Dissertation

submitted to the
Combined Faculties for the Natural Sciences and for Mathematics
of the Ruperto-Carola University of Heidelberg, Germany

for the degree of
Doctor of Natural Sciences

Put forward by
M.Sc. Daniel Berg
born in Simmern

Oral examination: July 27, 2018

Particle Filters for Nonlinear Data Assimilation

Referees:

Prof. Dr. Kurt Roth
Prof. Dr. Björn Malte Schäfer

Particle Filters for Nonlinear Data Assimilation: Environmental systems are nonlinear, multiscale and non-separable. Mathematical models describing these systems are typically high-dimensional and always have missing physics. Therefore, determining the system's state and its future development relies on in situ observations. Information from models and observations are combined using data assimilation methods, which are mainly developed for divergent systems as they arise from weather prediction. Applying them also to convergent systems requires modifications of these methods. I investigated the differences of data assimilation in convergent and divergent systems and found that parameter estimation is essential in convergent systems. In this work, I enhanced the particle filter, an ensemble-based data assimilation method. In contrast to other methods, the particle filter is able to handle nonlinear systems and to describe the resulting non-Gaussian probability density functions. However, for parameter estimation modifications of the resampling, i.e. the renewal of the ensemble, are necessary. I developed a resampling method that uses the weighted covariance information calculated from the ensemble to generate new particles. This method correlates observed with unobserved dimensions and can effectively estimate state and parameters in a convergent system. To be applicable in high-dimensional systems, particle filters need localisation. The introduced resampling allows localisation, which further increases the efficiency of the filter.

Partikelfilter für nichtlineare Datenassimilation: Umweltsysteme sind nichtlinear, vielskalig und nicht separierbar. Mathematische Modelle, die diese Systeme beschreiben, sind oft hochdimensional und besitzen nicht repräsentierte Physik. Die Bestimmung des Systemzustandes und dessen zukünftige Entwicklung benötigt daher insitu-Messungen. Informationen aus Modellen und Beobachtungen werden mit Hilfe von Datenassimilation kombiniert. Da die Methoden hierfür aus der Wettervorhersage hervorgehen, werden sie hauptsächlich für divergente Systeme entwickelt. Für eine Anwendung auf konvergente Systeme müssen die Methoden modifiziert werden. Ich habe die Unterschiede von konvergenten und divergenten Systemen im Bezug auf Datenassimilation untersucht, mit dem Ergebnis, dass Parameterschätzung in konvergenten Systemen essenziell ist. In dieser Arbeit habe ich den Partikelfilter, eine ensemblebasierte Methode für Datenassimilation, verbessert. Im Gegensatz zu anderen Methoden ist dieser für nichtlineare Systeme und die daraus resultierenden nicht-gaußschen Wahrscheinlichkeitsverteilungen geeignet. Für eine erfolgreiche Parameterschätzung ist es jedoch notwendig, das Resampling, also die Erneuerung des Ensembles, zu modifizieren. Für die Parameterschätzung habe ich ein Resampling entwickelt, das die gewichtete, aus dem Ensemble berechnete, Kovarianz verwendet, um neue Partikel zu erzeugen. Diese Methode korreliert unbeobachtete mit beobachteten Dimensionen und kann so den Zustand und die Parameter in einem konvergenten System effektiv schätzen. Um Partikelfilter in hochdimensionalen Systemen einsetzen zu können, benötigten sie Lokalisierung. Das eingeführte Resampling macht eine Lokalisierung des Partikelfilters möglich, was die Effizienz des Filters weiter erhöht.

Contents

1. Introduction

The study of environmental systems differs from the usual systems studied in the laboratory. In most natural science disciplines, the system is traditionally studied recursively towards fundamental laws as building blocks. However, this is not possible in environmental systems (Roth, 2017). "The ability to reduce everything to simple fundamental laws does not imply the ability to start from those laws and reconstruct the universe" (Anderson et al., 1972). Environmental systems are nonlinear, multiscale and non-separable (Roth, 2017). Mathematical models describing them are typically high-dimensional and they always have missing physics. This is either because there is no proper physical description or the description is computationally too expensive, but in most cases it is a combination of both (Van Leeuwen et al., 2015). To learn about a system's state and to predict its development, models are combined with in situ observations of the system.

Data assimilation combines the information of models with observations to determine the best knowledge of the system's state. All information is expressed by probability density functions (pdfs) and joined using Bayes' theorem. This requires a good representation of the uncertainties of both models and observations (*Liu and Gupta*, 2007; *Liu et al.*, 2012). The development of data assimilation methods is mainly driven in the field of atmospheric science. As the atmosphere is a divergent system, the distance between two states will increase exponentially in time, which is why data assimilation in meteorology focuses on state estimation (*Van Leeuwen et al.*, 2015).

Data assimilation is increasingly applied to convergent systems such as soil hydrology. This requires an adaptation of the methods used in atmospheric and related sciences for hydrology (*Liu and Gupta*, 2007). In soil hydrology, not only the state but also the parameters are typically ill-known. For this reason, it is mandatory to simultaneously estimate both, state and parameters. The ensemble Kalman filter (EnKF) (*Evensen*, 1994; *Burgers et al.*, 1998) is an often applied method for this purpose. It is based on the assumption of Gaussian pdfs, which greatly simplifies the theory and implementation. Because of its simplicity, it is probably the most applied data assimilation method in geosciences (*Van Leeuwen et al.*, 2015). Furthermore, including parameter estimation is operationally simple, as the covariance calculated from the ensemble transfers infor-

mation to parameter space. In hydrology, the EnKF is successfully used for state and parameter estimation (e.g. *Erdal et al.*, 2015; *Shi et al.*, 2015; *Bauser et al.*, 2016).

Nonlinear systems violate the assumption of Gaussian pdfs (*Harlim and Majda*, 2010; *DeChant and Moradkhani*, 2012; *Van Leeuwen et al.*, 2015) and the performance of the EnKF worsens (*Lei and Bickel*, 2011; *Liu et al.*, 2012; *Zhang et al.*, 2015). In soil hydrology, the dynamics of the Richards equation, which describes the unsaturated water flow in a porous medium, are generally dissipative and the Gaussian assumption is usually appropriate. However, jumps at layer boundaries, soliton-like fronts during strong infiltration and diverging potentials for strong evaporation deform the pdf and lead to non-Gaussianity.

Particle filters do not rely on Gaussian pdfs. They directly approximate the pdfs in Bayes' theorem using a weighted ensemble and are therefore able to represent non-Gaussian pdfs. The weights are updated in the analysis with each new observation. However, the particle filter is prone to filter degeneration, meaning that one ensemble member accumulates all the weight over time and the ensemble effectively collapses to this one member. *Gordon et al.* (1993) introduced resampling to particle filters. The idea of resampling is that after the analysis, ensemble members having large weights are replicated and ensemble members having small weights are dropped. This reduces the risk of filter degeneracy.

In the case of parameter estimation, resampling is challenging as it relies on a specified model error. On the one hand, parameters are often constant in time and a specification of a model error would be misleading. On the other hand, without an model error, resampled particles would stay identical and the filter would degenerate. Therefore, it is necessary to perturb the ensemble members in parameter space after resampling. *Moradkhani et al.* (2005a) suggested a random perturbation in parameter space. Further development was done by *Moradkhani et al.* (2012) and *Vrugt et al.* (2013), who used a Markov chain Monte Carlo (MCMC) method to generate new particles. *Abbaszadeh et al.* (2018) extended this method with an additional genetic algorithm. Even though the MCMC is accurate, it needs additional model runs, which is computationally expensive.

In high-dimensional systems, resampling alone is not sufficient to prevent particle filters from degeneration. Another challenge in nonlinear data assimilation with particle filters is the so-called curse of dimensionality (*Bengtsson et al.*, 2008; *Bickel et al.*, 2008; *Snyder et al.*, 2008): The necessary ensemble size to prevent filter degeneracy increases exponentially with the dimension of the system. In contrast to the particle filter, the EnKF has overcome this issue using localisation (*Houtekamer and Mitchell*, 1998, 2001; *Hamill et al.*, 2001). Localisation reduces the high-dimensional data assimilation problem to local problems and cuts off spurious long-range correlations, which occur due to a finite ensemble size. *Sny-*

der et al. (2015) mentions that localisation is the key idea that allows the EnKF to perform well for different high-dimensional geophysical systems. They conclude that also the particle filter needs localisation to be effective in high-dimensional systems. However, the transfer of localisation to particle filters is challenging (*Bengtsson et al.*, 2003). Calculating local weights results in a local resampling that introduces discontinuities (*Fearnhead and Künsch*, 2018). These discontinuities cause strong gradients, which typically lead to unphysical behaviour in geophysical models (*Van Leeuwen et al.*, 2015). Recently proposed localisation methods (e.g. *Poterjoy*, 2016; *Penny and Miyoshi*, 2016; *Robert and Künsch*, 2017) for particle filters show good results in numerical simulations. However, these methods lack the simplicity of the EnKF.

In this thesis, I investigated the challenges and difficulties of data assimilation in divergent and convergent systems using the Lorenz-96 model (*Lorenz*, 1996) and the Richards equation, respectively (Chapter 4). For state and parameter estimation with particle filters, I developed a new resampling technique that uses the weighted covariance matrix from the ensemble to generate new particles. Using the covariance, the unobserved parameters of the new particles are correlated to the observed state dimensions. The effectiveness of this method is demonstrated with a synthetic case – an unsaturated soil consisting of two homogeneous layers – by assimilating time domain reflectometry (TDR)-like measurements (Chapter 5). An advantage of the covariance resampling is that it can be localised similar to the EnKF. The localisation is tested successfully for the Lorenz-96 model (Chapter 6). The analysis of divergent and convergent systems shows the fundamental need for parameter estimation in convergent systems and the covariance resampling allows an efficient state and parameter estimation with particle filters.

2. Models

Mathematical models are used as an abstract representation of a physical system. They aim to describe the dynamics and the future development of system states.

The general representation of a physical system, as it used in this work, has four components: state, dynamics, forcing and parameters. The state is a set of quantities that characterise the system at one point in time such that future states can be calculated from it. The dynamics propagates the state forward in time and is often described by differential equations conditioned on the parameters and the forcing. The parameters, condense the subscale physics of the system and quantify the response of it to the described quantities. The forcing is the embedding in the larger system. It drives the system and couples it with its environment. It is usually an additional energy term or a boundary condition and can vary in space and time.

In the following, two models used in this thesis will be introduced. Section 2.1 describes a model for water movement in an unsaturated porous medium and Section 2.2 introduces the Lorenz-96 model that mimics some features of atmospheric circulation.

2.1. Soil hydrology

This section gives a brief introduction to water movement in soils and follows the description of *Roth* (2017). For a complete introdcution and an in-depth treatment please refer to *Hillel* (2003) or *Roth* (2017).

Soil can be described as a porous medium. This implies that its volume can be divided into a soil matrix and fully interconnected pores. Water movement inside this medium can be described as a two-phase flow, as long as the pores are not completely saturated with either water or air. In this work, the water movement is considered to be in the degenerate multiphase regime, where the air has an arbitrarily large mobility such that both phases are decoupled.

2.1.1. State

The state of the system is described by the volumetric water content θ $(-)$ and the matric potential ψ_m $(\mathrm{M\,L^{-1}\,T^2})$. The volumetric water content is an averaged macroscopic quantity and is defined for a representative elementary volume. This is the minimal averaging volume such that the averaged quantity is consistent and does not change significantly with small changes of the volume. The (volumetric) water content is the proportion of the total volume V $(\mathrm{L^3})$ occupied by water V_w $(\mathrm{L^3})$:

$$\theta = \frac{V_w}{V} \,. \tag{2.1}$$

The water is bound to the soil matrix through surface attraction and the matric potential is the energy required to remove the water from the soil. The matric potential is described as pressure difference between water and air $\psi_m = p_w - p_a$ and can be expressed as an equivalent head of water, the matric head h_m (L):

$$h_m = \frac{\psi_m}{\rho_w g} \,, \tag{2.2}$$

where ρ_w $(\mathrm{M\,L^{-3}})$ is the density of water and g $(\mathrm{L\,T^{-2}})$ is the gravitational acceleration.

2.1.2. Dynamics

In the degenerate multiphase regime, the water flux $\boldsymbol{j}$ $(\mathrm{L\,T^{-1}})$ is described by the Buckingham-Darcy law

$$\boldsymbol{j} = -\mathbf{K}^*(\theta)\left[\nabla\psi_m - \rho_w \boldsymbol{g}\right] \,, \tag{2.3}$$

with the hydraulic conductivity tensor $\mathbf{K}^*$ $(\mathrm{M\,L^3\,T})$, which is dependent on θ. The conductivity is usually assumed to be isotropic, such that it can be expressed as a scalar K^*.

This equation is only valid if the local equilibrium assumption holds, which requires that the time scales of the microscopic processes are much shorter than those of the macroscopic processes. Furthermore, the hydraulic conductivity has to be independent of the speed of the macroscopic development of the state.

The conservation of mass of the water phase, assuming incompressibility ($\rho_w =$ const), is

$$\partial_t \theta + \nabla \cdot \boldsymbol{j} = 0 \,. \tag{2.4}$$

Inserting Eq. (2.4) into Eq. (2.3) results in the Richards equation

$$\partial_t \theta - \nabla \cdot \left[K^*(\theta) \left[\nabla \psi_m - \rho_w \boldsymbol{g} \right] \right] . \tag{2.5}$$

Using the potential expressed per unit weight (Eq. (2.2)), Richards' equation Eq. (2.5) becomes

$$\partial_t \theta - \nabla \cdot \left[K(\theta) \left[\nabla h_m - 1 \right] \right] = 0 \, , \tag{2.6}$$

where the isotropic hydraulic conductivity K ($\mathrm{L\,T^{-1}}$) has the units of a velocity.

2.1.3. Parameters

To solve Richards' equation (Eq. (2.6)), it is necessary to know the dependency of the conductivity and the matric head on the water content. These relations represent the averaged subscale physics and are called material properties.

Different parametrisations for the material properties exist. One of them is the Mualem-van Genuchten parametrisation (*Mualem*, 1976; *Van Genuchten*, 1980). It is an often applied parametrisation and will be used in this thesis in its simplified form:

$$K(\Theta) = K_w \, \Theta^{\tau} \left[1 - \left[1 - \Theta^{n/[n-1]} \right]^{1-1/n} \right]^2 , \tag{2.7}$$

$$h_m(\Theta) = \frac{1}{\alpha} \left[\Theta^{-n/[n-1]} - 1 \right]^{1-1/n} , \tag{2.8}$$

with the saturation Θ $(-)$

$$\Theta := \frac{\theta - \theta_r}{\theta_s - \theta_r} . \tag{2.9}$$

The parameter θ_s $(-)$ is the saturated water content and θ_r $(-)$ is the residual water content. The matric head h_m is scaled with the parameter α ($\mathrm{L^{-1}}$) that can be related to the inverse air entry value. The parameter K_w ($\mathrm{L\,T^{-1}}$) is the saturated hydraulic conductivity, τ $(-)$ a tortuosity factor and n $(-)$ is a shape parameter that determines the sharpness of the air-entry. The parameter n can also associated with the width of the pore size distribution. Equation (2.7) and Eq. (2.8) describe the subscale physics with six parameters for a homogeneous soil layer.

An illustration of Eq. (2.7) and Eq. (2.8) is given in Fig. 2.1. Both functions are shown for sandy loam and loamy sand with parameters taken from *Carsel and Parrish* (1988) (see Table 2.1). Approaching residual water content, the

Table 2.1.: Mualem-van Genuchten parameters for loamy sand and sandy loam (*Carsel and Parrish*, 1988).

Parameter	Loamy sand	Sandy loam
θ_s [−]	0.41	0.41
θ_r [−]	0.057	0.065
τ [−]	0.5	0.5
n [−]	2.28	1.89
α [m]	−12.4	−7.5
K_w [m s^{-1}]	$4.00 \cdot 10^{-5}$	$1.23 \cdot 10^{-5}$

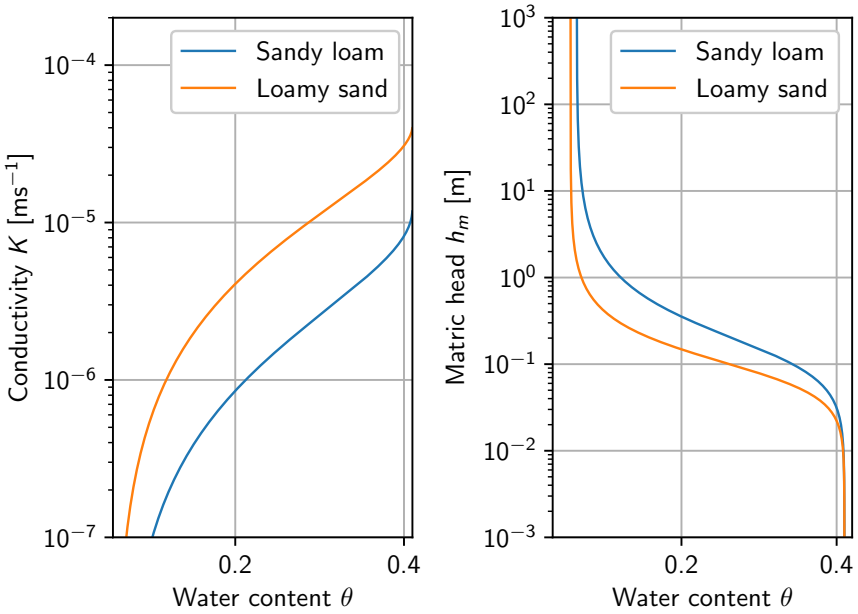

Figure 2.1.: Conductivity K and hydraulic head h_m as function of the water content θ using the Mualem-van Genuchten parametrisation for sandy loam (blue) and loamy sand (orange) using the parameters from *Carsel and Parrish* (1988) (see Table 2.1).

conductivity goes to zero while the hydraulic head diverges. Towards saturation, the conductivity approaches the saturated conductivity K_w while the hydraulic head converges to zero.

Combining Eq. (2.7) and Eq. (2.8) results in a conductivity function that depends on the hydraulic head:

$$K(h_m) = K_w \frac{\left[1 - (\alpha h_m)^{n-1}(1 + (\alpha h_m)^n)^{-1+1/n}\right]^2}{[1 + (\alpha h_m)^n]^{\tau(1-1/n)}} \,. \tag{2.10}$$

2.1.4. Forcing

The system is embedded in a larger environment. It is coupled to the atmosphere and the surrounding soil. In the mathematical model, the coupling is realised with boundary conditions that can vary in space and time.

The two most commonly used are the Dirichlet and the Neumann boundary condition. The Dirichlet boundary condition sets the value of the state (e.g. matric head) at the interface, while the Neumann boundary condition specifies the derivative of the same (e.g. flux). As a lower boundary, Dirichlet is often used to represent a ground water table. As long as the soil is not a limiting factor, a flux boundary is appropriate to describe infiltration and evaporation at the upper boundary. However, if the infiltration is too strong, the soil cannot absorb the water, which leads to ponding or surface runoff. If the system is near the residual water content, an evaporation flux boundary leads to a divergent matric head and the flux cannot be maintained by the soil. In both cases, a Dirichlet boundary condition should be used.

2.1.5. Properties

Infiltrations can lead to sharp fronts. This can be seen in the water content form of Richards' equation:

$$\partial_t \theta + \underbrace{\boldsymbol{V}(\theta) \cdot \nabla\theta}_{\text{advection}} - \underbrace{\nabla \cdot [D(\theta)\nabla\theta]}_{\text{dispersion}} = 0 \,, \tag{2.11}$$

with

$$V(\theta) := \frac{dK}{d\theta} \quad \text{and} \quad D(\theta) := \frac{dh_m(\theta)}{d\theta} K(\theta) \,. \tag{2.12}$$

With the previously described nonlinear parametrisation, the advection term is nonlinear. This will lead to the formation of shock fronts (a moving discontinuity

in θ), which are smoothed by the dispersion. This leads to a characteristic shape (soliton).

While the water content is discontinuous across interfaces between different materials, the hydraulic head is continuous. Therefore, Richards' equation is numerically solved in its potential form:

$$C(h_m)\,\partial_t h_m - \nabla \cdot [K(h_m)\,[\nabla h_m - 1]] = 0\,, \tag{2.13}$$

with the soil water capacity $C(h_m)$:

$$C(h_m) := \frac{\partial \theta}{\partial h_m}\,. \tag{2.14}$$

Using the Mualem-van Genuchten parametrisation, this equation is a nonlinear, degenerate elliptic-parabolic partial differential equation (*List and Radu*, 2016). From their mathematical perspective, the equation transitions from a parabolic to an elliptic partial differential equation for water contents close to the saturated water content. Because of this and the highly nonlinear material properties, it is difficult to solve equation Eq. (2.13) efficiently. A recent review (*Farthing and Ogden*, 2017) summarises the challenges and numerical methods for solving Richards' equation.

In this work, MuPhi (*Ippisch et al.*, 2006), a cell-centred finite-volume scheme with full-upwinding in space and an implicit Euler scheme in time that linearises the equation with an inexact Newton method with line search, was used to numerically solve Richards' equation.

2.2. Lorenz-96 model

The Lorenz-96 model (*Lorenz*, 1996) is an artificial model and cannot be derived from any dynamic equation (*Lorenz*, 2005). It can be interpreted as an unspecified scalar quantity x in a one dimensional atmosphere on a latitude circle on a linear lattice and was defined by *Lorenz* (1996) in the course of his study on predictability. The governing equation is a set of coupled ordinary differential equations:

$$\frac{dx_i}{dt} = (x_{i+1} - x_{i-2})x_{i-1} - x_i + F \quad \text{with} \quad i \in [1, 2, \ldots, J] \tag{2.15}$$

with constant forcing F, periodic boundaries ($x_{J+1} = x_1$) and dimension J. Even though it is not derived from physical principles, it shares certain properties of large atmospheric models (*Lorenz and Emanuel*, 1998). The quadratic terms

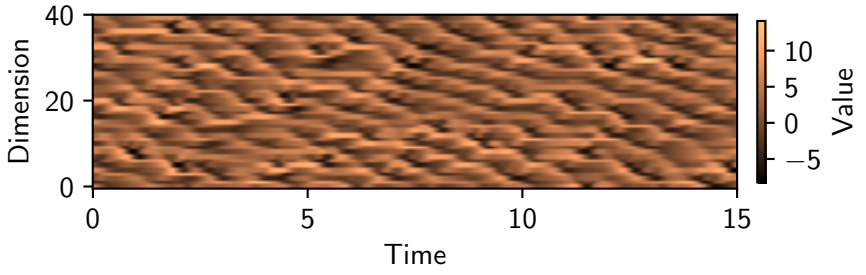

Figure 2.2.: State as a function of time of the 40-dimensional Lorenz-96 model for $F = 8$. The system is started on the attractor and shows irregular waves travelling in the direction opposite to the counting of the cells.

represent advection and conserve the total energy, while the linear term decreases the total energy comparable to dissipation. The constant F addresses external forcing and prevents the system's total energy from decaying to zero.

Lorenz (1996) scales the coefficients of the quadratic and linear terms to unity and assumes that this scaling corresponds to a time unit equal to 5 days. The system's behaviour is fully described by the forcing F and the total dimension J. Following the analysis of *Lorenz and Emanuel* (1998), the properties of the often used 40-dimensional case will be discussed. The system has three regimes:

- $F < 0.894$: the solution decays to F
- $0.894 < F < 4$: the system becomes periodic with regular waves
- $F > 4$: the system becomes chaotic with irregular waves (see Fig. 2.2)

The Lyapunov exponent quantifies the local speed of separation of two infinitesimal close trajectories. Analysing the averaged leading Lyapunov exponent for $F = 8$ results in a doubling time of 0.42 units or 2.1 days. This doubling time is similar to the one observed for large atmospheric models (cf. *Lorenz*, 1982; *Simmons et al.*, 1995; *Simmons and Hollingsworth*, 2002). However, the growth rate in a small interval can be considerably larger or smaller.

The Lorenz-96 model is a popular toy model to test new algorithms in data assimilation (e.g. *Fertig et al.*, 2007; *Li et al.*, 2009; *Anderson*, 2012) because it has similar properties as large-scale atmospheric models but is computational much more efficient.

3. Data assimilation

Parts of Section 3.3.3 are based on *Berg et al.* (2018).

Mathematical models are used to describe physical systems (see Chapter 2). Observations are another source of information. Data assimilation is used to combine these two sources of information to determine the best knowledge of the state. Observation and model information are incorporated into the current state to have an optimal estimate of the state of the system.

In data assimilation, information is treated as probability density function (pdf), therefore the quantification of uncertainty is an essential part (*Liu and Gupta*, 2007; *Liu et al.*, 2012). To incorporate an observation as additional information, it is necessary to combine the pdfs of the observation and the model. This is done using Bayes' theorem, which is the fundamental basis of all data assimilation methods:

$$P(\boldsymbol{\psi}|\boldsymbol{d}) = \frac{P(\boldsymbol{d}|\boldsymbol{\psi})P(\boldsymbol{\psi})}{P(\boldsymbol{d})}\,. \tag{3.1}$$

In Bayes' theorem, the prior $P(\boldsymbol{\psi})$, which contains all information of the state $\boldsymbol{\psi}$ before the observation $\boldsymbol{d}$, is multiplied with the observation pdf $P(\boldsymbol{d}|\boldsymbol{\psi})$ to determine the posterior $P(\boldsymbol{\psi}|\boldsymbol{d})$. This product is divided by the marginal distribution of the observation $P(\boldsymbol{d})$, which is the pdf of the observation before the actual measurement is made and not the pdf of the observation itself (*Van Leeuwen et al.*, 2015). Using the joint probability $P(\boldsymbol{\psi}, \boldsymbol{d})$, the marginal of the observation can be written as:

$$P(\boldsymbol{d}) = \int d\boldsymbol{\psi} P(\boldsymbol{\psi}, \boldsymbol{d}) = \int d\boldsymbol{\psi} P(\boldsymbol{d}|\boldsymbol{\psi})P(\boldsymbol{\psi})\,. \tag{3.2}$$

Thus, $P(\boldsymbol{d})$ is a normalisation constant such that the posterior integrates to one. This factor is hard to calculate directly and is mostly determined by normalisation.

The development of data assimilation methods is mainly driven in the area of weather forecasts (*Van Leeuwen et al.*, 2015). The methods can be separated into two main categories: variational methods and sequential methods. The variational methods calculate the highest mode, i.e. the most probable value, of the

posterior without an uncertainty estimate. The uncertainty can only be estimated under high computational cost.

Two main variational data assimilation methods exist: 3D-Var and 4D-Var. While 3D-Var uses the marginal distribution for every observation, 4D-Var calculates the posterior over a time window including several observations. Both methods assume Gaussian distributed prior and observation pdfs and are sophisticated to implement. Variational data assimilation methods are used in major numerical weather prediction (NWP) centres (*Buehner et al.*, 2010a; *Lorenc*, 2011). For more details about variational data assimilation please refer to e.g. *Daley* (1993) and *Kalnay* (2003).

Sequential data assimilation or filter methods, are primarily ensemble based. The two mainly used methods are the ensemble Kalman filter (EnKF) (*Evensen*, 1994; *Burgers et al.*, 1998) and the particle filter, including variations of both. They are sequential Monte Carlo methods and approximate the pdfs in Bayes' theorem using an ensemble of states. In the case of particle filters, the ensemble members are also called particles. The EnKF relies on the assumption of Gaussian pdfs, which makes it easy to implement. Because of its simplicity, it is probably the most applied data assimilation method in geosciences (*Van Leeuwen et al.*, 2015). It is used in meteorology (*Houtekamer and Zhang*, 2016) and is also applied in NWP centres (*Buehner et al.*, 2010b; *Miyoshi et al.*, 2010).

In nonlinear systems, the Gaussian assumption is violated (*Harlim and Majda*, 2010; *DeChant and Moradkhani*, 2012; *Van Leeuwen et al.*, 2015) and the performance of the EnKF worsens (*Lei and Bickel*, 2011; *Liu et al.*, 2012; *Zhang et al.*, 2015). This requires to overcome the limitation of Gaussian distributions when using EnKF,3D-Var or 4D-Var.

Particle filters instead do not rely on Gaussian pdfs. They directly approximate the pdfs and can therefore use non-Gaussian functions. However, so far they are only efficient for low-dimensional systems, An application to high-dimensional systems is still challenging.

This chapter introduces the sequential ensemble-based data assimilation methods: EnKF and particle filter. Section 3.1 starts with the basis of data assimilation methods, the recursive Bayes' theorem, and defines a sequential filter. Section 3.2 introduces the EnKF that is based on the assumption of Gaussian pdfs. Furthermore, a brief overview in the two extensions, localisation and inflation, is given. Section 3.3 follows with the general approach of importance sampling, which leads to the particle filter. Section 3.4 introduces the concept of state augmentation that allows state and parameter estimation using data assimilation methods. This chapter is closed with a brief description of the data assimilation software developed and used in this thesis (Section 3.5).

3.1. Recursive Bayesian estimation

This section describes the recursive Bayesian estimation and the assumptions leading to a sequential filter. Suppose that the state $\boldsymbol{\psi}$ with dimension N_{ψ} is propagated with a generic model equation:

$$\boldsymbol{\psi}^k = f(\boldsymbol{\psi}^{k-1}) + \boldsymbol{\beta}^k \,, \tag{3.3}$$

where $f(\cdot)$ is the deterministic part of the mathematical model, $\boldsymbol{\beta}^k$ is a stochastic model error and k denotes a discrete time step. Furthermore, assume a set of observations $\boldsymbol{d}$ with dimension $N_{\boldsymbol{d}}$ at discrete times. The sequence of these observations is defined as:

$$\boldsymbol{d}^{1:k} := (\boldsymbol{d}^1, \boldsymbol{d}^2, \ldots, \boldsymbol{d}^{k-1}, \boldsymbol{d}^k) \,. \tag{3.4}$$

The corresponding sequence of states, including the initial state $\boldsymbol{\psi}^0$ at time zero, is also given as:

$$\boldsymbol{\psi}^{0:k} := (\boldsymbol{\psi}^0, \boldsymbol{\psi}^1, \ldots, \boldsymbol{\psi}^{k-1}, \boldsymbol{\psi}^k) \,. \tag{3.5}$$

The discrete time can be related to a continuous time t_k, where t_k is the measurement time of observation $\boldsymbol{d}^k$. A continuous time also results in a continuous model equation.

Bayes' theorem (Eq. (3.1)) is used to combine the information from observations with the information of the model. Data assimilation aims to find the posterior distribution $P(\boldsymbol{\psi}^{0:k}|\boldsymbol{d}^{1:k})$ in Bayes' theorem. With the previous defined set of observations and states, the posterior is given by:

$$P(\boldsymbol{\psi}^{0:k}|\boldsymbol{d}^{1:k}) = \frac{P(\boldsymbol{d}^{1:k}|\boldsymbol{\psi}^{0:k})P(\boldsymbol{\psi}^{0:k})}{P(\boldsymbol{d}^{1:k})} \,. \tag{3.6}$$

In a stochastic point of view, the probability is propagated forward in time using a transition density

$$P(\boldsymbol{\psi}^{0:k}) = P(\boldsymbol{\psi}^k|\boldsymbol{\psi}^{0:k-1})P(\boldsymbol{\psi}^{0:k-1}) \,. \tag{3.7}$$

It is assumed that the development of the state is a Markov process, i.e. all information for the future state is contained in the present state without additional information from the past:

$$P(\boldsymbol{\psi}^k|\boldsymbol{\psi}^{0:k-1}) = P(\boldsymbol{\psi}^k|\boldsymbol{\psi}^{k-1}) \,. \tag{3.8}$$

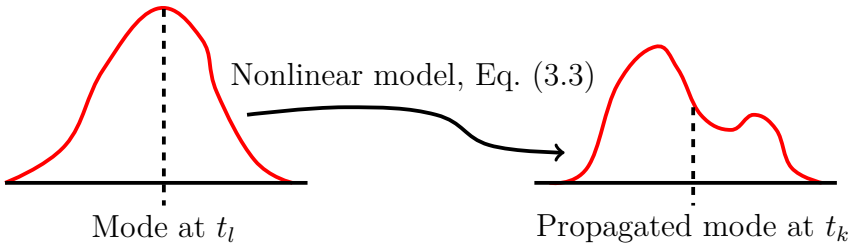

Figure 3.1.: Illustration of the problem of strong-constrained 4D-Var for nonlinear models. The mode calculated at time t_l will not be the correct mode at observation time t_k. Modified from *Van Leeuwen et al.* (2015).

Using this, Eq. (3.7) becomes

$$P(\boldsymbol{\psi}^{0:k}) = P(\boldsymbol{\psi}^k|\boldsymbol{\psi}^{k-1})P(\boldsymbol{\psi}^{0:k-1})\,, \tag{3.9}$$

which is a recursive formula that ends for the initial prior distribution $P(\boldsymbol{\psi}^0)$.

For the further derivation it is assumed that the observations are independent and only depend on the state at the current time. With this it is possible to write

$$P(\boldsymbol{d}^{1:k}|\boldsymbol{\psi}^{0:k}) = P(\boldsymbol{d}^k|\boldsymbol{\psi}^k)P(\boldsymbol{d}^{1:k-1}|\boldsymbol{\psi}^{0:k-1}) \tag{3.10}$$

and

$$P(\boldsymbol{d}^{1:k}) = P(\boldsymbol{d}^k)P(\boldsymbol{d}^{1:k-1})\,. \tag{3.11}$$

Using Eq. (3.9), Eq. (3.10) and Eq. (3.11), the posterior Eq. (3.6) becomes

$$P(\boldsymbol{\psi}^{0:k}|\boldsymbol{d}^{1:k}) = \frac{P(\boldsymbol{d}^k|\boldsymbol{\psi}^k)P(\boldsymbol{\psi}^k|\boldsymbol{\psi}^{k-1})}{P(\boldsymbol{d}^k)}P(\boldsymbol{\psi}^{0:k-1}|\boldsymbol{d}^{1:k-1})\,, \tag{3.12}$$

which defines a recursive formula for Bayes' theorem.

