution process is stationary ergodic. In case of spatial homogeneity the spatial correlation function becomes furthermore independent of x for N sufficiently large.

We also introduce a localization kernel $\mathcal{K}(x, x'; r_{\text{loc}})$ in order to define R -localization for an ESRF and the ETPF. The localization kernel can be as simple as

$$\mathcal{K}(x, x'; r_{\text{loc}}) = \begin{cases} 1 - \frac{1}{2}s & \text{for } s \leq 2, \\ 0 & \text{else,} \end{cases} \quad (2.50)$$

with

$$s := \frac{\min\{|x - x' - L|, |x - x'|, |x - x' + L|\}}{r_{\text{loc}}} \geq 0,$$

or a higher-order polynomial such as

$$\mathcal{K}(x, x'; r_{\text{loc}}) = \begin{cases} 1 - \frac{5}{3}s^2 + \frac{5}{8}s^3 + \frac{1}{2}s^4 - \frac{1}{4}s^5 & \text{for } s \leq 1, \\ -\frac{2}{3}s^{-1} + 4 - 5s + \frac{5}{3}s^2 + \frac{5}{8}s^3 - \frac{1}{2}s^4 + \frac{1}{12}s^5 & \text{for } 1 \leq s \leq 2, \\ 0 & \text{else.} \end{cases} \quad (2.51)$$

See Gaspari and Cohn (1999).

In order to compute the transformation matrix $S(x)$ for given x , we modify the k th diagonal entry in the measurement error covariance matrix $R \in \mathbb{R}^{N_y \times N_y}$ and define

$$\frac{1}{\tilde{r}_{kk}(x)} := \frac{\mathcal{K}(x, x_k; r_{\text{loc}, R})}{r_{kk}} \quad (2.52)$$

for $k = 1, \dots, N_y$. Given a localization radius $r_{\text{loc}, R} > 0$, this results in a matrix $\tilde{R}^{-1}(x)$ which replaces the R^{-1} in an ESRF and the ETPF.

More specifically, the LETKF is based on the following modifications to the ESRF. First one replaces (2.43) by

$$Q(x) = \left\{ I + \frac{1}{M-1} (A_y^f)^T \tilde{R}^{-1}(x) A_y^f \right\}^{-1}$$

and defines $D(x) = Q(x)^{1/2}$. Finally the localized transformation matrix $S(x)$ is given by

$$s_{ij}(x) = \frac{1}{M-1} q_{ij}(x) (\mathbf{y}_j^f - \bar{\mathbf{y}}^f) \tilde{R}^{-1}(x) (\mathbf{y}_{\text{obs}} - \bar{\mathbf{y}}^f) + d_{ij}(x), \quad (2.53)$$

which replaces (2.44). We mention that Anderson (2012) discusses practical methods for choosing the localization radius $r_{\text{loc},R}$ for EnKFs.

In order to extend the concept of R -localization to the ETPF, we also define the localized cost function

$$c_x(\mathbf{z}^f, \mathbf{z}^a) = \int_0^L \mathcal{K}(x, x'; r_{\text{loc},c}) \|\mathbf{z}^f(x') - \mathbf{z}^a(x')\|^2 dx' \quad (2.54)$$

with a localization radius $r_{\text{loc},c} \geq 0$, which can be chosen independently from the localization radius for the measurement error covariance matrix R .

The ETPF with R -localization can now be implemented as follows. At each spatial location $x \in [0, L)$ one determines the desired transformation matrix $S(x)$ by first computing the weights

$$w_i \propto e^{-\frac{1}{2} (h(\mathbf{z}_i^f) - \mathbf{y}_{\text{obs}})^T \tilde{R}^{-1}(x) (h(\mathbf{z}_i^f) - \mathbf{y}_{\text{obs}})} \quad (2.55)$$

and then minimizing the cost function

$$J(T) = \sum_{i,j=1}^M c_x(\mathbf{z}_i^f, \mathbf{z}_j^f) t_{ij} \quad (2.56)$$

over all admissible couplings. One finally sets $S(x) = MT^*$.

As discussed earlier any infinite-dimensional evolution equation such as the linear advection equation will be truncated in practice to a computational grid $x_j = j\Delta x$. The transform matrices $S(x)$ need then to be computed for each grid point only and the integral in (2.54) is replaced by a simple Riemann sum.

We mention that alternative filtering strategies for spatio-temporal processes have been proposed by Majda and Harlim (2012) in the context of turbulent systems. One of their strategies is to perform localization in spectral space in case of regularly spaced observations. Another spatial localization strategy for particle filters can be found in Rebeschini and van Handel (2013).

7 Applications

In this section we present some numerical results comparing the different LETFs for the chaotic Lorenz-63 (Lorenz, 1963) and Lorenz-96 (Lorenz, 1996) models. While the highly nonlinear Lorenz-63 model can be used to investigate the behavior

of different DA algorithms for strongly non-Gaussian distributions, the forty dimensional Lorenz-96 model is a prototype “spatially extended” system which demonstrates the need for localization in order to achieve skillful filter results for moderate ensemble sizes. We begin with the Lorenz-63 model.

We mention that theoretical results on the long time behavior of filtering algorithms for chaotic systems, such as the Lorenz-63 model, have been obtained, for example, by González-Tokman and Hunt (2013) and Law et al. (2013).

