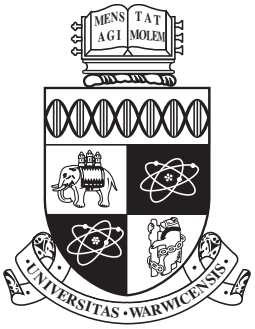

Bayes Linear Methods for Inverse Problems

by Jonas Latz

Dissertation MSc Scientific Computing

supervised by Prof. Andrew M. Stuart

second marker Dr. Claudia Schillings

August 2016

University of Warwick • Centre for Scientific Computing

Contents

Abstract

The objective of this report is the development and the examination of Bayes Linear methods in the Bayesian inverse problems framework. Firstly, this framework, along with that of the classical inverse problem, is motivated and briefly introduced. Thereafter, a rigorous presentation of the theory of both conditional expectations and Bayes Linear approximations for Hilbert space valued random variables is given, both of which focus on the best estimator property. These estimators will then be used as point estimators of the solution of a Bayesian inverse problem. However, due to the complexity of the examined model, the (mostly simulation based) derivation of the conditional expectation is relatively computationally expensive and the results of the less expensive Bayes Linear estimator are typically rather poor. Thus, the Ensemble Kalman Filter, which is an efficient method to approximate conditional expectations, is motivated as a sequential Bayes Linear strategy and presented in both the data assimilation and the Sequential Monte Carlo setting. With the aim of improving the efficiency of the Ensemble Kalman Filter by an adaptive step length, a gradient-free Wolfe-type condition is constructed based on the Bayes Linear estimator and discussed.

All presented techniques are numerically examined in several experiments. In particular, estimation results of analytical and simulation based Bayes Linear estimators are compared with the Ensemble Kalman Filter and with a Monte Carlo simulation (based on autonormalised importance sampling) of the conditional expectation. Furthermore, the Ensemble Kalman Filter both with and without the above-mentioned Bayes Linear line search are tested against each other. Finally, the Ensemble Kalman Filter is considered in a situation, where the posterior distribution is multi modal.

Es ist noch kein Meister vom Himmel gefallen.

No one is born a master.

Grandfather Frank H. Grewenig
(Master painter, 1949-2011)

Acknowledgements

Foremost, I would like to thank my supervisors Andrew M. Stuart and Claudia Schillings. They provided me with help, guidance and plenty of insightful contributions to my research. Their support was invaluable in allowing me to complete my dissertation and my successful introduction to the world of Bayesian inverse problems. Moreover, I would like to thank Claudia Schillings for providing me with her implementation of the Ensemble Kalman Filter for the groundwater flow inverse problem.

Additionally, I would like to thank David Albers and Matthew Levine from the Department of Biomedical Informatics of Columbia University New York for introducing me to their work, inviting me to collaborate with them and granting me access to their model implementations, MCMC simulation results and real world data.

I would also like to thank my fellow students Aurash Karimi and Dominic Richards for proofreading this dissertation.

Finally, I would like to thank my parents, grandparents and Konrad Adenauer Foundation for funding my studies at University of Warwick.

Jonas Latz

List of Tables

List of Figures

List of Abbreviations

a.e.	almost everywhere
a.s.	almost surely
BIP	Bayesian inverse problem
BLE	Bayes Linear estimator
$\widehat{\text{BLE}}$	empirical Bayes Linear estimator
BLUE	best linear unbiased estimator
dc	dominated convergence theorem
EnKF	Ensemble Kalman Filter for inverse problems
EnKFLS	Ensemble Kalman Filter for inverse problems with Bayes Linear line search
iid.	independent and identically distributed
Leb	Lebesgue measure
LM	Levenberg-Marquardt
MAP	Maximum-A-Posterior
MC	Monte Carlo
MCMC	Monte Carlo Markov Chain
pdf	probability density function
SD	standard deviation
SMC	Sequential Monte Carlo
UQ	Uncertainty Quantification

1. Introduction

Inverse problems appear in various scientific disciplines using mathematical models to describe the behaviour of some of their objects of interest. Such models are for instance used to reconstruct an object's behaviour and thus to enable a more accurate examination of the object or to predict its behaviour at some future point in time. However, due to the complexity and the lack of precise prior information of the considered objects, such models, examinations and predictions based on them tend to contain a great amount of uncertainty, which should be reduced to produce more reliable and accurate outputs. In order to do so, the inaccurate model has to be adjusted based on data that has been collected observing the object of interest. This process of incorporating data into the model is referred to as *inverse problem*.

More precisely, consider the following framework: some *model* G describes the behaviour of some object of interest. This model G is influenced by a *(model) parameter* u. The observable part of the object of interest is given by an *observation operator* $\mathcal{O}$ mapping the model output to the space of observations. Let y be some *observation* of this behaviour which is perturbed by some Gaussian noise η, i.e.

$$y := \mathcal{O} \circ G(u^{\text{true}}) + \eta, \tag{1.1}$$

where u^{true} is the *true underlying (model) parameter* describing the behaviour of the object of interest. Furthermore, $\mathcal{O} \circ G =: \mathcal{G}$ is called *forward response operator*.

Fundamental question in inverse problems: Is it possible to identify the true parameter u^{true} of the model G, given (noisy) observations y?

Precise definitions of an inverse problem are given in Section 1.1. Before moving on to that, the elliptic groundwater flow inverse problem is briefly introduced and discussed to illustrate and motivate the concept above. This groundwater flow inverse problem is considered for numerical experiments in Section 5.3 and within Uncertainty Quantification (UQ) thoroughly discussed in [48, Example 6.5] and the considered inverse problem e.g. in [11].

(1.2) Example (Groundwater flow)**.** Let $D \subseteq \mathbb{R}^2$ be an open, connected and bounded domain with Lipschitz boundary. The objective is to estimate the *permeability* of the system, which is given below by $\exp(u)$. The parameter u, referred to as *log-permeability* in this setting, is some Leb(2)-almost everywhere (Leb(2)-a.e.) bounded $\mathbb{R}$-valued function

on D, i.e. $u \in \mathcal{L}^\infty(D, \mathcal{B}D, \text{Leb}(2))$. Based on a predefined parameter u the model G returns the solution $\tilde{p}$ of the elliptic pde:

$$-\nabla \cdot \exp(u)\nabla p = f \text{ (on } D) \qquad (1.3)$$
$$p = 0 \text{ (on } \partial D)$$