At this point it is necessary to differentiate between smoothers and filters. A smoother calculates the marginal distribution $P(\boldsymbol{\psi}^l|\boldsymbol{d}^{1:k})$, a state in the past using observations up to the actual observation $\boldsymbol{d}^k$. This is done, e.g. for 4D-Var, which calculates the mode of the distribution at time l. For nonlinear systems, the mode at time l will not be the propagated mode at time k (*Van Leeuwen et al.*, 2015), which is illustrated in Fig. 3.1.

In contrast, a filter calculates the current state using the current observation $P(\boldsymbol{\psi}^k|\boldsymbol{d}^{1:k})$. This allows to assimilate new observations sequentially. The most prominent filters are the (ensemble) Kalman filter and the particle filter.

To calculate the distribution $P(\boldsymbol{\psi}^k|\boldsymbol{d}^{1:k})$, it is necessary to integrate over all past states such that that Eq. (3.12) only depends on $\boldsymbol{\psi}^k$:

$$P(\boldsymbol{\psi}^k|\boldsymbol{d}^{1:k}) = \frac{P(\boldsymbol{d}^k|\boldsymbol{\psi}^k)}{P(\boldsymbol{d}^k)} \int d\boldsymbol{\psi}^{k-1} P(\boldsymbol{\psi}^k|\boldsymbol{\psi}^{k-1}) P(\boldsymbol{\psi}^{k-1}|\boldsymbol{d}^{1:k-1}) \,. \tag{3.13}$$

The filter can be separated into a forecast (prediction) and an analysis (update) step. The forecast propagates the pdf forward in time to the next observation:

$$P(\boldsymbol{\psi}^k|\boldsymbol{d}^{1:k-1}) = \int d\boldsymbol{\psi}^{k-1} P(\boldsymbol{\psi}^k|\boldsymbol{\psi}^{k-1}) P(\boldsymbol{\psi}^{k-1}|\boldsymbol{d}^{1:k-1}) \,. \tag{3.14}$$

The resulting pdf is then updated in the analysis step with the new observation using

$$P(\boldsymbol{\psi}^k|\boldsymbol{d}^{1:k}) = \frac{P(\boldsymbol{d}^k|\boldsymbol{\psi}^k) P(\boldsymbol{\psi}^k|\boldsymbol{d}^{1:k-1})}{P(\boldsymbol{d}^k)} \,, \tag{3.15}$$

which results in the posterior pdf. The forecast and analysis steps are repeated with every new observation.

This section derived the recursive Bayes' theorem (Eq. (3.12)) and the sequential filter equation (Eq. (3.13)). The filter equation was derived assuming a Markov process for the state development and observations only dependent on the state at its measurement time. Without these assumptions, it is still possible to derive a recursive form of Bayes' theorem (see Appendix A.1). However, without the Markov assumption, it is impossible to define a sequential update, as it would be necessary to start the calculations form the beginning with each new observation (*Vetra-Carvalho et al.*, 2018).

3.2. Ensemble Kalman filter

The resulting integrals (Eq. (3.14) and Eq. (3.2)) are hard to calculate directly without using simplifying assumptions. Before discussing the general approach with Monte Carlo methods, it is assumed that all pdfs are Gaussian distributed.

With this assumption the analysis Eq. (3.15) becomes

$$P(\boldsymbol{\psi}^k|\boldsymbol{d}^{1:k}) \propto \underbrace{\exp\left[-\frac{1}{2}(\boldsymbol{d}^k - \mathbf{H}(\boldsymbol{\psi}^k))^\intercal \mathbf{R}^{-1} (\boldsymbol{d}^k - \mathbf{H}(\boldsymbol{\psi}^k))\right]}_{P(\boldsymbol{d}^k|\boldsymbol{\psi}^k)}$$
$$\cdot \underbrace{\exp\left[-\frac{1}{2}(\boldsymbol{\psi}^k - \boldsymbol{\psi}_\mathrm{f}^k)^\intercal \mathbf{P}_\mathrm{f}^{-1} (\boldsymbol{\psi}^k - \boldsymbol{\psi}_\mathrm{f}^k)\right]}_{P(\boldsymbol{\psi}^k|\boldsymbol{d}^{1:k-1})} \tag{3.16}$$
$$\propto \exp\left[-\frac{1}{2}(\boldsymbol{\psi}^k - \boldsymbol{\psi}_\mathrm{a}^k)^\intercal \mathbf{P}_\mathrm{a}^{-1} (\boldsymbol{\psi}^k - \boldsymbol{\psi}_\mathrm{a}^k)\right], \tag{3.17}$$

with the observation error covariance $\mathbf{R}$, the error covariance matrix $\mathbf{P}$ and the typically nonlinear observation operator $\mathbf{H}(\cdot)$, which maps the state from state space to observation space $\mathbf{H} : \mathbb{R}^{N_\psi} \to \mathbb{R}^{N_d}$. The indices 'a' and 'f' denote the analysis and the forecast, respectively.

The analysis state $\boldsymbol{\psi}_\mathrm{a}$ is derived by completing the squares such that the posterior is a Gaussian again. This constrains the observation operator to be linear. The analysis state becomes:

$$\boldsymbol{\psi}_\mathrm{a}^k = \boldsymbol{\psi}_\mathrm{f}^k + \mathbf{K}^k \left(\boldsymbol{d}^k - \mathbf{H}\boldsymbol{\psi}_\mathrm{f}^k\right) \tag{3.18}$$

with the Kalman gain

$$\mathbf{K}^k = \mathbf{P}_\mathrm{f}^k \mathbf{H}^\intercal (\mathbf{H}\mathbf{P}_\mathrm{f}^k \mathbf{H}^\intercal + \mathbf{R}^k)^{-1}. \tag{3.19}$$

The posterior covariance is given by

$$\mathbf{P}_\mathrm{a}^k = (\mathbf{1} - \mathbf{K}^k\mathbf{H})\mathbf{P}_\mathrm{f}^k. \tag{3.20}$$

For a linear model

$$\boldsymbol{\psi}^k = \mathbf{M}^{\mathbf{k-1}}\boldsymbol{\psi}^{k-1} + \boldsymbol{\beta}^k, \tag{3.21}$$

where $\mathbf{M}$ is a linear map and $\boldsymbol{\beta} \propto \mathcal{N}(\mathbf{0}, \mathbf{Q})$ a stochastic model error, the forecast of the error covariance Eq. (3.14) to the next observation is

$$\mathbf{P}_\mathrm{f}^k = \mathbf{M}^{k-1}\mathbf{P}^{k-1}(\mathbf{M}^{k-1})^\intercal + \mathbf{Q}^k. \tag{3.22}$$

This is the original Kalman filter (*Kalman*, 1960).

For nonlinear models, however, solving the high-dimensional integral Eq. (3.14) in a closed form is only possible for some special cases. Therefore, the integral is generally solved numerically using Monte Carlo methods. In the case of the

Kalman filter, the ensemble Kalman filter (EnKF) was introduced (*Evensen*, 1994; *Burgers et al.*, 1998), which approximates the Gaussian distributions by an ensemble of states. Each ensemble member is then propagated forward in time with the nonlinear model Eq. (3.3). The forecast error covariance $\mathbf{P}_\mathrm{f}^k$ is then given by the propagated ensemble:

$$\mathbf{P}_\mathrm{f}^k = \frac{1}{N-1}\sum_{i=0}^{N}(\boldsymbol{\psi}_{\mathrm{f},i}^k - \overline{\boldsymbol{\psi}_\mathrm{f}^k})^2\,, \tag{3.23}$$

with the ensemble size N, the forecast state of the i-th ensemble member $\boldsymbol{\psi}_{\mathrm{f},i}^k$ and the mean

$$\overline{\boldsymbol{\psi}_\mathrm{f}^k} = \frac{1}{N}\sum_{i=0}^{N}\boldsymbol{\psi}_{\mathrm{f},i}^k\,. \tag{3.24}$$

In the analysis, the Kalman gain Eq. (3.19) is applied to every ensemble member. However, it is necessary to add a realisation of the observation error to the observation for each ensemble member. This changes Eq. (3.18) to

$$\boldsymbol{\psi}_{\mathrm{a},i}^k = \boldsymbol{\psi}_{\mathrm{f},i}^k + \mathbf{K}^k\left(\boldsymbol{d}^k - \mathbf{H}\boldsymbol{\psi}_{\mathrm{f},i}^k + \boldsymbol{\epsilon}_i^k\right) \quad \text{with} \quad \boldsymbol{\epsilon}_i^k \propto \mathcal{N}(\mathbf{0}, \mathbf{R}^k) \tag{3.25}$$

such that the update of the error covariance (Eq. (3.20)) is still correct (*Burgers et al.*, 1998). For sequential data assimilation, the process of forecast and analysis is iterated for every new observation. This cycle is visualised in Fig. 3.2 for the EnKF.

In high-dimensional systems only a small ensemble size $\mathcal{O}(100)$ can be afforded, which leads to rank deficient matrices (*Van Leeuwen et al.*, 2015; *Houtekamer and Zhang*, 2016). To counteract this and issues caused by non Gaussian distributions, methods like inflation (e.g. *Anderson and Anderson*, 1999; *Bauser et al.*, 2018; *Gharamti*, 2018), localisation (*Houtekamer and Mitchell*, 1998, 2001; *Hamill et al.*, 2001) and damping (*Franssen and Kinzelbach*, 2008; *Wu and Margulis*, 2011) are introduced to the EnKF.

3.2.1. Localisation

The EnKF approximates the covariance matrix by an ensemble (Eq. (3.23)). In high-dimensional systems, the ensemble size N is generally much smaller than the system dimension $N_{\boldsymbol{\psi}}$ and often also much smaller than the number of observations $N_{\boldsymbol{d}}$. However, the ensemble can only provide $N-1$ directions in phase space for the calculation of the covariance matrix, which leads to a covariance matrix

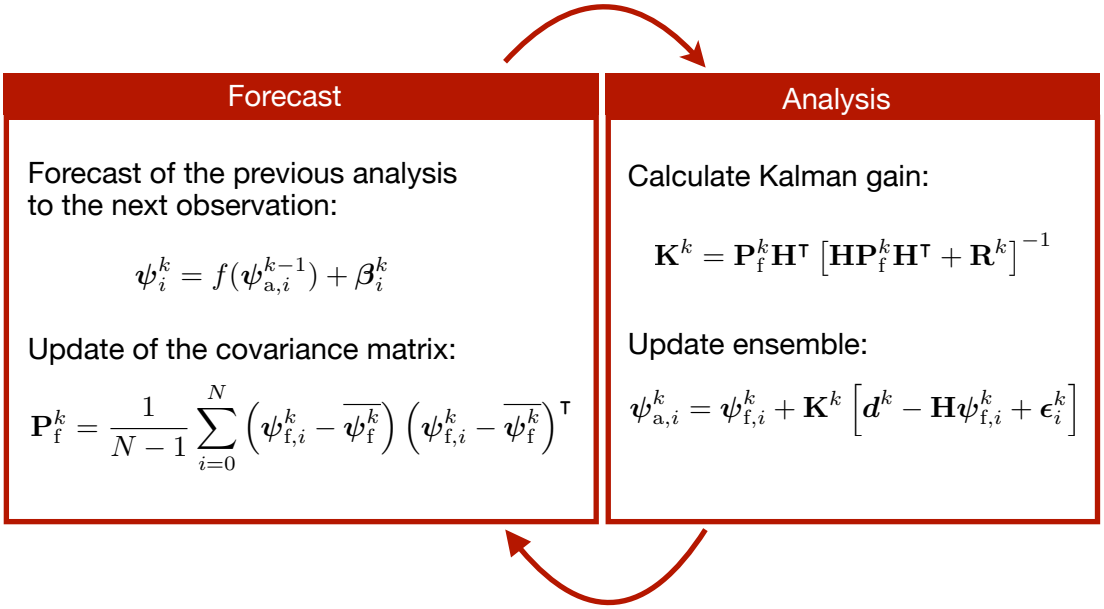

Figure 3.2.: Illustration of the forecast and analysis cycle of the ensemble Kalman filter. The forecast (Eq. (3.14)) of the error covariance is realised by propagating the ensemble with the nonlinear model equation (Eq. (3.3)) and then calculating the propagated error covariance using the ensemble (Eq. (3.23)). In the analysis (Eq. (3.15)), the ensemble is updated using the Kalman gain (Eq. (3.19)) and the newest observation $\boldsymbol{d}^k$. Modified from *Welch and Bishop* (2006).

that does not have full rank. This issue is called the rank (deficiency) problem (*Lorenc*, 2003; *Houtekamer and Zhang*, 2016).

In physical systems, the correlation between two points in space generally decreases with increasing distance. However, due to the rank problem, the noise of the covariance estimates becomes larger than the actual signal for increasing distances (*Hamill et al.*, 2001), which leads to spurious long-range correlations.

This problem is addressed using localisation. Localisation reduces the high-dimensional data assimilation problem to several local problems. Each of these local problems has $N-1$ local directions of the ensemble available. This is done by introducing a localisation matrix ρ that reduces the covariance for increasing distance. The localisation matrix is multiplied with the error covariance matrix $\mathbf{P}_\mathrm{f}$:

$$\mathbf{P}_\mathrm{f}^k \quad \rightarrow \quad \rho \circ \mathbf{P}_\mathrm{f}^k\,, \tag{3.26}$$

where $\circ$ is the entrywise product (Hadamard product). This changes the Kalman gain (Eq. (3.19)) to

$$\mathbf{K}^k = (\rho \circ \mathbf{P}_\mathrm{f}^k)\mathbf{H}^\intercal \left[\mathbf{H}(\rho \circ \mathbf{P}_\mathrm{f}^k)\mathbf{H}^\intercal + \mathbf{R}^k\right]^{-1} . \tag{3.27}$$

An often used localisation matrix to relate the distance $|z| = |i - j|$ between two points (i, j) with a length scale c is the fifth-order polynomial Gaspari and Cohn (*Gaspari and Cohn*, 1999) function

$$\rho_{i,j}(a) = \begin{cases} -\frac{1}{4}a^5 + \frac{1}{2}a^4 + \frac{5}{8}a^3 - \frac{5}{3}a^2 + 1 & 0 \leq a \leq 1 \\ \frac{1}{12}a^5 - \frac{1}{2}a^4 + \frac{5}{8}a^3 + \frac{5}{3}a^2 - 5a + 4 - \frac{2}{3}a^{-1} & 1 \leq a \leq 2 \\ 0 & \text{else} \end{cases} \tag{3.28}$$

with $a = |z|/c$. This function approximates a Gaussian with half-width c and has a compact support i.e. becomes 0 for $|z| > 2c$. For high-dimensional applications, localisation is an essential component of the EnKF (*Hamill et al.*, 2001; *Anderson*, 2012) to provide efficient data assimilation with feasible ensemble sizes.

3.2.2. Inflation

Localisation alone cannot compensate the effects of a finite ensemble size completely. This leads to underestimated error covariances (*van Leeuwen*, 1999). Unrepresented model errors in the forecast (*Whitaker and Hamill*, 2012; *Houtekamer and Zhang*, 2016) and the propagation of the ensemble with a nonlinear model equation (Eq. (3.3)) increase this effect. In nonlinear systems, the Gaussian assumption is violated (*Harlim and Majda*, 2010; *DeChant and Moradkhani*, 2012), which leads to biases that result in an underestimated error covariance (*Lei and Bickel*, 2011). If the ensemble spread becomes too small, the Kalman filter degenerates. This means that new observations no longer lead to a change in the states (Kalman gain $\mathbf{K} \approx \mathbf{0}$), which causes data assimilation to fail.

Inflation methods focus on increasing the ensemble spread again and alleviate the problem. They can be separated into additive inflation, multiplicative inflation and relaxation methods. A summary of these methods can be found in the EnKF review of *Houtekamer and Zhang* (2016). Newest developments focus on adaptive inflation methods (*Bauser et al.*, 2018; *Gharamti*, 2018).

3.3. Particle filter

3.3.1. Importance Sampling

This subsection discusses the importance sampling and follows *Van Leeuwen et al.* (2015). The Gaussian assumption does not hold in nonlinear systems (*Harlim and Majda*, 2010; *Moradkhani et al.*, 2012) and the performance of the EnKF worsens (*Lei and Bickel*, 2011; *Liu et al.*, 2012; *Zhang et al.*, 2015). Without loss of generality, a scalar random variable x is used in the following, to simplify the notation. Instead of assuming Gaussian distributions, the integrals can be approximated with importance sampling. The idea of importance sampling is to calculate the variable of interest by drawing samples of the distribution $P(x)$

$$\overline{f(x)} = \int dx f(x) P(x) \,. \tag{3.29}$$

If it is possible to sample directly from the distribution $P(x)$, the distribution can be approximated by N samples $x_i \sim P(x)$

$$P(x) \simeq \frac{1}{N} \sum_{i=1}^{N} \delta(x - x_i) \,, \tag{3.30}$$

with the Dirac delta distribution $\delta(x - x_i)$. Then the expectation value can be calculated with

$$\overline{f(x)}_N = \sum_{i=1}^{N} \int dx f(x) \delta(x - x_i) = \frac{1}{N} \sum_{i=1}^{N} f(x_i) \,. \tag{3.31}$$

However, drawing samples from an arbitrary distribution can be difficult and inefficient. Therefore, it can be preferable to draw samples from a proposal distribution $\Pi(x)$ from which it is easy to draw, e.g. a Gaussian or a uniform distribution. The expectation value then reads

$$\overline{f(x)} = \int dx f(x) \frac{P(x)}{\Pi(x)} \Pi(x) \,, \tag{3.32}$$

which is actually a multiplication by one and does not change the mean $\overline{f(x)}$. Instead using $P(x)$, samples x_i are drawn from the proposal $\Pi(x)$

$$\Pi(x) \simeq \frac{1}{N} \sum_{i=1}^{N} \delta(x - x_i) \,, \tag{3.33}$$

This leads to:

$$\overline{f(x)}_N = \sum_{i=1}^{N} \int dx f(x) \frac{P(x)}{\Pi(x)} \delta(x - x_i) \tag{3.34}$$

$$= \sum_{i=1}^{N} f(x_i) \frac{P(x_i)}{\Pi(x_i)}\,, \tag{3.35}$$

such that the distribution $P(x)$ is approximated by a weighted ensemble

$$P(x) \simeq \sum_{i=1}^{N} \delta(x - x_i) \frac{P(x_i)}{\Pi(x_i)} = \sum_{i=1}^{N} w_i\, \delta(x - x_i)\,, \tag{3.36}$$

with weights $w_i = \frac{P(x_i)}{\Pi(x_i)}$. The samples x_i are drawn from $\Pi(x)$. It is important that the support of the proposal is larger than the support of the original distribution, otherwise a division by zero would occur.

3.3.2. Standard particle filter

Particle filters are based on importance sampling. The posterior is approximated by a weighted ensemble. The ensemble members are also called particles. The posterior at the previous time step is given as a weighted ensemble

$$P(\boldsymbol{\psi}^k | \boldsymbol{d}^{1:k-1}) = \sum_{i=1}^{N} w_i^{k-1} \delta(\boldsymbol{\psi}^{k-1} - \boldsymbol{\psi}_i^{k-1})\,. \tag{3.37}$$

Inserting this equation in the filter equation (Eq. (3.13)) results in:

$$P(\boldsymbol{\psi}^k | \boldsymbol{d}^{1:k}) = \frac{P(\boldsymbol{d}^k | \boldsymbol{\psi}^k)}{P(\boldsymbol{d}^k)} \sum_{i=1}^{N} \int d\boldsymbol{\psi}^{k-1} P(\boldsymbol{\psi}^k | \boldsymbol{\psi}^{k-1}) w_i^{k-1} \delta(\boldsymbol{\psi}^{k-1} - \boldsymbol{\psi}_i^{k-1}) \tag{3.38}$$

$$= \sum_{i=1}^{N} w_i^{k-1} \frac{P(\boldsymbol{d}^k | \boldsymbol{\psi}^k)}{P(\boldsymbol{d}^k)} P(\boldsymbol{\psi}^k | \boldsymbol{\psi}_i^{k-1})\,. \tag{3.39}$$

The ensemble is propagated forward in time using the nonlinear model Eq. (3.3). The transition density then becomes

$$P(\boldsymbol{\psi}^k | \boldsymbol{\psi}_i^{k-1}) = \delta(\boldsymbol{\psi}^k - \boldsymbol{\psi}_i^k)\,, \tag{3.40}$$

where $\boldsymbol{\psi}_i^k$ is the propagated particle including the stochastic model error. Using this equation and reformulating it such that the posterior $P(\boldsymbol{\psi}^k | \boldsymbol{d}^{1:k})$ is represented as a weighted ensemble

$$P(\boldsymbol{\psi}^k | \boldsymbol{d}^{1:k}) = \sum_{i=1}^{N} w_i^k \delta(\boldsymbol{\psi}^k - \boldsymbol{\psi}_i^k) \tag{3.41}$$

the weights are then updated with

$$w_i^k = w_i^{k-1} \frac{P(\boldsymbol{d}^k|\boldsymbol{\psi}_i^k)}{P(\boldsymbol{d}^k)} \,. \tag{3.42}$$

It is difficult to calculate the pdf $P(\boldsymbol{d}^k)$ directly. However, since the sum of all weights has to be one, the pdf can be obtained in a normalisation process

$$\sum_{i=0}^{N} w_i^k \stackrel{!}{=} 1 \quad \Rightarrow \quad P(\boldsymbol{d}^k) = \sum_{i=0}^{N} w_i^{k-1} P(\boldsymbol{d}^k|\boldsymbol{\psi}_i^k) \,. \tag{3.43}$$

Particle filters directly represent the posterior by a weighted ensemble, which allows to have arbitrarily distributions. However, they tend to filter degeneracy, which is also referred to as filter impoverishment. The variance of the weights increases and after several analysis steps, one particle gets all statistical information as its weight becomes increasingly large. The remaining particles only get a small weight such that the ensemble effectively collapses to this one particle. Because of this, the filter does not react to new observations and follows the particle with the large weight. This situation is illustrated in Fig. 3.3.

An estimate for the number of significant ensemble members is given by the effective sample size (*Kong et al.*, 1994; *Doucet*, 1998):

$$N_{\text{eff}} = \frac{1}{\sum_{i=1}^{N} w_i^2} \,. \tag{3.44}$$

For example, if one particle accumulates all the weight $N_{\text{eff}} = 1$, the ensemble is effectively described only by this particle.

3.3.3. Resampling

The problem of filter degeneration caused a rare usage of particle filters in data assimilation (*van Leeuwen*, 2009). *Gordon et al.* (1993) introduced resampling to particle filters, a technique that reduces the variance in the weights and has the potential to prevent filter degeneracy. The idea of resampling is that after the analysis, particles with large weights are replicated and particles with small weights are dropped. This helps that the regions with high weighted particles are represented better by the ensemble, which alleviates the degeneracy of the filter. The forecast and analysis cycle for a particle filter with resampling is illustrated in Fig. 3.4.

There are many different resampling algorithms (see *van Leeuwen* (2009) for a summary). In the following, two techniques will be presented: the residual resampling and the stochastic universal resampling.

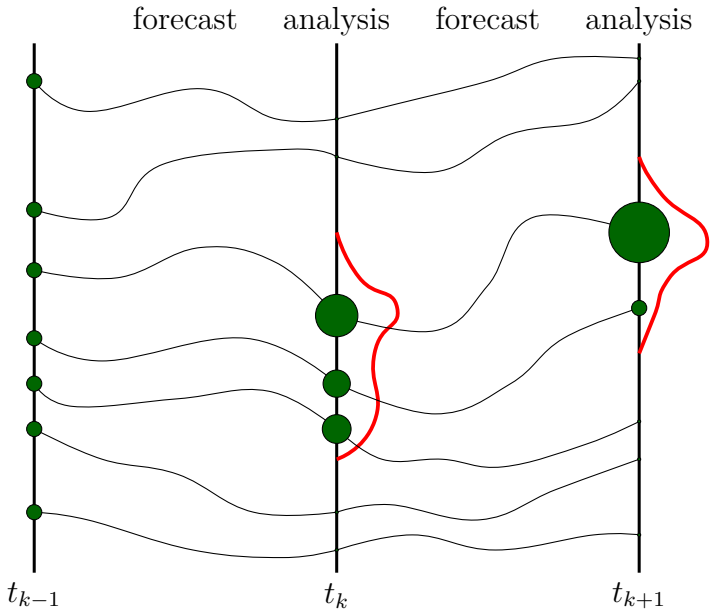

Figure 3.3.: Illustration of a particle filter with filter degeneration. The weights are shown as green dots. The weight of a particle corresponds to the size of the dot. At time t_{k-1} all particles have the same weight. The ensemble is propagated forward in time from t_{k-1} to t_k. The particles are weighted (Eq. (3.42)) according to the observation pdf $P(\boldsymbol{d}^k|\boldsymbol{\psi}_i^k)$ (red). In this example, three particles get a significant weight and four particles get a negligible weight. The ensemble is propagated to the next observation time t_{k+1}. After the weighting at t_{k+1} one particle has a large weight and the effective sample size (Eq. (3.44)) is $N_{\text{eff}} \approx 1$. The filter is degenerated. Modified from *van Leeuwen* (2009).

Residual resampling The residual resampling was introduced by *Liu and Chen* (1998). In this method all weights are multiplied with the ensemble size N. The floor of the resulting numbers $z_i = \lfloor Nw_i \rfloor$ gives the number of copies of a particle i. Then these integer parts are subtracted from Nw_i. The necessary remaining particles to have an ensemble size of N again are drawn randomly with probabilities proportional to the residual weights

$$r_i = Nw_i - z_i \,. \tag{3.45}$$

After the resampling step, all weights are set to N^{-1}.

Stochastic universal resampling The stochastic universal resampling (*Kitagawa*, 1996) can be summarised as follows (see also Fig. 3.5): All weights are aligned on an interval $[0, 1]$. A random number in the interval $[0, N^{-1}]$ is drawn

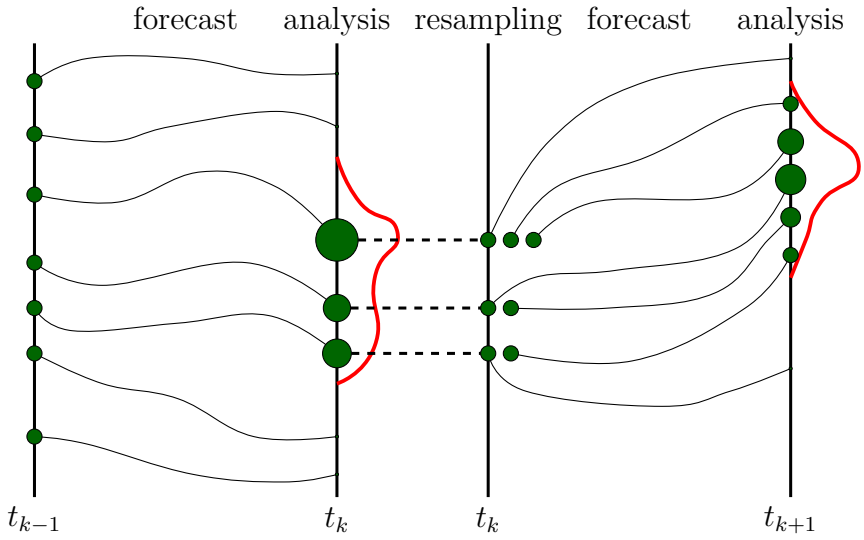

Figure 3.4.: Illustration of a particle filter with resampling. The weights are shown as green dots. The weight of a particle corresponds to the size of the dot. At time t_{k-1} all particles have the same weight. The ensemble is propagated forward in time from t_{k-1} to t_k. The particles are weighted (Eq. (3.42)) according to the observation pdf $P(\boldsymbol{d}^k|\boldsymbol{\psi}_i^k)$ (red). In this example, three particles get a large weight and four particles get a negligible weight. The resampling duplicates particles with large weight and drops the other particles. After the resampling, all particles are equal weighted again. The identical particles become distinguishable during the next forecast due to the stochastic model error (Eq. (3.3)). Modified from *van Leeuwen* (2009).

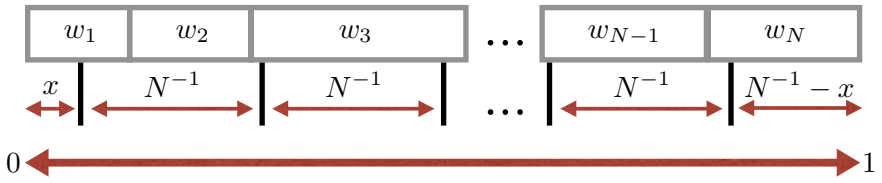

Figure 3.5.: Illustration of the stochastic universal resampling process. A random number x is drawn from a uniform distribution in the interval $[0, N^{-1}]$. The endpoint of this number indicates the first particle. Then, N^{-1} is added $(N-1)$-times to this random number, where every endpoint is a particle of the new ensemble. In this illustration, particle one is chosen once, particle two is not chosen and particle three is chosen twice. This way some particles are replicated and other particles are dropped. If the model does not have a stochastic model error, it is necessary to perturb the new particles, otherwise they would be identical and the filter would degenerate.

from a uniform distribution. This number points to the first particle of the new ensemble selected by the corresponding weight. Then, N^{-1} is added $(N-1)$-times to x. Each endpoint selects the corresponding particle for the new ensemble. This way, some particles get duplicated and some particles are dropped. With this approach, particles with a weight smaller than N^{-1} can be chosen maximally once, whereas a weight larger than N^{-1} guarantees that the particle is at least chosen once. If all particles have equal weights, no particle is dropped. The result is a new set of N particles. After the resampling step, all weights are set to N^{-1}. The stochastic universal resampling has the lowest sampling noise compared to other resampling methods (*van Leeuwen*, 2009).

3.3.4. Optimal proposal

Similar to the importance sampling, a proposal distribution Π can be introduced in Eq. (3.39):

$$P(\boldsymbol{\psi}^k|\boldsymbol{d}^{1:k}) = \sum_{i=1}^{N} w_i^{k-1} \frac{P(\boldsymbol{d}^k|\boldsymbol{\psi}^k)}{P(\boldsymbol{d}^k)} \frac{P(\boldsymbol{\psi}^k|\boldsymbol{\psi}_i^{k-1})}{\Pi(\boldsymbol{\psi}^k|\boldsymbol{\psi}_i^{k-1}, \boldsymbol{d}^k)} \Pi(\boldsymbol{\psi}^k|\boldsymbol{\psi}_i^{k-1}, \boldsymbol{d}^k)\,. \tag{3.46}$$

The difference is now, that a sample $\boldsymbol{\psi}_i^k$ is now generated from the proposal

$$\Pi(\boldsymbol{\psi}_i^k|\boldsymbol{\psi}_i^{k-1}, \boldsymbol{d}^k) = \delta(\boldsymbol{\psi}^k - \boldsymbol{\psi}_i^k)\,. \tag{3.47}$$

This can be a direct sample from the proposal or can be used to propagate the ensemble with a different model equation (*van Leeuwen*, 2010). Representing the posterior as a weighted ensemble, the weight update Eq. (3.42) changes to:

$$w_i^k = \frac{P(\boldsymbol{d}^k|\boldsymbol{\psi}_i^k)}{P(\boldsymbol{d}^k)} \frac{P(\boldsymbol{\psi}_i^k|\boldsymbol{\psi}_i^{k-1})}{\Pi(\boldsymbol{\psi}_i^k|\boldsymbol{\psi}_i^{k-1}, \boldsymbol{d}^k)} w_i^{k-1}\,. \tag{3.48}$$

The weights now have an additional term that includes the proposal and the original transition density.

With this extension, an optimal proposal can be defined that minimises the variance in the weights (*Doucet*, 1998; *Doucet et al.*, 2000)

$$\Pi(\boldsymbol{\psi}^k|\boldsymbol{\psi}_i^{k-1}, \boldsymbol{d}^k) = P(\boldsymbol{\psi}^k|\boldsymbol{\psi}_i^{k-1}, \boldsymbol{d}^k)\,. \tag{3.49}$$

The optimal proposal can reduce the necessary ensemble size significantly. However, sampling from the optimal proposal is typically not possible since the proposal distribution is generally unknown because of the conditioning on the future observation.

3.3.5. Curse of Dimensionality

Resampling alleviates filter degeneration of the standard particle. However, resampling is not sufficient to avoid filter degeneration in high-dimensional systems. Based on the work of *Bengtsson et al.* (2008) and *Bickel et al.* (2008), *Snyder et al.* (2008) showed that the necessary ensemble size N to avoid filter degeneration increases exponentially with the variance of the logarithmic weights τ^2. This variance τ^2 can be regarded as an effective dimension that depends not only on the state but also on the prior and the observations (*Snyder et al.*, 2008). The so-called curse of dimensionality is a challenging issue for nonlinear data assimilation in high-dimensional systems with particle filters.

Van Leeuwen et al. (2015) show in an illustrative example that an increasing observation dimension results in filter degeneration. This argumentation will be shown in the following:

Assume a set of $N_{\boldsymbol{d}}$ observations and two particles at one particular time step. One particle is always 0.1σ away from the observations and the other particle is always 0.2σ away, where σ is the standard deviation of each Gaussian distributed observation. The weights of the particles are calculated using Eq. (3.42):

$$w_1 = A \exp\left[-\frac{1}{2}(\boldsymbol{d} - \mathbf{H}(\boldsymbol{\psi}_1)^\intercal \mathbf{R}^{-1}(\boldsymbol{d} - \mathbf{H}(\boldsymbol{\psi}_1)\right] = A \exp\left[-0.005 N_{\boldsymbol{d}}\right] \tag{3.50}$$

for the first particle and

$$w_1 = A \exp\left[-\frac{1}{2}(\boldsymbol{d} - \mathbf{H}(\boldsymbol{\psi}_2)^\intercal \mathbf{R}^{-1}(\boldsymbol{d} - \mathbf{H}(\boldsymbol{\psi}_2)\right] = A \exp\left[-0.02 N_{\boldsymbol{d}}\right] \tag{3.51}$$

for the second particle. The ratio of the weights is

$$\frac{w_2}{w_1} = \exp\left[-0.015 N_{\boldsymbol{d}}\right] . \tag{3.52}$$

For 1000 independent observations, the ratio is

$$\frac{w_2}{w_1} = \exp\left[-15\right] \approx 3 \cdot 10^{-7} . \tag{3.53}$$

In this case, the first particle has a negligible weight compared to the second one, even though both particles are close to the observation. This example shows that the number of observations is crucial for filter degeneration.