7.1 Lorenz-63 model

The Lorenz-63 model is given by the differential equation (2.7) with state variable $z = (x, y, z)^T \in \mathbb{R}^3$, right-hand side

$$f(z) = \begin{pmatrix} \sigma(y - x) \\ x(\rho - z) - y \\ xy - \beta z \end{pmatrix},$$

and parameter values $\sigma = 10$, $\rho = 28$, and $\beta = 8/3$. The resulting ODE (2.7) is discretized in time by the implicit midpoint method (Ascher, 2008), i.e.,

$$z^{n+1} = z^n + \Delta t f(z^{n+1/2}), \quad z^{n+1/2} = \frac{1}{2}(z^{n+1} + z^n) \quad (2.57)$$

with step-size $\Delta t = 0.01$. Let us abbreviate the resulting map by Ψ_{IM} . Then the dynamical system (2.6) is defined as

$$\Psi = \Psi_{\text{IM}}^{[12]}.$$

In other words observations are assimilated every 12 time-steps. We only observe the x variable with a Gaussian measurement error of variance $R = 8$.

We used different ensemble sizes from 10 to 80 as well as different inflation factors ranging from 1.0 to 1.12 by increments of 0.02 for the EnKF and rejuvenation parameters ranging from 0 to 0.4 by increments of 0.04 for the ETPF. Note that a rejuvenation parameter of $h = 0.4$ corresponds to an inflation factor $\alpha = \sqrt{1 + h^2} \approx 1.0770$.

The following variant of the ETPF with localized cost function has also been implemented. We first compute the importance weights w_i of a given observation. Then each component of the state vector is updated using only the distance in that component in the cost function $J(T)$. For example, the x_i^f components of the forecast ensemble members $z_i^f = (x_i^f, y_i^f, z_i^f)^T$, $i = 1, \dots, M$, are updated according to

$$x_i^a = M \sum_{i=1}^M x_i^f t_{ij}^*$$

with the coefficients $t_{ij}^* \geq 0$ minimizing the cost function

$$J(T) = \sum_{i,j=1}^M t_{ij} |x_i^f - x_j^f|^2$$

subject to (2.36). We use ETPF_R0 as the shorthand form for this method. This variant of the ETPF is of special interest from a computational point of view since the linear transport problem in \mathbb{R}^3 reduces to three simple one-dimensional problems.

The model is run over $N = 20,000$ assimilation steps after discarding 200 steps to lose the influence of the initial conditions. The resulting *root-mean-square errors* averaged over time (RMSEs)

$$\text{RMSE} = \frac{1}{N} \sum_{n=1}^N \sqrt{\|\bar{z}^{n,a} - z_{\text{ref}}^n\|^2}$$

are reported in Fig. 2 a)-c). We dropped the results for the ETPF and ETPF_R0 with ensemble size $M = 10$ as they indicated strong divergence. We see that the

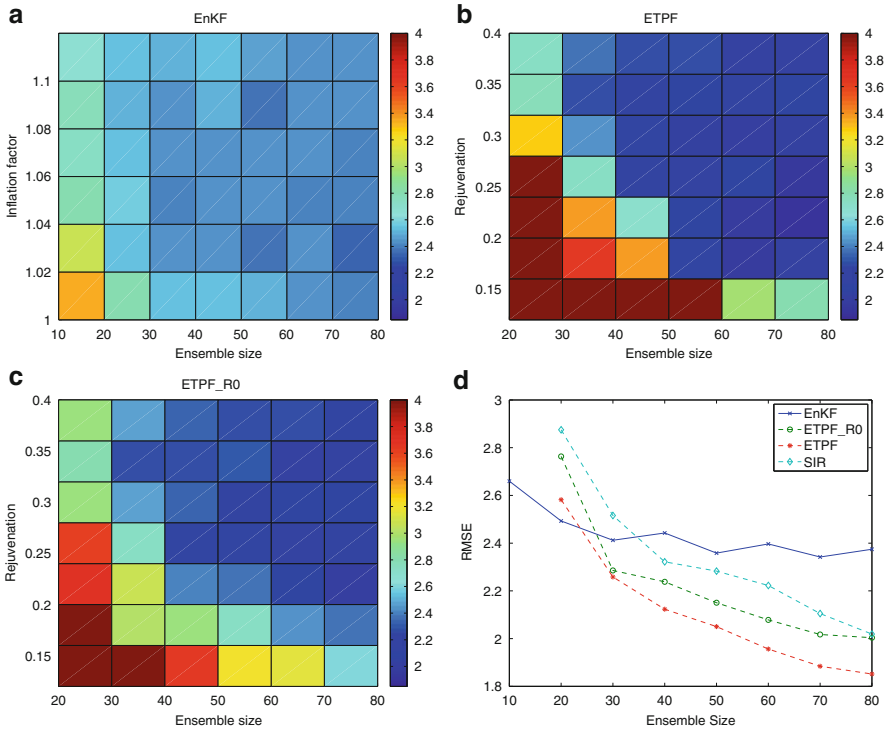


Fig. 2 a)-c): Heatmaps showing the RMSEs for different parameters for the EnKF, ETPF, and ETPF_R0 for the Lorenz-63 model. d): RMSEs for different ensemble sizes using ‘optimal’ inflation factors and rejuvenation.

EnKF produces stable results while the other filters are more sensitive to different choices for the rejuvenation parameter. However, with increasing ensemble size and ‘optimal’ choice of parameters the ETPF and the ETPF_R0 outperform the EnKF which reflects the biasedness of the EnKF.