The observation operator $\mathcal{O}$ is given by $\tilde{p} \mapsto (\tilde{p}(x_m) : m \in \{1, ..., M\}) \in \mathbb{R}^M$, where $x \in D^M$. Observations are given by $y_m := \tilde{p}(x_m) + \eta_m$ $(m \in \{1, ..., M\})$, where $(\eta_i : i \in \{1, ..., M\}) \sim \text{N}(0, \Gamma)$ is an independent and identically distributed (iid.) family of Gaussian random variables. Intuitively, the observations are noisy values of a set of specific points of the solution $\tilde{p}$ of the pde (1.3). •

There are different ways to approach inverse problems. Presented are the classical and the Bayesian way, since the methods to solve inverse problems discussed here are on the borderline between these two approaches. The classical approach consists of a least squares minimisation problem. This and its regularisation are briefly introduced in Subsection 1.1.1, based on [48, Section 6.1]. Thereafter in Subsection 1.1.2, some foundations of Bayesian statistics are discussed, which then lead to the Bayesian approach's underlying idea and a sketched derivation of the general posterior density, which is considered to be the solution of this class of inverse problems. Last, the connection of the (Bayesian) Maximum-A-Posteriori estimator and (classical) regularised least squares problem is demonstrated. Foundations of Bayesian statistics are given in [44, Section 4.1], [48, Section 2.4] and [8, Chapter 2]. The Bayesian approach to inverse problems is rigorously discussed in [10], [46] and [48, Section 6.2].

Bayes Linear is a collection of approximate methods in Bayesian statistics which are applied here in the Bayesian inverse problem framework. The fundamental idea is based on the optimality property of conditional expectations in expected square error loss. Whilst conditional expectations are given by some general square integrable function, Bayes Linear adjusts one's expectation using an affine function, which is also optimal in expected square error loss. This affine function can then be derived given only expected values and (co-)variances of the considered random variables. This report introduces conditional expectations rigorously in a Hilbert space setting (cf. [32], [12]) in Section 2.1 before constructing the Bayes Linear estimator in the same setting in Section 2.2. The conditional expectations in this setting are similarly defined in [12, Section V.1] and [32]. The Bayes Linear estimator in a Hilbert space setting is given in [14, Section 4.1] while the finite dimensional setting is fundamentally discussed in [19] and the generalised polynomial Bayes estimator in [18]. A motivational and application-related introduction with a thorough

discussion of properties of several Bayes Linear related methods is given in the textbook of Goldstein and Wooff [21]. Relevant properties and diverse examples of the Bayes Linear estimator are considered in Section 2.3. Thereafter, past work on Bayes Linear approaches for inverse problems (that does not include the Ensemble Kalman Filter) is discussed in Section 2.4, especially the methods of [20], [49] and [9] are briefly introduced, and direct Bayes Linear approximations of inverse problem solutions are given, that are either derived analytically or with Monte Carlo simulated parameters.

The Ensemble Kalman Filter for Inverse Problems as proposed by Iglesias, Law and Stuart [28] is derived in Chapter 3. It is constructed by applying the typical data assimilation Ensemble Kalman Filter, which is introduced in Section 3.1, as it is suggested by Evensen [15], respectively [16], to some auxiliary discrete data assimilation problem. The Ensemble Kalman Filter has been applied successfully in the Bayesian inverse problem framework as a robust and often computationally cheap (effective given small ensemble sizes) alternative to MCMC (cf. [43, Chapter 1]), e.g. in [14], [43], and [28]. The connection of Bayes Linear and the Ensemble Kalman Filter for inverse problems is demonstrated, as it is presented by [14, 4.1] and [37, Subsection 3.2.1] and also mentioned by [43], which refer to Bayes Linear as a Gaussian-type approximation. In the remainder of the report, one then concentrates on the version of the Ensemble Kalman Filter for inverse problems that is motivated by Sequential Monte Carlo (SMC) and given in [43]. Evidence is presented in Chapter 3 that this filter can be understood as a sequential Bayes Linear method.

In Chapter 4, a Bayes Linear approximation is used to construct a backtracking line search method for the EnKF. This is motivated by the Levenberg-Marquardt approach proposed by Iglesias in [26] and [27] and also by the typical backtracking line search in some gradient descent method, as presented in [38], [50] and [51]. Both of these methods are briefly introduced in Sections 4.1 and 4.2, before the Bayes Linear line search is proposed in Section 4.3.

Results of numerical experiments are presented in Chapter 5. Considered are a linear inverse problem [17], [34], [36], a cubic inverse problem, which is similarly given in [46, p. 461, Example 2.2], the groundwater flow inverse problem mentioned in Example (1.2) and an inverse problem induced by a glucose-insulin model, introduced in [47] and discussed in [31]. Tested and discussed are the numerical performances of a pure Bayes Linear approach, the Ensemble Kalman Filter and the Ensemble Kalman Filter with Bayes linear (backtracking) line search. They are compared in different settings with each other, and Monte Carlo/Monte Carlo Markov Chain solutions. These settings are specified in the introduction of Chapter 5. The theoretical and numerical results are finally discussed in Chapter 6 and suggestions for future research are given.

1.1. Notation and Problem Setting

The following paragraph lists several notations and the underlying setting. Assume that this presented setting holds throughout the entire report. Notations and definitions that are not given here are given in appendices A and B or in the standard literature cited there. Let $p \in [1, \infty)$. The *space of $\mathcal{L}^p$-functions* mapping from some measure space $(\Omega, \mathcal{A}, \mu)$ to $\mathbb{R}$ is given by $\mathcal{L}^p(\Omega, \mathcal{A}, \mu) := \{f : (\Omega, \mathcal{A}) \to (\mathbb{R}, \mathcal{B}) : \int_\Omega |f|^p \mathrm{d}\mu < \infty\}$, where this symbol is used to refer to both the function space and the space of equivalence classes of μ-a.e. equal functions. Hence, specific Banach or Hilbert-space properties, that are then only true for the equivalence classes, are mentioned specifically. The *norm* of $\mathcal{L}^p(\Omega, \mathcal{A}, \mu)$ is $\|\cdot\|_{\mathcal{L}^p} : \mathcal{L}^p(\Omega, \mathcal{A}, \mu) \to [0, \infty), f \mapsto \left(\int_\Omega |f|^p \mathrm{d}\mu\right)^{1/p}$. Furthermore, $\mathcal{L}^2(\Omega, \mathcal{A}, \mu)$ is a Hilbert space with *inner product* $[\mathcal{L}^2(\Omega, \mathcal{A}, \mu)]^2 \ni (f, g) \mapsto \int_\Omega fg \mathrm{d}\mu := \langle f, g\rangle_{\mathcal{L}^2(\Omega,\mathcal{A},\mu)} := \langle f, g\rangle_{\mathcal{L}^2} \in \mathbb{R}$.