Different methods based on the freedom of the proposal density (Section 3.3.4) were suggested to lift the curse of dimensionality, e.g. the implicit particle filter (*Chorin and Tu*, 2009; *Morzfeld et al.*, 2012; *Chorin and Morzfeld*, 2013) and the equal weights particle filter (*van Leeuwen*, 2010, 2011; *van Leeuwen and Ades*,

2013). *Snyder et al.* (2015) showed that the optimal proposal can increase the efficiency significantly but still needs an exponentially increasing ensemble size to avoid filter degeneracy. Only the equal weights filter, further studied by *Ades and van Leeuwen* (2013, 2015a,b), is partly not covered by the study because the proposal becomes sharper with increasing ensemble size (*Snyder et al.*, 2015).

Localisation for particle filters is another approach to lift the curse of dimensionality. *Metref et al.* (2014) and *Snyder et al.* (2015) argued that localisation (Section 3.2.1) causes the EnKF to be successful in high-dimensional systems with ensemble size $\mathcal{O}(100)$. They conclude that localisation is also necessary for high-dimensional particle filtering to limit the influence of the observations and locally update the ensemble. The curse of dimensionality and particle filters with localisation will be revisited in Chapter 6.

3.4. State and parameter estimation

The state and parameter estimation is realised by state augmentation. The original state $\boldsymbol{\psi}$ is extended with the parameters $\boldsymbol{p}$ to an augmented state

$$\boldsymbol{u} = \begin{pmatrix} \boldsymbol{\psi} \\ \boldsymbol{p} \end{pmatrix} . \tag{3.54}$$

The model equation (Eq. (3.3)) changes to

$$\boldsymbol{u}^k = \begin{pmatrix} \boldsymbol{\psi}^k \\ \boldsymbol{p}^k \end{pmatrix} = \begin{pmatrix} f_{\boldsymbol{\psi}}\left(\boldsymbol{\psi}^{k-1}\right) + \boldsymbol{\beta}_{\boldsymbol{\psi}}^k \\ f_{\boldsymbol{p}}\left(\boldsymbol{p}^{k-1}\right) + \boldsymbol{\beta}_{\boldsymbol{p}}^k \end{pmatrix} , \tag{3.55}$$

with the deterministic models $f_{\boldsymbol{\psi}}$ and $f_{\boldsymbol{p}}$, as well as, the stochastic model errors $\boldsymbol{\beta}_{\boldsymbol{\psi}}$ and $\boldsymbol{\beta}_{\boldsymbol{p}}$ for state and parameters, respectively. Parameters are often constant in time $f_{\boldsymbol{p}}\left(\boldsymbol{p}^{k-1}\right) = \boldsymbol{p}^{k-1}$.

With the augmented state, parameter estimation is straightforward for the EnKF. The Kalman gain (Eq. (3.19)) transfers the information from observed to unobserved dimensions using the error covariance (Eq. (3.23)). The augmented state can also be used to incorporate nonlinear observation operators to the EnKF.

For particle filters, however, the transfer from state to combined state and parameter estimation is challenging. Since parameters are often constant in time, a specification of a model error $\boldsymbol{\beta}_{\boldsymbol{p}}$ would be misleading. However, without a model error, the variety of the ensemble decreases with every resampling since the resampled particles are identical copies of high weighted particles and will stay identical during the forward propagation. This leads to filter degeneration after a few analysis steps.

Parameter estimation for particle filters will be revisited in Chapter 5 with the description of a new resampling method.

3.5. Software – KnoFu

During the thesis I developed the data assimilation software KnoFu (Knowledge Fusion). It is designed to provide a flexible framework for the development of new ideas in the context of data assimilation for the purpose of knowledge fusion. Knowledge fusion is the idea of a consistent aggregation of all information pertinent to some observed reality (*Bauser et al.*, 2016).

One part of the flexibility is the augmented state (Section 3.4). The concept of the augmented state is directly mapped to the software, which allows state and parameter estimation. It also gives the ability to extend the augmented state to estimate other model components as well.

The software contains different ordinary differential equation (ODE) solvers, for instance a fourth order Runge-Kutta method. This has the advantage that new ODE models can be easily defined. Different ODE models are part of the package, including the Lorenz-63 model (*Lorenz*, 1963) and the Lorenz-96 model (Section 2.2), which are often used toy models to test data assimilation methods.

To perform data assimilation for soil hydrology, KnoFu is coupled to MuPhi (*Ippisch et al.*, 2006) in cooperation with *Jaumann* (2018), and DORiE (DUNE-Operated Richards equation solving Environment) (*Riedel et al.*, 2018), which is based on DUNE (Distributed and Unified Numerics Environment) (*Blatt et al.*, 2016). The coupling of DORiE was done by L. Riedel.

For data assimilation, KnoFu provides the ensemble Kalman filter (EnKF), the ensemble square root filter (EnSRF) (*Tippett et al.*, 2003) and the standard particle filter. The EnKF and the EnSRF, including localisation, damping and several inflation methods were implemented to the framework by *Bauser* (2018). The particle filter contains different resampling methods including the covariance resampling and the localised version of the covariance resampling.

KnoFu is part of an ongoing PhD-project by L. Riedel. So far, it was used for two papers (*Bauser et al.*, 2018; *Berg et al.*, 2018), a finished PhD project (*Bauser*, 2018), several finished master theses, a bachelor thesis and two project practicals (C. Pixius, M. Weiler). Table 3.1 lists the titles and authors of the bachelor and master theses.

The results in Chapter 4, Chapter 5 and Chapter 6 are all produced using this software.

Table 3.1.: Finished master (MSc) and bachelor (BSc) theses using KnoFu.

Author	Title, Year	Type
K. Mosny	Enhancing the Particle Filter by Markov-Chain-Monte-Carlo methods – A performance assessment, 2017	MSc
F. Oraschewski	Estimating effective soil hydraulic parameters for two-dimensional heterogeneous media using the Ensemble Kalman Filter, 2018	BSc
L. Riedel	Advancing Data Assimilation of Soil Water Flow with a DG Based Richards Solver - Merging the KnoFu and DORiE Software Packages, 2017	MSc
F. Riexinger	Food Webs: Augmented State Estimation and Model Enhancement using the Ensemble Kalman Filter, 2017	MSc
M. Walch	A Study On Stochastic Nonlinear Hyperbolic Partial Differential Equations with Bayesian Filtering Methods, 2017	MSc

4. Convergent and divergent systems

Data assimilation methods are mainly developed in meteorology. They are therefore designed to meet the challenges in the atmosphere – a divergent system. In a divergent system, close states will inevitably drift apart, even if the system is described by a perfect model (*Kalnay*, 2003). This gives an upper limit to the predictability in divergent systems (*Lorenz*, 1982).

From the meteorological perspective, the major limit for long time predictions is an unknown initial condition because its error will increase exponentially. Therefore, data assimilation in operational weather forecasting primarily focusses on state estimation (*Reichle*, 2008; *Van Leeuwen et al.*, 2015).

Data assimilation methods are also increasingly applied to convergent systems such as soil hydrology. In a convergent system, nearby trajectories will stay nearby and even become more similar. If the model to describe such a convergent system is perfect, this results in a high predictability (*Lorenz*, 1996). An error in the initial state will decay towards the truth after some transient phase. However, this is only true if the parameters are perfectly known and the model is correct, which is usually not the case for environmental systems. In many systems, the parameters are ill- or unknown. This can lead to a bias and the state converges to a wrong state. The convergence of the system also favours filter degeneration. This difference to divergent systems, where the ensemble spread increases exponentially, makes the direct transfer of data assimilation methods from divergent systems to convergent systems challenging.

The developed methods are often adapted to prevent filter degeneracy. For example, in the field of hydrology, *Shi et al.* (2014) and *Rasmussen et al.* (2015) kept the ensemble spread constant. However, this does not directly address the issue. *Liu and Gupta* (2007) and *Liu et al.* (2012) pointed out that, in contrast to atmospheric science, parameter estimation is an important part in hydrology. They discussed the importance of correct error representation for successful data assimilation.

This chapter investigates the differences and challenges of ensemble data assimilation methods in divergent and convergent systems. The divergent case is illustrated using the Lorenz-96 model (Section 2.2). For the convergent case, a

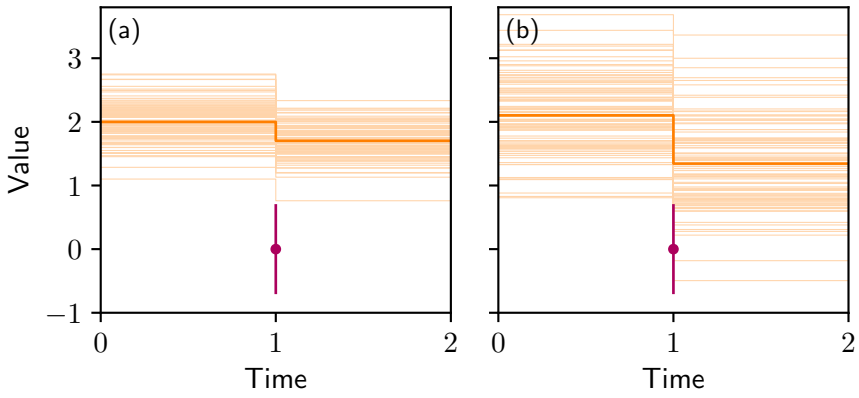

Figure 4.1.: Illustration of the Kalman update (Eq. (3.25)) of the ensemble mean (orange) and the ensemble (light orange) for one observation (purple) for different ensemble spreads (a): $\sigma^2 = 0.1$ (b): $\sigma^2 = 0.4$. The uncertainty of the observation is $\sigma^2_{\text{Obs}} = 0.5$ and the ensemble size is $N = 100$.

soil hydrological system (Section 2.1) is used.

4.1. Case study

For this study the ensemble Kalman filter (EnKF) is used. In contrast to the particle filter, it has the advantage that it does not require a model error. Therefore, it can be applied in a perfect model scenario. On the downside, the EnKF assumes Gaussian pdfs in Bayes' theorem. However, this is not limiting here, since the behaviour of the spread is examined and the best possible representation of the pdf is not necessary.

The behaviour of the EnKF is briefly outlined, in order to understand the cases in this chapter qualitatively. Figure 4.1 shows the behaviour of the EnKF for different ensemble spreads for one assimilation step in one dimension. The observation has a value $d = 0$ and the observation error is $\sigma^2_{\text{Obs}} = 0.5$. In Fig. 4.1a the ensemble has a variance of $\sigma^2 = 0.1$, which leads to a much smaller correction than in Fig. 4.1b, where the ensemble spread is $\sigma^2 = 0.4$. For an accurate observation with equal ensemble spread, the correction of the ensemble members is larger than for an inaccurate observation.

For the following cases, an ensemble size of $N = 100$ was chosen. The EnKF is used without any extensions. Furthermore, the model error in Eq. (3.3) is set to zero.

The cases will be analysed with the variance of the ensemble in one state dimension σ^2_{dim} and the mean variance of the ensemble over all dimensions σ^2. The variance in one dimension j is:

$$\sigma^2_{\text{dim},\,j} = \frac{1}{N-1}\sum_{i=1}^{N}\left(\boldsymbol{\psi}_{i,j} - \overline{\boldsymbol{\psi}_j}\right)^2 \tag{4.1}$$

with the ensemble mean

$$\overline{\boldsymbol{\psi}_j} = \frac{1}{N}\sum_{i=1}^{N}\boldsymbol{\psi}_{i,j}\,. \tag{4.2}$$

The mean variance is defined as:

$$\sigma^2 = \frac{1}{N_{\boldsymbol{\psi}}}\sum_{j=1}^{N_{\boldsymbol{\psi}}}\sigma^2_{\text{dim,j}}\,, \tag{4.3}$$

where $N_{\boldsymbol{\psi}}$ is the state dimension. In the following the index j will be omitted.

4.2. Divergent system

The data assimilation behaviour for a divergent system is demonstrated with a 40-dimensional Lorenz-96 model. The model is solved using a fourth order Runge-Kutta method with a time step of $\Delta t = 0.01$.

To generate an initial state for the data assimilation run, the model was run until time 2000 with an initial state $x_i = 4.0\ \forall i \in [1, 2, \ldots, 39]$ and $x_{40} = 4.001$, with the typical value $F = 8$ for the forcing parameter (*Lorenz and Emanuel*, 1998). The final state of this run is used as the initial state for the data assimilation. This ensures that the state is on the attractor without the initial transient phase.

The observations are generated by a forward run, using the true value and perturbing it with a Gaussian with zero mean and a standard deviation of $\sigma_{\text{obs}} = 1.0$. Observations are generated in each dimension at 5 different times with an observation interval of $\Delta t_{\text{Obs}} = 0.5$. This observation interval is chosen rather large, compared to other studies (e.g. *Nakano et al.*, 2007; *van Leeuwen*, 2010; *Poterjoy*, 2016), to ensures a large divergence of the system.

In the following, two different cases for the assimilation run are illustrated. For both cases, the initial ensemble is generated by perturbing the true initial state with a Gaussian $\mathcal{N}(0, 1)$. In the first case, the ensemble is propagated with the true model that was used to generate the observations. In the second case, the

same observations are assimilated but the ensemble is propagated with a model that has a changed model parameter $F = 10$, to show the behaviour of a divergent system for an unrepresented parameter error.

In Appendix A.2 three additional cases are shown. The numbering of the cases from this chapter will be continued. Divergent case 3 (DC3, Appendix A.2.1) propagates the ensemble with $F = 6$. In divergent case 4 (DC4, Appendix A.2.2) the ensemble is propagated with $F = 10$. The observation interval is reduced to $\Delta t_{\text{Obs}} = 0.05$. This intervalwd is often used in other studies (e.g. *Nakano et al.*, 2007; *van Leeuwen*, 2010; *Poterjoy*, 2016). Due to the frequent observations, the system cannot develop its full divergent behaviour and the Kalman gain keeps the variance at a low level. In divergent case 5 (DC5, Appendix A.2.3), observation interval is also $\Delta t_{\text{Obs}} = 0.05$ but the parameter error is represented by the ensemble.

4.2.1. Divergent case 1 (DC1) – true parameter

In this case, the ensemble is propagated with the same parameter ($F = 8$) as the observations are generated. Figure 4.2a shows the time development of one state dimension of the Lorenz-96 model and Fig. 4.2b shows σ^2_{dim} and σ^2. The ensemble has a sufficient ensemble spread such that the EnKF is able to correct the states and to follow the truth.

The mean ensemble spread σ^2 increases exponentially between two observations. At each observation, the EnKF updates the ensemble and the variance decreases.

The behaviour of σ^2_{dim} differs from the behaviour of σ^2. During the forecast, σ^2_{dim} sometimes increases, decreases or stays approximately constant. This happens because the Lorenz-96 model is bounded and has, therefore, convergent and divergent directions. A local linear stability analysis shows that the Lorenz-96 model has 13 positive Lyapunov exponents (*Lorenz and Emanuel*, 1998). So 13 directions of the eigenbasis are divergent, while the other 27 directions are convergent.

4.2.2. Divergent case 2 (DC2) – wrong parameter

In this case, the ensemble is propagated with a wrong parameter of $F = 10$ instead of $F = 8$. The Lorenz-96 model with $F = 10$ has 14 positive Lyapunov exponents instead of 13 as the truth. Furthermore, the doubling time, the time the distance between two nearby states doubles, decreases from 0.42 time units to 0.3 units (*Lorenz*, 1996), which also decreases the predictability of the system. So the ensemble is propagated with a model that is more divergent than the synthetic

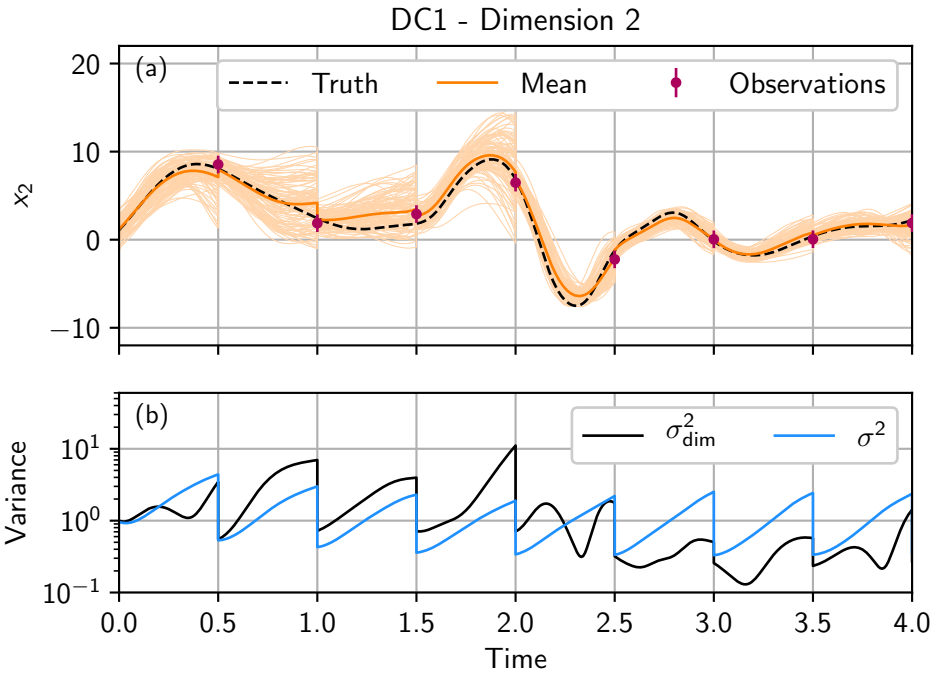

Figure 4.2.: Divergent case 1 (DC1): the ensemble is propagated with the same parameter as the truth and the state is estimated. (a): The ensemble mean (orange) and the ensemble (light orange) in the data assimilation run for state dimension 2. The observations (purple), generated from the truth (black dashed line), and the ensemble are propagated with the same model parameter $F = 8$. (b): Mean variance σ^2 (light blue) of the ensemble over all dimensions and variance σ^2_{dim} (black) of state dimension 2.

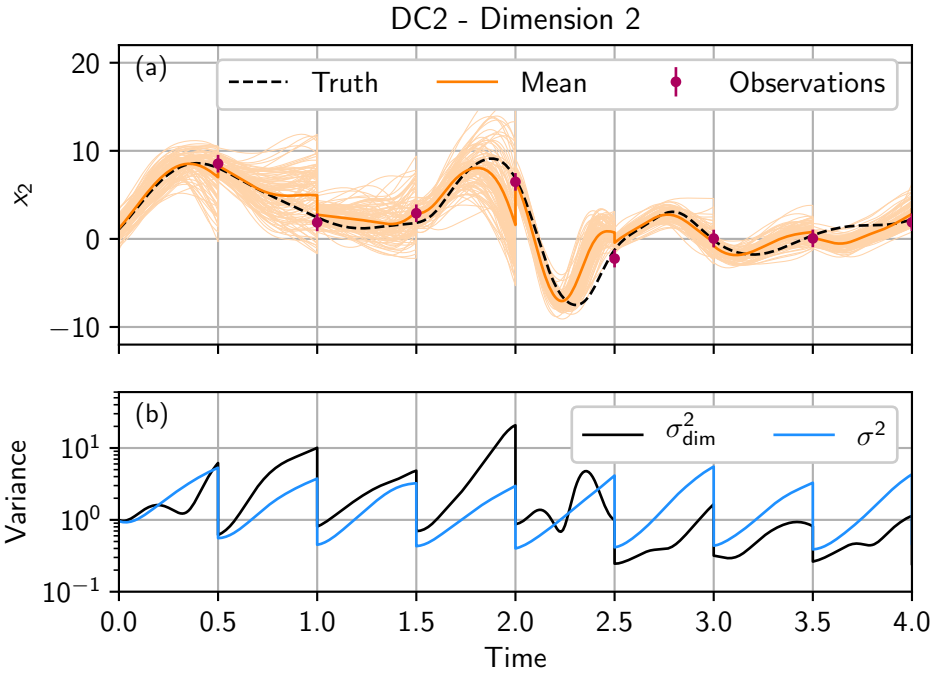

Figure 4.3.: Divergent case 2 (DC2): the ensemble is propagated with $F = 10$ instead of $F = 8$ and the state is estimated. (a): The ensemble mean (orange) and the ensemble (light orange) in the data assimilation run for state dimension 2. The observations (purple) and the truth (black dashed line) are generated with $F = 8$. (b): Mean variance σ^2 (light blue) of the ensemble over all dimensions and variance σ^2_{dim} (black) of the dimension 2.

truth, from which the observations are taken. This can be seen in Fig. 4.3b. Compared to DC1 (Fig. 4.2b), the σ^2 and σ^2_{dim} increase faster and the ensemble spread reaches higher values. Propagating the ensemble with this different model leads to a larger deviation of the ensemble mean from the truth (see Figure 4.2a) than in DC1. However, since the ensemble spread increases faster, the Kalman gain is larger and the corrections of the EnKF stronger, which compensates the wrong parameter and leads to a worsened but still good estimation of the truth.

The behaviour is also tested for a smaller value of $F = 6$ (see Appendix A.2.1). In DC3 the system is more convergent and the corrections of the EnKF are smaller. Sometimes the ensemble does not represent the truth anymore but the filter does not degenerate.

4.3. Convergent system

For the case studies of data assimilation in a convergent system, a homogeneous soil of loamy sand is used with an extent of 1 m.

Following Chapter 2, the description will be separated into dynamics, parameters, forcing and state. The model of the considered system is described in Section 2.1 in more detail.

Dynamics: The dynamics in an unsaturated porous medium can be described by the Richards equation (Eq. (2.6)). The resolution for the numerical solver MuPhi (*Ippisch et al.*, 2006) is set to 1 cm, which results in a 100-dimensional water content state.

Parameters: For these case studies the Mualem-van Genuchten parametrisation (Eqs. (2.7) and (2.8)) is used. For the true trajectories and the observations, parameters by *Carsel and Parrish* (1988) for loamy sand are used (see Table 2.1).

Forcing: A Dirichlet condition with zero potential (groundwater table) is used as the lower boundary. The upper boundary condition is chosen as one infiltration over the whole observation time with a flux of $5 \cdot 10^{-7}$ m s^{-1}.

State: Initially, the system is in equilibrium and will be forced by the boundary condition. The initial state is shown in Fig. 4.4. Four time domain reflectometry (TDR)-like observations are generated equidistantly at the positions $(0.2, 0.4, 0.6, 0.8)$ m. The observation error is chosen to be $\sigma_{\text{Obs}} = 0.007$ (e.g. *Jaumann and Roth*, 2017). Observations are taken hourly for a duration of 30 h.

To generate the initial ensemble, the equilibrium state is perturbed by a correlated multivariate Gaussian. The main diagonal of the covariance matrix is 0.003^2. The off-diagonal entries are determined by multiplying the variance on the main diagonal with the Gaspari and Cohn function (Eq. (3.28)) using a length scale of $c = 10$ cm. The distance for the Gaspari and Cohn function is is the distance of the off-diagonal entry from the main diagonal. This ensures a correlated initial state, which increases the diversity of the ensemble. If instead uncorrelated Gaussian random numbers with zero mean were used, the dissipative nature of the system would lead to a fast dissipation of the perturbation

In the following, four cases will be studied. In convergent case 1 (CC1), an interpolated state is used as the initial state for the ensemble and the ensemble is propagated with the same model as the truth without any estimation. In the

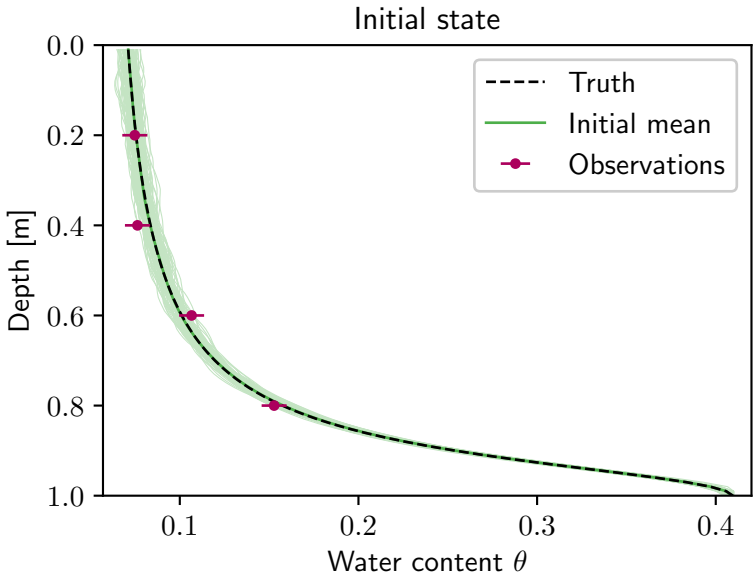

Figure 4.4.: Initial state for convergent case 2, 3 and 4. The ensemble (light green) with 100 ensemble members is generated by perturbing the initial truth (black dashed line), which is used for the observations (purple), with a spatially correlated Gaussian.

other three cases, the initial ensemble shown in Fig. 4.4 is used. The ensemble is propagated with a changed parameter n. Convergent case 2 (CC2) estimates the state but not the parameter, convergent case 3 (CC3) represents the error in the parameter and convergent case 4 (CC4) estimates the state and the parameter simultaneously.

4.3.1. Convergent case 1 (CC1) – no estimation

In this case, an interpolated initial condition is used for the ensemble. The observations at time zero are interpolated linearly. Additionally, the saturated water content for loamy sand, which is 0.41, is taken as the lower boundary. The approximated state is used as the ensemble mean for the EnKF. This state is then perturbed with the same correlated Gaussian as in Fig. 4.4. The initial ensemble represents the uncertainty of the water content in most parts (see Fig. 4.5).

The ensemble is propagated with the true model. The temporal development of the water content in 20 cm depth, the position of the uppermost observation, is shown in Fig. 4.6a. The initially broad ensemble slowly collapses to the truth.

The ensemble variance in this depth σ^2_{dim} and the ensemble variance over all

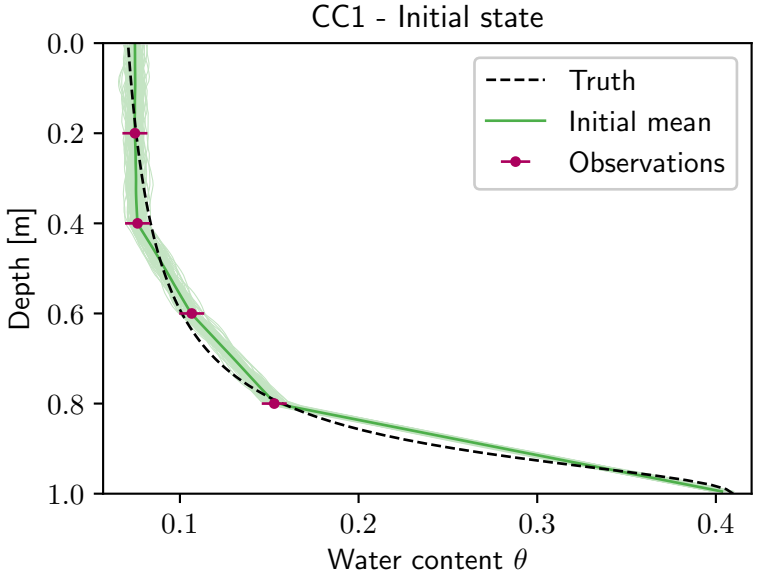

Figure 4.5.: Initial state for convergent case 1 (CC1). Observations (purple) at time zero are connected linearly and set constant towards the upper boundary. For the lower boundary, the saturated water content $\theta_r = 0.41$ of loamy sand is used for the interpolation. The ensemble (light green) with 100 ensemble members is generated by perturbing the interpolated state with a spatially correlated Gaussian. The initial truth, which is used for the observations, is shown as a black dashed line.

dimensions σ^2 are shown in Fig. 4.6b. The variance σ^2_{dim} increases when the infiltration front reaches 20 cm. Because of the nonlinear conductivity function (Eq. (2.7)), the different initial water contents lead to a different arrival time of the infiltration front at the first observation position. This leads to an increase in the ensemble spread.

For increasing water content, the ensemble collapses faster since the conductivity increases (see Fig. 2.1), which leads to a faster convergence to the truth. The variance σ^2 decreases slowly and will reach zero after some time. If the system would have been started with a higher water content instead of equilibrium, this collapse would have been faster.

This case shows that the predictability for a perfect convergent system is infinite. After a transient phase, the states converge to the truth as discussed by *Kalnay* (2003). However, in real systems, the model parameters are uncertain, which leads to difficulties.

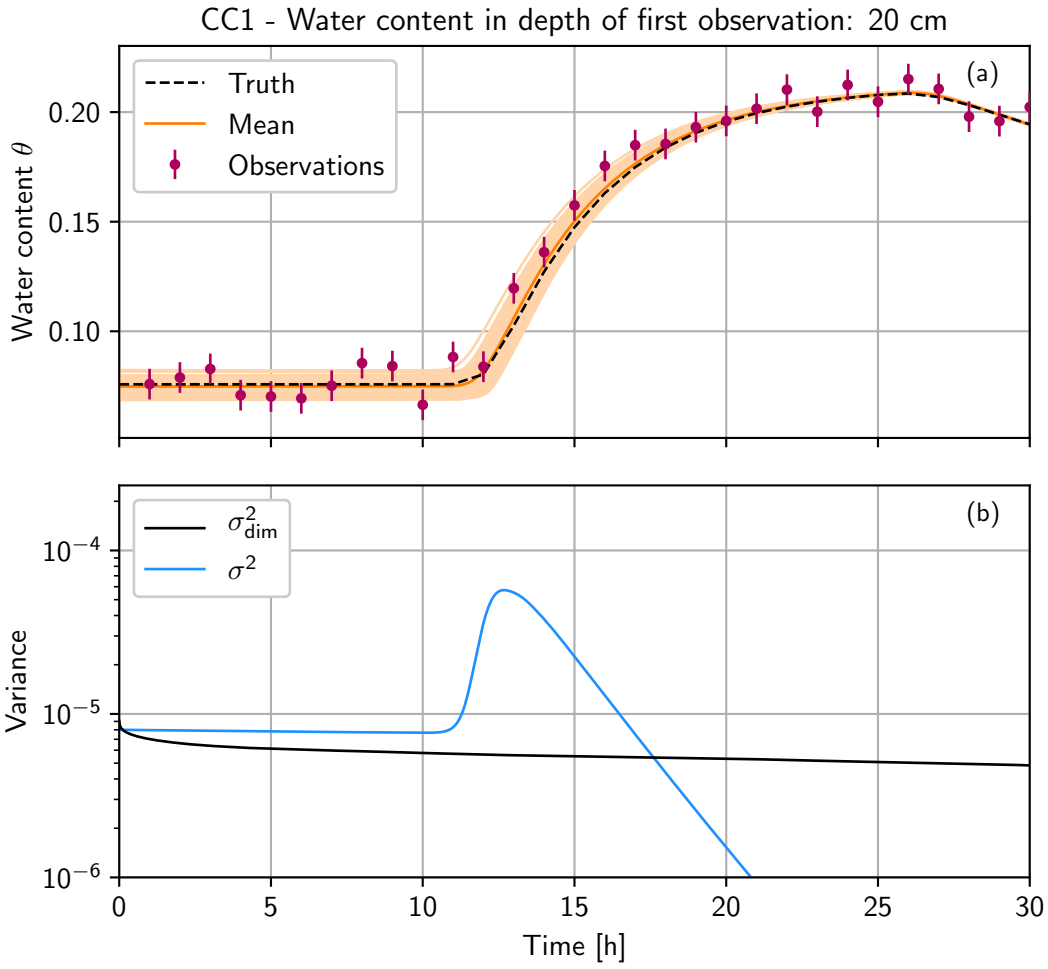

Figure 4.6.: Convergent case 1 (CC1): forward run without data assimilation, using an interpolated initial condition (see Fig. 4.5). (a): The ensemble mean (orange) and the ensemble (light orange) during the forward run at the depth of the uppermost observation (20 cm). The truth, which is used to generate the observations (purple), is shown as a black dashed line. (b): Mean variance σ^2 (light blue) of the ensemble over all dimensions and variance σ^2_{dim} (black) at the depth of 20 cm.

4.3.2. Convergent case 2 (CC2) – state estimation, wrong parameter

In this case, the state is estimated but the ensemble is propagated with a wrong parameter for n. Instead of $n_{\text{true}} = 2.28$, n is chosen to be $n = 2.68$. In Fig. 4.7a the temporal development of the water content in the depth of the uppermost observation (20 cm) is shown. A larger n results in a higher conductivity. The infiltration front reaches the depth of 20 cm earlier for the ensemble than for the truth. The EnKF tries to correct the ensemble but fails because the variance of the ensemble is too small and cannot represent the truth. Due to the convergent system, σ^2 decreases constantly while σ^2_{dim} decreases fast to zero after the infiltration front reaches the depth of 20 cm (see Fig. 4.7b). This convergence leads to a false trust in the model and the filter degenerates.