Figure 2 d) shows the RMSEs for each ensemble size using the parameters that yield the lowest RMSE. Here we see again that the stability of the EnKF leads to good results even for very small ensemble sizes. The downside is also evident: While the ETPFs fail to track the reference solution as well as the EnKF for very small ensemble sizes a small increase leads to much lower RMSEs. The asymptotic consistent ETPF outperforms the ETPF_R0 for large ensemble sizes but is less stable otherwise. We also included RMSEs for the SIR filter with rejuvenation parameters chosen from the same range of values as for the ETPFs. Although not shown here, this range seems to cover the ‘optimal’ choice for the rejuvenation parameter. The comparison with the EnKF is as expected: for small ensemble sizes the SIR performs worse but beats the EnKF for larger ensemble sizes due to its asymptotic consistency. However, the equally consistent ETPF yields lower RMSEs throughout for the ensemble sizes considered here. Interestingly, the SIR only catches up with the inconsistent but computationally cheap ETPF_R0 for the largest ensemble size in this experiment. We mention that the RMSE drops to around 1.4 with the SIR filter with an ensemble size of 1000.

At this point we note that the computational burden increases considerably for the ETPF for larger ensemble sizes due to the need of solving increasingly large linear transport problems. See the discussion from Section 5.3.

7.2 Lorenz-96 model

Given a periodic domain $x \in [0, L]$ and N_z equally spaced grid-points $x_j = j\Delta x$, $\Delta x = L/N_z$, we denote by u_j the approximation to $z(x)$ at the grid points x_j , $j = 1, \dots, N_z$. The following system of differential equations

$$\frac{du_j}{dt} = -\frac{u_{j-1}u_{j+1} - u_{j-2}u_{j-1}}{3\Delta x} - u_j + F, \quad j = 1, \dots, 40, \quad (2.58)$$

is due to Lorenz (1996) and is called the Lorenz-96 model. We set $F = 8$ and apply periodic boundary conditions $u_j = u_{j+40}$. The state variable is defined by $z = (u_1, \dots, u_{40})^T \in \mathbb{R}^{40}$. The Lorenz-96 model (2.58) can be seen as a coarse spatial approximation to the PDE

$$\frac{\partial u}{\partial t} = -\frac{1}{2} \frac{\partial(u)^2}{\partial x} - u + F, \quad x \in [0, 40/3],$$

with mesh-size $\Delta x = 1/3$ and $N_z = 40$ grid points. The implicit midpoint method (2.57) is used with a step-size of $\Delta t = 0.005$ to discretize the differential

equations (2.58) in time. Observations are assimilated every 22 time-steps and we observe every other grid point with a Gaussian measurement error of variance $R = 8$. The large assimilation interval and variance of the measurement error are chosen because of a desired non-Gaussian ensemble distribution.

We used ensemble sizes from 10 to 80, inflation factors from 1.0 to 1.12 with increments of 0.02 for the EnKF and rejuvenation parameters between 0 and 0.4 with increments of 0.05 for the ETPFs.

As mentioned before, localization is required and we take (2.51) as our localization kernel. For each value of M we fixed a localization radius $r_{\text{loc},R}$ in (2.52). The particular choices can be read off of the following table:

M	10	20	30	40	50	60	70	80
$r_{\text{loc},R}^{\text{EnKF}}$	2	4	6	6	7	7	8	8
$r_{\text{loc},R}^{\text{ETPF}}$	1	2	3	4	5	6	6	6

These values have been found by trial and error and we do not claim that these values are by any means ‘optimal’.

As for localization of the cost function (2.56) for the ETPF we used the same kernel as for the measurement error and implemented different versions of the localized ETPF which differ in the choice of the localization radius: ETPF_R1 corresponds to the choice of $r_{\text{loc},c} = 1$ and ETPF_R2 is used for the ETPF with $r_{\text{loc},c} = 2$. As before we denote the computationally cheap version with cost function $c_{x_j}(z^f, z^a) = |u_j^f - u_j^a|^2$ at grid point x_j by ETPF_R0.

The localization kernel and the localization radii $r_{\text{loc},c}$ are not chosen by any optimality criterion but rather by convenience and simplicity. A better kernel or localization radii may be derived from looking at the time averaged spatial correlation coefficients (2.49) as shown in Fig. 3. Our kernel gives higher weights




Fig. 3 Time averaged spatial correlation between solution components depending on their distance.