Let $k, M \in \mathbb{N} := \{1, 2, 3, ...\}$. The *$k$-dimensional Lebesgue measure* is denoted by $\mathrm{Leb}(k) : \mathcal{B}^k \to [0, \infty]$. The *parameter space* X is a separable Hilbert space, either given by $X := \mathcal{L}^2(D, \mathcal{B}D, \mathrm{Leb}(k))$, where $D \subseteq \mathbb{R}^k$ is open and bounded, or $X := \mathbb{R}^k$, and its respective Euclidean inner product. Given is also the *data space* $Y := \mathbb{R}^M$, which is the separable Hilbert space with inner product $Y^2 \ni (x, y) \mapsto x^T A y =: \langle x, y\rangle_A =: \langle x, y\rangle_Y$, defined for some symmetric, strictly positive definite matrix $A \in \mathbb{R}^{M\times M}$. The probabilistic framework of this report is based on some probability triple $(\Theta, \mathcal{F}, \mathbb{P})$. Furthermore, $\mathcal{X}$ and $\mathcal{Y}$ are the Hilbert spaces of *square integrable X-valued and Y-valued random variables*.

$$\begin{aligned}
\mathcal{X} &:= \mathcal{L}^2(\Theta, \mathcal{F}, \mathbb{P}; X) := \mathcal{L}^2(\Theta, \mathcal{F}, \mathbb{P}) \otimes X \\
&:= \left\{ f : (\Theta, \mathcal{F}) \to X \text{ (}\mathbb{P}\text{-measurable)} : \int \|f\|_X^2 \mathrm{d}\mathbb{P}(\theta) < \infty \right\}, \\
\mathcal{Y} &:= \mathcal{L}^2(\Theta, \mathcal{F}, \mathbb{P}; Y) := \mathcal{L}^2(\Theta, \mathcal{F}, \mathbb{P}) \otimes Y \\
&:= \left\{ f : (\Theta, \mathcal{F}) \to Y \text{ (}\mathbb{P}\text{-measurable)} : \int \|f\|_A^2 \mathrm{d}\mathbb{P}(\theta) < \infty \right\}.
\end{aligned}$$

The inner products of $\mathcal{X}$ and $\mathcal{Y}$ are given by $\langle f, g\rangle_{\mathcal{X}} := \int \langle f(\theta), g(\theta)\rangle_X \mathrm{d}\mathbb{P}(\theta)$ and $\langle f, g\rangle_{\mathcal{Y}} := \int \langle f(\theta), g(\theta)\rangle_Y \mathrm{d}\mathbb{P}(\theta)$.

Let $X := \mathbb{R}^k$. $x \in \mathcal{X}$ is *(non-degenerate) Gaussian (distributed)*, if some strictly positive definite and symmetric matrix $\Gamma \in \mathbb{R}^{k\times k}$ and a vector $m \in \mathbb{R}^k$ exist, such that $x^{\#}\mathbb{P} =: \mathrm{N}(m, \Gamma) \ll \mathrm{Leb}(k)$ and

$$\frac{\mathrm{dN}(m, \Gamma)}{\mathrm{dLeb}(k)} = \frac{1}{(2\pi)^{k/2}(\det \Gamma)^{1/2}} \exp(-\tfrac{1}{2}\langle \cdot - m, \Gamma^{-1}(\cdot - m)\rangle_X) =: \gamma_{m,\Gamma} \quad (\mathrm{Leb}(k)\text{-a.e.}).$$

In general, if $X \in \{\mathbb{R}^k, \mathcal{L}^2(D, \mathcal{B}D, \mathrm{Leb}(k))\}$, some random variable $x \in \mathcal{X}$ is *Gaussian (distributed)*, if for any $T \in X^*$, some $m \in \mathbb{R}$ and $\sigma^2 \in [0, \infty)$ exist, such that $\langle T|x\rangle \sim \mathrm{N}(m, \sigma^2)$. In this case $\mathrm{N}(m, 0) := \delta_m$ is the *Dirac measure* concentrated on m, i.e. $\langle T|x\rangle = m$ ($\mathbb{P}$-a.s.). x Gaussian is denoted by $x \sim \mathrm{N}(m_0, \mathcal{C}_0)$, where the vector or function $m_0 := \int x \mathrm{d}\mathbb{P}$ is the *mean* of x and the operator $\mathcal{C}_0 : X \to X, \varphi \mapsto \mathcal{C}_0\varphi$, defined by

$$\mathcal{C}_0\varphi(\omega) := \langle c(\cdot, \omega), \varphi\rangle_X \quad (\omega \in D \text{ resp. } \{1, ..., k\}),$$

is the *covariance operator* of x and

$$c(\omega, \omega') := \int (x(\omega, \theta) - m_0(\omega))(x(\omega', \theta) - m_0(\omega'))\mathrm{d}\mathbb{P}(\theta) \quad (\omega, \omega' \in D \text{ resp. } \{1, ..., k\}),$$

is the *covariance function* of x.

The forward response operator is typically a continuous function $\mathcal{G} : X \to Y$, the unknown parameter u is an element of X and the data $y \in Y$ is given by an evaluation of $\mathcal{G}(u) + \eta$, where $\eta \sim R = \mathrm{N}(0, \Gamma)$ is a centred Gaussian random variable with some symmetric and strictly positive definite covariance matrix $\Gamma \in \mathbb{R}^{M\times M}$ with probability density function $\rho : \mathbb{R}^m \to [0, \infty)$. The strict positive definiteness of Γ is generally not necessary in the inverse problem setting. However, this report exclusively considers the noisy setting with strictly positive definite covariance.

All notations are well-defined (cf. [48, Sections 2.7 and 3.5] and [6, Sections 1.2 and 2.2]).

1.1.1. Classical Approach

Due to stochastic perturbation and potential further complexity, it is typically neither possible nor reasonable to 'solve' the equation $y = \mathcal{G}(u)$ - this problem is ill-posed (or: not well-posed) in the sense of Hadamard [22, p. 50]. This ill-posedness can for instance be induced by the equation being underdetermined ($\dim X > \dim Y$) or the data being out of the range of the forward response operator ($y \notin \mathrm{im}(\mathcal{G}) \subseteq Y$). Thus, the classical approach to inverse problems is the minimisation of the *least squares problem*

$$\min_{u\in X} \|\mathcal{G}(u) - y\|_Y^2 \tag{1.4}$$

and its minimiser $\widehat{u} := \mathrm{argmin}_{u\in X}\|\mathcal{G}(u) - y\|_Y^2$ is then the *solution of the inverse problem*. However, the functional $\|\mathcal{G}(\cdot) - y\|_Y^2$ is not necessarily convex. Given that case, it might have several local minima, which implies that standard methods of numerical optimisation

are possibly not able to find a global minimum or to converge.