This case illustrates that a wrong parameter in a convergent system can lead to filter degeneration. This is in direct contrast to DC2 (Section 4.2.2) and DC3 (Appendix A.2.1), where the filter is still able to estimate the state. However, if the observation interval in the Lorenz-96 model is too short such that it cannot develop its full divergent behaviour, the filter will also degenerate for a wrong parameter (see Appendix A.2.2).

4.3.3. Convergent case 3 (CC3) – state estimation, represented error

In this case, the parameter error is represented with the ensemble but not estimated. Each ensemble member has a different parameter n. The parameters are Gaussian distributed with $\mathcal{N}(2.68, 0.4^2)$ such that the truth lies within one standard deviation.

Figure 4.8a shows the temporal development of the water content in a depth of 20 cm. The infiltration front reaches this depth at different times due to the different parameter n for each ensemble member. This increases the variance in the ensemble both, in this depth and overall (see Fig. 4.8b). The variance increases rapidly between the observations, similar to the divergent cases. This way, the ensemble spread stays large enough such that the EnKF can correct the states. The ensemble can follow and represent the truth. This behaviour can also be observed for the Lorenz-96 model with a short observation interval (see Appendix A.2.3).

Representing the model error adds a divergent part to the convergent model. This allows the EnKF to correct the state and follow the truth. However, the predictability of the system decreases since each ensemble member converges to a

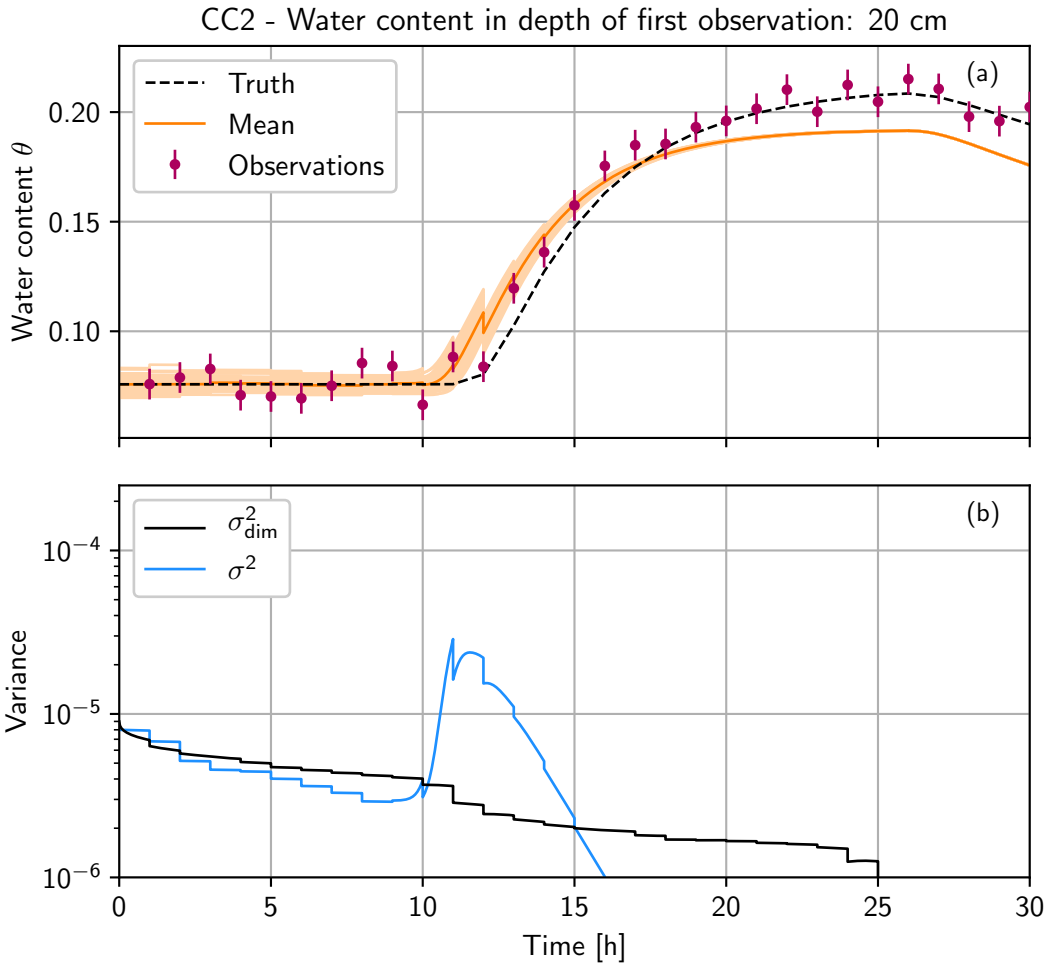

Figure 4.7.: Convergent case 2 (CC2): the ensemble is propagated with $n = 2.68$ instead of $n_{\text{true}} = 2.28$ and the state is estimated. (a): The ensemble mean (orange) and the ensemble (light orange) in the data assimilation run at the depth of the uppermost observation (20 cm). The truth, which is used to generate the observations (purple), is shown as a black dashed line. (b): Mean variance σ^2 (light blue) of the ensemble over all dimensions and variance σ^2_{dim} (black) at the depth of 20 cm.

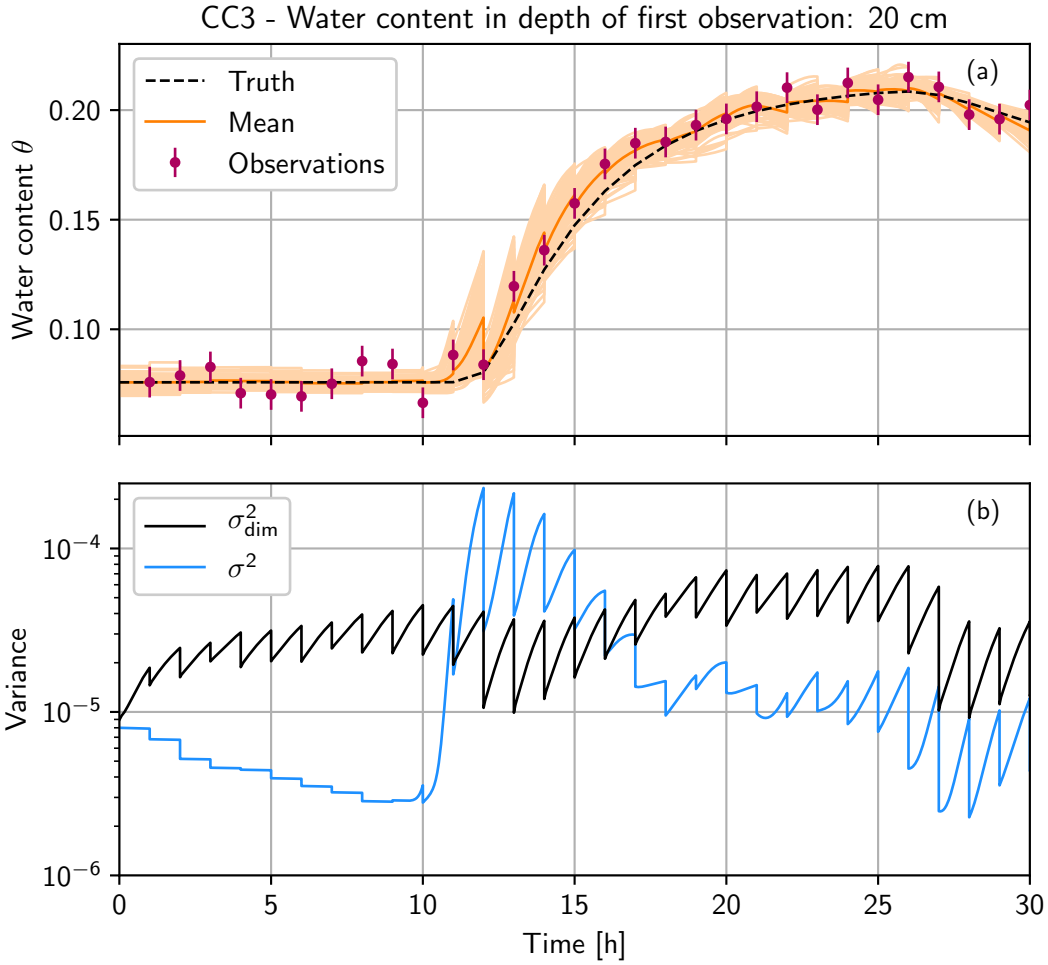

Figure 4.8.: Convergent case 3 (CC3): the parameter error is represented by the ensemble using $\mathcal{N}(2.68, 0.4^2)$ and the state is estimated. (a): The ensemble mean (orange) and the ensemble (light orange) in the data assimilation run at the depth of the uppermost observation (20 cm). The truth, which is used to generate the observations (purple), is shown as a black dashed line. (b): Mean variance σ^2 (light blue) of the ensemble over all dimensions and variance σ^2_{dim} (black) at the depth of 20 cm.

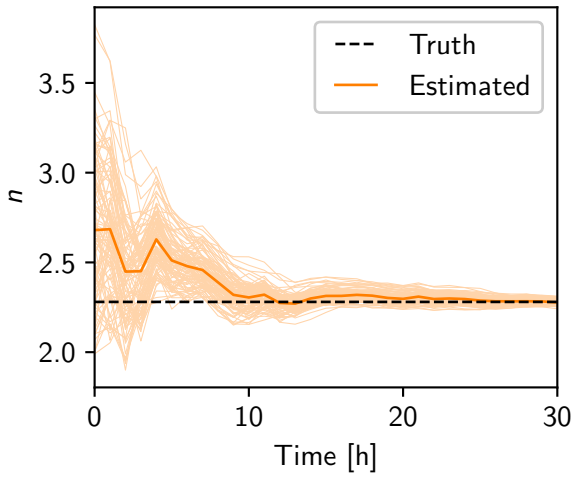

Figure 4.9.: Estimation of parameter n in convergent case 4. The ensemble mean is shown in orange and the ensemble in light orange. The truth is a black dashed line.

different fixed point apart from the truth. To increase the predictability, parameter estimation is necessary.

4.3.4. Convergent case 4 (CC4) – state and parameter estimation

In this case, the error in the parameter n is not only represented but also estimated. Therefore, an augmented state $\boldsymbol{u}$ (Section 3.4) is used. The model equation (Eq. (3.3)) thus changes to Eq. (3.55) with a constant model for the parameters

$$\boldsymbol{p}^k = \boldsymbol{p}^{k-1} \tag{4.4}$$

and Richards' equation as $f_{\psi}(\cdot)$. Note that both stochastic model errors of Eq. (3.55) are set to zero. The initial parameter set is Gaussian distributed with $\mathcal{N}(2.68, 0.4^2)$ such that the truth is located within one standard deviation.

The estimation of n is shown in Fig. 4.9. The ensemble converges to the truth in a fast way because only one parameter is estimated, so every deviation from the truth is mainly caused by this parameter.

The mean variance σ^2 increases (see Fig. 4.10b) because at the beginning, the parameter has not been sufficiently improved such that the ensemble members still have different n. This leads to a divergent ensemble in state space during

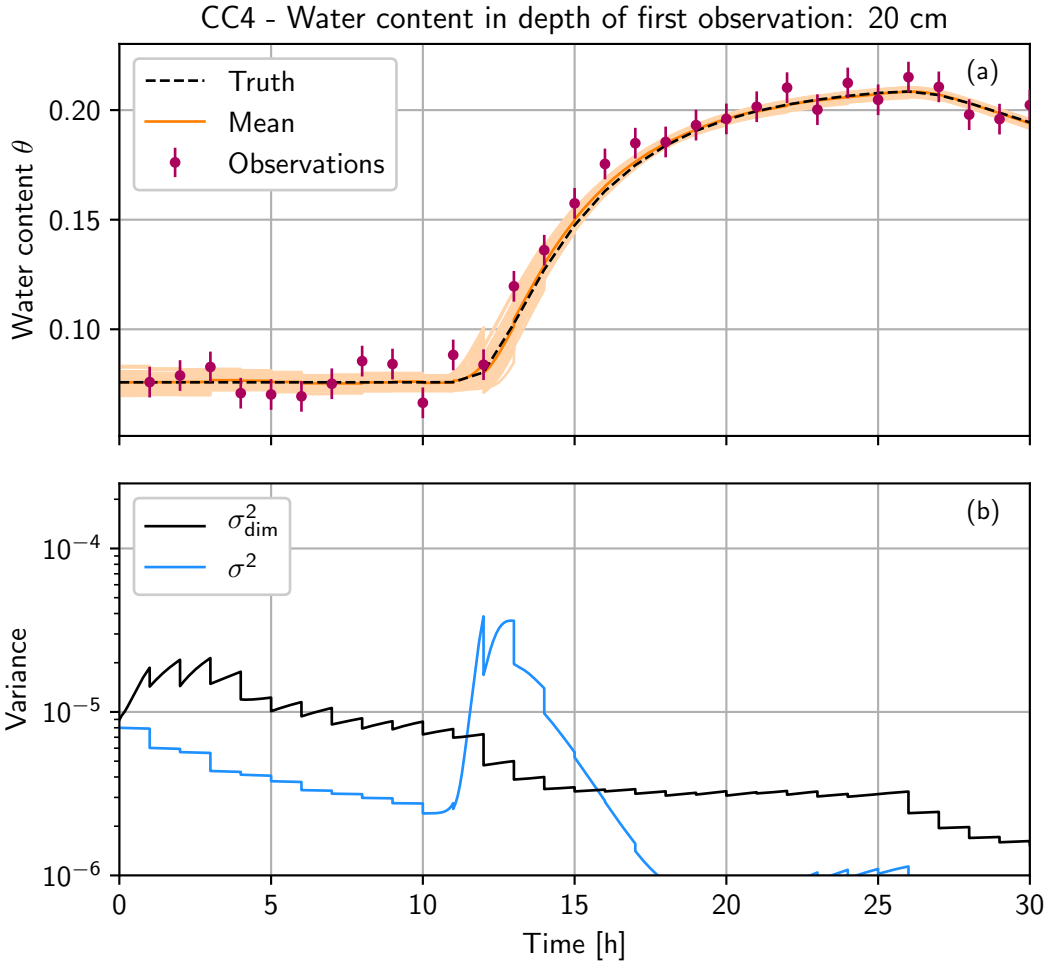

Figure 4.10.: Convergent case 4 (CC4): simultaneous state and parameter estimation. (a): The ensemble mean (orange) and the ensemble (light orange) in the data assimilation run at the depth of the uppermost observation (20 cm). The truth, which is used to generate the observations (purple), is shown as a black dashed line. (b): Mean variance σ^2 (light blue) of the ensemble over all dimensions and variance σ^2_{dim} (black) at the depth of 20 cm.

the infiltration similar to CC3 (Section 4.3.3). While the parameter is estimated, the variance of the ensemble decreases fast and the convergent property of the system becomes dominant.

The temporal development of the water content in a depth of 20 cm is shown in Fig. 4.10a. In contrast to CC3 (Section 4.3.3), the corrections of the EnKF to the state are much smaller. The mean of the parameter n comes close to the true value and the uncertainty of it decreases. This causes the forward propagation to come close to the true model as well. The propagation with an almost correct model supports the state estimation because of the convergent nature of the system that forces the state to the true value.

4.4. Summary and discussion

For the divergent Lorenz-96 system, the EnKF is able to estimate the state for the true model as well as for the case with a wrong parameter. In a divergent system, the volume of the prior in phase space increases during forward propagation (*Evensen*, 1994). For the EnKF, this is directly connected to the ensemble spread, which increases rapidly between the observations. This prevents a collapse of the ensemble even in the presence of an unrepresented parameter error. However, if the observation interval is too small such that the system cannot develop its full divergent behaviour, the EnKF leads to a decrease in the ensemble spread such that the filter collapses in the case of an unrepresented parameter error.

In a convergent soil hydrological system, the volume of the prior distribution decreases during forward propagation such that the prior becomes more certain during the forward propagation even without an observation or data assimilation.

For a perfect model, the predictability and state estimation in a convergent model are trivial. The initial ensemble will converge to the truth after some time, even with a rough approximation. In this case, data assimilation is not necessary.

In the case of wrong parameters the situation is different. The ensemble converges to a wrong state, the filter degenerates and data assimilation fails. Increasing the ensemble size can only improve the performance marginally since all ensemble members will converge to the same fixed point.

Representing the parameter error by assigning each ensemble member a different parameter, increases the divergence of the system and the filter is able to estimate the state again. Between the observations, the ensemble spread increases rapidly because the ensemble members diverge to different fixed points apart from the truth. This results in a finite predictability. By representing the parameter error, Richards' equation behaves more similar to the Lorenz-96 model. However, in the case of the Lorenz-96 model, it is possible to estimate the state without rep-

resenting the parameter error. In the case of a convergent system, it is necessary to represent the parameter error, otherwise the ensemble collapses.

To increase the predictability of the system again, it is necessary to not only represent but also estimate the wrong parameters. If the parameters are estimated to the true values, the convergent property of the system supports the state estimation. With the true parameters the system converges to the true state and the predictability increases. Therefore, parameter estimation is important for convergent systems. For divergent systems parameter estimation also increases the predictability but the predictability is limited by divergent system dynamics.

For the application of data assimilation to real data, there is still the problem of unrepresented model errors. In hydrology, the model error is typically ill-known (*Li and Ren*, 2011) and can vary both in space and time, which can lead to filter degeneracy and wrong parameters. In this case, inflation methods (see Section 3.2.2), can help to increase the ensemble spread and avoid filter degeneracy. For example, *Zhang et al.* (2017) keep a constant ensemble spread for the parameters to provide a sufficient spread in state space. However, keeping the spread constant prevents the EnKF from reducing the prediction uncertainty. *Zhang et al.* (2017) describes that they accept this disadvantage in order to prevent filter degeneration. This procedure is similar to the case shown in Section 4.3.3. Even in the absence of a real parameter error, keeping an ensemble spread in the parameter space adds a divergent component to the system that results in an increased ensemble spread in state space.

Adaptive inflation methods (e.g. *Bauser et al.*, 2018; *Gharamti*, 2018), are more sophisticated to ensure a sufficient ensemble spread. However, if the parameters are estimated during a time interval where the model is correct or almost correct, an estimation during an time interval where the model cannot describe the dynamics leads to biased parameters. In those cases it is useful to apply a closed-eye period (*Bauser et al.*, 2016). In the closed-eye period, the parameters are kept constant during the time of unrepresented dynamics and only the state is estimated. If then the model can describe the dynamics correct again, the state fits to the observations and the parameters do not contain a bias through unrepresented model dynamics.

If state and parameters are already converged to the truth and a case of unrepresented model errors occurs, the ensemble is already too small and will not be able to incorporate the new observations anymore. This collapse is difficult to prevent by inflation or the closed-eye period. In a divergent system, this will not occur because the system dynamics prevents this case and inflation methods will work more effectively, as they do not have to work against the dynamics.

Data assimilation methods require a divergent part in the model to keep a sufficient ensemble spread and prevent filter degeneration. While in divergent

systems this is inherent to the system, in convergent systems it has to be added. This makes the transfer to convergent systems challenging. Methods to increase the ensemble spread artificially or the representation of errors adds a divergent part to the system and avoids filter degeneracy. In the case of parameter errors, estimating the parameter is helpful for state estimation.

The difference and challenges of data assimilation were highlighted for the Lorenz-96 model and Richards' equation. I expect that the results of this chapter can be transferred to other convergent and divergent systems.

5. Covariance resampling

This chapter is based on *Berg et al.* (2018).

In soil hydrology both, state and parameters, are typically ill-known and have to be estimated. The ensemble Kalman filter (EnKF) (Section 3.2) is a popular data assimilation method in hydrology. It has the advantage of using the ensemble covariance to correlate dimensions with observations to unobserved dimensions. The EnKF with parameter estimation is successfully applied to several hydrological systems. *Moradkhani et al.* (2005b) used the EnKF for a rainfall-runoff model and *Chen and Zhang* (2006) for saturated flow modelling. For systems based on Richards' equation, the EnKF is mostly applied in synthetic studies (e.g. *Wu and Margulis*, 2011; *Song et al.*, 2014; *Erdal et al.*, 2015; *Shi et al.*, 2015; *Man et al.*, 2016). However, some applications to real data exist (e.g. *Li and Ren*, 2011; *Bauser et al.*, 2016; *Botto et al.*, 2018).

Nonlinear systems violate the EnKF assumption of Gaussian probability density functions (pdfs) (*Harlim and Majda*, 2010; *DeChant and Moradkhani*, 2012). The dynamics of Richards' equation is generally dissipative and the Gaussian assumption is appropriate. However, jumps at layer boundaries, soliton-like fronts during strong infiltration and diverging potentials for strong evaporation deform the pdf and lead to non-Gaussianity, which results in poorer performance of the EnKF (*Lei and Bickel*, 2011; *Liu et al.*, 2012; *Zhang et al.*, 2015). In this case, the pdf requires higher statistical moments to be described correctly. Particle filters (Section 3.3) can accomplish this task.

Parameter estimation with particle filters is challenging. Resampling (Section 3.3.3) generates identical copies of particles with large weights. During the forecast, the stochastic model error is added to each particle, which makes them distinguishable. Parameters are typically constant in time. Thus it is difficult to define a stochastic model error. Resampled particles stay identical and the filter degenerates.

If a model does not have a model error, which also happens because the error unknown and therefore set to zero (*Poterjoy*, 2016), it is necessary to perturb the copied particles after resampling. Even in the presence of a model error it can be useful to perturb the particles. For example, if the model error is ill-known or structurally incorrect, it can help to guarantee a sufficient ensemble spread and

diversity. In hydrology, the model error is typically ill-known and can vary both, in space and time.

To apply the particle filter for parameter estimation, *Moradkhani et al.* (2005a) suggested to perturb the parameters using Gaussian noise with zero mean after the resampling step. They used an unweighted variance of the ensemble modified with a damping factor such that the ensemble spread did not become too large.

This method or similar has been used for state and parameter estimation land surface models (*Qin et al.*, 2009; *Plaza et al.*, 2012), rainfall-runoff models (*Weerts and El Serafy*, 2006) and soil hydrology (*Montzka et al.*, 2011; *Manoli et al.*, 2015). However, with only a rough initial guess of the initial state, perturbing only the parameters does not guarantee a sufficient ensemble spread in state space for small water contents (see Section 4.3.3) and the filter can degenerate.

Additional development of resampling for parameter estimation was done by *Moradkhani et al.* (2012) and *Vrugt et al.* (2013). They used a Markov chain Monte Carlo (MCMC) method to generate new particles. This method was further used by e.g. *Yan et al.* (2015) and *Zhang et al.* (2017). The latter compared the performance of this method to an EnKF and the particle filter presented by *Moradkhani et al.* (2005a) and found that the performance of the filters were similar with slight advantages for the EnKF. While the MCMC is accurate, it is also expensive, as it requires additional model runs. To increase the efficiency, *Abbaszadeh et al.* (2018) additionally combined it with a generic algorithm.

This chapter introduces the covariance resampling, a resampling method that generates new particles using the ensemble covariance. This method conserves the first two statistical moments in the limit of large numbers while partly maintaining the structure of the pdf in the retained ensemble. Using the covariance, the unobserved parameters of the new particles are correlated to the observed state dimensions. The particle filter with covariance resampling (PFCR) is able to estimate state and parameters in case of a difficult initial condition without additional model evaluations.

5.1. Method

Different approaches where suggested to perturb the particles using the ensemble covariance, however without application for parameter estimation. *Pham* (2001) proposed to sample new particles by perturbing the identical particles using a Gaussian with the (damped) ensemble covariance matrix as covariance. *Xiong et al.* (2006) sampled the whole ensemble from a Gaussian using the first two moments specified by the ensemble (full covariance information), which neglects the particle filter ability to use non-Gaussian distributions.

The covariance resampling neither perturbs the duplicated states nor draws a complete new ensemble. It uses the stochastic universal resampling (Section 3.3.3) to choose the ensemble members that are kept. For this step, other resampling methods like the residual resampling (Section 3.3.3) can be equally used. In this work, the stochastic universal resampling is used because it has the lowest sampling noise (*Van Leeuwen et al.*, 2015).

Instead of duplicating the particles and setting the weights to N^{-1}, the weight of the particles is changed to

$$w_i = \frac{z}{N} \quad \text{with} \quad i \in \{1, 2, .., N'\} \ , \tag{5.1}$$

where the particle i is chosen z-times and N' is the number of kept particles. This ensures that the estimated statistical moments are conserved.

The total ensemble reduces to N'. To have N ensemble members again, $N - N'$ new particles have to be generated. These particles are sampled from a Gaussian $\mathcal{N}(\overline{\boldsymbol{u}}, \mathbf{P})$ with the weighted mean

$$\overline{\boldsymbol{u}} = \sum_{i=1}^{N} w_i \, \boldsymbol{u}_i \ . \tag{5.2}$$

and the weighted covariance

$$\mathbf{P} = \frac{1}{1 - \sum_{i=1}^{N} w_i^2} \sum_{i=1}^{N} w_i \left[\boldsymbol{u}_i - \overline{\boldsymbol{u}}\right] \left[\boldsymbol{u}_i - \overline{\boldsymbol{u}}\right]^\intercal \ , \tag{5.3}$$

where the factor $\frac{1}{1-\sum_{i=1}^{N} w_i^2}$ is Bessel's correction for a weighted covariance. The mean and the covariance of the ensemble are calculated using the updated weights before resampling (Eq. (3.42)).

Sampling only the dropped particles from a Gaussian conserves the first two statistical moments in the limit of large numbers with the advantage that information of the non-Gaussian pdf is partly conserved in the retained ensemble (see Section 5.2). A weight of N^{-1} is assigned to each of the new particles, which results in a sum of all weights being larger than one. Therefore, it is necessary to renormalise the weights. The whole resampling process is illustrated in Fig. 5.1.

New particles are generated using a Cholesky decomposition of the covariance matrix. The calculation of the covariance from the ensemble can result in small numerical errors that have to be regularised, otherwise the decomposition would fail. For details about the generation of new particles and regularisation of the covariance matrix see Appendix A.4.

Pham (2001) and *Moradkhani et al.* (2005a) introduced a a tuning parameter to modify the covariance matrix or the variance, respectively. They used the tuning

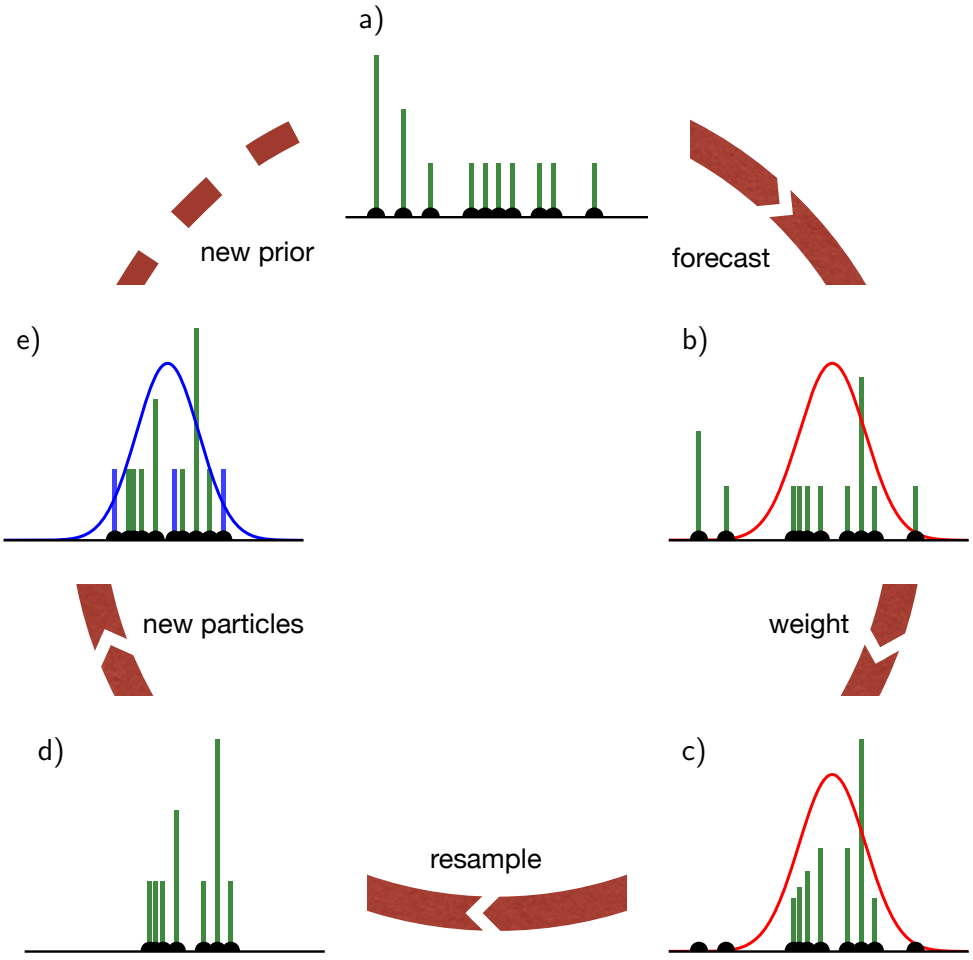

Figure 5.1.: Illustration of the PFCR. The green bars show the weight of each ensemble member (ten in this example) and the black dots their position. (a) The prior is represented through the ensemble. (b): The ensemble is propagated to the next observation (depicted as Gaussian, red curve). (c): The particles are weighted according to the observation. At this point, some particles have already negligible weight. (d): The universal resampling drops particles with low weight (three in this example). Instead of adding new particles at the same position, only the weights of the kept particles are changed. If a particle is resampled k-times, it will get the weight $k\,N^{-1}$. The ensemble size is reduced and new particles have to be added to conserve the ensemble size and avoid filter degeneration. (e): The new particles are drawn from the full covariance of the ensemble (Eq. (5.3)) and their weight is set to N^{-1}. Since new particles with weights are added to the ensemble, it is necessary to renormalise the weights to one. This results in the posterior or the next prior, respectively.

factor to reduce the amplitude of the perturbation. For the covariance resampling, I also introduce a tuning parameter. The covariance matrix is modified by a multiplicative factor $\boldsymbol{\gamma}$

$$\mathbf{P}' = (\boldsymbol{\gamma}\,\boldsymbol{\gamma}^{\mathsf{T}}) \circ \mathbf{P} \;, \tag{5.4}$$

where $\circ$ is the entrywise product (Hadamard product).

If the model dynamics does not support a sufficient spread for the ensemble, the perturbation of the covariance resampling has to be large enough to prevent the ensemble from collapsing. For example, in the case of constant parameters it is beneficial to chose a factor larger than one to prevent filter degeneracy.

5.2. Non-Gaussian distribution

To visualise the behaviour of the covariance resampling for non-Gaussian distributions, it is illustrative to look at the analysis step for a one dimensional case. The covariance resampling is compared with the EnKF and the standard particle filter. This study follows the comparison of deterministic and stochastic Kalman filters by *Lawson and Hansen* (2004).

Since the goal of this section is to investigate the behaviour of the algorithms for non-Gaussian distributions, a bimodal distribution is chosen as the prior. The bimodal distribution is constructed with two equally probable Gaussian distributions:

$$P(x) = \frac{1}{2\sqrt{2\pi}} \left(\exp\left[-\frac{1}{2}(x-\overline{x})^2 \right] + \exp\left[-\frac{1}{2}(x+\overline{x})^2 \right] \right) , \tag{5.5}$$

where $\overline{x}$ is the offset of each peak from zero. In the following example $\overline{x}$ is chosen as $\overline{x} = 4$. For the calculation of the Kalman gain (Eq. (3.19)), the bimodal prior with zero mean and a variance of $\sigma^2_{\text{Prior}} = 17$, is equal to a Gaussian with these two statistical moments ($\mathcal{N}(0, \sigma^2_{\text{Prior}})$).

Using Eq. (5.5) a rather larger ensemble ($N = 5000$) is generated. This ensemble size is used to observe the behaviour of the filters without large statistical noise. The prior distribution, including the generated ensemble, is shown in Fig. 5.2.

The analysis is calculated for an observation d at $d = 3.5$, with three different observation errors σ^2_{Obs}. The observation errors are chosen as $\sigma^2_{\text{Obs}} = \sigma^2_{\text{Prior}}/2$, $\sigma^2_{\text{Obs}} = \sigma^2_{\text{Prior}}$ and $\sigma^2_{\text{Obs}} = 2\sigma^2_{\text{Prior}}$. Figure 5.3 shows the resulting analytical posterior distributions calculated using Bayes' theorem (Eq. (3.1)) and the posterior ensemble of the EnKF, the particle filter and the PFCR.

The particle filter without resampling can sample the posterior accurately for all three cases. However, the variance of the weights increases and after some

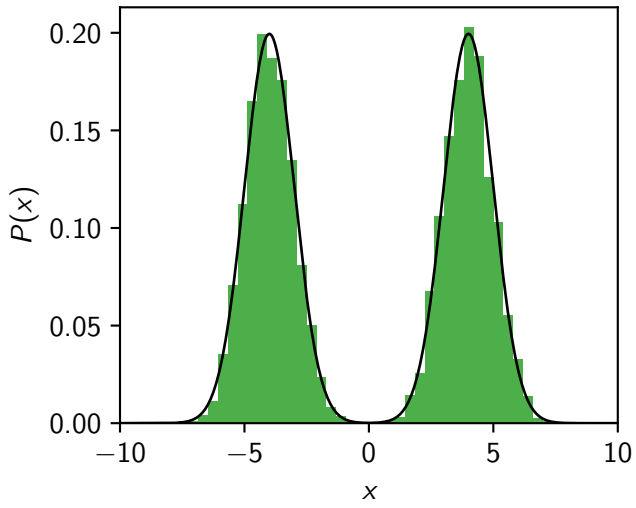

Figure 5.2.: Prior bimodal distribution $P(x)$ (black, Eq. (5.5)) and histogram of the 5000 ensemble members (green).

assimilation cycles the particle filter would degenerate. For example, in the case of high accuracy observation (left), the effective sample size (Eq. (3.44)) reduces from the forecast $N_{\text{eff}} = 5000$ to $N_{\text{eff}} \approx 2454$ in the analysis.