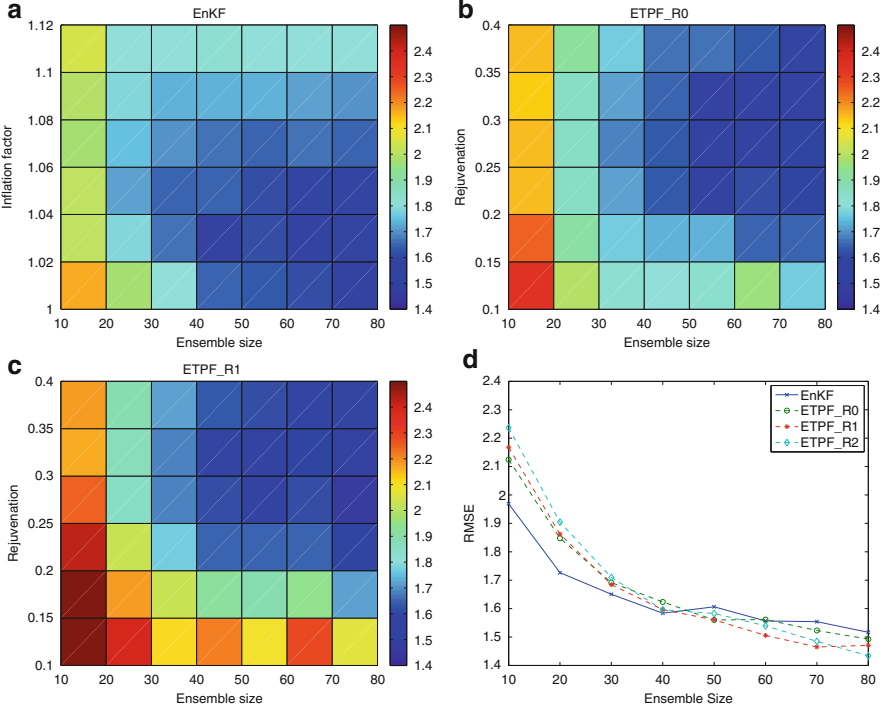


Fig. 4 a)-c): Heatmaps showing the RMSEs for different parameters for the EnKF, ETPF_R0, and ETPF_R1 for the Lorenz-96 model. d): RMSEs for different ensemble sizes using ‘optimal’ inflation factors and rejuvenation.

to components closer to the one to be updated, even though the correlation with the immediate neighbor is relatively low.

The model is run over $N = 10,000$ assimilation steps after discarding 500 steps to lose the influence of the initial conditions. The resulting time averaged RMSEs are displayed in Fig. 4. We dropped the results for the smallest rejuvenation parameters as the filters showed strong divergence. Similar to the results for the Lorenz-63 model the EnKF shows the most stable overall performance for various parameters but fails to keep up with the ETPFs for higher ensemble sizes, though the difference between the different filters is much smaller than for the Lorenz-63 system. This is no surprise since the Lorenz-96 system does not have the highly non-linear dynamics of the Lorenz-63 system which causes the involved distributions to be strongly non-Gaussian. The important point here is that the ETPF as a particle filter is able to compete with the EnKF even for small ensemble sizes. Traditionally, high dimensional systems required very high ensemble sizes for particle filters to perform reasonably well. Hundreds of particles are necessary for the SIR to be even close to the true state.

8 Historical comments

The notion of data assimilation has been coined in the field of meteorology and more widely in the geosciences to collectively denote techniques for combining computational models and physical observations in order to estimate the current state of the atmosphere or any other geophysical process. The perhaps first occurrence of the concept of data assimilation can be found in the work of Richardson (1922), where observational data needed to be interpolated onto a grid in order to initialize the computational forecast process. With the rapid increase in computational resolution starting in the 1960s, it became quickly necessary to replace simple data interpolation by an optimal combination of first guess estimates and observations. This gave rise to techniques such as the *successive correction method*, *nudging*, *optimal interpolation*, and *variational least square techniques* (3D-Var and 4D-Var). See Daley (1993); Kalnay (2002) for more details.

Leith (1974) proposed *ensemble* (or Monte Carlo) *forecasting* as an alternative to conventional single forecasts. However ensemble forecasting did not become operational before 1993 due to limited computer resources (Kalnay, 2002). The availability of ensemble forecasts subsequently lead to the invention of the EnKF by Evensen (1994) with a later correction by Burgers et al. (1998) and many subsequent developments, which have been summarized in Evensen (2006). We mention that the analysis step of an EnKF with perturbed observations is closely related to a method now called *randomized likelihood method* (Kitanidis, 1995; Oliver, 1996; Oliver et al., 1996).

In a completely independent line of research the problem of optimal estimation of stochastic processes from data has led to the theory of *filtering* and *smoothing*, which started with the work of Wiener (1948). The state space approach to filtering of linear systems gave rise to the celebrated *Kalman filter* and more generally to the *stochastic PDE formulations* of Zakai and Kushner-Stratonovitch in case of continuous-time filtering. See Jazwinski (1970) for the theoretical developments up to 1970. Monte Carlo techniques were first introduced to the filtering problem by Handschin and Mayne (1969), but it was not until the work of Gordon et al. (1993) that the SMCM became widely used (Doucet et al., 2001). The *McKean interacting particle approach* to SMCMs has been pioneered by del Moral (2004). The theory of particle filters for time-continuous filtering problems is summarized in Bain and Crisan (2009).

Standard SMCMs suffer from the *curse of dimensionality* in that the necessary number of ensemble members M increases exponentially with the dimension N_z of state space (Bengtsson et al., 2008). This limitation has prevented SMCMs from being used in meteorology and the geosciences. On the other hand, it is known that EnKFs lead to inconsistent estimates which is problematic when multimodal forecast distributions are to be expected. Current research work is therefore focused on a theoretical understanding of EnKFs and related sequential assimilation techniques (see, for example, González-Tokman and Hunt (2013); Law et al. (2013)), extensions of particle filters/SMCMs to PDE models (see, for

example, Morzfeld and Chorin (2012); van Leeuwen and Ades (2012); Beskov et al. (2013); Metref et al. (2013)), and Bayesian inference on function spaces (see, for example, Stuart (2010); Cotter et al. (2009); Dashti et al. (2013)) and hybrid variational methods such as by, for example, Bonavita et al. (2012); Clayton et al. (2013).