Regularisation is a typical way to make the least squares problem (1.4) more approachable for numerical methods. Several regularisation methods are proposed in [48, Section 6.1]. This report only considers the *quadratic regularised* least squares problem. This is given by

$$\min_{u \in H} \|\mathcal{D}^{-\frac{1}{2}}(\mathcal{G}(u) - y)\|_Y^2 + \|u - m\|_H^2, \tag{1.5}$$

where H is some Banach space embedded in X, $\mathcal{D} \in \mathbb{R}^{M \times M}$ is symmetric and strictly positive definite matrix, and $m \in H$. There are different ways to choose $\mathcal{D}, H$ and m. One possible choice is given below in Corollary (1.12), which is derived within the Bayesian approach to inverse problems.

1.1.2. Bayesian Approach

The Bayesian approach to inverse problems is, as its name implies, based on Bayesian statistics. Therefore, some foundations of Bayesian statistics in finite dimensions are summarised first.

The objective is to gain information about some parameter $\theta \in X := \mathbb{R}^k$ given data $\widehat{z} \in Y$, which in this case, can be interpreted as a realisation of a random variable $z \in \mathcal{Y}$. The *likelihood* of $\widehat{z}$ given this parameter θ is given by some function $L : Y \times X \to [0, \infty)$, that is proportional to the probability density function of z that is parametrised by θ. Let $\theta \sim \mu_0$ be itself a random variable distributed according to a *prior distribution* μ_0 with *probability density function (pdf)* f_0. This prior distribution reflects previous knowledge about θ. *Bayes' rule* is then applied to adjust the prior knowledge by the information that is given by the data $\widehat{z}$. This adjustment of the prior knowledge is then given by the density of the *posterior distribution* f^{post}, which in this case, is the density of the conditional distribution of θ given $z = \widehat{z}$ and can be derived using Bayes' rule.

$$\text{(Bayes' rule)} \qquad f^{\text{post}}(\theta|\widehat{z}) \propto L(\widehat{z}|\theta) \cdot f_0(\theta)$$

Conditional distributions are rigorously introduced in Section 2.1. The relation of a conditional distribution to prior and likelihood as given in Bayes' rule and an intuitive understanding of conditioning in probability theory is sufficient to proceed within the current subsection.

Even if Bayesian Statistics only considers probability distributions to describe parameters,

it is sometimes convenient to also have point estimates for parameters. Similarly to a Maximum Likelihood estimator in a classical inference setting, the *Maximum-A-Posteriori (MAP)* estimator is given by

$$\theta^{\mathrm{MAP}} :\in \operatorname{argmax}_{\theta \in X} f^{\mathrm{post}}(\theta|\widehat{z}). \tag{1.6}$$

A similar further point estimate, that is used within Bayesian statistics, often called *Bayes Estimator*, is the conditional expectation of θ given that $z = \widehat{z}$. These are discussed in Section 2.1.

The regularised classical approach appears to be a reasonable way to solve an inverse problem. However, the complexity of the forward response operator, multiple minima of the target functional and a potential dependence of the parameters on each other, implies that gaining more information than just a point estimate of the solution as given in the classical approach is reasonable. The *Bayesian approach to inverse problems* or a *Bayesian inverse problem (BIP)* considers an inverse problem in the presented framework of Bayesian Statistics. It is composed in the following way: let u, which is still the unknown parameter, be a random variable $u \sim \mu_0$, where μ_0 is the prior distribution of u and u is independent of the noise η. The likelihood of y, that is the likelihood that the random variable $\mathcal{G}(\widehat{u})+\eta$ equals y (or η equals $\mathcal{G}(\widehat{u}) - y$), given some parameter $\widehat{u} \in X$, is given by

$$L(y|\widehat{u}) := \exp(-\tfrac{1}{2}\|\mathcal{G}(\widehat{u}) - y\|^2_{\Gamma^{-1}}), \tag{1.7}$$

since $\eta \sim R = \mathrm{N}(0, \Gamma)$. The solution of the Bayesian approach to inverse problems is the posterior distribution of u, which is again the conditional distribution of u given that $\mathcal{G}(u) + \eta = y$, and denoted by $\mu^y := \mathbb{P}(u \in \cdot|\mathcal{G}(u) + \eta = y)$. However, the parameter space X is a possibly infinite dimensional function space, i.e. the probability distribution μ_0 of u is not given by a probability density function with respect to the Lebesgue measure, and this setting is not covered by the given version of (Bayes' rule). The next theorem deals with this case.

(1.8) Theorem (Bayes' rule in the Inverse Problem setting)**.** Let $\Phi : X \times Y \to \mathbb{R}$, $(u, y) \mapsto \Phi(u; y) = \frac{1}{2}\|\mathcal{G}(u)-y\|^2_{\Gamma^{-1}}$ be $(\mu_0 \otimes R)$-measurable, the function $-\Phi$ is called *log-likelihood* and Φ is called *potential*. Furthermore, assume that

$$Z(y) := \int \exp\left(-\Phi(u; y)\right) \mathrm{d}\mu_0(u) > 0 \ (y \in Y, R\text{−a.s.}). \tag{1.9}$$

Then, the conditional probability distribution of u given $\mathcal{G}(u) + \eta = y$ exists and $\mathbb{P}(u \in$

$\cdot|\mathcal{G}(u)+\eta=y) =: \mu^y \ll \mu_0$. The Radon-Nikodym derivative of μ^y with respect to μ_0 is

$$\frac{\mathrm{d}\mu^y}{\mathrm{d}\mu_0}(t) = \frac{1}{Z(y)}\exp\left(-\Phi(t;y)\right) \quad (t\in X, \mu_0\text{−a.s.}, y\in Y, R\text{−a.s.}). \tag{1.10}$$

Therefore,

$$\mu^y = \frac{1}{Z(y)}\exp\left(-\Phi(\cdot;y)\right)\mu_0. \tag{1.11}$$ •

Proof. The proof requires measure theoretic considerations that are not given in this report. The background material and a rigorous derivation are presented in [10] and [48, Section 6.2]. □

The last corollary in this section proposes a connection between the regularised least squares approach and the Bayesian approach. It shows how prior knowledge can be used in combination with a point estimator.

(1.12) Corollary. Let $\mu_0 = \mathrm{N}(m_0, \mathcal{C}_0)$. The MAP estimator in the Bayesian approach to inverse problems is given by

$$u^{\mathrm{MAP}} \in \operatorname{argmax}_{u\in H} -\|\mathcal{G}(u)-y\|^2_{\Gamma^{-1}} - \|u-m_0\|^2_H, \tag{1.13}$$

where H is the Cameron-Martin space of μ_0 on X and $\|\cdot\|_H := \|\mathcal{C}_0^{-\frac{1}{2}}\cdot\|_X$. The map estimator is equivalent to the solution of some regularised least squares problem as given in Equation (1.5). •

Proof. The equivalence of regularised least squares problem and the maximisation problem is obvious.