In the case of an accurate observation $\sigma^2_{\text{Obs}} = \sigma^2_{\text{Prior}}/2$ (left), the observation determines the right mode of the bimodal prior as the posterior. The EnKF corrects the states towards the observation and the posterior becomes approximately Gaussian but with a mean shifted to lower values and a larger variance compared to the truth. The PFCR can describe the posterior excellently. It does not need to shift the ensemble members of the left mode to the observation like the EnKF, instead almost 50 % of the ensemble members are dropped and resampled. This resampling is effective because the covariance resampling samples from a Gaussian distribution and the posterior is approximately Gaussian.

In the case of a less accurate observation $\sigma^2_{\text{Obs}} = \sigma^2_{\text{Prior}}$ (middle), the observation information is less dominant and the bimodal structure of the prior is partly visible. For the EnKF, the ensemble members of the left mode are shifted towards the observations such that the posterior becomes unimodal. The bimodal structure cannot be described by the EnKF properly. The PFCR can sample both posterior peaks with the retained ensemble. The new particles, however, are generated with a mean that lies in between both peaks, such that some of the new particles are located between the modes.

For $\sigma^2_{\text{Obs}} = 2\sigma^2_{\text{Prior}}$ (right), the posterior is similar to the prior with a less likely

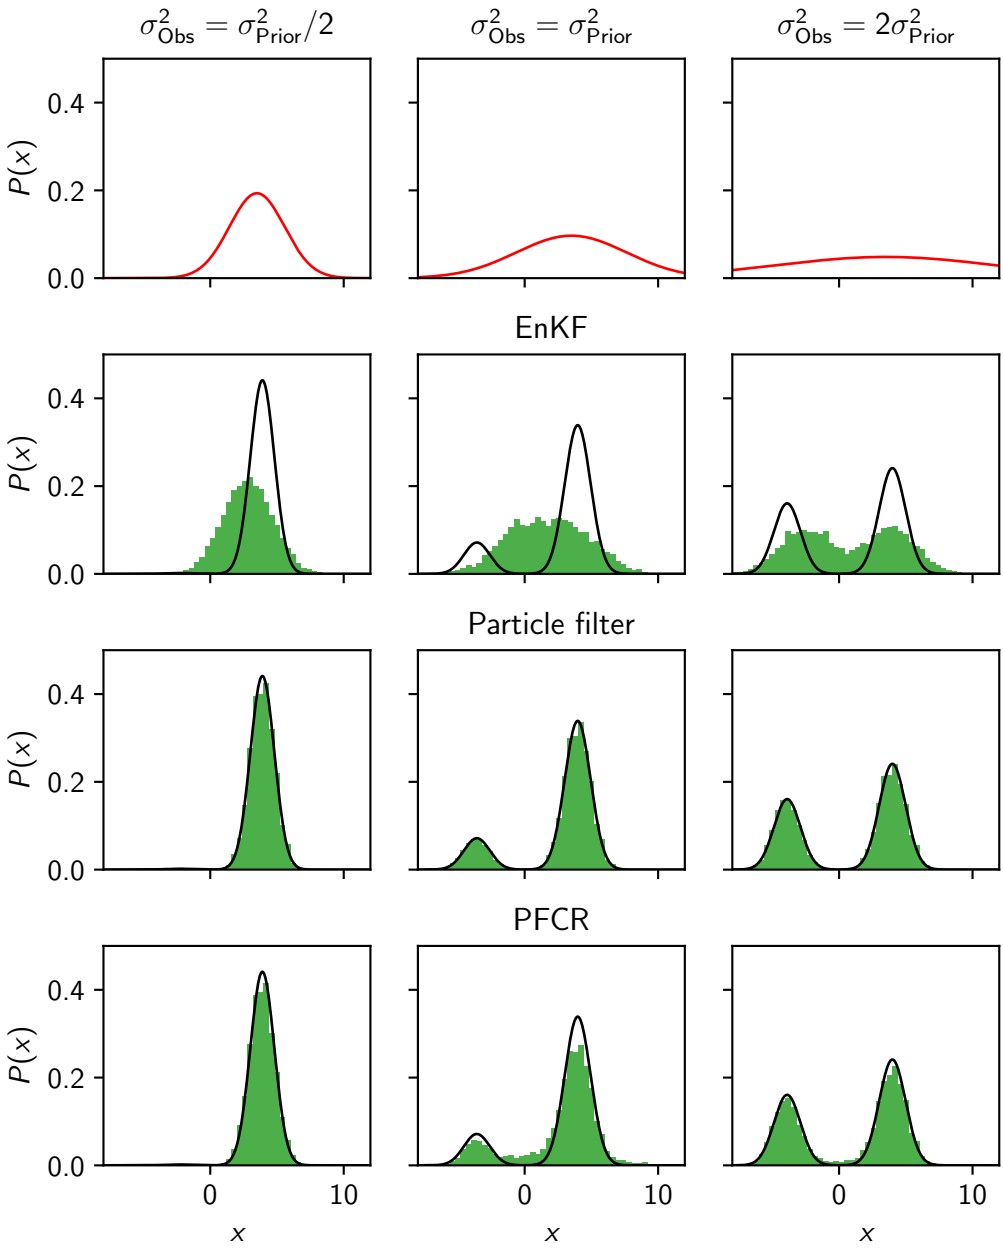

Figure 5.3.: Analysis ensemble (green) for EnKF (second row), particle filter (third row) and PFCR (bottom row) for a bimodal prior (Fig. 5.2), for observations (top row, red) with different observation errors: $\sigma^2_{\text{Obs}} = \sigma^2_{\text{Prior}}/2$ (left column), $\sigma^2_{\text{Obs}} = \sigma^2_{\text{Prior}}$ (middle column) and $\sigma^2_{\text{Obs}} = 2\sigma^2_{\text{Prior}}$ (right column). The true posterior (black) is calculated with Bayes' theorem. In case of the PFCR, 49.7 % (left), 32.6 % (middle) and 10.0 % (right) particles are resampled.

left mode. The large uncertainty makes it impossible to confidently decide in which peak the truth lies. The EnKF retains the bimodal distribution of the prior because the Kalman gain becomes small if the observation error is large, which leads to a small correction in the analysis step. The PFCR describes the posterior almost as accurate as the standard particle filter. The particles have approximately equal weights because of the large observation error. Therefore, only 10.0 % of the particles are resampled which reduces the effect that particles are generated in between the peaks as in the case $\sigma^2_{\text{Obs}} = \sigma^2_{\text{Prior}}$ (middle).

The particle filter can sample the posterior distribution accurately, however the filter would degenerate without resampling. The PFCR is superior in approximating the posterior distribution compared to the EnKF for all cases. Additional experiments with different distributions are shown in Appendix A.3.

5.3. Case study

In Section 5.2 one analysis step was examined. This section describes the settings for a full data assimilation run with state and parameter estimation in a hydrological case.

The covariance resampling is tested for a one-dimensional unsaturated porous medium with two homogeneous layers. The system has a vertical extent of 1 m with the layer boundary in the middle at 50 cm.

Following Chapter 2, the description is separated into dynamics, parameters, forcing and state. The model of the considered system is described in Section 2.1 in more detail.

Dynamics: The dynamics in an unsaturated porous medium can be described by the Richards equation (Eq. (2.6)). The resolution for the numerical solver MuPhi (*Ippisch et al.*, 2006) is set to 1 cm, which results in a 100-dimensional water content state.

Parameters: In this case study, the Mualem-van Genuchten parametrisation (Eqs. (2.7) and (2.8)) is used, where the parameters α, n and K_w are estimated for each layer. The conductivity function (Eq. (2.10)) incorporates all estimated parameters.

For the true trajectories and the observations, parameters for loamy sand (upper layer, layer 1) and sandy loam (lower layer, layer 2) are taken from *Carsel and Parrish* (1988). The true parameter values are given in Table 2.1. In the following, the parameters will have an index representing their corresponding layer.

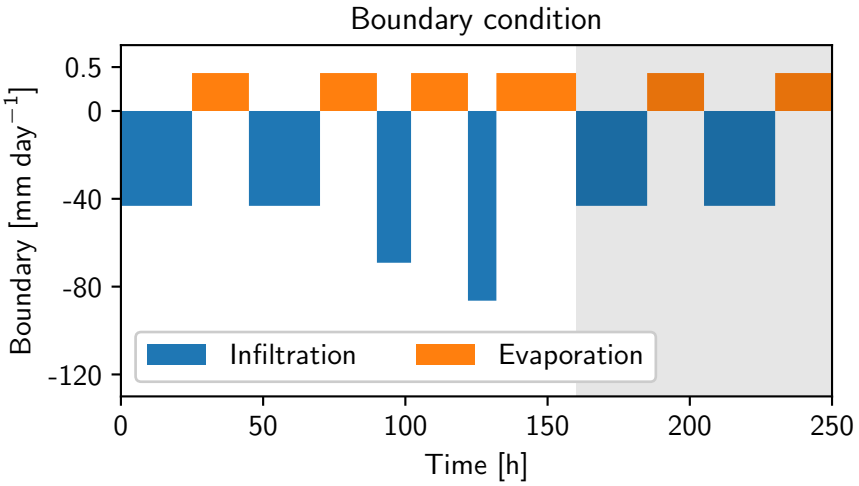

Figure 5.4.: Upper boundary condition for the data assimilation case. Four rain events (blue) each followed by a dry period (orange) are used for the test of the data assimilation run. After this run, two additional rain events and dry periods are used in a free run to test the assimilation results (grey background). Note the different axes for infiltration and evaporation.

Since state and parameters are estimated, an augmented state $\boldsymbol{u}$ (Section 3.4) is used. The model equation (Eq. (3.3)) thus changes to Eq. (3.55) with a constant model for the parameters

$$\boldsymbol{p}^k = \boldsymbol{p}^{k-1} \tag{5.6}$$

and Richards' equation for $f_\psi(\cdot)$. Note that both stochastic model errors of Eq. (3.55) are set to zero.

Forcing: A Dirichlet condition with zero potential (groundwater table) is used as the lower boundary. The upper boundary condition is chosen as a flux boundary (Neumann), representing four rain events with increasing intensity and dry periods in between (see Fig. 5.4). Using infiltrations with an increasing intensity has the advantage that the system has more time to adjust the parameters. This way, less observations are necessary to resolve the infiltration front and the information of the observations can be incorporated in the state and the parameters. The stronger infiltration front at the end of the assimilation run is used to test the robustness of the estimation.

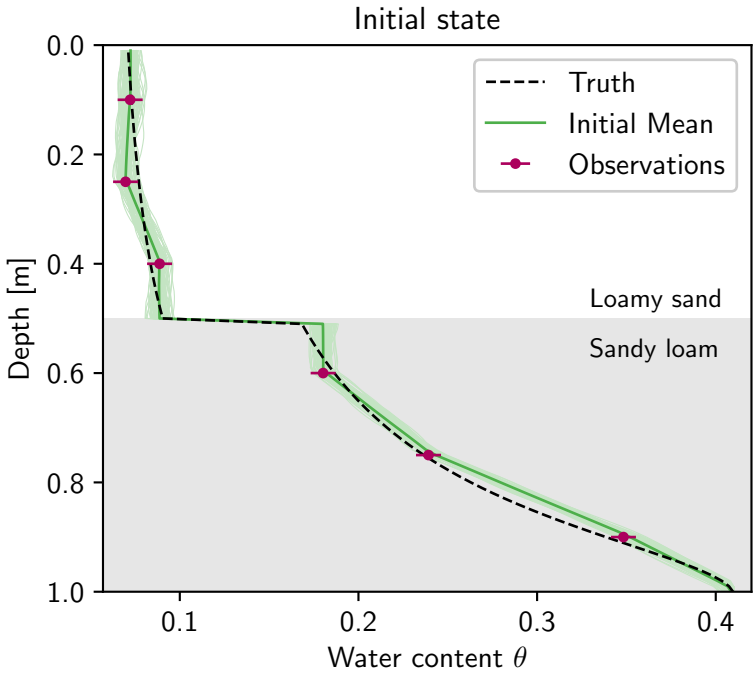

Figure 5.5.: Initial state for the data assimilation run. Observations (purple) at time zero are connected linearly and set constant towards the layer and upper boundary. For the lower boundary, the saturated water content $\theta_r = 0.41$ of sandy loam is used for the interpolation. The ensemble (light green) with 100 ensemble members is generated by perturbing the interpolated state using a spatially correlated Gaussian. The initial truth that is used for the observations (purple) is shown as a black dashed line.

State: Initially, the system is in equilibrium and will be forced by the boundary condition. The initial state is depicted in Fig. 5.5. Six time domain reflectometry (TDR)-like observations are taken equidistantly in each layer at the positions $(0.1, 0.25, 0.3)$ m for layer 1 and $(0.6, 0.75, 0.9)$ m for layer 2. The measurement error is chosen to be $\sigma_{\text{Obs}} = 0.007$ (e.g. *Jaumann and Roth*, 2017). Observations are taken hourly for a duration of 160 h.

For the initial state of the data assimilation, the observations at time zero are used. The measured water content is interpolated linearly between the measurements and kept constant towards the boundary. Additionally, the saturated water content for sandy loam, which is 0.41, is taken as the lower boundary. The approximated state is used as the ensemble mean for the particle filter. This procedure is a viable option for real data although it represents a rather crude approximation of the real initial condition.

Table 5.1.: True Mualem-van Genuchten parameters and range of the uniformly distributed initial guess. The true values are taken from *Carsel and Parrish* (1988).

Parameter	Truth	Lower	Upper
n_1 [−]	2.28	2.2	3.5
n_2 [−]	1.89	1.8	3.2
α_1 [m]	-12.4	-14	-12
α_2 [m]	-7.5	-10.5	-6.5
$\log(K_{w,1})$, K_w in [m s^{-1}]	-4.40	-7	-4
$\log(K_{w,2})$ K_w in [m s^{-1}]	-4.91	-7.5	-4

The approximated state is perturbed by a correlated multivariate Gaussian. The main diagonal of the covariance matrix is 0.003^2. The variance is chosen such that the ensemble represents the uncertainty of the water content in most parts (see Fig. 5.5). The off-diagonal entries are determined by the following two steps: (i) All covariances between the two layers are set to zero to ensure no correlations across the layer boundary. (ii) The remaining entries are the variance of the main diagonal multiplied with the Gaspari and Cohn function (Eq. (3.28)) The distance for the Gaspari and Cohn function is is the distance of the off-diagonal entry from the main diagonal and a length scale of $c = 10\,\text{cm}$ is used. This way, the water content is only correlated in the range of 20 cm.

The use of the covariance increases the diversity of the ensemble and also helps to avoid degeneration. Using uncorrelated Gaussian random numbers with zero mean would result in a fast degeneration of the particle filter caused by the dissipative nature of the system. The perturbation would simply dissipate and the ensemble would collapse.

The initial parameters for the ensemble are uniformly distributed. The ranges of the uniform distributions are given in Table 5.1. Note that the logarithm of the saturated conductivity K_w is used for the estimation, so K_w spans three orders of magnitudes.

The data assimilation run is performed with 100 ensemble members. The multiplicative factor Eq. (5.4) is set to:

$$\boldsymbol{\gamma} = \begin{bmatrix} \mathbf{1}_{100} \\ \mathbf{1.2}_6 \end{bmatrix} , \tag{5.7}$$

where the index denotes the dimension. The covariance in the 100-dimensional state space is unchanged. For the parameter space a factor of 1.2 is used to compensate the missing dynamics.

After the assimilation, the final ensemble is used to run a forecast without data assimilation. In this run two additional infiltration events and evaporation periods are used to test the forecasting ability of the estimated states and parameters (see Fig. 5.4) . The mean of this forward run is calculated from the ensemble with the weights of the last assimilation step.

5.4. Results

The development of the parameters is depicted in Fig. 5.6. The saturated conductivity $K_{w,1}$ (Fig. 5.6a) can be estimated quickly because the infiltration front induces dynamics in the first layer which is strongly dependent on K_w. Instead $K_{w,2}$ (Fig. 5.6b) is not sensitive to the dynamics in the first hours, as the infiltration front did not reach the second layer yet. At around 46 h, the infiltration front reaches the first observation position in the second layer and the estimation for $K_{w,2}$ improves quickly.

If dynamic is induced in the system, the ensemble spread in the water content space increases because of different parameter sets. This makes the particles and their corresponding parameter sets distinguishable and parameter estimation possible. The parameters n_1 and n_2 (Fig. 5.6c and d) as well as α_2 (Fig. 5.6f) can be estimated well. One exception is α_1 (Fig. 5.6e). This parameter is insensitive to the observations. The effect of α on the trajectory of the ensemble members is limited to a small region next to the layer boundary. Further away, wrong values can be compensated by n. The parameter α_1 jitters randomly around a value slightly larger than the truth.

In a synthetic case, the estimation of α_1 can be improved easily by having observations directly next to the boundary. This way a different value for α_1 would have a direct influence on the trajectory and α_1 would become sensitive to the observations. However, in reality it is not feasible to change the measurement position or measure directly next to the layer interface. I decided to retain these positions to illuminate an often encountered practical difficulty.

To see the effect of the parameters on the forward propagation, it is necessary to have a closer look at the conductivity function Eq. (2.10). This function is used for the model forward propagation and many parameter sets can effectively describe the same situation in a limited regime of the hydraulic head. The function is shown in Fig. 5.7 for both layers and reveals that the difference between the truth and the estimated parameters is small.

The final water content state agrees with the synthetic truth in a narrow band (see Fig. 5.8). The estimated state is slightly biased to higher water contents. It was checked that the direction of the bias is a random effect and is different

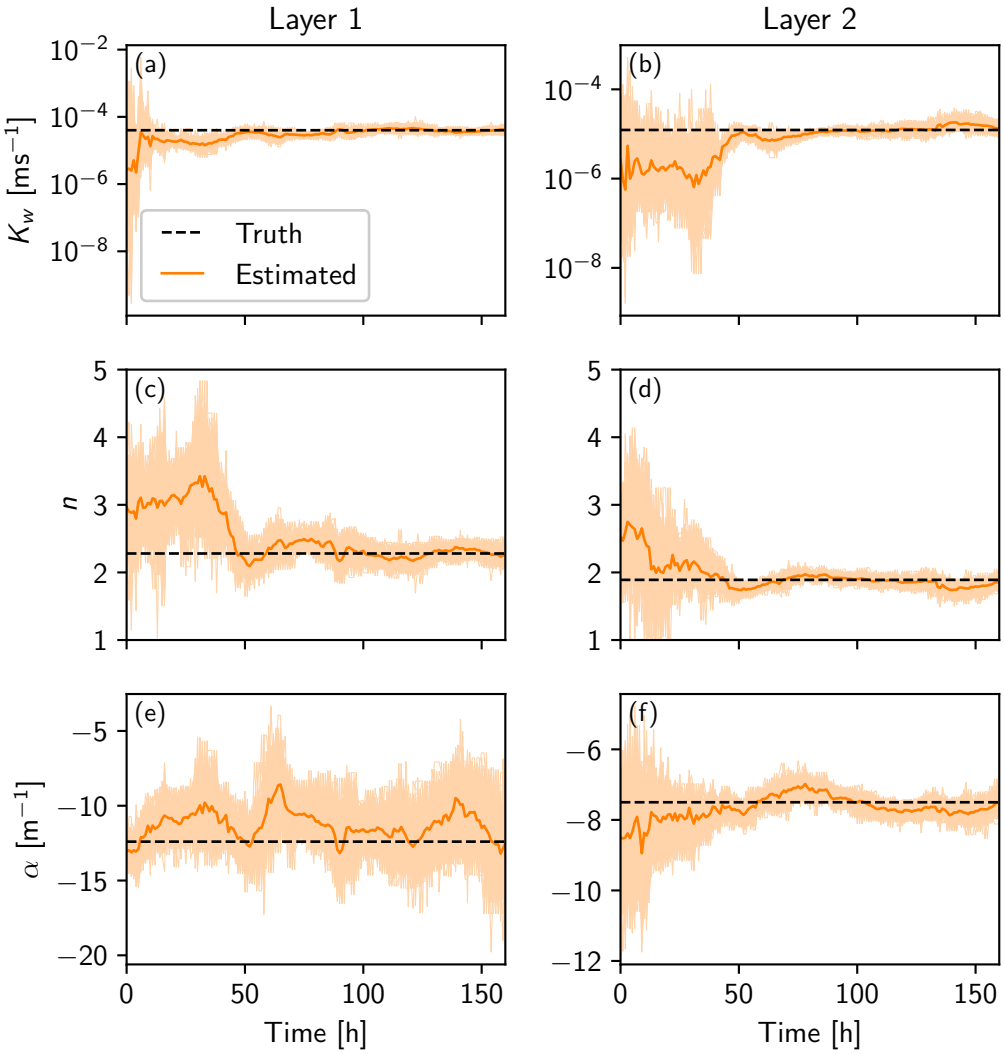

Figure 5.6.: Estimation of K_w, n and α of the Mualem-van Genuchten parametrisation (Eqs. (2.7) and (2.8)) for both layers ((a): $K_{w,1}$, (b): $K_{w,2}$, (c): n_1, (d): n_2, (e): α_1, (f): α_2). The ensemble mean is shown in orange and the ensemble in light orange. The truth is a black dashed line.

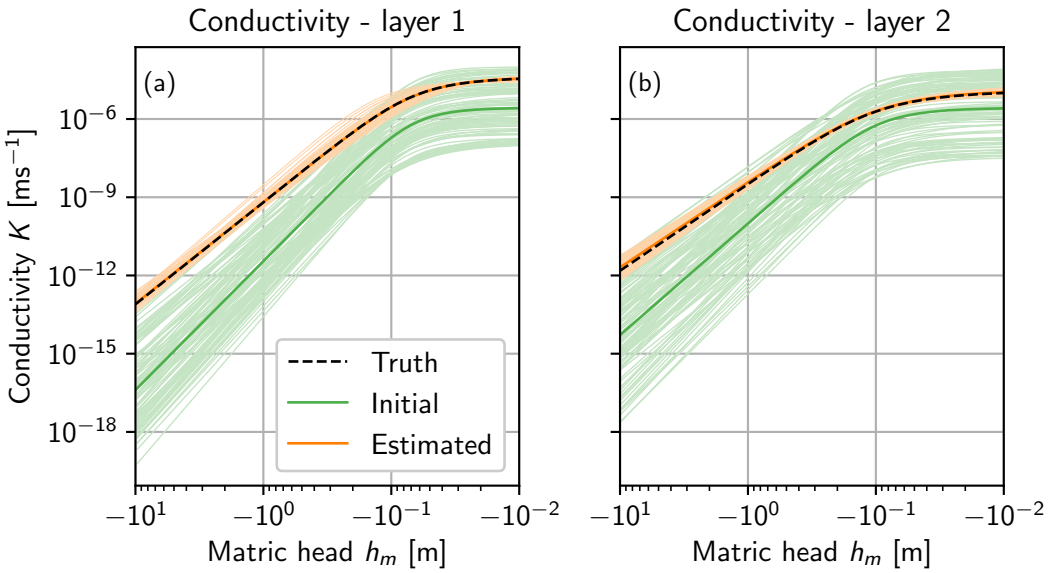

Figure 5.7.: Conductivity function $K(h_m)$ (Eq. (2.10)) for (a): layer 1 and (b): layer 2. In Eq. (2.10), all estimated parameters are represented. The initial ensemble and the mean are shown in light green and green, respectively. The truth (black dashed line) is almost congruent with the estimated mean (orange), such that only the final ensemble (light orange) is visible.

for different seeds. The observation of a bias instead is caused by long-range correlations of the model. In this case, the system has started to relax after the last infiltration and a higher water content in one part results in a higher water content in the rest of the layer. The ensemble spread next to the layer boundary is caused by the large spread of α_1.

To analyse the ensemble, it is useful to take a closer look at the effective sample size (Eq. (3.44)) and the number of resampled particles. Figure 5.9 shows the effective sample size and the number of new particles over time. The effective sample size drops every time new information is available and the number of resampled particles increases. For times $t < 15\,\mathrm{h}$, the effective sample size drops to small values. The infiltration front propagates through the first layer, which leads to a large ensemble spread caused by unknown parameters and only a few particles have a significant weight. The filter assimilates the information from the observations and resamples particles with low weight. This improves the state and parameters and leads to an increasing effective sample size until the infiltration front reaches the second layer ($t \approx 46\,\mathrm{h}$). The effective sample size decreases rapidly because the parameters in the second layer are still unknown and lead to

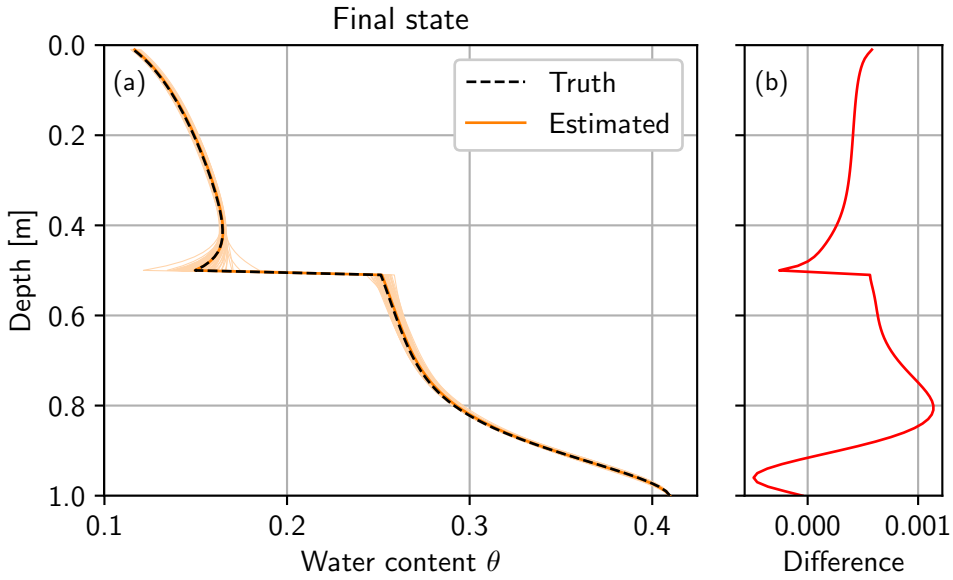

Figure 5.8.: (a): Final water content state after the assimilation run. The truth (black dashed line) is almost congruent with the estimated mean (orange), such that only the final ensemble (light orange) is visible. The final ensemble with the corresponding weights is used to start a free forward run afterwards. (b): The difference (red) of the estimated water content and the synthetic truth lies in a narrow band, with a small bias to larger water contents.

a large ensemble spread again. This variation of the effective sample size occurs every time the filter gets new information that leads to a discrepancy between states of the particles and the observations.

The effective sample size is a crucial parameter for the covariance resampling. If it drops to low values the filter can degenerate because effectively too few particles contribute to the weighted covariance (Eq. (5.3)) and the covariance information becomes insignificant.

For further analysis, the RMSE is calculated based on the difference of the true water content and the weighted mean at every observation time. This includes also the unobserved dimensions. The RMSE shows a similar behaviour as the parameters and the effective sample size (see Fig. 5.10). During the first infiltration, the state and the parameters are improved, the RMSE becomes smaller. Then the infiltration front reaches the boundary interface. The parameters of the second layer and α_1 are still wrong and diverse. This leads to a spread of the ensemble by the system dynamics and also a shift of the mean away from the truth. The parameters in the second layer are estimated and the state is corrected, which

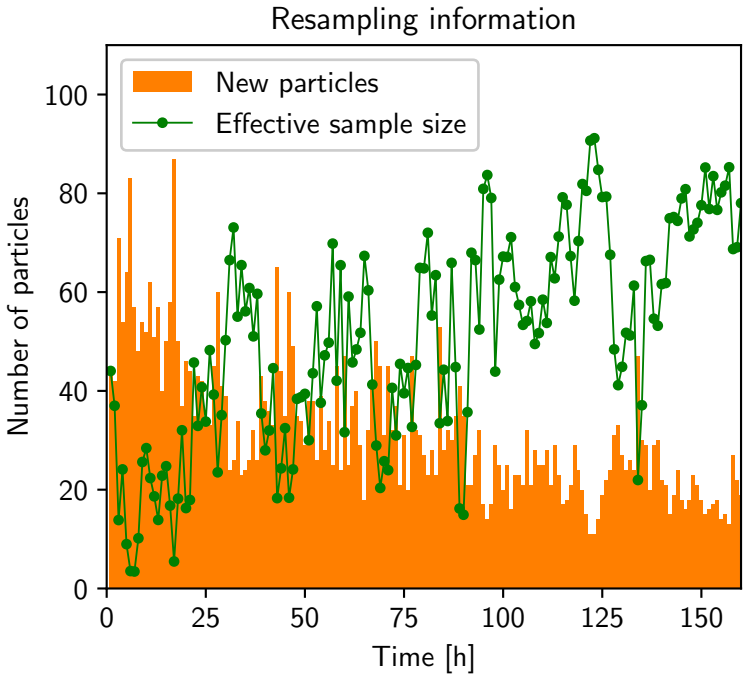

Figure 5.9.: Amount of particles that are resampled (orange) and the effective ensemble size (green dots). The lines connecting the dots are for better visibility of the time dependent behaviour. The effective sample size increases while the number of resampled particles decreases. Every infiltration reduces the effective sample size and leads to more resampled particles.

leads to a decrease in the RMSE. The increase for the next infiltration events becomes smaller since state and parameters are already improved.

The forward run without data assimilation shows that the RMSE oscillates in a narrow range. These oscillations are part of the unobserved space next to the boundary and are mainly caused by the wrong value of α for the first layer. In the beginning, the RMSE stays small, but when the infiltration front reaches the boundary of the two layers, the mean state and the truth begin to deviate from each other (limited to the boundary interface). After the infiltration front passed, the state starts to equilibrate and is increasingly defined by the whole parameter set, which has a certain distance to the true equilibrium.

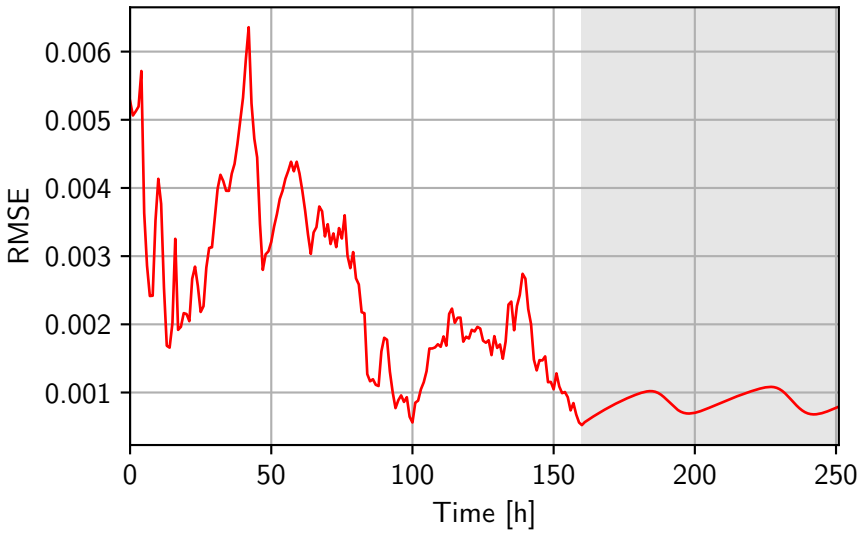

Figure 5.10.: The RMSE of the water content calculated between the truth and the estimated mean of the forward propagated ensemble. The mean includes observed and unobserved dimensions. After 160 hours, the free run starts (grey background). During this forward run, the RMSE is about 10^{-3}. For the assimilation and the free run, the RMSE increases with each infiltration.

5.5. Discussion

In this chapter, I introduced a resampling method for particle filters that uses the covariance information of the ensemble to generate new particles and effectively avoids filter degeneracy. The covariance connects information between observed and unobserved dimensions. This has some similarity to the ensemble Kalman filter (EnKF) but the covariance resampling partly maintains information of the non-Gaussian pdf in the retained ensemble. This was shown for bimodal prior distribution in an one dimensional example for one analysis step. In this example, the covariance resampling was superior to the EnKF in representing the posterior distribution.

The method was also tested in a synthetic one-dimensional unsaturated porous medium with two homogeneous layers. Even with just a rough initial guess, a broad parameter range and only 100 ensemble members, the estimation shows excellent results. After the assimilation, the results are verified in a free run with the final results. Even tough the RMSE of the water content includes the unobserved state dimensions, it stays in a narrow range (RMSE is about 10^{-3}) during the forecast. With every infiltration, the RMSE shows excursions caused

by a wrong value of parameter α in the first layer that results in a wrong state near the layer boundary during the infiltration.

Transferring the information to the unobserved dimensions helps the filter to not degenerate when only a rough initial guess is available. The states and parameters are both altered actively. For the used initial condition, perturbing the parameters only (*Moradkhani et al.*, 2005a), can lead to filter degeneracy because the state is only changed by the dynamics of the system, which does not guarantee a sufficient ensemble spread for low water contents (see Section 4.3.3). Compared to the particle filter with MCMC resampling (*Moradkhani et al.*, 2012; *Vrugt et al.*, 2013), the covariance resampling presented in this study has the advantage that it does not need additional model runs to generate new particles. However, the covariance resampling has to calculate the covariance matrix and perform a Cholesky decomposition every assimilation step. Similar to localisation for the ensemble Kalman filter (*Houtekamer and Mitchell*, 2001; *Hamill et al.*, 2001), it is possible to localise the covariance in the resampling to increase the efficiency.

The effective sample size (Eq. (3.44)) is a crucial parameter for this method. The covariance resampling needs a sufficient effective sample size, otherwise the calculation of the covariance matrix (Eq. (5.3)) becomes inaccurate and the filter may degenerate. In such a situation, the filter can be improved by resetting the weights to N^{-1} or increasing the ensemble size. In our example this was not necessary because the effective sample size was critical only for single assimilation steps.