A historical account of optimal transportation can be found in Villani (2009). The work of McCann (1995) provides the theoretical link between the classic linear assignment problem and the Monge-Kantorovich problem of coupling PDFs. The ETPF is a computational procedure for approximating such couplings using importance sampling and linear transport instead.

9 Summary and Outlook

We have discussed various ensemble/particle-based algorithms for sequential data assimilation in the context of LETFs. Our starting point was the McKean interpretation of Feynman-Kac formulae. The McKean approach requires a coupling of measures which can be discussed in the context of optimal transportation. This approach leads to the ETPF when applied in the context of SMCMs. We have furthermore discussed extensions of LETFs to spatially extended systems in form of R -localization.

The presented work can be continued along several lines. First, one may replace the empirical forecast measure

$$\pi_{\text{emp}}^f(z) := \frac{1}{M} \sum_{i=1}^M \delta(z - z_i^f), \quad (2.59)$$

which forms the basis of SMCMs and the ETPF, by a *Gaussian mixture*

$$\pi_{\text{GM}}^f(z) := \frac{1}{M} \sum_{i=1}^M n(z; z_i^f, B), \quad (2.60)$$

where $B \in \mathbb{R}^{N_z \times N_z}$ is a given covariance matrix and

$$n(z; m, B) := \frac{1}{(2\pi)^{N_z/2} |B|^{1/2}} e^{-\frac{1}{2}(z-m)^T B^{-1}(z-m)}.$$

Note that the empirical measure (2.59) is recovered in the limit $B \rightarrow 0$. While the weighted empirical measure

$$\pi_{\text{emp}}^a(z) := \sum_{i=1}^M w_i \delta(z - z_i^f)$$

with weights given by (2.35) provides the analysis in case of an empirical forecast measure (2.59) and an observation y_{obs} , a Gaussian mixture forecast PDF (2.60) leads to an analysis PDF in form of a weighted Gaussian mixture provided the forward operator $h(z)$ is linear in z . This fact allows one to extend the ETPF to Gaussian mixtures. See Reich and Cotter (2015) for more details. Alternative implementations of Gaussian mixture filters can, for example, be found in Stordal et al. (2011) and Frei and Künsch (2013).

Second, one may factorize the likelihood function $\pi_{Y^{1:N}}(y^{1:N}|z^{0:N})$ into $L > 1$ identical copies

$$\hat{\pi}_{Y^{1:N}}(y^{1:N}|z^{0:N}) := \frac{1}{(2\pi)^{N_y N/2} |R/L|^{N/2}} \prod_{n=1}^N e^{-\frac{1}{2L} (h(z^n) - y^n)^T R^{-1} (h(z^n) - y^n)},$$

i.e.,

$$\pi_{Y^{1:N}}(y^{1:N}|z^{0:N}) = \prod_{l=1}^L \hat{\pi}_{Y^{1:N}}(y^{1:N}|z^{0:N})$$

and one obtains a sequence of L “incremental” Feynman-Kac formulae. Each of these formulae can be approximated numerically by any of the methods discussed in this review. For example, one obtains the continuous EnKF formulation of Bergemann and Reich (2010) in the limit $L \rightarrow \infty$ in case of an ESRF. We also mention the continuous Gaussian mixture ensemble transform filter (Reich, 2012). An important advantage of an incremental approach is the fact that the associated weights (2.35) remain closer to the uniform reference value $1/M$ in each iteration step. See also related methods such as *running in place* (RIP) (Kalnay and Yang, 2010), the iterative EnKF approach of Bocquet and Sakov (2012); Sakov et al. (2012), and the embedding approach of Beskov et al. (2013) for SMCs.

Third, while this paper has been focused on discrete time algorithms, most of the presented results can be extended to differential equations with observations arriving continuously in time such as

$$dy_{\text{obs}}(t) = h(z_{\text{ref}}(t))dt + \sigma dW(t),$$

where $W(t)$ denotes standard Brownian motion and $\sigma > 0$ determines the amplitude of the measurement error. The associated marginal densities satisfy the Kushner-Stratonovitch stochastic PDE (Jazwinski, 1970). Extensions of the McKean approach to continuous-in-time filtering problems can be found in Crisan and Xiong (2010) and Yang et al. (2013). We also mention the continuous-in-time formulation of the EnKF by Bergemann and Reich (2012). More generally, a reformulation of LETFs in terms of continuously-in-time arriving observations is of the abstract form

$$dz_j = f(z_j)dt + \sum_{i=1}^M z_i ds_{ij}, \quad j = 1, \dots, M. \quad (2.61)$$

Here $S(t) = \{s_{ij}(t)\}$ denotes a matrix-valued stochastic process which depends on the ensemble $\{z_i(t)\}$ and the observations $y_{\text{obs}}(t)$. In other words, (2.61) leads to a particular class of interacting particle systems and we leave further investigations of its properties for future research. We only mention that the continuous-in-time EnKF formulation of Bergemann and Reich (2012) leads to

$$ds_{ij} = \frac{1}{M-1} (y_i - \bar{y}) \sigma^{-1} (dy_{\text{obs}} - y_j dt + \sigma^{1/2} dW_j),$$

where the $W_j(t)$'s denote standard Brownian motion, $y_j = h(z_j)$, and $\bar{y} = \frac{1}{M} \sum_{i=1}^M y_i$. See also Amezcua et al. (2014) for related reformulations of ESRFs.