Let $X := \mathbb{R}^k$. The MAP estimator is the solution of the maximisation problem

$$\max_{u\in X} \log L(y|u) + \log f_0(u),$$

where f_0 is the probability density function of $\mu_0 = \mathrm{N}(m_0, \mathcal{C}_0)$. The target function is given by

$$\begin{aligned}\log L(y|u) + \log f_0(u) &= \log \rho(\mathcal{G}(u)-y) + \log f_0(u)\\ &= -\frac{1}{2}\|\mathcal{G}(u)-y\|^2_{\Gamma^{-1}} - \frac{1}{2}\|\mathcal{C}_0^{-\frac{1}{2}}(u-m_0)\|^2_X + c,\end{aligned}$$

where $c\in\mathbb{R}$ does not depend on u. The summand c and the prefactors are negligible.

The case where the given parameter space is infinite dimensional is handled in [10] and [28]. □

To conclude the introduction an inverse problem including the connection of the classical and Bayesian approach to solve it is presented graphically in Figure 1.

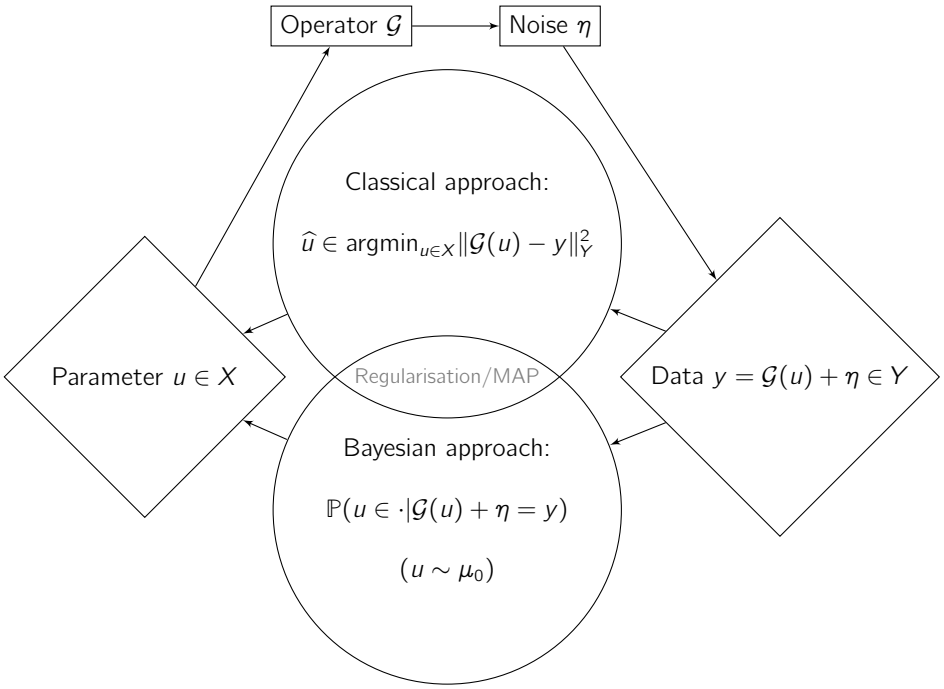

Figure 1 – Classical and Bayesian approach to inverse problems

2. Bayes Linear Statistics

Bayes Linear Statistics is a collection of methods to linearly approximate the conditional expectation of some square integrable random variable given another square integrable random variable. Section 2.1 presents the theory of conditional expectations for $\mathbb{R}^k$- and $\mathcal{L}^2$-valued random variables and especially their optimal estimator property. The subsection is concluded with examples illustrating ways additional information is used by conditional expectations to adjust the expectation one has concerning some random variable. Section 2.2 describes how the Bayes Linear estimator is constructed for $\mathbb{R}^k$- and $\mathcal{L}^2$-valued random variables. Thereafter, further statistical properties of the Bayes Linear estimator and examples are considered in Section 2.3 before proceeding to direct Bayes Linear methods for inverse problems in Section 2.4. Discussed is also the approach of Goldstein and his collaborators in [9], [20], and [49].

The content of this chapter is based on the textbooks of Diestel and Uhl [12, Chapters III

and V], Ash and Doléans-Dade [2, Chapters 4 and 5], Billingsley [3, Chapter 16], Bogachev [6, Chapter 2], Bühlmann and Gisler [8, Chapters 2 and 3], and Goldstein and Wooff [21, Chapters 2 and 6] and the papers of Krug [32] and Bochner [5]. Basic principles in measure theory, (Bochner)-integration and probability theory and standard definitions left out here can be looked up in [2], [3], [12] or other standard textbooks on (vector) measure and probability theory. Some relevant definitions and results are also introduced in appendices A and B.

(2.1) Assumptions. Let $\mathcal{C} \subseteq \mathcal{F}$ be a σ-algebra containing Θ and $(\Theta', \mathcal{F}')$ some further measurable space, fulfilling the condition $\{\theta'\} \in \mathcal{F}'$ for any $\theta' \in \Theta'$. Let $x \in \mathcal{X}, y \in \mathcal{Y}$ be square integrable X- resp. Y-valued random variables and $z : (\Theta, \mathcal{F}) \to (\Theta', \mathcal{F}')$ be some other random variable. •

2.1. Conditional Expectation

Let Assumptions (2.1) hold. The overall objective of this chapter is to find an optimal way to approximate the stochastic behaviour of the random variable x, i.e. to use given *additional information* to adjust one's expectation concerning x.

Without any *additional information*, the expected value of x can be seen as the best possible approximation of the unknown x. *Best* refers here to expected square error loss. This property is briefly discussed in Remark (2.12).

(2.2) Definition (Expected Value)**.** Let Assumptions (2.1) hold. The *expected value* of x is given by

$$\mathbb{E}[x] := \int x(\theta) \mathrm{d}\mathbb{P}(\theta) \overset{(B.10)}{=} \int \xi \mathrm{d}(x^{\#}\mathbb{P})(\xi). \qquad \bullet$$

The expected value is a constant approximation of x with respect to the probability space. Better, non-constant approximations of x can be obtained using additional information, that is either:

1. Some other random variable z depending on x
2. A σ-algebra $\mathcal{C}$ of events one has prior information about, i.e. it is known whether C contains $\theta \in \Theta$ for each $C \in \mathcal{C}$. θ is a realisation of the probability space $(\Theta, \mathcal{F}, \mathbb{P})$.

Consider the case where the given additional information is a second random variable z and its observed realisation.