Different parameter sets can approximately describe the same conductivity function (Eq. (2.10)) in a certain matric head regime. Model dynamics is necessary to differentiate between those sets. If the infiltration covers only a small regime, the conductivity function is only significant in the observed range and can differ from the truth otherwise. This is also reflected in the chosen boundary condition. Starting with infiltrations with low intensity but longer duration helps the filter to explore the water content range slowly and the observations can resolve the infiltration front.

The covariance resampling connects observed with unobserved dimensions to effectively estimate parameters and prevent filter degeneracy. It conserves the first two statistical moments in the limit of large numbers, while partly maintaining the structure of the pdf in the retained ensemble. The method is able to estimate state and parameters in case of a difficult initial condition without additional model evaluations and using a rather small ensemble size.

6. Localised covariance resampling

The particle filter with covariance resampling (PFCR) is able to effectively estimate state and parameters in soil hydrology (see Chapter 5). However, the particle filter and also the PFCR suffer from the curse of dimensionality (Section 3.3.5).

Snyder et al. (2015) state that localisation is the key idea that allows the ensemble Kalman filter (EnKF) to perform well for different geophysical systems. Localisation utilises a common property in geophysical systems that state variables with a large spatial distance are nearly independent. They conclude that also particle filters need localisation to be effective in high-dimensional systems. This conclusion is supported by the recent analysis of *Morzfeld et al.* (2017).

Bengtsson et al. (2003) investigated the possibilities for localisation in particle filters and note the complications due to a lack of spatial smoothness. The calculation of local weights results in a local resampling that introduces discontinuities (*Fearnhead and Künsch*, 2018). These discontinuities cause strong gradients, which typically lead to unphysical behaviour in geophysical models (*Van Leeuwen et al.*, 2015).

To circumvent the problem of local resampling, *Lei and Bickel* (2011) proposed the nonlinear ensemble adjustment filter (NLEAF). The NLEAF utilises a particle filter to correct the bias of the EnKF, which occurs for nonlinear data assimilation. This allows to use the localisation of the EnKF and avoid resampling completely. Another method is the local ensemble Kalman particle filter (LEnKPF) introduced by *Robert and Künsch* (2017), which is based on the ensemble Kalman particle filter (EnKPF) (*Frei and Künsch*, 2013). The EnKPF makes a continuous transition between the EnKF and a particle filter update. *Robert and Künsch* (2017) proposed two approaches to localise the EnKPF. One calculates local analysis pdfs for each dimension, using only observations close to this dimension and ignoring dependencies between the dimensions. This approach still has discontinuities. The second approach divides the observations into blocks, which are assimilated sequentially. Only state dimensions that are directly influenced by the observations in one block are updated. The remaining dimensions are divided into the ones correlated with the updated dimensions and those which are not. The correlated ones are updated such that potential discontinuities to the uncorrelated dimensions are smoothed out. This is repeated until

all blocks are assimilated.

Other suggestions are local particle filters that rely on transformations similar to the ensemble transform Kalman filter (ETKF) (*Bishop et al.*, 2001) or its localised form (*Hunt et al.*, 2007). The transformation is used to transform the forecast ensemble to an analysis ensemble. The transform interpretation of the particle filter has been explored by *Reich* (2013) and *Metref et al.* (2014). Localised particle filters based on the transformation of the ensemble are the local particle filter (LPF) (*Penny and Miyoshi*, 2016) and the nonlinear ensemble transform filter (NETF) (*Tödter and Ahrens*, 2015; *Tödter et al.*, 2016). The LPF augments the particle filter with an observation-space localisation approach, which calculates an independent analysis at each grid point. The local weights are interpolated to smooth the transition between neighbouring points to prevent large gradients. The NETF is comparable to the NLEAF by *Lei and Bickel* (2011). Using a transformation, both construct an analysis ensemble from the forecast ensemble with a mean and a covariance that exactly match the particle filter estimates. The difference between the NETF and the NLEAF is that the NETF uses a deterministic update, while the NLEAF uses a stochastic update.

Poterjoy (2016) calculated local weights for each dimension and localises these weights with a defined correlation function. The applied resampling merges resampled particles with prior particles in the localisation region. The resulting ensemble is corrected for higher order moments with probability mapping. The filter is successfully tested for different high-dimensional systems (*Poterjoy and Anderson*, 2016; *Poterjoy et al.*, 2017).

This chapter introduces localisation for the covariance resampling. This localisation is similar to the localisation of the EnKF, is easy to implement and mathematically simple compared to the other suggested filters. It also has the advantage that including parameter estimation is straightforward. The localised covariance resampling is tested with a divergent 40-dimensional Lorenz-96 model.

6.1. Method

The covariance resampling can be localised analogous to localisation for the EnKF (Section 3.2.1). The weighted covariance matrix $\mathbf{P}$ (Eq. (5.3)) that is used to generate new particles after resampling is multiplied with a localisation matrix ρ

$$\mathbf{P}_{\text{localised}} = \rho \circ \mathbf{P}\,, \tag{6.1}$$

where $\circ$ is the entrywise product (Hadamard product).

By multiplying the covariance matrix with the localisation matrix, it is possible to cut off the spurious correlations that emerge from small ensemble sizes. How-

ever, in contrast to the EnKF, it does not reduce the high-dimensional problem to a local problem. In the EnKF, the forecast covariance is localised and the Kalman gain for the analysis is then calculated using this covariance (see Eq. (3.27)). In the case of covariance resampling, the analysis weights are already calculated and the localisation is used to generate localised new particles.

6.2. Case study

The localised covariance resampling is tested with a 40-dimensional Lorenz-96 model (see Section 2.2). The model is solved with a fourth order Runge Kutta scheme with a time step of $\Delta t = 0.01$.

To generate an initial state for the data assimilation run, the model is run until time 2000 with an initial state $x_i = 4.0 \; \forall i \in [1, 2, \ldots, 39]$ and $x_{40} = 4.001$ using the typical value for the forcing parameter $F = 8$ (*Lorenz and Emanuel*, 1998). The final state of this run is used as the initial state for the data assimilation. This ensures that the state is on the attractor without the initial transient phase. For all data assimilation runs, the initial ensemble is generated by perturbing this final state with a Gaussian $\mathcal{N}(0, 1)$.

The observations are generated by a forward run. Every second dimension is observed using the true value, which is perturbed using a Gaussian with zero mean and a standard deviation of $\sigma_{\text{obs}} = 1.0$. The observation interval is equal to the typical forecast length of $\Delta t_{\text{Obs}} = 0.05$ (corresponding to 6 h (*Lorenz*, 1996)). The chosen values are comparable to those used in other studies (e.g. *Nakano et al.*, 2007; *van Leeuwen*, 2010; *Lei and Bickel*, 2011). The observations are generated for 200 times.

For small effective sample sizes, the perturbation of the covariance resampling is small, which can lead to filter degeneration. This can be alleviated by increasing the perturbation using the tuning parameter $\boldsymbol{\gamma}$ (Eq. (5.4)). Since the Lorenz-96 model is a divergent system, a small increase is sufficient. In this chapter, the factor is chosen to be $\boldsymbol{\gamma} = 1.15$.

To determine the minimum ensemble size to prevent filter degeneration, the RMSE is calculated over all dimensions and time steps, including the unobserved dimensions. Because of the stochastic nature of the PFCR, the RMSE and therefore the minimum ensemble size are strongly seed dependent. The simulations are performed for 40 different seeds and, similar to *Kirchgessner et al.* (2014), the MRMSE is defined as the mean of the RMSE over all realisations to reduce statistical noise.

In the following, three cases will be regarded to test the covariance localisation. In case 1, the PFCR is used to assimilate the observations. In case 2, localisation

is used with different localisation radii and ensemble sizes to find the optimal localisation radius. In case 3, this localisation radius ($c = 4.5$) is used to compare the efficiency of the localised covariance resampling with case 1.

The particle filter's performance is dependent on the tails of the observation pdf, which can lead to degeneration. This can be alleviated by increasing the observation error for the calculation of the weights (*van Leeuwen*, 2003). Therefore, all three cases are performed for the actual observation error and a larger assumed observation error to investigate the behaviour of the particle filter for a broader distribution. The generation of the observations is still done with $\sigma_{\text{Obs}} = 1$ but for the calculation of the weights an observation error $\sigma_{\text{assumed}} = 1.6 \cdot \sigma_{\text{Obs}}$ is assumed. This broadens the Gaussian distribution and leads to higher weights for distant particles. The factor 1.6 was determined after a rough analysis of the RMSE for one ensemble size and different σ_{assumed}. For too large values, the RMSE of the estimated mean increases despite the convergence of the filter since the observation information is assumed to be uncertain. For too small factors, the increase of the probability in the tails is too small to result in a significant effect on the weights.

A thought experiment to illustrate this behaviour: Assume a Gaussian observation with a small error, such that the probability is approximately a delta distribution. This distribution correctly describes the observation error and the true value. The EnKF would directly correct the state to the observation and the truth would be estimated. In contrast, the particle filter would degenerate because the observation has an infinitesimally small volume in phase space and the probability that a particle lies in this observation is close to zero. This leads to degeneracy despite an almost perfect observation.

The localisation matrix for case 2 (Section 6.3.2) and case 3 (Section 6.3.3) is constructed using the Gaspari-Cohn function (Eq. (3.28)). The localisation matrix should resemble the real correlation structure of the used model. Following *Anderson* (2007) and *Anderson* (2009), the distance $|z|$ for the localisation matrix is chosen as

$$|z| = \min(|k - i|, 40 - |k - i|)\,, \tag{6.2}$$

where i and k are the dimensional indices. This distance function reflects the periodic boundary of the Lorenz-96 model.

This localisation matrix is illustrated in Fig. 6.1 for a localisation radius of $c = 7$. The periodic structure is well recognisable for dimensions around one and 40.

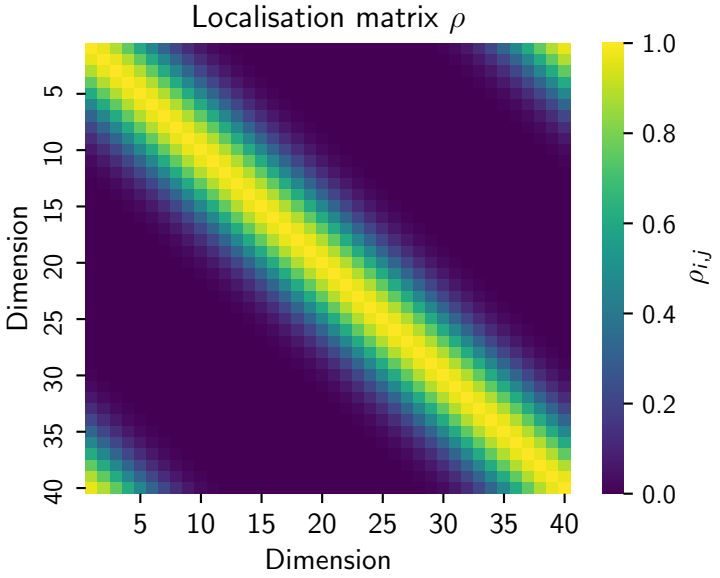

Figure 6.1.: Localisation matrix ρ using the Gaspari-Cohn function, the distance function $z = \min(|k-i|, 40-|k-i|)$ and a localisation radius of $c = 7$ in 40 dimensions.

6.3. Results

In all case studies the MRMSE is calculated from 40 realisations. To compare the cases for convergence, I defined the minimum ensemble size to be the ensemble size for which the MRMSE stays smaller than the true observation noise ($\sigma_{\text{Obs}} = 1$).

6.3.1. Case 1 – no localisation

In this case, the PFCR is applied without using localisation. Figure 6.2 shows the mean (red) and median (blue) of the RMSE and the corresponding 70 %-quantile (darker blue) and the 100 %-quantile (light blue) for a different number of particles.

In the case of of $\sigma_{\text{assumed}} = \sigma_{\text{Obs}}$ (Fig. 6.2a), the minimum ensemble size is approximately 1400. For this ensemble size, the 70 %-quantile stays smaller than 1 but still some realisations degenerate as the 100 %-quantile stays almost constant. For more than approximately 1100 particles, the median drops rapidly, but the mean decreases much slower. The median and mean deviate from each other also for ensemble sizes larger than 1500 because at least 4 realisations still have an RMSE larger than one (see Fig. 6.3a).

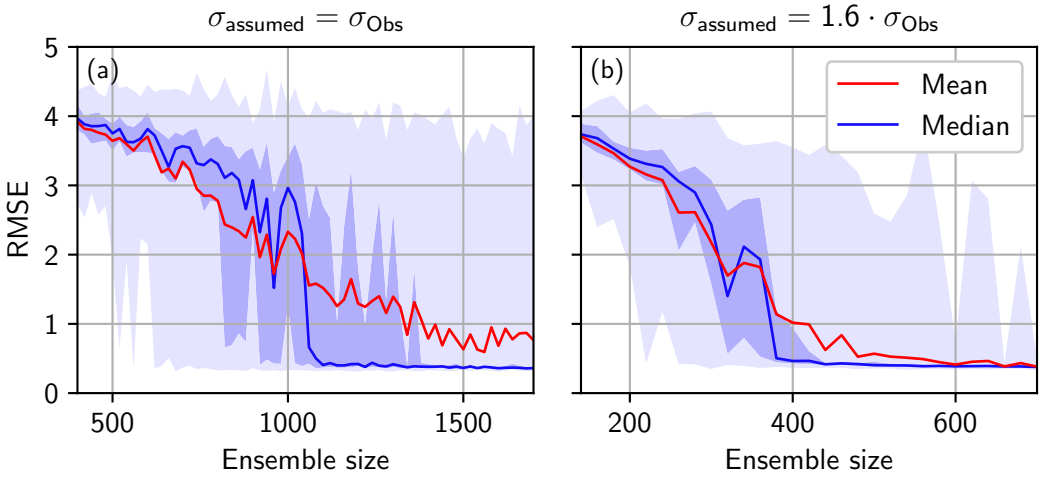

Figure 6.2.: Mean and median of the RMSE for different ensemble sizes for the PFCR without localisation using an assumed observation error of (a): $\sigma_{\text{assumed}} = \sigma_{\text{Obs}}$ and (b): $\sigma_{\text{assumed}} = 1.6 \cdot \sigma_{\text{Obs}}$. The blue areas represent the 70 %-quantile (darker blue) and the 100 %-quantile (light blue), respectively. Note the different scaling of the x-axes.

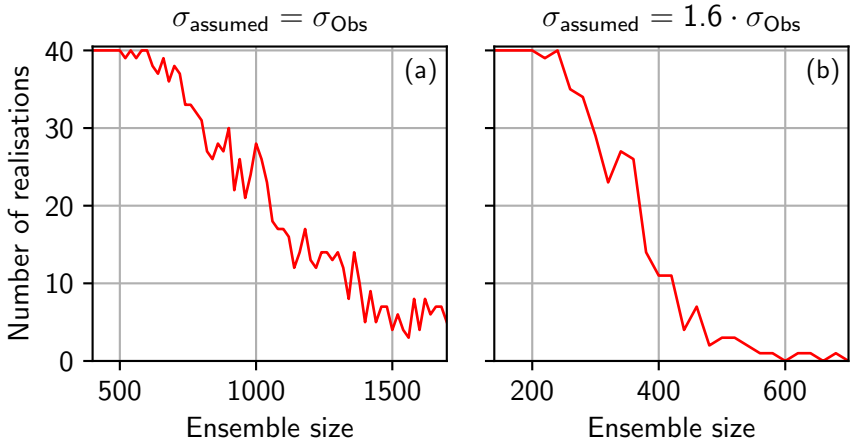

Figure 6.3.: Amount of realisations with an RMSE larger than the true observation error $\sigma_{\text{Obs}} = 1.0$ for the case of the PFCR without localisation using an assumed observation error of (a): $\sigma_{\text{assumed}} = \sigma_{\text{Obs}}$ and (b): $\sigma_{\text{assumed}} = 1.6 \cdot \sigma_{\text{Obs}}$. Note the different scaling of the x-axes.

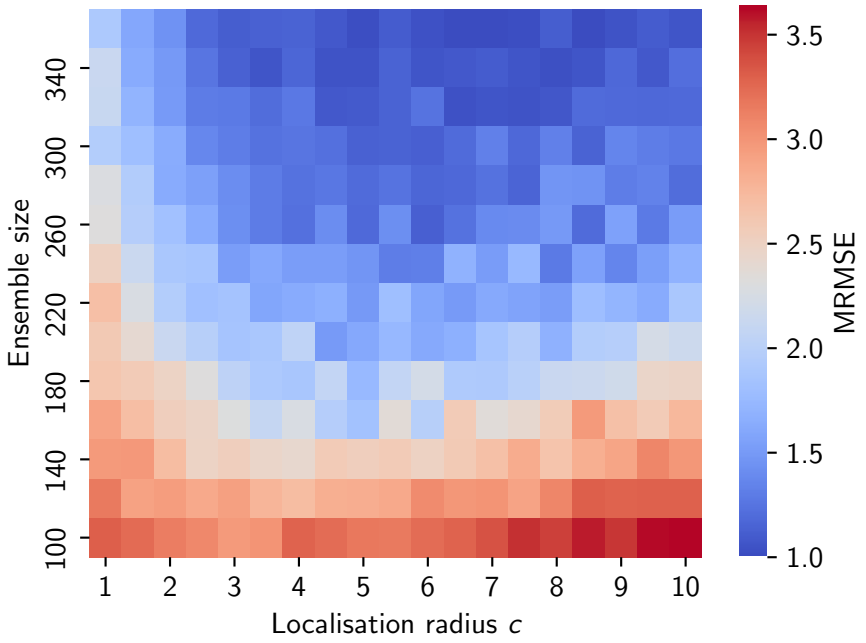

Figure 6.4.: MRMSE for different ensemble sizes and localisation radii c using the localised PFCR.

In the case of $\sigma_{\text{assumed}} = 1.6 \cdot \sigma_{\text{Obs}}$ (Fig. 6.2b), the efficiency increases significantly. The minimum ensemble size is approximately 420. For larger ensemble sizes, the mean and median are almost equal. For ensemble sizes larger than 480 only one realisation has an RMSE larger than one (see Fig. 6.3b). The 100 %-quantile shows that the RMSE of this one realisation is also decreasing.

6.3.2. Case 2 – determine localisation radius

In this case, the localised PFCR is used. The simulations are performed with varied localisation radii and ensemble sizes to investigate the behaviour of the localisation and to determine an optimal localisation radius. The observation error is $\sigma_{\text{Obs}} = 1$. In the case of a larger assumed observation error $\sigma_{\text{assumed}} = 1.6 \cdot \sigma_{\text{Obs}}$ the result for the localisation radii is similar but shifted towards smaller ensemble sizes.

In Fig. 6.4 the MRMSE is shown for different localisation radii and ensemble sizes. The smallest MRMSE can be found for localisation radii in the interval $[3.5, 5.5]$. For smaller localisation radii, the necessary ensemble size to keep the MRMSE constant increases because the reduction of the covariance by the local-

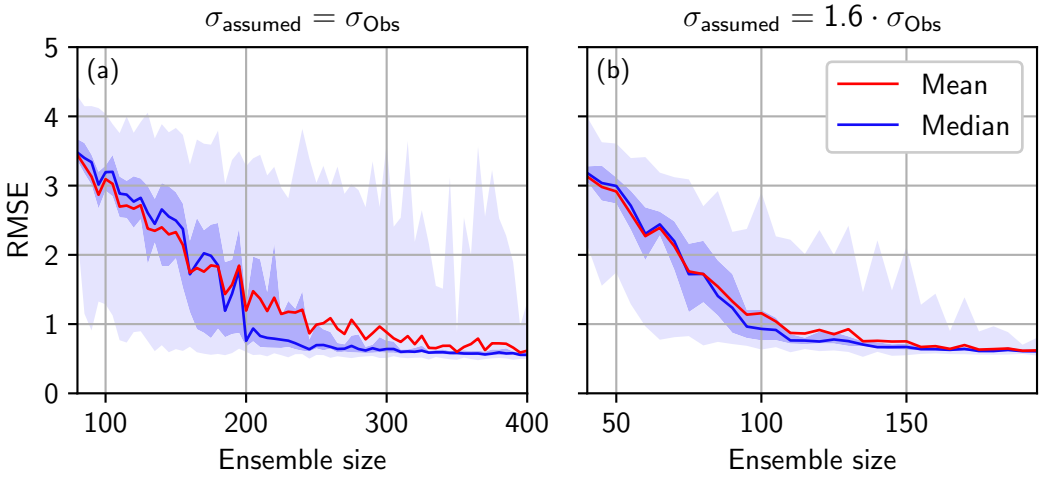

Figure 6.5.: Mean and median of the RMSE for different ensemble sizes using the PFCR with localisation. A localisation radius of $c = 4.5$ is used. The observation error is assumed to be (a): $\sigma_{\text{assumed}} = \sigma_{\text{Obs}}$ and (b): $\sigma_{\text{assumed}} = 1.6 \cdot \sigma_{\text{Obs}}$. The blue areas represent the 70 %-quantile (darker blue) and the 100 %-quantile (light blue), respectively. Note the different scaling of the x-axes.

isation matrix is too strong. For example, in the case of a localisation radius of $c = 1$, only the covariance of the directly adjacent dimension is used to generate a new particle and this covariance is multiplied by a factor of 0.5. This results in new particles that are almost randomly perturbed in each dimension.

For localisation radii larger than 5.5, the necessary ensemble size to keep the MRMSE constant increases with increasing localisation radii because the effects of the spurious correlation are increasing again. The increase of the necessary ensemble size is slower. In this case, if no knowledge is available it is better to choose a larger localisation radius than one that is too small since it can be compensated using an increased ensemble size.

6.3.3. Case 3 – localisation

In this case, the PFCR is used with the localisation. As the MRMSE is smallest for localisation radii in the interval $[3.5, 5.5]$ (cf. Fig. 6.4), the localisation radius is chosen to be $c = 4.5$. This radius is similar to those used for the EnKF e.g. $c = 4$ (*Anderson*, 2007, 2009). Figure 6.5 shows the mean (red) and median (blue) of the RMSE and the corresponding 70 %-quantile (darker blue) and the 100 %-quantile (light blue) for different ensemble sizes.

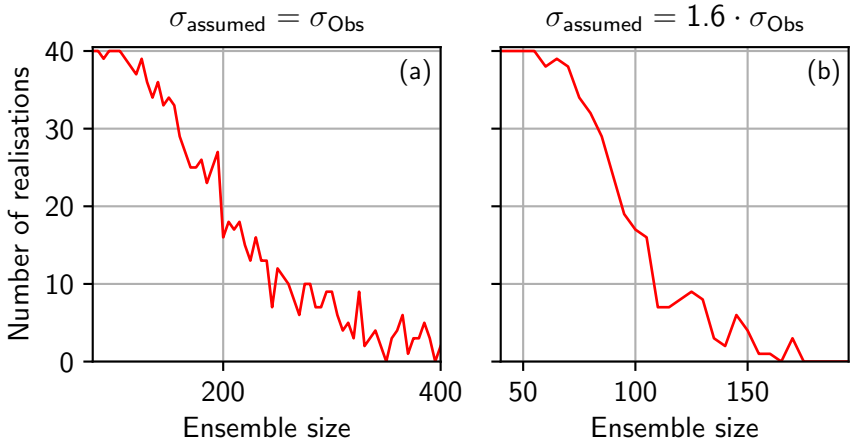

Figure 6.6.: Amount of realisations with an RMSE larger than the true observation error $\sigma_{\text{Obs}} = 1.0$ for the case of the PFCR using localisation. A localisation radius of $c = 4.5$ is used. The observation error is assumed as (a): $\sigma_{\text{assumed}} = \sigma_{\text{Obs}}$ and (b): $\sigma_{\text{assumed}} = 1.6 \cdot \sigma_{\text{Obs}}$. Note the different scaling of the x-axes.

In the case of of $\sigma_{\text{assumed}} = \sigma_{\text{Obs}}$ (Fig. 6.5a), the minimum ensemble size is approximately 280. Without localisation five times more particles are necessary for the same result (see Fig. 6.2a). Furthermore, the 70 %-quantile converges to the median much faster and the 100 %-quantile starts to converge. Mean and median are closer to each other since the number of realisations with an RMSE larger than 1 drops fast for increasing ensemble sizes and sometimes reaches zero (see Fig. 6.6a).

Similar to the case without localisation, the efficiency increases for an assumed observation error of $\sigma_{\text{assumed}} = 1.6 \cdot \sigma_{\text{Obs}}$. The minimum ensemble size is approximately 110 (Fig. 6.5b). The convergence of the realisations improves similar to the case without localisation. The RMSE of the 100 %-quantile decreases and for ensemble sizes larger than 180 no realisation has an RMSE larger than 1 (see Fig. 6.6b).

Fig. 6.7 compares the MRMSE for the case with localisation (dashed lines), without localisation (solid lines), as well as, for $\sigma_{\text{assumed}} = \sigma_{\text{Obs}}$ (red) and a $\sigma_{\text{assumed}} = 1.6 \cdot \sigma_{\text{Obs}}$ (blue). Localisation significantly reduces the necessary ensemble size. For the case $\sigma_{\text{assumed}} = \sigma_{\text{Obs}}$ the efficiency increased by a factor of 5 (compare red lines) and for the case $\sigma_{\text{assumed}} = 1.6 \cdot \sigma_{\text{Obs}}$ by a factor of approximately 3.8 (compare blue lines). Using a larger assumed observation error increases the efficiency for the case without localisation by factor of approximately 3.3 (compare solid lines). For the case with localisation the necessary ensemble

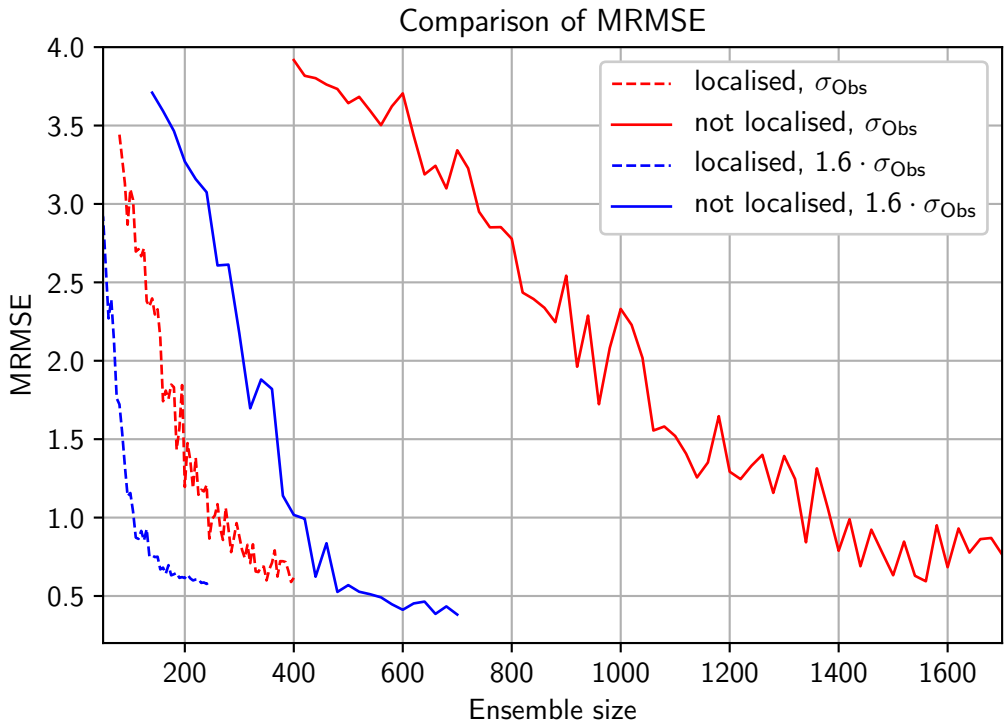

Figure 6.7.: Comparison of the MRMSE for the case of the PFCR with localisation (dashed lines), without localisation (solid lines), as well as with an assumed observation error of $\sigma_{\text{assumed}} = \sigma_{\text{Obs}}$ (red) and $\sigma_{\text{assumed}} = 1.6 \cdot \sigma_{\text{Obs}}$ (blue).

size reduces by a factor of approximately 2.5 (compare dashed lines).

6.4. Summary and discussion

This chapter introduced a localisation method for the PFCR. Using the covariance resampling (introduced in Chapter 5) allows to define a localisation for the particle filter, which is very similar to the one already applied to the EnKF (Section 3.2.1). The spurious long-range correlations that originate from small ensemble sizes are cut off by the localisation matrix. However, the problem is not divided into small subproblems as suggested by *Snyder et al.* (2015) because the localisation is applied after the weighting step and only the newly sampled particles are localised. This is in contrast to the EnKF, where the localisation is applied to the forecast covariance and the analysis is calculated with an already localised covariance.

The localised covariance resampling still suffers from the curse of dimension-

ality. Not all particles end up close to the observation such that the effective sample size decreases to very small values and the covariance matrix is effectively described by only a few particles. This can also be seen in the strong dependence on the width of the observation distribution. Assuming an observation error of 1.6-times the original error increased the efficiency of the PFCR by a factor of 2.5 with localisation and a factor of 3.3 without localisation (see Fig. 6.7).

The efficiency of the PFCR increased significantly using localisation. For the case without localisation five times more particles are necessary to obtain a converging result. The PFCR with localisation is able assimilate the Lorenz-96 model with 280 particles.

It is difficult to compare this outcome with the results of other studies because the filters have been tested for different models or the observations were generated differently. *Robert and Künsch* (2017) used a model based on a modified shallow water equation (*Würsch and Craig*, 2014) and *Tödter and Ahrens* (2015) used a 80-dimensional Lorenz-96 model with Laplace distributed observations. Studies based on a 40-dimensional Lorenz-96 model with Gaussian distributed observation errors obtained good results for ensemble sizes of 40 (*Penny and Miyoshi*, 2016) to 200 (*Poterjoy*, 2016) particles. However, *Penny and Miyoshi* (2016) used randomly distributed observations such that observations can occur between two dimensions. They also applied Gaussian noise in the forecast and the observation error was $\sigma_{\text{Obs}} = 0.5$ compared to $\sigma_{\text{Obs}} = 1.0$ in this study. The more accurate observations increase the predictability for an estimated state but are also more prone to filter degeneration. In this case, the initial state is more important. The most comparable study is possibly *Poterjoy* (2016), who used randomly distributed observations like *Penny and Miyoshi* (2016). The performance of the PFCR with localisation is in the same order of magnitude.

The decisive advantage of the localised covariance resampling is the more accessible mathematics and the resulting much easier implementation compared to other methods. While the other filters have not been tested for simultaneous state and parameter estimation, the PFCR has been demonstrated successfully for both tasks (as shown in Chapter 5), which can also be done using localisation.

7. Summary and conclusion

Three topics were addressed in this thesis: first, the differences and challenges of data assimilation in convergent and divergent systems. Second, parameter estimation with particle filters, for which a new resampling method has been introduced and third, a localisation approach for the new resampling technique.

This work recognises the fundamental differences between divergent and convergent systems for data assimilation. These differences are studied using the divergent Lorenz-96 model and the convergent Richards equation. Successful data assimilation in the presence of a model error, in this study represented by an incorrect parameter, relies on unstable dimensions. These unstable dimensions keep the ensemble spread large enough to assimilate new observations and prevent filter degeneracy. This behaviour is inherent to divergent models such as the Lorenz-96 model, which allows state estimation despite an incorrect parameter. In a convergent system, the ensemble converges to a wrong state due to the incorrect parameter and the filter degenerates. Therefore, the convergent model must be transformed into one with divergent dimensions. This can be done by increasing the ensemble spread artificially using inflation methods or by representing parameter errors with the ensemble using an augmented state. The latter case leads to a diverging ensemble members and thus, to a larger ensemble spread. However, this reduces the predictability of the system. To increase the predictability again, it is necessary to estimate the parameters. In summary, a good representation of model errors and parameter estimation is essential for the success of data assimilation methods in convergent systems.

In soil hydrology, parameters are typically ill-known and need to be estimated. Parameter estimation for particle filters is challenging, since it is necessary to perturb the ensemble members in parameter space to avoid them being identical after resampling, which leads to filter degeneracy. In this thesis, I developed a new resampling method, the covariance resampling. Instead of duplicating particles, new particles are generated using the weighted covariance calculated from the ensemble. The covariance matrix correlates observed with unobserved dimensions to effectively estimate parameters and to prevent filter degeneracy. The covariance resampling conserves the first two statistical moments in the limit of large numbers, while also partly retaining the overall structure of the posterior distribution in the kept ensemble. The effective sample size is a crucial parameter

for the covariance resampling, as for small effective ensemble sizes, the calculation of the covariance matrix becomes inaccurate and the filter may degenerate. The method is successfully demonstrated for state and parameter estimation in a synthetic one-dimensional unsaturated porous medium with two homogeneous layers.

Particle filters suffer from the curse of dimensionality. Localisation is the key to successfully apply particle filters to high-dimensional nonlinear systems (*Snyder et al.*, 2015; *Morzfeld et al.*, 2017). However, the direct transfer of localisation from the EnKF to particle filters is challenging because of discontinuities resulting from local resampling. An advantage of the covariance resampling is that it allows localisation analogous to the one used for the EnKF. However, the localised covariance resampling still suffers from the curse of dimensionality because the localisation is applied to the analysis covariance and not to the forecast covariance as in the EnKF. The localisation reduces the influence of spurious long-range correlations and for the test case of a Lorenz-96 model it increases the efficiency of the particle filter with covariance resampling by a factor of five. This reduce the minimum ensemble size for convergence to 280. It is difficult to compare the performance with other studies because the filters have been tested on different models or the observations were generated differently. Studies using the Lorenz-96 models show good results for ensemble sizes of 40 (*Penny and Miyoshi*, 2016) to 200 (*Poterjoy*, 2016) ensemble members but are still hardly comparable. Both studies used randomly distributed observations and *Penny and Miyoshi* (2016) additionally used more accurate observations. Compared to other localisation methods introduced to the particle filter (e.g. *Poterjoy*, 2016; *Penny and Miyoshi*, 2016; *Robert and Künsch*, 2017), the localised covariance resampling has the advantage of more accessible mathematics and therefore easier implementation. In addition, the particle filter with covariance resampling has been successfully demonstrated for state and parameter estimation and can be extended by localisation.