Acknowledgements We would like to thank Yann Brenier, Dan Crisan, Mike Cullen, and Andrew Stuart for inspiring discussions on ensemble-based filtering methods and the theory of optimal transportation.

References

- J. Amezcua, E. Kalnay, K. Ide, and S. Reich. Ensemble transform Kalman-Bucy filters. *Quarterly J. Royal Meteor. Soc.*, 140:995–1004, 2014.
- J.L. Anderson. An adaptive covariance inflation error correction algorithm for ensemble filters. *Tellus*, 59A:210–224, 2007.
- J.L. Anderson. A non-Gaussian ensemble filter update for data assimilation. *Monthly Weather Review*, 138:4186–4198, 2010.
- J.L. Anderson. Localization and sampling error correction in ensemble Kalman filter data assimilation. *Mon. Wea. Rev.*, 140:2359–2371, 2012.
- U.M. Ascher. *Numerical Methods for Evolutionary Differential Equations*. SIAM, Philadelphia, MA, 2008.
- A. Bain and D. Crisan. *Fundamentals of Stochastic Filtering*, volume 60 of *Stochastic modelling and applied probability*. Springer-Verlag, New-York, 2009.
- T. Bengtsson, P. Bickel, and B. Li. Curse of dimensionality revisited: Collapse of the particle filter in very large scale systems. In *IMS Lecture Notes - Monograph Series in Probability and Statistics: Essays in Honor of David F. Freedman*, volume 2, pages 316–334. Institute of Mathematical Sciences, 2008.
- K. Bergemann and S. Reich. A localization technique for ensemble Kalman filters. *Q. J. R. Meteorological Soc.*, 136:701–707, 2010.
- K. Bergemann and S. Reich. An ensemble Kalman-Bucy filter for continuous data assimilation. *Meteorolog. Zeitschrift*, 21:213–219, 2012.
- A. Beskov, D. Crisan, A. Jasra, and N. Whiteley. Error bounds and normalizing constants in sequential Monte Carlo in high dimensions. *Adv. App. Probab.*, 2013. to appear.
- M. Bocquet and P. Sakov. Combining inflation-free and iterative ensemble Kalman filters for strongly nonlinear systems. *Nonlin. Processes Geophys.*, 19:383–399, 2012.
- M. Bonavita, L. Isaksen, and E. Holm. On the use of EDA background error variances in the ECMWF 4D-Var. *Q. J. Royal. Meteor. Soc.*, 138:1540–1559, 2012.
- G. Burgers, P.J. van Leeuwen, and G. Evensen. On the analysis scheme in the ensemble Kalman filter. *Mon. Wea. Rev.*, 126:1719–1724, 1998.
- R.E. Caflisch. Monte Carlo and quasi-Monte Carlo methods. In *Acta Numerica*, volume 7, pages 1–49. Cambridge University Press, 1988.

- R.E. Caflisch, W. Morokoff, and A.B. Owen. Valuation of mortgage backed securities using Brownian bridges to reduce effective dimension. *Journal of Computational Finance*, pages 27–46, 1997.
- A.J. Chorin, M. Morzfeld, and X. Tu. Implicit filters for data assimilation. *Comm. Appl. Math. Comp. Sc.*, 5:221–240, 2010.
- A.M. Clayton, A.C. Lorenc, and D.M. Barker. Operational implementation of a hybrid ensemble/4D-Var global data assimilation system at the Met Office. *Q. J. Royal. Meteo. Soc.*, 139:1445–1461, 2013.
- S.L. Cotter, M. Dashti, J.C. Robinson, and A.M. Stuart. Bayesian inverse problems for functions and applications to fluid mechanics. *Inverse Problems*, 25:115008, 2009.
- D. Crisan and J. Xiong. Approximate McKean-Vlasov representation for a class of SPDEs. *Stochastics*, 82:53–68, 2010.
- M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300, Lake Tahoe, Nevada, 2013.
- R. Daley. *Atmospheric Data Analysis*. Cambridge University Press, Cambridge, 1993.
- M. Dashti, K.J.H. Law, A.M. Stuart, and J. Voss. MAP estimators and posterior consistency. *Inverse Problems*, 29:095017, 2013.
- P. del Moral. *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Springer-Verlag, New York, 2004.
- A. Doucet and A.M. Johansen. A tutorial on particle filtering and smoothing: fifteen years later. In D. Crisan and B. Rozovskii, editors, *Oxford Handbook of Nonlinear Filtering*, pages 656–704, 2011.
- A. Doucet, N. de Freitas, and N. Gordon (eds.). *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, Berlin Heidelberg New York, 2001.
- G. Evensen. Sequential data assimilation with a nonlinear quasigeostrophic model using Monte Carlo methods for forecasting error statistics. *J. Geophys. Res.*, 99:10143–10162, 1994.
- G. Evensen. *Data Assimilation. The Ensemble Kalman Filter*. Springer-Verlag, New York, 2006.
- M. Frei and H.R. Künsch. Mixture ensemble Kalman filters. *Computational Statistics and Data Analysis*, 58:127–138, 2013.
- G. Gaspari and S.E. Cohn. Construction of correlation functions in two and three dimensions. *Q. J. Royal Meteorological Soc.*, 125:723–757, 1999.
- G.H. Golub and Ch.F. van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, 3rd edition, 1996.
- C. González-Tokman and B.R. Hunt. Ensemble data assimilation for hyperbolic systems. *Physica D*, 243:128–142, 2013.
- N.J. Gordon, D.J. Salmon, and A.F.M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEEE Proceedings F on Radar and Signal Processing*, 140:107–113, 1993.
- Th.M. Hamill, J.S. Whitaker, and Ch. Snyder. Distance-dependent filtering of background covariance estimates in an ensemble Kalman filter. *Mon. Wea. Rev.*, 129:2776–2790, 2001.
- J.E. Handschin and D.Q. Mayne. Monte Carlo techniques to estimate the conditional expectation in multi-stage non-linear filtering. *Internat. J. Control*, 1:547–559, 1969.
- P.L. Houtekamer and H.L. Mitchell. A sequential ensemble Kalman filter for atmospheric data assimilation. *Mon. Wea. Rev.*, 129:123–136, 2001.
- P.L. Houtekamer and H.L. Mitchell. Ensemble Kalman filtering. *Q. J. Royal Meteorological Soc.*, 131:3269–3289, 2005.
- B.R. Hunt, E.J. Kostelich, and I. Szunyogh. Efficient data assimilation for spatialtemporal chaos: A local ensemble transform Kalman filter. *Physica D*, 230:112–137, 2007.
- T. Janjić, D. McLaughlin, S.E. Cohn, and M. Verlaan. Conservation of mass and preservation of positivity with ensemble-type Kalman filter algorithms. *Mon. Wea. Rev.*, 142:755–773, 2014.
- A.H. Jazwinski. *Stochastic Processes and Filtering Theory*. Academic Press, New York, 1970.
- J. Kaipio and E. Somersalo. *Statistical and Computational Inverse Problems*. Springer-Verlag, New York, 2005.
- E. Kalnay. *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, 2002.