(2.3) Theorem and Definition (Conditional Expectation given $z = \cdot$)**.** Let Assumptions (2.1) hold. A $z^{\#}\mathbb{P}$-a.s. unique $z^{\#}\mathbb{P}$-measurable function $g : (\Theta', \mathcal{F}') \to X$ exists, such that the *Radon-Nikodym-equation*:

$$\int_{\{z \in F'\}} x \mathrm{d}\mathbb{P} = \int_{F'} g(\theta) \mathrm{d}(z^{\#}\mathbb{P})(\theta) \tag{2.4}$$

holds in the Hilbert space X for each $F' \in \mathcal{F}'$. The function $g =: \mathbb{E}[x|z = \cdot]$ is the *(factorised) conditional expectation of x given $z = \cdot$*. Additionally, let

$$\mathbb{E}[x|z] := \mathbb{E}[x|z = \cdot] \circ z := g(z) : (\Theta, \mathcal{F}) \to X$$

be a random variable called the *conditional expectation of x given z*. •

Here $\mathbb{E}[x|z = \tilde{z}]$ is the expected value of x knowing that the event $\{z = \tilde{z}\}$ occurrs (for some $\tilde{z} \in \Theta'$), which is (in this general setting) defined even if $\mathbb{P}(z = \tilde{z}) = 0$. Alternatively $\mathbb{E}[x|z]$ is by definition a random variable which can be considered as an approximation of x based on z. Another interpretation of $\mathbb{E}[x|z]$ is given in Corollary (2.8).

Before proving the theorem above, the conditional expectation given some σ-algebra $\mathcal{C}$ is considered, along with some further result.

(2.5) Theorem and Definition (Conditional Expectation given a sub-σ-algebra)**.** Let Assumptions (2.1) hold. A $\mathbb{P}$-measurable and $\mathbb{P}$-a.s. unique function $h : (\Theta, \mathcal{C}) \to X$ exists such that the *Radon-Nikodym-equation*

$$\int_C h \mathrm{d}\mathbb{P} = \int_C x \mathrm{d}\mathbb{P} \tag{2.6}$$

holds in the Hilbert space X for each $C \in \mathcal{C}$. $h =: \mathbb{E}[x|\mathcal{C}]$ is called *conditional expectation of x given $\mathcal{C}$*. •

Proof. The proof only considers the case $X = \mathbb{R}$. The general case can be proven similarly, but requires to introduce vector measures and the Radon-Nikodym theory for vector measures. The proof is similar in this more general case since X is a Hilbert space and the theorem of Radon-Nikodym holds for Hilbert spaces. [12, p. 122-123, Lemma 3, Theorem 4 and p. 100, Corollary 4 (von Neumann)]

Without loss of generality assume that $x \geq 0$ $\mathbb{P}$-a.s.. (Otherwise, consider two cases $x_+ := \max\{x, 0\}$, $x_- := \max\{-x, 0\}$). Hence, the map

$$\widehat{\mu} : \mathcal{C} \to [0, \infty], C \mapsto \int_C x \mathrm{d}\mathbb{P}$$

is a measure with $\mathbb{P}$-density x. By the Theorem of Radon-Nikodym (B.8), $\widehat{\mu} \ll \mathbb{P}|_{\mathcal{C}}$ and also by Radon-Nikodym, a $\mathbb{P}$-a.s. unique map $h := \frac{\mathrm{d}\widehat{\mu}}{\mathrm{d}\mathbb{P}|_{\mathcal{C}}} : (\Theta, \mathcal{C}) \to (X, \mathcal{B}X)$ exists and fulfils Equation (2.6). □

The following Lemma, which is a result about measurable functions and stated similarly in [32, p. 4, Lemma 3] and [2, p. 216, 5.4.2], is required to prove Theorem (2.3).

(2.7) Lemma. Let Assumptions (2.1) hold. Furthermore, let $x : (\Theta, \sigma(z)) \to X$ be $\mathbb{P}$-measurable. Some $z^{\#}\mathbb{P}$-a.s. unique $z^{\#}\mathbb{P}$-measurable function $j : (\Theta', \mathcal{F}') \to X$ exists, such that: $x = j \circ z$. •

Proof. In the following, the existence of the map j is proven and j is constructed. Then its measurability is shown and, last, its $z^{\#}\mathbb{P}$-a.s. uniqueness is discussed. The functions x, j, and z and the construction of j are visualised in Figure 2.

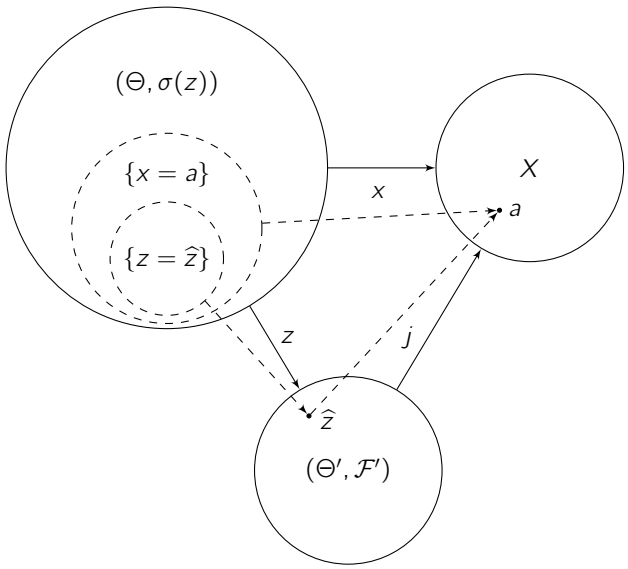

Figure 2 – Maps x, z, j and the construction of j as given in Lemma (2.7)

By assumption j has to be a function mapping $\Theta' \ni \widehat{z} \mapsto x(\theta) \in X$ for all $\theta \in \{z = \widehat{z}\}$, which implies that x has to be constant on any level set $\{z = \cdot\}$ of z. Since $\mathcal{F}'$ contains the singleton $\{\theta'\}$ for any $\theta' \in \Theta'$, $\sigma(z)$ contains any level set $\{z = \cdot\}$ of z. $\sigma(z)$ also contains any level set $\{x = \cdot\}$ of x: let $a \in X$. Either x is simple, then it is given by a finite partition $\mathcal{P} \subseteq \mathcal{F}$ of Θ and a vector $(\xi_P : P \in \mathcal{P}) \in X^{\mathcal{P}}$: $x = \sum_{P \in \mathcal{P}} \xi_P \mathbb{1}_P$, then $\{x = a\} = \bigcup_{P \in \mathcal{P}: \xi_P = a} P \in \mathcal{F}$ or x is the limit of a sequence of simple functions; in which case an antitone sequence $(F_n : n \in \mathbb{N}) \in \mathcal{F}^{\mathbb{N}}$ can be chosen such that $\mathcal{F} \ni \bigcap_{n=1}^{\infty} F_n = \{x = a\}$. The sequence $(F_n : n \in \mathbb{N})$ could for instance be given by $F_n = \bigcup_{P \in \mathcal{P}_n, P \cap \{x = a\} \neq \emptyset} P$, where $\mathcal{P}_n$ is the partition of Θ, that is used to define the n-th simple function in the sequence x is based on. Therefore, every level set of x is in $\sigma(z)$, which implies that it is either empty

or the non-empty preimage $z^{-1}[F']$ of z of some set $F' \in \mathcal{F}'$. This non-empty preimage is then either a level set of z or a superset of some level set of z, since $\sigma(z)$ contains no non-empty strict subsets of level sets of z, due to the construction of initial σ-algebras. Clearly, x is constant on its level-sets. Therefore x is indeed constant on any level set of z.