With the new methods for resampling and localisation introduced in this thesis, an important step towards nonlinear data assimilation with particle filters for large convergent systems has been made. This opens the path to handle also highly nonlinear systems that are hard for the EnKF. Examples of such systems include soil water flow, as demonstrated in this thesis, but also transport of reactive solutes, surface runoff processes and interaction networks in ecosystems.

A. Appendix

A.1. Recursive Bayes' theorem

In this section the recursive Bayes' theorem is derived without the assumption of independent observations (Eq. (3.10)) and without the assumption of Markovian state development (Eq. (3.8)) made in Section 3.1.

First some general relations are defined for the random variables $\boldsymbol{a}$, $\boldsymbol{b}$ and $\boldsymbol{c}$, such that they can be used for the further derivation. The joint probability is:

$$P(\boldsymbol{a}, \boldsymbol{b}) = P(\boldsymbol{a}|\boldsymbol{b})P(\boldsymbol{b}) = P(\boldsymbol{b}|\boldsymbol{a})P(\boldsymbol{a})\,. \tag{A.1}$$

Using Eq. (A.1), the following relations can be defined:

$$P(\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}) = P(\{\boldsymbol{a}, \boldsymbol{b}\}, \boldsymbol{c}) = P(\boldsymbol{a}, \boldsymbol{b}|\boldsymbol{c})P(\boldsymbol{c}) \tag{A.2}$$

$$P(\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}) = P(\boldsymbol{a}, \{\boldsymbol{b}, \boldsymbol{c}\}) = P(\boldsymbol{a}|\boldsymbol{b}, \boldsymbol{c})P(\boldsymbol{b}, \boldsymbol{c}) = P(\boldsymbol{a}|\boldsymbol{b}, \boldsymbol{c})P(\boldsymbol{b}|\boldsymbol{c})P(\boldsymbol{c})\,, \tag{A.3}$$

which directly leads to:

$$P(\boldsymbol{a}, \boldsymbol{b}|\boldsymbol{c}) = P(\boldsymbol{a}|\boldsymbol{b}, \boldsymbol{c})P(\boldsymbol{b}|\boldsymbol{c})\,. \tag{A.4}$$

The starting point of the derivation is Bayes' theorem for a set of discrete observations and states (Eq. (3.6)). The observations $\boldsymbol{d}^{1:k}$ can be sub grouped such that the observation pdfs change to:

$$P(\boldsymbol{d}^{1:k}|\boldsymbol{\psi}^{0:k}) = P(\boldsymbol{d}^{1:k-1}, \boldsymbol{d}^{k}|\boldsymbol{\psi}^{0:k}) \tag{A.5}$$

$$\overset{\text{Eq. (A.4)}}{=} P(\boldsymbol{d}^{k}|\boldsymbol{d}^{1:k-1}, \boldsymbol{\psi}^{0:k})P(\boldsymbol{d}^{1:k-1}|\boldsymbol{\psi}^{0:k}) \tag{A.6}$$

and

$$P(\boldsymbol{d}^{1:k}) = P(\boldsymbol{d}^{1:k-1}, \boldsymbol{d}^{k}) \tag{A.7}$$

$$\overset{\text{Eq. (A.4)}}{=} P(\boldsymbol{d}^{k}|\boldsymbol{d}^{1:k-1})P(\boldsymbol{d}^{1:k-1})\,. \tag{A.8}$$

Inserting Eq. (A.6) and Eq. (A.8) in Eq. (3.6) results in:

$$P(\boldsymbol{\psi}^{0:k}|\boldsymbol{d}^{1:k}) = \frac{P(\boldsymbol{d}^{k}|\boldsymbol{d}^{1:k-1}, \boldsymbol{\psi}^{0:k})P(\boldsymbol{d}^{1:k-1}|\boldsymbol{\psi}^{0:k})P(\boldsymbol{\psi}^{0:k})}{P(\boldsymbol{d}^{k}|\boldsymbol{d}^{1:k-1})P(\boldsymbol{d}^{1:k-1})}\,. \tag{A.9}$$

Using Bayes' theorem (Eq. (3.1)) for $P(\boldsymbol{d}^{1:k-1}|\boldsymbol{\psi}^{0:k})$:

$$P(\boldsymbol{d}^{1:k-1}|\boldsymbol{\psi}^{0:k}) = \frac{P(\boldsymbol{\psi}^{0:k}|\boldsymbol{d}^{1:k-1})P(\boldsymbol{d}^{1:k-1})}{P(\boldsymbol{\psi}^{0:k})}\,, \tag{A.10}$$

Eq. (A.9) becomes

$$P(\boldsymbol{\psi}^{0:k}|\boldsymbol{d}^{1:k}) = \frac{P(\boldsymbol{d}^{k}|\boldsymbol{d}^{1:k-1},\boldsymbol{\psi}^{0:k})P(\boldsymbol{\psi}^{0:k}|\boldsymbol{d}^{1:k-1})}{P(\boldsymbol{d}^{k}|\boldsymbol{d}^{1:k-1})}\,. \tag{A.11}$$

Sub grouping the forecast pdf $P(\boldsymbol{\psi}^{0:k}|\boldsymbol{d}^{1:k-1})$, it becomes:

$$P(\boldsymbol{\psi}^{0:k}|\boldsymbol{d}^{1:k-1}) \quad = \quad P(\boldsymbol{\psi}^{0:k-1},\boldsymbol{\psi}^{k}|\boldsymbol{d}^{1:k-1}) \tag{A.12}$$

$$\overset{\text{Eq. (A.4)}}{=} P(\boldsymbol{\psi}^{k}|\boldsymbol{d}^{1:k-1},\boldsymbol{\psi}^{1:k-1})P(\boldsymbol{\psi}^{0:k-1}|\boldsymbol{d}^{1:k-1}) \tag{A.13}$$

Inserting Eq. (A.13) in Eq. (A.11) results in the general recursive Bayes' theorem:

$$P(\boldsymbol{\psi}^{0:k}|\boldsymbol{d}^{1:k}) = \frac{P(\boldsymbol{d}^{k}|\boldsymbol{d}^{1:k-1},\boldsymbol{\psi}^{0:k})P(\boldsymbol{\psi}^{k}|\boldsymbol{d}^{1:k-1},\boldsymbol{\psi}^{1:k-1})P(\boldsymbol{\psi}^{0:k-1}|\boldsymbol{d}^{1:k-1})}{P(\boldsymbol{d}^{k}|\boldsymbol{d}^{1:k-1})}\,. \tag{A.14}$$

The difference of the recursive Bayes' theorem (Eq. (3.12)) in Section 3.1 using the assumption of independent observations and the assumption of Markovian state development to (Eq. (3.12)) is:

$$\frac{P(\boldsymbol{d}^{k}|\boldsymbol{\psi}^{k})P(\boldsymbol{\psi}^{k}|\boldsymbol{\psi}^{k-1})}{P(\boldsymbol{d}^{k})} \quad \rightarrow \quad \frac{P(\boldsymbol{d}^{k}|\boldsymbol{d}^{1:k-1},\boldsymbol{\psi}^{0:k})P(\boldsymbol{\psi}^{k}|\boldsymbol{d}^{1:k-1},\boldsymbol{\psi}^{1:k-1})}{P(\boldsymbol{d}^{k}|\boldsymbol{d}^{1:k-1})}\,. \tag{A.15}$$

The pdfs are now dependent on the previous observations and states, such that it is impossible to define a sequential filter. For every new observation, it would be necessary to propagate the ensemble from the initial prior $P(\boldsymbol{\psi}^{0})$ again (*Vetra-Carvalho et al.*, 2018).

A.2. Divergent system

This section shows three additional runs for the study of data assimilation in convergent and divergent systems (Chapter 4). These three cases are using the divergent Loren-96 model (Section 2.2). The numbering of Chapter 4 will be continued. Divergent case 3 (DC3) is an addition to the investigation of the influence of a wrong parameter DC2 (Section 4.2.2). Instead of $F = 10$ the ensemble is propagated with $F = 6$, while the observations are still taken with $F = 8$. In divergent case 4 (DC4) and divergent case 5 (DC5) the observation interval is reduced to $\Delta t_{\text{Obs}} = 0.05$, to study the behaviour in a more linear regime of the Lorenz-96 model.

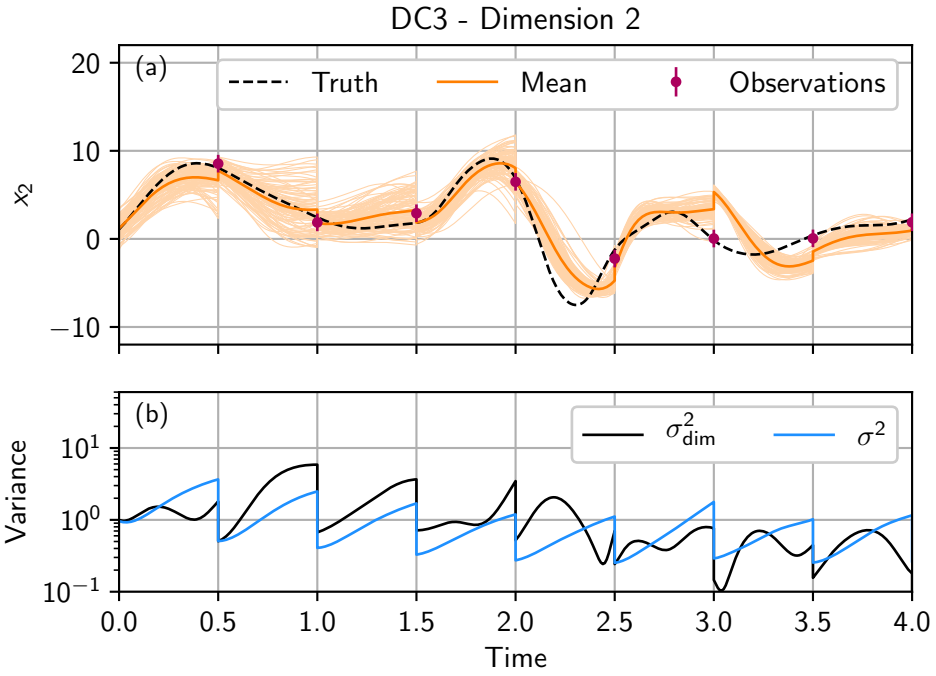

Figure A.1.: Divergent case 3 (DC3): the ensemble is propagated with $F = 6$ instead of $F = 8$ and the state is estimated. (a): The ensemble mean (orange) and the ensemble (light orange) in the data assimilation run for state dimension 2. The observations (purple) are generated with the true model run using $F = 8$. The truth is shown as a black dashed line. (b): Mean variance σ^2 (light blue) of the ensemble over all dimensions and variance σ^2_{dim} (black) of dimension 2.

A.2.1. Divergent case 3 (DC3) – wrong parameter

The observations and the settings are the same as in DC2 Section 4.2.2, except the ensemble is propagated with a forcing parameter of $F = 6$. This results in a less divergent forward propagation. The ensemble spread, σ^2_{dim} and σ^2, (see Fig. A.1b) stays lower than in the case of $F = 10$ (see Fig. 4.3b). This leads to smaller corrections of the EnKF and sometimes the truth is not covered by the ensemble anymore (see Fig. 4.3a). However, due to the divergent system dynamics the filter is able to follow the truth.

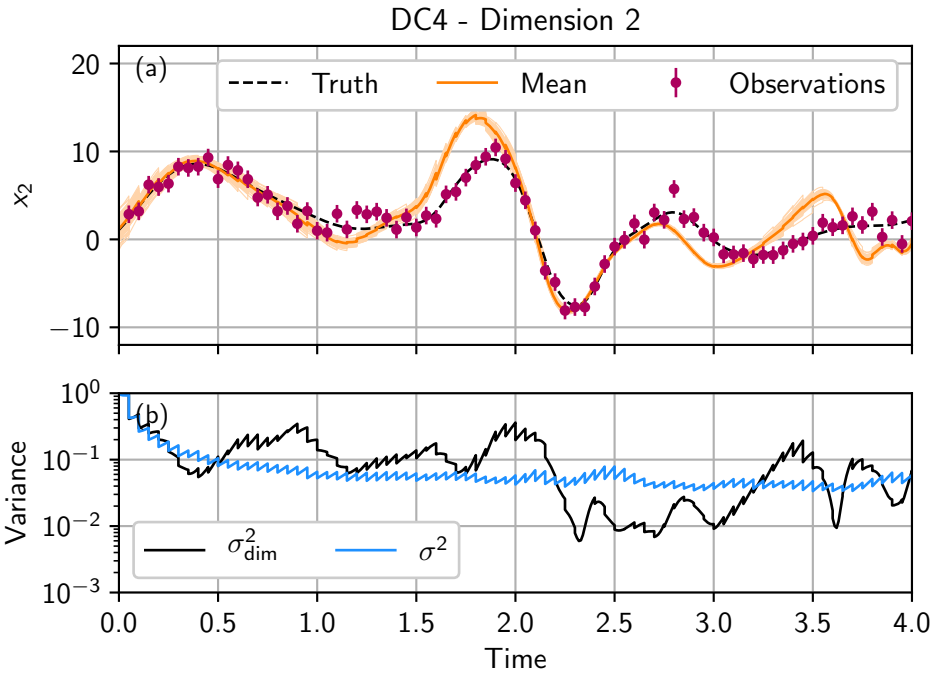

Figure A.2.: Divergent case 4 (DC4): the observation interval is reduced to $\Delta t_{\text{Obs}} = 0.05$, the ensemble is propagated with $F = 10$ instead of $F = 8$ and the state is estimated. (a): The ensemble mean (orange) and the ensemble (light orange) in the data assimilation run for state dimension 2. The observations (purple) are generated with the true model run using $F = 8$. The truth is shown as a black dashed line. (b): Mean variance σ^2 (light blue) of the ensemble over all dimensions and variance σ^2_{dim} (black) of dimension 2.

A.2.2. Divergent case 4 (DC4) – small interval, wrong parameter

In this case, the observation interval is reduced to $\Delta t_{\text{Obs}} = 0.05$. The high frequency of analysis steps, where the Kalman update reduces the variance in the ensemble, prevents the Lorenz-96 model to develop enough divergence to increase the ensemble spread sufficiently (see Fig. A.2b). Although the ensemble is propagated with a more divergent model ($F = 10$), the divergence is not sufficient and the filter degenerates (see Fig. A.2a).

A.2.3. Divergent case 5 (DC5) – small interval, represented error

In this case, the parameter error is represented by the ensemble. Each ensemble member has a different parameter F. The forcing parameter F is drawn from a Gaussian $\mathcal{N}(8, 2^2)$. The true value lies within one standard deviation. Representing the error increases the variance only slightly (Fig. A.3b) compared to DC4 (Appendix A.2.2 and the minimal variance in the dimension has higher values. This helps the filter enough such that it does not degenerate (Fig. A.3a).

Representing the parameter error in the case of frequent observations prevents filter degeneracy. This is similar to the convergent CC3 in Section 4.3.3.

A.3. Experiments with different distributions

This section extends the study in Section 5.2 and illustrates the behaviour of the EnKF and the particle filter with covariance resampling (PFCR) for two additional prior distributions. The first prior is a Gaussian and the second a lognormal distribution.

For all cases 5000 ensemble member are used such that the behaviour of the filter can be shown without statistical noise. The analysis step is calculated for an observation at $d = 3.5$ with three different observation errors $\sigma^2_{\text{Obs}} = \sigma^2_{\text{Prior}}/2$, $\sigma^2_{\text{Obs}} = \sigma^2_{\text{Prior}}$ and $\sigma^2_{\text{Obs}} = 2\sigma^2_{\text{Prior}}$. The variance σ^2_{Prior} is $\sigma^2_{\text{Prior}} = 17$. The particle filter without resampling is not shown in both cases because it samples the posterior almost perfectly, however it would degenerate after several assimilation steps.

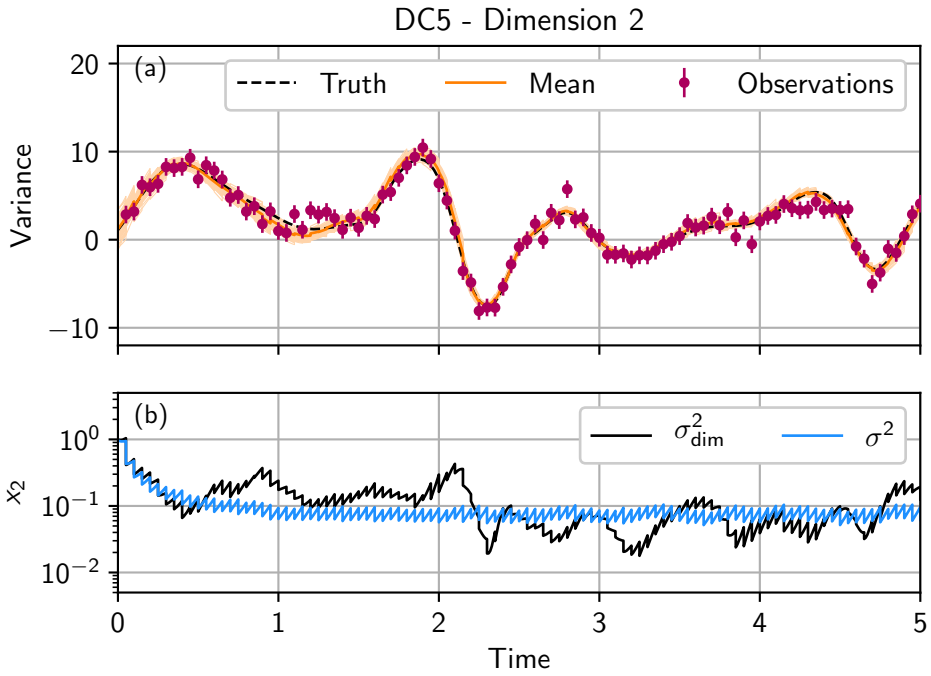

Figure A.3.: Divergent case 5 (DC5): the observation interval is reduced to $\Delta t_{\text{Obs}} = 0.05$, the parameter error is represented by the ensemble using $\mathcal{N}(8, 2^2)$ and the state is estimated. (a): The ensemble mean (orange) and the ensemble (light orange) in the data assimilation run for state dimension 2. The observations (purple) are generated with the true model run using $F = 8$. The truth is shown as a black dashed line. (b): Mean variance σ^2 (light blue) of the ensemble over all dimensions and variance σ^2_{dim} (black) of dimension 2.

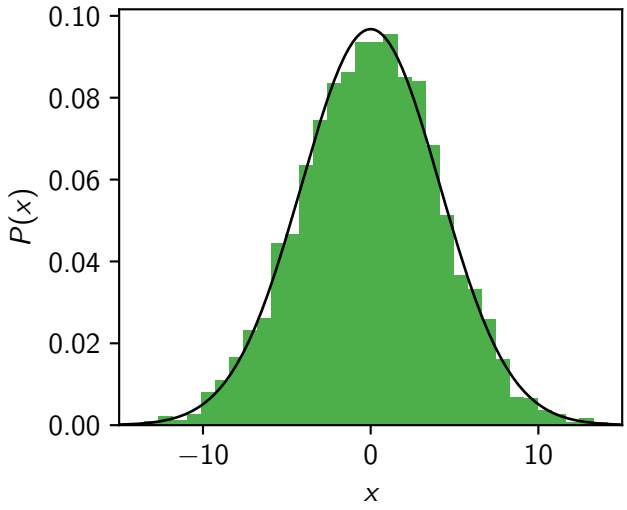

Figure A.4.: Prior Gaussian distribution $P(x)$ (black, $\mathcal{N}(0, 17)$) and histogram of the 5000 ensemble members (green).

A.3.1. Gaussian prior

The prior Gaussian distribution is chosen such that it has the same mean and variance as the bimodal prior (Eq. (5.5)) in Section 5.2. So the initial ensemble is sampled from $\mathcal{N}(0, 17)$. The prior with the initial ensemble is shown in Fig. A.4.

Fig. A.5 shows the analysis for the EnKF and the PFCR. With a Gaussian prior the EnKF assumptions are valid and the representation of the posterior is in all cases excellent. The PFCR shows similar results. Since the posterior is Gaussian and the covariance resampling conserves the first two statistical moments in the limit of large ensemble sizes, newly generated particles fit well to the posterior.

A.3.2. Lognormal prior

In this section the lognormal distribution is used as a prior. This distribution has some similarity with those observed in soil hydrology for an initially Gaussian distributed ensemble in state and parameter space. An initial Gaussian becomes skewed during an infiltration and is limited in the range. The minimal possible water content is the residual water content θ_r and the maximum value is the saturated water content θ_r. However, the exact shape was not analysed in detail such that the lognormal distribution is used at this point for easier calculations.

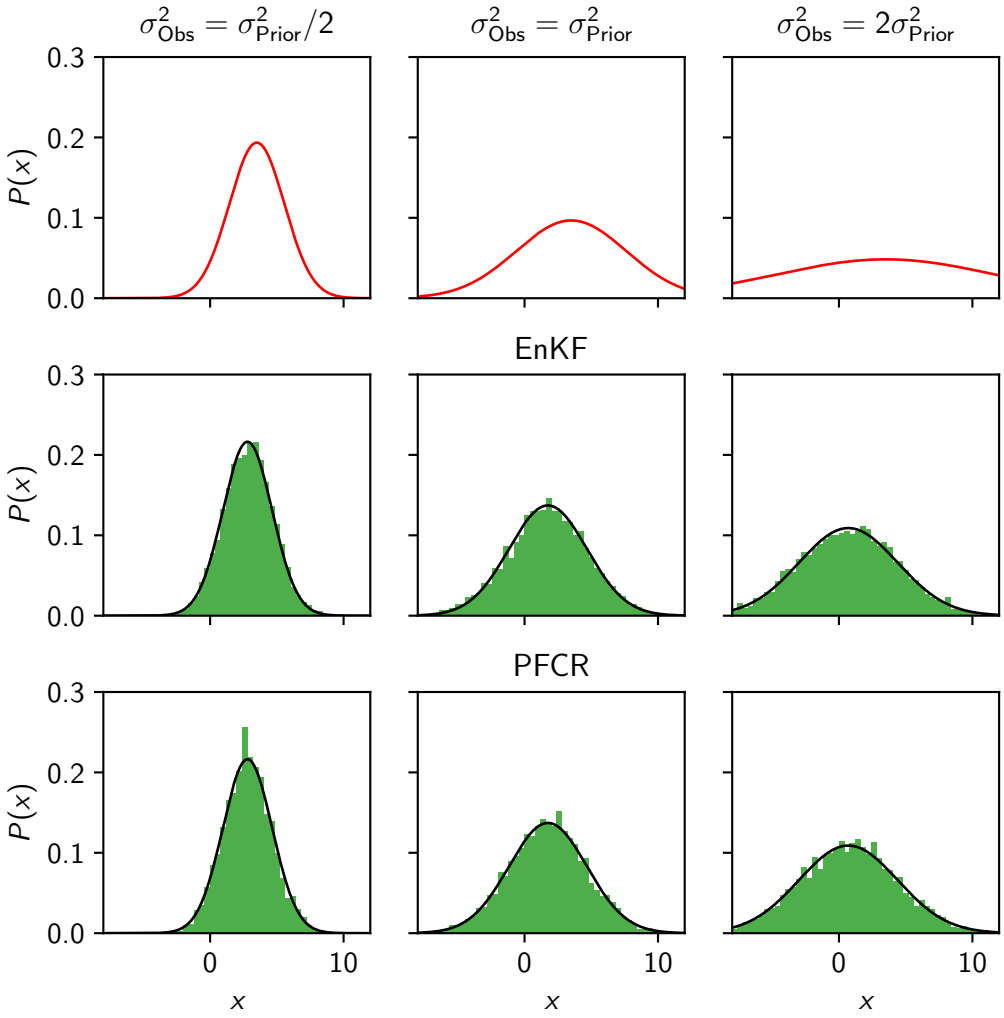

Figure A.5.: Analysis ensemble (green) for EnKF and PFCR for a Gaussian prior (Fig. A.4), for observations (red) with different observation errors: $\sigma^2_{\text{Obs}} = \sigma^2_{\text{Prior}}/2$ (left column), $\sigma^2_{\text{Obs}} = \sigma^2_{\text{Prior}}$ (middle column) and $\sigma^2_{\text{Obs}} = 2\sigma^2_{\text{Prior}}$ (right column). The true posterior (black) is calculated using Bayes' theorem. In case of the PFCR, 47.5 % (left), 23.8 % (middle) and 8.8 % (right) particles are resampled.

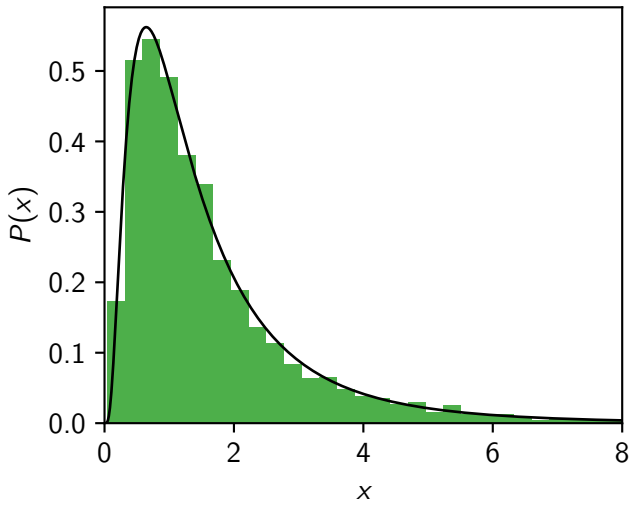

Figure A.6.: Prior lognormal distribution $P(x)$ (black, Eq. (A.16)) and histogram of the 5000 ensemble members (green).

The lognormal distribution is defined as

$$P(x) = \begin{cases} \frac{1}{\sqrt{2\pi}\lambda x} \exp\left[-\frac{(\ln(x)-\mu)^2}{2\lambda^2}\right] & x > 0 \\ 0 & x \leq 0 \end{cases} \tag{A.16}$$

where μ and λ are the mean and standard deviation of the natural logarithm of x. The values for are chosen as $\mu = 0.2$ and $\lambda = 0.8$. For those values the EnKF's capabilities in representing the posterior decreases strongly. For a less skewed function ($\lambda \lesssim 0.5$) the EnKF is capable in representing the posterior since the function becomes more Gaussian. The prior with the chosen parameters and the generated ensemble is shown in Fig. A.6

Fig. A.5 shows the analysis for the EnKF and the PFCR. The EnKF shifts the maximum of the distribution to the observation. This results in a bias towards the observation. Except for the inaccurate observation $\sigma^2_{\text{Obs}} = 2\sigma^2_{\text{prior}}$, parts of the prior are conserved since the Kalman gain (Eq. (3.19)) becomes small for large observation errors.

The PFCR can sample all posteriors with only small deviations because maximal 16.2 % of the particles were resampled, such that the prior distribution is conserved in the retained ensemble.

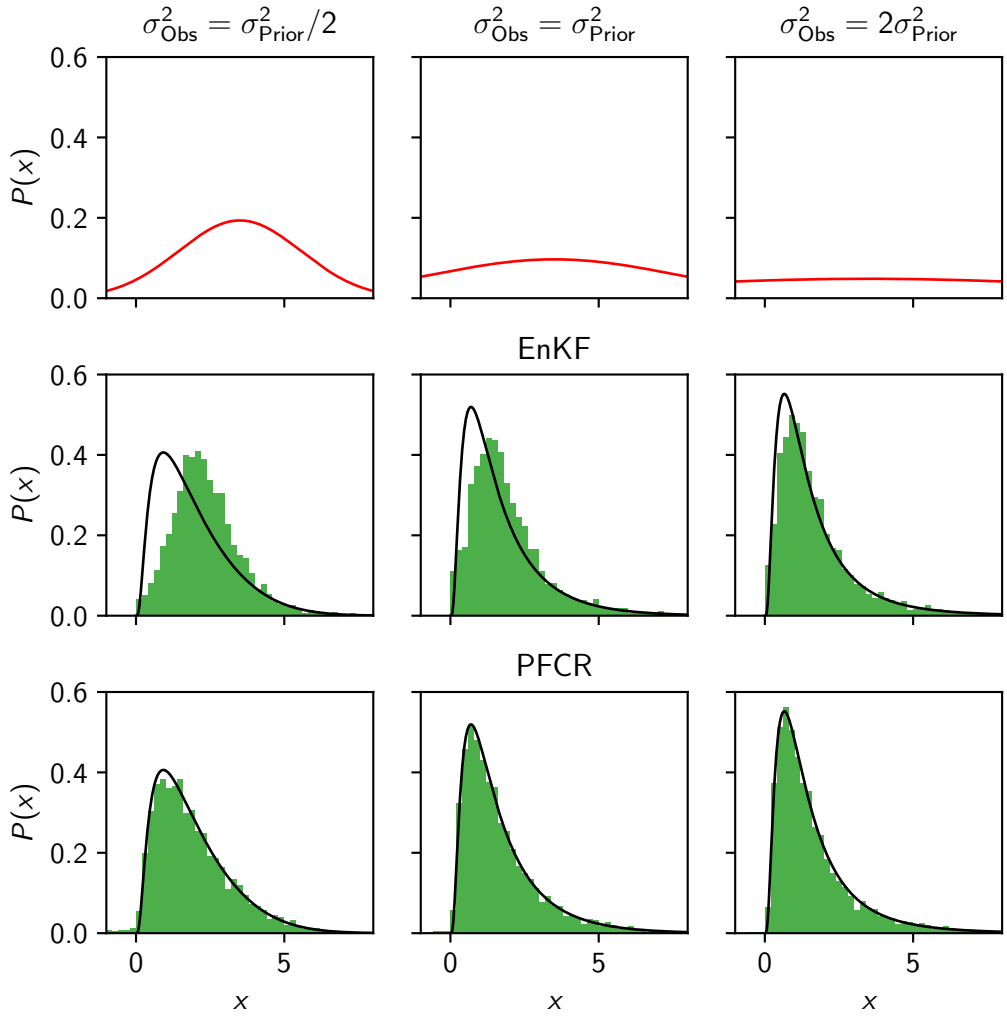

Figure A.7.: Analysis ensemble (green) for EnKF and PFCR for a lognormal prior (Fig. A.6), for observations (red) with different observation errors: $\sigma^2_{\text{Obs}} = \sigma^2_{\text{Prior}}/2$ (left column), $\sigma^2_{\text{Obs}} = \sigma^2_{\text{Prior}}$ (middle column) and $\sigma^2_{\text{Obs}} = 2\sigma^2_{\text{Prior}}$ (right column). The true posterior (black) is calculated using Bayes' theorem. In case of the PFCR, 16.2 % (left), 4.2 % (middle) and 0.6 % (right) particles are resampled.

A.4. Generation of correlated random numbers

This section is based on *Berg et al.* (2018).

A.4.1. Cholesky Decomposition

Correlated random numbers are generated using the Cholesky decomposition. We use the LDLT decomposition which is part of the *Eigen3* library (*Guennebaud, Jacob et al.*, 2010). Decomposing the covariance matrix $\mathbf{Q}$ leads to

$$\mathbf{Q} = \mathbf{L}\mathbf{D}\mathbf{L}^\intercal \,, \tag{A.17}$$

where $\mathbf{D}$ is a diagonal matrix and $\mathbf{L}$ is a lower unit triangular matrix. The LDLT form of the decomposition is related to the LLT-form by

$$\mathbf{Q} = \mathbf{L}'\mathbf{L}'^\intercal \quad \text{with} \quad \mathbf{L}' := \mathbf{L}\mathbf{D}^{\frac{1}{2}} \,. \tag{A.18}$$

To draw a random vector $\mathbf{y}$ from a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \mathbf{Q})$ with mean $\boldsymbol{\mu}$, we first generate a normal distributed ($\mathcal{N}(\mathbf{0}, \mathbf{I})$) random vector $\boldsymbol{x}$. This vector is multiplied with $\mathbf{L}'$ and the mean $\boldsymbol{\mu}$ is added:

$$\boldsymbol{y} = \mathbf{L}'\boldsymbol{x} + \boldsymbol{\mu} \tag{A.19}$$

To verify that this gives the correct result the covariance matrix of $\boldsymbol{y}$ is calculated:

$$\overline{(\boldsymbol{y} - \boldsymbol{\mu})(\boldsymbol{y} - \boldsymbol{\mu})^\intercal} = \overline{\mathbf{L}'\boldsymbol{x}(\mathbf{L}'\boldsymbol{x})^\intercal} = \mathbf{L}'\overline{\boldsymbol{x}\boldsymbol{x}^\intercal}\mathbf{L}'^\intercal = \mathbf{L}'\mathbf{I}\mathbf{L}'^\intercal = \mathbf{Q} \tag{A.20}$$

yields $\mathbf{Q}$ as required.

A.4.2. Regularisation of the ensemble covariance matrix

The calculation of the Cholesky decomposition (LDLT-version) is only possible if the matrix is not indefinite. Mathematically, a covariance matrix has to be positive semidefinite:

$$\begin{aligned} \boldsymbol{v}^\intercal\mathbf{Q}\boldsymbol{v} &= \boldsymbol{v}^\intercal \sum (\boldsymbol{y}_i - \boldsymbol{\mu})(\boldsymbol{y}_i - \boldsymbol{\mu})^\intercal \boldsymbol{v} && \text{(A.21)} \\ &= \sum \boldsymbol{v}^\intercal (\boldsymbol{y}_i - \boldsymbol{\mu})(\boldsymbol{y}_i - \boldsymbol{\mu})^\intercal \boldsymbol{v} && \text{(A.22)} \\ &= \sum (\boldsymbol{v}^\intercal (\boldsymbol{y}_i - \boldsymbol{\mu}))^2 \geq 0 \quad \text{with} \quad \boldsymbol{v} \in \mathbb{R}^d \,, && \text{(A.23)} \end{aligned}$$

but the covariance matrix calculated from our ensemble is occasionally indefinite. The reason for the covariance matrix being indefinite is a numerical error in the

calculation of this matrix. In fact, the calculation of the eigenvalues λ results in negative values in the order of $\mathcal{O}(10^{-20})$.