- E. Kalnay and S.-C. Yang. Accelerating the spin-up of ensemble Kalman filtering. *Quart. J. Roy. Meteor. Soc.*, 136:1644–1651, 2010.
- P.K. Kitanidis. Quasi-linear geostatistical theory for inverting. *Water Resources Research*, 31: 2411–2419, 1995.
- K.J.H. Law, A. Shukla, and A.M. Stuart. Analysis of the 3DVAR filter for the partially observed Lorenz’63 model. *Discrete and Continuous Dynamical Systems A*, 2013. to appear.
- T.S. Lee and D. Mumford. Hierarchical Bayesian inference in the visual cortex. *J. Opt. Soc. Am. A*, 20:1434–1448, 2003.
- J. Lei and P. Bickel. A moment matching ensemble filter for nonlinear and non-Gaussian data assimilation. *Mon. Weath. Rev.*, 139:3964–3973, 2011.
- C.E. Leith. Theoretical skills of Monte Carlo forecasts. *Mon. Weath. Rev.*, 102:409–418, 1974.
- T. Lelièvre, M. Rousset, and G. Stoltz. *Free Energy Computations - A Mathematical Perspective*. Imperial College Press, London, 2010.
- J.M. Lewis, S. Lakshmivarahan, and S.K. Dhall. *Dynamic Data Assimilation: A Least Squares Approach*. Cambridge University Press, Cambridge, 2006.
- J.S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer-Verlag, New York, 2001.
- D.M. Livings, S.L. Dance, and N.K. Nichols. Unbiased ensemble square root filters. *Physica D*, 237:1021–1028, 2008.
- E.N. Lorenz. Deterministic non-periodic flows. *J. Atmos. Sci.*, 20:130–141, 1963.
- E.N. Lorenz. Predictability: A problem partly solved. In *Proc. Seminar on Predictability*, volume 1, pages 1–18, ECMWF, Reading, Berkshire, UK, 1996.
- A. Majda and J. Harlim. *Filtering Complex Turbulent Systems*. Cambridge University Press, Cambridge, 2012.
- R.J. McCann. Existence and uniqueness of monotone measure-preserving maps. *Duke Mathematical Journal*, 80:309–323, 1995.
- H.P. McKean. A class of Markov processes associated with nonlinear parabolic equations. *Proc. Natl. Acad. Sci. USA*, 56:1907–1911, 1966.
- S. Metref, E. Cosme, C. Snyder, and P. Brasseur. A non-Gaussian analysis scheme using rank histograms for ensemble data assimilation. *Nonlinear Processes in Geophysics*, 21:869–885, 2013.
- T. Miyoshi. The Gaussian approach to adaptive covariance inflation and its implementation with the local ensemble transform Kalman filter. *Mon. Wea. Rev.*, 139:1519–1535, 2011.
- K.W. Morton and D.F. Mayers. *Numerical Solution of Partial Differential Equations*. Cambridge University Press, Cambridge, 2nd edition, 2005.
- M. Morzfeld and A.J. Chorin. Implicit particle filtering for models with partial noise and an application to geomagnetic data assimilation. *Nonlinear Processes in Geophysics*, 19:365–382, 2012.
- M. Morzfeld, X. Tu, E. Atkins, and A.J. Chorin. A random map implementation of implicit filters. *J. Comput. Phys.*, 231:2049–2066, 2012.
- T.A. El Moselhy and Y.M. Marzouk. Bayesian inference with optimal maps. *J. Comput. Phys.*, 231:7815–7850, 2012.
- L. Nerger, T. Janjic Pfander, J. Schröter, and W. Hiller. A regulated localization scheme for ensemble-based Kalman filters. *Quarterly J. Royal Meteor. Soc.*, 138:802–812, 2012.
- J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer-Verlag, New York, 2nd edition, 2006.
- D.S. Oliver. On conditional simulation to inaccurate data. *Math. Geology*, 28:811–817, 1996.
- D.S. Oliver, N. He, and A.C. Reynolds. Conditioning permeability fields on pressure data. Technical report, presented at the 5th European Conference on the Mathematics of Oil Recovery, Leoben, Austria, 1996.
- I. Olkin and F. Pukelsheim. The distance between two random vectors with given dispersion matrices. *Linear Algebra and its Applications*, 48:257–263, 1982.
- E. Ott, B.R. Hunt, I. Szunyogh, A.V. Zimin, E.J. Kostelich, M. Corazza, E. Kalnay, D.J. Patil, and J.A. Yorke. A local ensemble Kalman filter for atmospheric data assimilation. *Tellus*, A 56: 415–428, 2004.