$x : (\Theta, \sigma(z)) \to X$ is a $\mathbb{P}$-measurable function, i.e. x is either simple or the limit of a series of simple functions. Let $F \in \sigma(z)$, i.e. $F = z^{-1}[F']$ for some $F' \in \mathcal{F}'$ and $x = \mathbb{1}_F$. Choose $j := \mathbb{1}_{F'}$, then $j \circ z = \mathbb{1}_{F'} \circ z = \mathbb{1}_{F'}(z) = \mathbb{1}_{z^{-1}[F']} = \mathbb{1}_F = x$. Let $\mathcal{P} \subseteq \sigma(z)$ and $\xi \in X^{\mathcal{P}}$ be given as above such that $x = \sum_{P \in \mathcal{P}} \xi_P \mathbb{1}_P$. Let P' be chosen for each P to be the maximal set that fulfils $P = z^{-1}[P']$. Then $\{P' : P \in \mathcal{P}\}$ is a partition of Θ'. Furthermore, $j := \sum_{P \in \mathcal{P}} \xi_P \mathbb{1}_{P'}$ implies, that $j \circ z := \sum_{P \in \mathcal{P}} \xi_P \mathbb{1}_{P'}(z) = \sum_{P \in \mathcal{P}} \xi_P \mathbb{1}_{z^{-1}[P']} = x$. Let now $(x_n : \Theta \to X, n \in \mathbb{N})$ be a sequence of simple functions with $x_n \to x$ as $n \to \infty$. Let $(j_n : \Theta' \to X, n \in \mathbb{N})$ the sequence of simple functions as constructed before. Then $j_n \circ z = x_n$ $(n \in \mathbb{N})$. Set then $j := \lim_{n \to \infty} j_n$, wherever $(j_n : n \in \mathbb{N})$ converges and 0 otherwise. I.e. j is $z^{\#}\mathbb{P}$-measurable, whenever $x : (\Theta, \sigma(z)) \to X$ is $\mathbb{P}$-measurable, by the definition of measurability.

The function j, as it is given in the first paragraph, is uniquely defined on $\mathrm{im}(z)$. Furthermore, $\{z \in \Theta' \backslash \mathrm{im}(z)\} = \emptyset$, therefore $z^{\#}\mathbb{P}(\Theta' \backslash \mathrm{im}(z))$ is defined, even if $\Theta' \backslash \mathrm{im}(z) \notin \mathcal{F}'$, and $z^{\#}\mathbb{P}(\Theta' \backslash \mathrm{im}(z)) = \mathbb{P}(z \in \Theta' \backslash \mathrm{im}(z)) = \mathbb{P}(\emptyset) = 0$. $j|_{\Theta' \backslash \mathrm{im}(z)}$ can be chosen arbitrarily. □

The preceding Lemma is given in [32, p. 4, Lemma 3] in a slightly different setting. The function j is uniquely defined instead of only $z^{\#}\mathbb{P}$-a.s. unique, given the further assumption that $\mathrm{im}(z) = \Theta'$. Even if this is a stronger result, the given extra assumption does not seem to be a good match in the Bayesian inverse problem framework. Furthermore, it is stated in the standard notation of $(\mathcal{F}\text{-}\mathcal{F}')$-measurable functions, rather than Bochner's Definition (B.3) (cf. also [5], [12, Section II.1]) notation of $\mathbb{P}$-measurability.

As outlined above, the result of Lemma (2.7) is crucial to prove Theorem (2.3). Thus, one can now proceed to this proof.

Proof of Theorem (2.3). Let $C \in \sigma(z)$ and $C' \in \mathcal{F}'$ be some measurable set, such that $C = z^{-1}[C']$. By Theorem (2.5), a $\mathbb{P}$-measurable, $\mathbb{P}$-a.s. unique function $h : (\Theta, \sigma(z)) \to X$ exists, such that $\int_C h \mathrm{d}\mathbb{P} = \int_C x \mathrm{d}\mathbb{P}$ holds for any $C \in \sigma(z))$. Furthermore, Lemma (2.7) implies, that a $z^{\#}\mathbb{P}$-a.s. uniquely defined $z^{\#}\mathbb{P}$-measurable function $g : (\Theta', \mathcal{F}') \to X$

exists, such that $h = g \circ z$. Then:

$$\int_{z^{-1}[C']} x \mathrm{d}\mathbb{P} = \int_{z^{-1}[C']} h \mathrm{d}\mathbb{P} = \int_{z^{-1}[C']} g \circ z \mathrm{d}\mathbb{P} \overset{(B.10)}{=} \int_{C'} g \mathrm{d}z^{\#}\mathbb{P},$$

which is equivalent to Definition (2.3). □

Even though the proof of Theorem (2.3) applies Theorem (2.5), the two different conditional expectations seem to handle two entirely different structures and can be thought of as two distinct ways to approximate x. The next two corollaries are well-known results (cf. e.g. [2]) that link these two structures and lead to the conclusion, that considering only the $z = \cdot$ case is for *the most part* (cf. Remark (2.10)) sufficient.