For this purpose, the identity matrix $\mathbf{I}$, which is multiplied by a scalar $\lambda_{\max}$, is added to the covariance matrix. The value of $\lambda_{\max}$ is in the order of magnitude of the largest negative eigenvalue of $\mathbf{Q}$. Thus, the regularised covariance matrix reads

$$\mathbf{Q}_{\text{Reg.}} = \mathbf{Q} + \lambda_{\max}\mathbf{I}\,. \tag{A.24}$$

In the experiments of Chapter 5 and Chapter 6, the smallest variance on the main diagonal of the covariance matrix is still 16 orders of magnitude larger than $\lambda_{\max}$ such that the influence of this correction is negligible and does not change the results.

Bibliography

Abbaszadeh, P., H. Moradkhani, and H. Yan, Enhancing hydrologic data assimilation by evolutionary particle filter and Markov chain Monte Carlo, *Advances in Water Resources*, *111*, 192 – 204, doi:10.1016/j.advwatres.2017.11.011, 2018.

Ades, M., and P. J. van Leeuwen, An exploration of the equivalent weights particle filter, *Quarterly Journal of the Royal Meteorological Society*, *139*(672), 820–840, doi:10.1002/qj.1995, 2013.

Ades, M., and P. J. van Leeuwen, The effect of the equivalent-weights particle filter on dynamical balance in a primitive equation model, *Monthly Weather Review*, *143*(2), 581–596, doi:10.1175/MWR-D-14-00050.1, 2015a.

Ades, M., and P. J. van Leeuwen, The equivalent-weights particle filter in a high-dimensional system, *Quarterly Journal of the Royal Meteorological Society*, *141*(687), 484–503, doi:10.1002/qj.2370, 2015b.

Anderson, J. L., An adaptive covariance inflation error correction algorithm for ensemble filters, *Tellus A*, *59*(2), 210–224, doi:10.1111/j.1600-0870.2006.00216.x, 2007.

Anderson, J. L., Spatially and temporally varying adaptive covariance inflation for ensemble filters, *Tellus A*, *61*(1), 72–83, doi:10.1111/j.1600-0870.2008.00361.x, 2009.

Anderson, J. L., Localization and sampling error correction in ensemble Kalman filter data assimilation, *Monthly Weather Review*, *140*(7), 2359–2371, doi:10.1175/MWR-D-11-00013.1, 2012.

Anderson, J. L., and S. L. Anderson, A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts, *Monthly Weather Review*, *127*(12), 2741–2758, doi:10.1175/1520-0493(1999)127<2741:AMCIOT>2.0.CO;2, 1999.

Bauser, H. H., Knowledge Fusion in soil hydrology, dissertation, Heidelberg University, 2018.

Bauser, H. H., S. Jaumann, D. Berg, and K. Roth, EnKF with closed-eye period – towards a consistent aggregation of information in soil hydrology, *Hydrology and Earth System Sciences*, *20*(12), 4999–5014, doi:10.5194/hess-20-4999-2016, 2016.

Bauser, H. H., D. Berg, O. Klein, and K. Roth, Inflation method for ensemble Kalman filter in soil hydrology, *Hydrology and Earth System Sciences Discussions*, *2018*, 1–18, doi:10.5194/hess-2018-74, in review, 2018.

Bengtsson, T., C. Snyder, and D. Nychka, Toward a nonlinear ensemble filter for high-dimensional systems, *Journal of Geophysical Research: Atmospheres*, *108*(D24), doi:10.1029/2002JD002900, 8775, 2003.

Bengtsson, T., P. Bickel, and B. Li, *Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems*, *Collections*, vol. Volume 2, pp. 316–334, Institute of Mathematical Statistics, Beachwood, Ohio, USA, doi:10.1214/193940307000000518, 2008.

Berg, D., H. H. Bauser, and K. Roth, Covariance resampling for particle filter – state and parameter estimation for soil hydrology, *Hydrology and Earth System Sciences Discussions*, *2018*, 1–20, doi:10.5194/hess-2018-121, in review, 2018.

Bickel, P., B. Li, and T. Bengtsson, *Sharp failure rates for the bootstrap particle filter in high dimensions*, *Collections*, vol. Volume 3, pp. 318–329, Institute of Mathematical Statistics, Beachwood, Ohio, USA, doi:10.1214/074921708000000228, 2008.

Bishop, C. H., B. J. Etherton, and S. J. Majumdar, Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects, *Monthly Weather Review*, *129*(3), 420–436, doi:10.1175/1520-0493(2001)129<0420:ASWTET>2.0.CO;2, 2001.

Blatt, M., et al., The distributed and unified numerics environment, version 2.4, *Archive of Numerical Software*, *4*(100), 13–29, 2016.

Botto, A., E. Belluco, and M. Camporese, Multi-source data assimilation for physically-based hydrological modeling of an experimental hillslope, *Hydrology and Earth System Sciences Discussions*, *2018*, 1–26, doi:10.5194/hess-2018-18, in review, 2018.

Buehner, M., P. L. Houtekamer, C. Charette, H. L. Mitchell, and B. He, Intercomparison of variational data assimilation and the ensemble Kalman filter for global deterministic NWP. Part I: Description and

single-observation experiments, *Monthly Weather Review*, *138*(5), 1550–1566, doi:10.1175/2009MWR3157.1, 2010a.

Buehner, M., P. L. Houtekamer, C. Charette, H. L. Mitchell, and B. He, Intercomparison of variational data assimilation and the ensemble Kalman filter for global deterministic NWP. Part II: One-month experiments with real observations, *Monthly Weather Review*, *138*(5), 1567–1586, doi:10.1175/2009MWR3158.1, 2010b.

Burgers, G., P. J. van Leeuwen, and G. Evensen, Analysis scheme in the ensemble Kalman filter, *Monthly Weather Review*, *126*(6), 1719–1724, doi:10.1175/1520-0493(1998)126<1719:ASITEK>2.0.CO;2, 1998.

Carsel, R. F., and R. S. Parrish, Developing joint probability distributions of soil water retention characteristics, *Water Resources Research*, *24*(5), 755–769, doi:10.1029/WR024i005p00755, 1988.

Chen, Y., and D. Zhang, Data assimilation for transient flow in geologic formations via ensemble Kalman filter, *Advances in Water Resources*, *29*(8), 1107 – 1122, doi:10.1016/j.advwatres.2005.09.007, 2006.

Chorin, A. J., and M. Morzfeld, Conditions for successful data assimilation, *Journal of Geophysical Research: Atmospheres*, *118*(20), 11,522–11,533, doi:10.1002/2013JD019838, 2013.

Chorin, A. J., and X. Tu, Implicit sampling for particle filters, *Proceedings of the National Academy of Sciences*, *106*(41), 17,249–17,254, doi:10.1073/pnas.0909196106, 2009.

Daley, R., *Atmospheric data analysis*, 2, Cambridge university press, 1993.

DeChant, C. M., and H. Moradkhani, Examining the effectiveness and robustness of sequential data assimilation methods for quantification of uncertainty in hydrologic forecasting, *Water Resources Research*, *48*(4), doi:10.1029/2011WR011011, 2012.

Doucet, A., On sequential simulation-based methods for Bayesian filtering, 1998.

Doucet, A., S. Godsill, and C. Andrieu, On sequential monte carlo sampling methods for bayesian filtering, *Statistics and Computing*, *10*(3), 197–208, doi:10.1023/A:1008935410038, 2000.

Erdal, D., M. Rahman, and I. Neuweiler, The importance of state transformations when using the ensemble Kalman filter for unsaturated flow modeling: Dealing with strong nonlinearities, *Advances in Water Resources*, *86*, 354 – 365, doi:10.1016/j.advwatres.2015.09.008, 2015.

Evensen, G., Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, *Journal of Geophysical Research: Oceans*, *99*(C5), 10,143–10,162, doi:10.1029/94JC00572, 1994.

Farthing, M. W., and F. L. Ogden, Numerical solution of Richards' equation: A review of advances and challenges, *Soil Science Society of America Journal*, doi:10.2136/sssaj2017.02.0058, 2017.

Fearnhead, P., and H. R. Künsch, Particle filters and data assimilation, *Annual Review of Statistics and Its Application*, *5*(1), 421–449, doi:10.1146/annurev-statistics-031017-100232, 2018.

Fertig, E. J., J. Harlim, and B. R. Hunt, A comparative study of 4D-VAR and a 4D ensemble Kalman filter: perfect model simulations with Lorenz-96, *Tellus A*, *59*(1), 96–100, doi:10.1111/j.1600-0870.2006.00205.x, 2007.

Franssen, H. J. H., and W. Kinzelbach, Real-time groundwater flow modeling with the ensemble Kalman filter: Joint estimation of states and parameters and the filter inbreeding problem, *Water Resources Research*, *44*(9), doi:10.1029/2007WR006505, 2008.

Frei, M., and H. R. Künsch, Bridging the ensemble Kalman and particle filters, *Biometrika*, *100*(4), 781–800, doi:10.1093/biomet/ast020, 2013.

Gaspari, G., and S. E. Cohn, Construction of correlation functions in two and three dimensions, *Quarterly Journal of the Royal Meteorological Society*, *125*(554), 723–757, doi:10.1002/qj.49712555417, 1999.

Gharamti, M. E., Enhanced adaptive inflation algorithm for ensemble filters, *Monthly Weather Review*, *146*(2), 623–640, doi:10.1175/MWR-D-17-0187.1, 2018.

Gordon, N. J., D. J. Salmond, and A. F. M. Smith, Novel approach to nonlinear/non-Gaussian Bayesian state estimation, *IEE Proceedings F - Radar and Signal Processing*, *140*(2), 107–113, doi:10.1049/ip-f-2.1993.0015, 1993.

Guennebaud, G., B. Jacob, et al., Eigen v3, `http://eigen.tuxfamily.org`, 2010.

Hamill, T. M., J. S. Whitaker, and C. Snyder, Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter, *Monthly Weather Review*, *129*(11), 2776–2790, doi:10.1175/1520-0493(2001)129<2776:DDFOBE>2.0.CO;2, 2001.

Harlim, J., and A. J. Majda, Catastrophic filter divergence in filtering nonlinear dissipative systems, *Commun. Math. Sci.*, *8*(1), 27–43, 2010.

Hillel, D., *Introduction to environmental soil physics*, first ed., Academic Press, Burlington, doi:10.1016/B978-012348655-4/50025-3, 2003.

Houtekamer, P. L., and H. L. Mitchell, Data assimilation using an ensemble Kalman filter technique, *Monthly Weather Review*, *126*(3), 796–811, doi:10.1175/1520-0493(1998)126<0796:DAUAEK>2.0.CO;2, 1998.

Houtekamer, P. L., and H. L. Mitchell, A sequential ensemble Kalman Filter for atmospheric data assimilation, *Monthly Weather Review*, *129*(1), 123–137, doi:10.1175/1520-0493(2001)129<0123:ASEKFF>2.0.CO;2, 2001.

Houtekamer, P. L., and F. Zhang, Review of the Ensemble kalman filter for atmospheric data assimilation, *Monthly Weather Review*, *144*(12), 4489–4532, doi:10.1175/MWR-D-15-0440.1, 2016.

Hunt, B. R., E. J. Kostelich, and I. Szunyogh, Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter, *Physica D: Nonlinear Phenomena*, *230*(1), 112 – 126, doi:10.1016/j.physd.2006.11.008, data Assimilation, 2007.

Ippisch, O., H.-J. Vogel, and P. Bastian, Validity limits for the van Genuchten–Mualem model and implications for parameter estimation and numerical simulation, *Advances in Water Resources*, *29*(12), 1780 – 1789, doi:10.1016/j.advwatres.2005.12.011, 2006.

Jaumann, S., Estimation of soil water dynamics based on hydrogeophysical measurements, dissertation, Heidelberg University, 2018.

Jaumann, S., and K. Roth, Effect of unrepresented model errors on estimated soil hydraulic material properties, *Hydrology and Earth System Sciences*, *21*(9), 4301–4322, doi:10.5194/hess-21-4301-2017, 2017.

Kalman, R. E., A new approach to linear filtering and prediction problems, *Journal of basic Engineering*, *82*(1), 35–45, doi:10.1115/1.3662552, 1960.

Kalnay, E., *Atmospheric modeling, data assimilation and predictability*, Cambridge university press, 2003.

Kirchgessner, P., L. Nerger, and A. Bunse-Gerstner, On the choice of an optimal localization radius in ensemble Kalman filter methods, *Monthly Weather Review*, *142*(6), 2165–2175, doi:10.1175/MWR-D-13-00246.1, 2014.

Kitagawa, G., Monte Carlo filter and smoother for non-Gaussian nonlinear state space models, *Journal of Computational and Graphical Statistics*, *5*(1), 1–25, doi:10.1080/10618600.1996.10474692, 1996.

Kong, A., J. S. Liu, and W. H. Wong, Sequential imputations and Bayesian missing data problems, *Journal of the American Statistical Association*, *89*(425), 278–288, doi:10.1080/01621459.1994.10476469, 1994.

Lawson, W. G., and J. A. Hansen, Implications of stochastic and deterministic filters as ensemble-based data assimilation methods in varying regimes of error growth, *Monthly Weather Review*, *132*(8), 1966–1981, doi:10.1175/1520-0493(2004)132<1966:IOSADF>2.0.CO;2, 2004.

Lei, J., and P. Bickel, A moment matching ensemble filter for nonlinear non-gaussian data assimilation, *Monthly Weather Review*, *139*(12), 3964–3973, doi:10.1175/2011MWR3553.1, 2011.

Li, C., and L. Ren, Estimation of unsaturated soil hydraulic parameters using the ensemble Kalman filter, *Vadose Zone Journal*, *10*(4), 1205, doi:10.2136/vzj2010.0159, 2011.

Li, H., E. Kalnay, and T. Miyoshi, Simultaneous estimation of covariance inflation and observation errors within an ensemble Kalman filter, *Quarterly Journal of the Royal Meteorological Society*, *135*(639), 523–533, doi:10.1002/qj.371, 2009.

List, F., and F. A. Radu, A study on iterative methods for solving Richards' equation, *Computational Geosciences*, *20*(2), 341–353, doi:10.1007/s10596-016-9566-3, 2016.

Liu, J. S., and R. Chen, Sequential Monte Carlo methods for dynamic systems, *Journal of the American Statistical Association*, *93*(443), 1032–1044, doi:10.1080/01621459.1998.10473765, 1998.

Liu, Y., and H. V. Gupta, Uncertainty in hydrologic modeling: Toward an integrated data assimilation framework, *Water Resources Research*, *43*(7), doi:10.1029/2006WR005756, w07401, 2007.

Liu, Y., et al., Advancing data assimilation in operational hydrologic forecasting: progresses, challenges, and emerging opportunities, *Hydrology and Earth System Sciences*, *16*(10), 3863–3887, doi:10.5194/hess-16-3863-2012, 2012.

Lorenc, A. C., The potential of the ensemble Kalman filter for NWP—a comparison with 4D-Var, *Quarterly Journal of the Royal Meteorological Society*, *129*(595), 3183–3203, doi:10.1256/qj.02.132, 2003.

Lorenc, A. C., Developments of variational data assimilation, in *ECMWF Seminars*, 2011.

Lorenz, E. N., Deterministic nonperiodic flow, *Journal of the Atmospheric Sciences*, *20*(2), 130–141, doi:10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2, 1963.

Lorenz, E. N., Atmospheric predictability experiments with a large numerical model, *Tellus*, *34*(6), 505–513, doi:10.1111/j.2153-3490.1982.tb01839.x, 1982.

Lorenz, E. N., Predictability: A problem partly solved, in *Proc. Seminar on predictability*, vol. 1, 1996.

Lorenz, E. N., Designing chaotic models, *Journal of the Atmospheric Sciences*, *62*(5), 1574–1587, doi:10.1175/JAS3430.1, 2005.

Lorenz, E. N., and K. A. Emanuel, Optimal sites for supplementary weather observations: Simulation with a small model, *Journal of the Atmospheric Sciences*, *55*(3), 399–414, doi:10.1175/1520-0469(1998)055<0399:OSFSWO>2.0.CO;2, 1998.

Man, J., W. Li, L. Zeng, and L. Wu, Data assimilation for unsaturated flow models with restart adaptive probabilistic collocation based Kalman filter, *Advances in Water Resources*, *92*, 258 – 270, doi:10.1016/j.advwatres.2016.03.016, 2016.

Manoli, G., M. Rossi, D. Pasetto, R. Deiana, S. Ferraris, G. Cassiani, and M. Putti, An iterative particle filter approach for coupled hydro-geophysical inversion of a controlled infiltration experiment, *Journal of Computational Physics*, *283*(Supplement C), 37 – 51, doi:10.1016/j.jcp.2014.11.035, 2015.

Metref, S., E. Cosme, C. Snyder, and P. Brasseur, A non-Gaussian analysis scheme using rank histograms for ensemble data assimilation, *Nonlinear Processes in Geophysics*, *21*(4), 869–885, doi:10.5194/npg-21-869-2014, 2014.

Miyoshi, T., Y. Sato, and T. Kadowaki, Ensemble Kalman filter and 4D-Var intercomparison with the japanese operational global analysis and prediction system, *Monthly Weather Review*, *138*(7), 2846–2866, doi:10.1175/2010MWR3209.1, 2010.

Montzka, C., H. Moradkhani, L. Weihermüller, H.-J. H. Franssen, M. Canty, and H. Vereecken, Hydraulic parameter estimation by remotely-sensed top soil moisture observations with the particle filter, *Journal of Hydrology*, *399*(3), 410 – 421, doi:10.1016/j.jhydrol.2011.01.020, 2011.

Moradkhani, H., K. Hsu, H. Gupta, and S. Sorooshian, Uncertainty assessment of hydrologic model states and parameters: Sequential data assimilation using the particle filter, *Water Resources Research*, *41*(5), doi:10.1029/2004WR003604, 2005a.

Moradkhani, H., S. Sorooshian, H. V. Gupta, and P. R. Houser, Dual state–parameter estimation of hydrological models using ensemble Kalman filter, *Advances in Water Resources*, *28*(2), 135 – 147, doi:10.1016/j.advwatres.2004.09.002, 2005b.

Moradkhani, H., C. M. DeChant, and S. Sorooshian, Evolution of ensemble data assimilation for uncertainty quantification using the particle filter-Markov chain Monte Carlo method, *Water Resources Research*, *48*(12), doi:10.1029/2012WR012144, 2012.

Morzfeld, M., X. Tu, E. Atkins, and A. J. Chorin, A random map implementation of implicit filters, *Journal of Computational Physics*, *231*(4), 2049 – 2066, doi:10.1016/j.jcp.2011.11.022, 2012.

Morzfeld, M., D. Hodyss, and C. Snyder, What the collapse of the ensemble Kalman filter tells us about particle filters, *Tellus A: Dynamic Meteorology and Oceanography*, *69*(1), 1283,809, doi:10.1080/16000870.2017.1283809, 2017.

Mualem, Y., A new model for predicting the hydraulic conductivity of unsaturated porous media, *Water Resources Research*, *12*(3), 513–522, doi:10.1029/WR012i003p00513, 1976.

Nakano, S., G. Ueno, and T. Higuchi, Merging particle filter for sequential data assimilation, *Nonlinear Processes in Geophysics*, *14*(4), 395–408, doi:10.5194/npg-14-395-2007, 2007.

Penny, S. G., and T. Miyoshi, A local particle filter for high-dimensional geophysical systems, *Nonlinear Processes in Geophysics*, *23*(6), 391–405, doi:10.5194/npg-23-391-2016, 2016.

Pham, D. T., Stochastic methods for sequential data assimilation in strongly nonlinear systems, *Monthly Weather Review*, *129*(5), 1194–1207, doi:10.1175/1520-0493(2001)129<1194:SMFSDA>2.0.CO;2, 2001.

Plaza, D. A., R. De Keyser, G. J. M. De Lannoy, L. Giustarini, P. Matgen, and V. R. N. Pauwels, The importance of parameter resampling for soil moisture data assimilation into hydrologic models using the particle filter, *Hydrology and Earth System Sciences*, *16*(2), 375–390, doi:10.5194/hess-16-375-2012, 2012.

Poterjoy, J., A localized particle filter for high-dimensional nonlinear systems, *Monthly Weather Review*, *144*(1), 59–76, doi:10.1175/MWR-D-15-0163.1, 2016.

Poterjoy, J., and J. L. Anderson, Efficient assimilation of simulated observations in a high-dimensional geophysical system using a localized particle filter, *Monthly Weather Review*, *144*(5), 2007–2020, doi:10.1175/MWR-D-15-0322.1, 2016.

Poterjoy, J., R. A. Sobash, and J. L. Anderson, Convective-scale data assimilation for the weather research and forecasting model using the local particle filter, *Monthly Weather Review*, *145*(5), 1897–1918, doi:10.1175/MWR-D-16-0298.1, 2017.

Qin, J., S. Liang, K. Yang, I. Kaihotsu, R. Liu, and T. Koike, Simultaneous estimation of both soil moisture and model parameters using particle filtering method through the assimilation of microwave signal, *Journal of Geophysical Research: Atmospheres*, *114*(D15), doi:10.1029/2008JD011358, d15103, 2009.

Rasmussen, J., H. Madsen, K. H. Jensen, and J. C. Refsgaard, Data assimilation in integrated hydrological modeling using ensemble Kalman filtering: evaluating the effect of ensemble size and localization on filter performance, *Hydrology and Earth System Sciences*, *19*(7), 2999–3013, doi:10.5194/hess-19-2999-2015, 2015.

Reich, S., A nonparametric ensemble transform method for bayesian inference, *SIAM Journal on Scientific Computing*, *35*(4), A2013–A2024, doi:10.1137/130907367, 2013.

Reichle, R. H., Data assimilation methods in the earth sciences, *Advances in Water Resources*, *31*(11), 1411 – 1418, doi:10.1016/j.advwatres.2008.01.001, hydrologic Remote Sensing, 2008.

Riedel, L., D. Häfner, O. Klein, and F. Riexinger, DORiE – The DUNE Operated Richards Equation Solving Environment, `http://ts.iup.uni-heidelberg.de/research/terrestrial-systems/dorie/`, accessed: 2018-04-22, 2018.

Robert, S., and H. R. Künsch, Localizing the ensemble Kalman particle filter, *Tellus A: Dynamic Meteorology and Oceanography*, *69*(1), 1282,016, doi:10.1080/16000870.2017.1282016, 2017.

Roth, K., *Physics of Terrestrial Systems*, vol. 0.2, Institute of Environmental Physics, Heidelberg University, Germany, 2017.

Shi, L., X. Song, J. Tong, Y. Zhu, and Q. Zhang, Impacts of different types of measurements on estimating unsaturated flow parameters, *Journal of Hydrology*, *524*, 549 – 561, doi:10.1016/j.jhydrol.2015.01.078, 2015.

Shi, Y., K. J. Davis, F. Zhang, C. J. Duffy, and X. Yu, Parameter estimation of a physically based land surface hydrologic model using the ensemble Kalman filter: A synthetic experiment, *Water Resources Research*, *50*(1), 706–724, doi:10.1002/2013WR014070, 2014.

Simmons, A. J., and A. Hollingsworth, Some aspects of the improvement in skill of numerical weather prediction, *Quarterly Journal of the Royal Meteorological Society*, *128*(580), 647–677, doi:10.1256/003590002321042135, 2002.

Simmons, A. J., R. Mureau, and T. Petroliagis, Error growth and estimates of predictability from the ECMWF forecasting system, *Quarterly Journal of the Royal Meteorological Society*, *121*(527), 1739–1771, doi:10.1002/qj.49712152711, 1995.

Snyder, C., T. Bengtsson, P. Bickel, and J. Anderson, Obstacles to high-dimensional particle filtering, *Monthly Weather Review*, *136*, 4629, doi:10.1175/2008MWR2529.1, 2008.

Snyder, C., T. Bengtsson, and M. Morzfeld, Performance bounds for particle filters using the optimal proposal, *Monthly Weather Review*, *143*(11), 4750–4761, doi:10.1175/MWR-D-15-0144.1, 2015.

Song, X., L. Shi, M. Ye, J. Yang, and I. M. Navon, Numerical comparison of iterative ensemble Kalman filters for unsaturated flow inverse modeling, *Vadose Zone Journal*, *13*(2), doi:doi:10.2136/vzj2013.05.0083, 2014.

Tippett, M. K., J. L. Anderson, C. H. Bishop, T. M. Hamill, and J. S. Whitaker, Ensemble square root filters, *Monthly Weather Review*, *131*(7), 1485–1490, doi:10.1175/1520-0493(2003)131<1485:ESRF>2.0.CO;2, 2003.

Tödter, J., and B. Ahrens, A second-order exact ensemble square root filter for nonlinear data assimilation, *Monthly Weather Review*, *143*(4), 1347–1367, doi:10.1175/MWR-D-14-00108.1, 2015.

Tödter, J., P. Kirchgessner, L. Nerger, and B. Ahrens, Assessment of a nonlinear ensemble transform filter for high-dimensional data assimilation, *Monthly Weather Review*, *144*(1), 409–427, doi:10.1175/MWR-D-15-0073.1, 2016.

Van Genuchten, M. T., A closed-form equation for predicting the hydraulic conductivity of unsaturated soils, *Soil science society of America journal*, *44*(5), 892–898, doi:10.2136/sssaj1980.03615995004400050002x, 1980.

van Leeuwen, P. J., Comment on "Data assimilation using an ensemble Kalman filter technique", *Monthly Weather Review*, *127*(6), 1374–1377, doi:10.1175/1520-0493(1999)127<1374:CODAUA>2.0.CO;2, 1999.

van Leeuwen, P. J., A variance-minimizing filter for large-scale applications, *Monthly Weather Review*, *131*(9), 2071–2084, doi:10.1175/1520-0493(2003)131<2071:AVFFLA>2.0.CO;2, 2003.

van Leeuwen, P. J., Particle filtering in geophysical systems, *Monthly Weather Review*, *137*, 4089, doi:10.1175/2009MWR2835.1, 2009.

van Leeuwen, P. J., Nonlinear data assimilation in geosciences: Fan extremely efficient particle filter, *Quarterly Journal of the Royal Meteorological Society*, *136*(653), 1991–1999, doi:10.1002/qj.699, 2010.

van Leeuwen, P. J., Efficient nonlinear data-assimilation in geophysical fluid dynamics, *Computers & Fluids*, *46*(1), 52 – 58, doi:10.1016/j.compfluid.2010.11.011, 10th ICFD Conference Series on Numerical Methods for Fluid Dynamics (ICFD 2010), 2011.

van Leeuwen, P. J., and M. Ades, Efficient fully nonlinear data assimilation for geophysical fluid dynamics, *Computers & Geosciences*, *55*, 16 – 27, doi:10.1016/j.cageo.2012.04.015, 2013.

Van Leeuwen, P. J., Y. Cheng, and S. Reich, *Nonlinear data assimilation*, vol. 2, Springer, doi:10.1007/978-3-319-18347-3, 2015.

Vetra-Carvalho, S., P. J. van Leeuwen, L. Nerger, A. Barth, M. U. Altaf, P. Brasseur, P. Kirchgessner, and J.-M. Beckers, State-of-the-art stochastic data assimilation methods for high-dimensional non-gaussian problems, *Tellus A: Dynamic Meteorology and Oceanography*, *70*(1), 1445,364, doi:10.1080/16000870.2018.1445364, 2018.

Vrugt, J. A., C. J. ter Braak, C. G. Diks, and G. Schoups, Hydrologic data assimilation using particle Markov chain Monte Carlo simulation: Theory, concepts and applications, *Advances in Water Resources*, *51*(Supplement C), 457 – 478, doi:10.1016/j.advwatres.2012.04.002, 35th Year Anniversary Issue, 2013.

Weerts, A. H., and G. Y. H. El Serafy, Particle filtering and ensemble Kalman filtering for state updating with hydrological conceptual rainfall-runoff models, *Water Resources Research*, *42*(9), doi:10.1029/2005WR004093, w09403, 2006.

Welch, G., and G. Bishop, An introduction to the Kalman filter, *University of North Carolina at Chapel Hill, Department of Computer Science*, *TR 95-041*, 2006.

Whitaker, J. S., and T. M. Hamill, Evaluating methods to account for system errors in ensemble data assimilation, *Monthly Weather Review*, *140*(9), 3078–3089, doi:10.1175/MWR-D-11-00276.1, 2012.

Wu, C.-C., and S. A. Margulis, Feasibility of real-time soil state and flux characterization for wastewater reuse using an embedded sensor network data assimilation approach, *Journal of Hydrology*, *399*(3), 313 – 325, doi:10.1016/j.jhydrol.2011.01.011, 2011.

Würsch, M., and G. C. Craig, A simple dynamical model of cumulus convection for data assimilation research, *Meteorologische Zeitschrift*, (5), 483–490, doi:10.1127/0941-2948/2014/0492, 2014.

Xiong, X., I. M. Navon, and B. Uzunoglu, A note on the particle filter with posterior Gaussian resampling, *Tellus A*, *58*(4), 456–460, doi:10.1111/j.1600-0870.2006.00185.x, 2006.

Yan, H., C. M. DeChant, and H. Moradkhani, Improving soil moisture profile prediction with the particle filter-Markov chain Monte Carlo method, *IEEE Transactions on Geoscience and Remote Sensing*, *53*(11), 6134–6147, doi:10.1109/TGRS.2015.2432067, 2015.

Zhang, D., H. Madsen, M. E. Ridler, J. C. Refsgaard, and K. H. Jensen, Impact of uncertainty description on assimilating hydraulic head in the MIKE SHE distributed hydrological model, *Advances in Water Resources*, *86*, 400 – 413, doi:10.1016/j.advwatres.2015.07.018, 2015.

Zhang, H., H.-J. Hendricks Franssen, X. Han, J. A. Vrugt, and H. Vereecken, State and parameter estimation of two land surface models using the ensemble Kalman filter and the particle filter, *Hydrology and Earth System Sciences*, *21*(9), 4927–4958, doi:10.5194/hess-21-4927-2017, 2017.

Bibliography of own publications

Parts of the results presented in this dissertation have been submitted and are currently under review. I wrote the following publication as first author and identical wording occurs in this thesis:

Berg, D. , H. H. Bauser, and K. Roth, Covariance resampling for particle filter — state and parameter estimation for soil hydrology, *Hydrology and Earth System Sciences Discussions*, *2018*, 1–20, doi:10.5194/hess-2018-121, in review, 2018.

Used in this thesis and cited

Bauser, H. H., S. Jaumann, **D. Berg**, and K. Roth, EnKF with closed-eye period – towards a consistent aggregation of information in soil hydrology, *Hydrology and Earth System Sciences*, *20* (12), 4999–5014, doi:10.5194/hess-20-4999-2016, 2016.

Bauser, H. H., **D. Berg**, O. Klein, and K. Roth, Inflation method for ensemble Kalman filter in soil hydrology,*Hydrology and Earth System Sciences Discussions*, *2018*, 1–18, doi:10.5194/hess-2018-74, in review, 2018.

Publications neither used nor cited in this thesis

Peer-reviewed

Bartelmann, M., F. Fabis, **D. Berg**, E. Kozlikin, R. Lilow, and C. Viermann, A microscopic, non-equilibrium, statistical field theory for cosmic structure formation, *New Journal of Physics*, *18* (4), 043,020, doi:10.1088/1367- 2630/18/4/043020, 2016.

Non peer-reviewed

Bartelmann, M., F. Fabis, **D. Berg**, E. Kozlikin, R. Lilow, and C. Viermann, Nonequilibrium statistical field theory for classical particles: Non-linear structure evolution with first-order interaction, *ArXiv e-prints*, arXiv:1411.1502, 2014a.

Bartelmann, M., F. Fabis, **D. Berg**, E. Kozlikin, R. Lilow, and C. Viermann, Nonequilibrium statistical field theory for classical particles: Linear and mildly nonlinear evolution of cosmological density power spectra, *ArXiv e-prints*, arXiv:1411.1153, 2014b.

Fabis, F., **D. Berg**, E. Kozlikin, and M. Bartelmann, Non-equilibrium statistical field theory for classical particles: Initially correlated grand canonical ensembles, *ArXiv e-prints*, arXiv:1412.2572, 2014.

Acknowledgements

First of all I would like to thank my advisor Prof. Dr. Kurt Roth for the opportunity to do this thesis, productive discussions, and his guidance and support during the last years. I also want to thank Prof. Dr. Peter Bastian for co-supervising my dissertation and Prof. Dr. Björn Malte Schäfer, who immediately agreed to be available as second referee.

Another thank goes to all group members of the TS-CCEES work group for funny lunch times and inspiring conversations. Additionally, I want to acknowledge the support and discussions especially in the final phase of my thesis: Hannes Bauser, Sabrina Ebenhoch, Lisa Hantschel, Angelika Klein, Philipp Kreyenberg, Karsten Mosny, Peter Stromberger and Juliane Weber.

A special thanks goes to Hannes Bauser for the fruitful discussions and the good collaboration. I also want to thank my office colleagues for the great atmosphere and the evening gaming sessions, which led to various interdisciplinary discussions. I really enjoyed this office as a working place.

Furthermore, I am grateful for my girlfriend Juliane Weber. She supported me during the whole thesis and also encouraged me in the final phase. Last but not least I thank my family and friends for the support during my studies and life.