- O. Pele and M. Werman. Fast and robust earth mover's distances. In *Computer Vision, 2009 IEEE 12th international conference*, pages 460–467, 2009.
- D.T. Pham. Stochastic methods for sequential data assimilation in strongly nonlinear systems. *Mon. Wea. Rev.*, 129:1194–1207, 2001.
- P. Rebescini and R. van Handel. Can local particle filters beat the curse of dimensionality? *arXiv:1301.6585*, 2013.
- S. Reich. A dynamical systems framework for intermittent data assimilation. *BIT Numer Math*, 51: 235–249, 2011.
- S. Reich. A Gaussian mixture ensemble transform filter. *Q. J. R. Meteorolog. Soc.*, 138:222–233, 2012.
- S. Reich. A nonparametric ensemble transform method for Bayesian inference. *SIAM J. Sci. Comput.*, 35:A2013–A2024, 2013a.
- S. Reich. A guided sequential Monte Carlo method for the assimilation of data into stochastic dynamical systems. In *Recent Trends in Dynamical Systems*, volume 35 of *Springer Proceedings in Mathematics and Statistics*, pages 205–220, 2013b.
- S. Reich and C. J. Cotter. Ensemble filter techniques for intermittent data assimilation. In M. Cullen, Freitag M. A., S. Kindermann, and R. Scheichl, editors, *Large Scale Inverse Problems. Computational Methods and Applications in the Earth Sciences*, volume 13 of *Radon Ser. Comput. Appl. Math.*, pages 91–134. Walter de Gruyter, Berlin, 2013.
- S. Reich and C.J. Cotter. *Probabilistic Forecasting and Bayesian Data Assimilation*. Cambridge University Press, Cambridge, 2015.
- L.F. Richardson. *Weather Prediction by Numerical Processes*. Cambridge University Press, Cambridge, 1922.
- Ch. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer-Verlag, New York, 2nd edition, 2004.
- J.C. Robinson. *Infinite-Dimensional Dynamical Systems*. Cambridge University Press, Cambridge, 2001.
- P. Sakov, D. Oliver, and L. Bertino. An iterative EnKF for strongly nonlinear systems. *Mon. Wea. Rev.*, 140:1988–2004, 2012.
- S. Särkkä. *Bayesian Filtering and Smoothing*. Cambridge University Press, Cambridge, 2013.
- A.S. Stordal, H.A. Karlsen, G. Nævdal, H.J. Skaug, and B. Vallés. Bridging the ensemble Kalman filter and particle filters: The adaptive Gaussian mixture filter. *Comput. Geosci.*, 15:293–305, 2011.
- G. Strang. *Introduction to Applied Mathematics*. Wellesley-Cambridge Press, 2nd edition, 1986.
- A.M. Stuart. Inverse problems: a Bayesian perspective. In *Acta Numerica*, volume 17, pages 451–559. Cambridge University Press, Cambridge, 2010.
- A. Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, Philadelphia, 2005.
- M.K. Tippett, J.L. Anderson, G.H. Bishop, T.M. Hamill, and J.S. Whitaker. Ensemble square root filters. *Mon. Wea. Rev.*, 131:1485–1490, 2003.
- P.J. van Leeuwen. Nonlinear data assimilation in the geosciences: An extremely efficient particle filter. *Q.J.R. Meteorolog. Soc.*, 136:1991–1996, 2010.
- P.J. van Leeuwen and M. Ades. Efficient fully nonlinear data assimilation for geophysical fluid dynamics. *Computers and Geosciences*, 47:in press, 2012.
- C. Villani. *Topics in Optimal Transportation*. American Mathematical Society, Providence, Rhode Island, NY, 2003.
- C. Villani. *Optimal Transportation: Old and New*. Springer-Verlag, Berlin Heidelberg, 2009.
- X. Wang, C.H. Bishop, and S.J. Julier. Which is better, an ensemble of positive-negative pairs or a centered spherical simplex ensemble? *Mon. Wea. Rev.*, 132:1590–1505, 2004.
- N. Wiener. *Extrapolation, Interpolation and Smoothing of Stationary Time Series: With Engineering Applications*. MIT Press, Cambridge, MA, 1948.
- T. Yang, P.G. Mehta, and S.P. Meyn. Feedback particle filter. *IEEE Trans. Automatic Control*, 58: 2465–2480, 2013.

Nonlinear Data Assimilation

Van Leeuwen, P.J.; Cheng, Y.; Reich, S.

2015, XII, 118 p. 19 illus., 15 illus. in color., Softcover

ISBN: 978-3-319-18346-6