(2.8) Corollary. Let Assumptions (2.1) hold. Then

$$\mathbb{E}[x|z] = \mathbb{E}[x|\sigma(z)] \ (\mathbb{P}\text{-a.s.}) \qquad \bullet$$

Proof. Let $A' \in \sigma(z)$ and choose $A \in \mathcal{F}'$ such that $A' := z^{-1}[A]$. One has to prove, that $\mathbb{E}[x|z] := \mathbb{E}[x|z = \cdot] \circ z$ fulfils the Radon-Nikodym-equation (2.6) for A':

$$\int_{A'} \mathbb{E}[x|z] \mathrm{d}\mathbb{P} = \int_{\{z \in A\}} (\mathbb{E}[x|z = \cdot] \circ z)(\theta) \mathrm{d}\mathbb{P}(\theta) \overset{(B.10)}{=} \int_{A} \mathbb{E}[x|z = \cdot] \mathrm{d}(z^{\#}\mathbb{P}) \overset{(2.3)}{=} \int_{A'} x \mathrm{d}\mathbb{P}.$$

□

(2.9) Corollary. Let Assumptions (2.1) hold. A measurable space $(\widehat{\Theta}, \widehat{\mathcal{F}})$ and a random variable $\widehat{y} : (\Theta, \mathcal{F}) \to (\widehat{\Theta}, \widehat{\mathcal{F}})$ can be chosen such that $\mathbb{E}[x|\mathcal{C}] = \mathbb{E}[x|\widehat{y}]$ ($\mathbb{P}$-a.s.). •

Proof. Set for instance: $\widehat{\Theta} := \Theta$ and $\widehat{y} := \mathrm{id}_{\Theta} : (\Theta, \mathcal{F}) \to (\Theta, \mathcal{C})$. Then: $\sigma(\widehat{y}) = \mathcal{C}$. □

(2.10) Remark. 1. As already stated above, the preceding corollaries imply that one can sufficiently only consider the case of $\mathbb{E}[x|z = \cdot]$, since this conditional expectation can be used to construct $\mathbb{E}[x|\mathcal{C}]$. However, the Bayes Linear estimator is defined under further assumptions (where the square integrable y is considered instead of z), where depending on the given probability space $(\Theta, \mathcal{F}, \mathbb{P})$ it might be not possible to prove Corollary (2.9) in the way depicted. This report does not cover any application of $\mathbb{E}[x|\mathcal{C}]$ within the Bayesian Inverse Problem setting. Thus, this loss of generality is negligible here.

2. Let $X := \mathbb{R}^k$. Up to this point, the whole section could have been derived with weaker assumptions on x. In Assumptions (2.1), x is assumed to be square integrable, but merely the existence of $\mathbb{E}[x]$ is sufficient, the expected value could be even infinite, which is discussed in [2, Chapter 5] and [3, Chapter 6]. The theorems stated above are equivalent

in the more general case. However, the given assumptions are essential in the remaining part of the section. •

Besides the Radon-Nikodym-Equation, the conditional expectation of x given $z = \cdot$ can be characterised as a best approximation of x in the expected square error loss sense. This is not only essential to define the Bayes Linear estimator, but also used as definition of the conditional expectation by some authors (cf. [37, 2.3.2 Definition 9] and [35, Definition 4.45, Theorem 4.52], where both definitions are given, but their equivalence is not discussed). The equivalence is stated and proven in the theorem below. The case where X is finite dimensional is also considered in [4, p. 90, 3.2.6].

(2.11) Theorem. Let Assumptions (2.1) hold and let $g \in \mathcal{J} := \mathcal{L}^2(\Theta', \mathcal{F}', z^{\#}\mathbb{P}; X)$ be some measurable function. The following statements are equivalent.

1. $\begin{cases} g = \mathbb{E}[x|z = \cdot] \ (\mathbb{P}\text{-a.s.}), & \text{if } X = \mathbb{R}^k, \\ g = \mathbb{E}[x|z = \cdot] \ (\mathbb{P}\text{-a.s., Leb}(k)\text{-a.e.}), & \text{if } X = \mathcal{L}^2(D, \mathcal{B}D, \text{Leb}(k)). \end{cases}$

2. g minimises $\min\limits_{\widehat{g} \in \mathcal{J}} \mathbb{E}[\|\widehat{g}(z) - x\|_X^2]$ globally. •

Proof. By Lemma (A.6), the minimality of 2 is implied, if

$$0 = \langle g(z) - x, f(z) \rangle_{\mathcal{X}} = \mathbb{E}[\langle g(z) - x, f(z) \rangle_X] \quad (f \in \mathcal{J}).$$

Let $f \in \mathcal{J}$ be chosen arbitrarily. Then,

$$\begin{aligned} \mathbb{E}[\langle g(z) - x, f(z) \rangle_X] &= \int \langle g(z(\theta)), f(z(\theta)) \rangle_X \mathrm{d}\mathbb{P}(\theta) - \int \langle x(\theta), f(z(\theta)) \rangle_X \mathrm{d}\mathbb{P}(\theta) \\ &\overset{(B.10)}{=} \int \langle g, f \rangle_X \mathrm{d}(z^{\#}\mathbb{P}) - \int \langle x, f(z) \rangle_X \mathrm{d}\mathbb{P}. \end{aligned}$$

Let $(\Omega, \mathcal{A}, \mu)$ be a measure space, which is either $(\Omega, \mathcal{A}, \mu) := (D, \mathcal{B}D, \text{Leb}(k))$ or $(\Omega, \mathcal{A}, \mu) := (\{1, ..., k\}, 2^{\{1,...,k\}}, \#)$ and therefore in both cases σ-finite. By definition, $X = \mathcal{L}^2(\Omega, \mathcal{A}, \mu)$ in each case.

$(1 \Rightarrow 2)$: Consider $f : \Theta' \times \Omega \to \mathbb{R}, (\theta', \omega) \mapsto f(\theta', \omega) := \mathbb{1}_{F'}(\theta')$, for some $F' \in \mathcal{F}'$.

$$\begin{aligned}
\mathbb{E}[\langle g(z) - x, \mathbb{1}_{F'}(z)\rangle_X] &= \int \langle g(\theta'), \mathbb{1}_{F'}(\theta')\rangle_X \mathrm{d}(z^{\#}\mathbb{P})(\theta') - \int \langle x(\theta), \mathbb{1}_{F'}(z(\theta))\rangle_X \mathrm{d}\mathbb{P}(\theta) \\
&= \int_{F'} \int g(\theta', \omega)\mathrm{d}\mu(\omega)\mathrm{d}(z^{\#}\mathbb{P})(\theta') - \int_{\{z \in F'\}} \int x(\theta, \omega)\mathrm{d}\mu(\omega)\mathrm{d}\mathbb{P}(\theta) \\
&= \int \left(\int_{F'} g(\theta', \omega)\mathrm{d}(z^{\#}\mathbb{P})(\theta') - \int_{\{z \in F'\}} x(\theta, \omega)\mathrm{d}\mathbb{P}(\theta) \right) \mathrm{d}\mu(\omega) = 0,
\end{aligned}$$

by Fubini's theorem. Consider the definition of $\mathbb{P}$-measurability. The linearity of the integral implies that all simple functions f also fulfil this condition. A $\mathbb{P}$-measurable function, that is not simple, is the limit of a sequence of simple functions. Given such a function f, all simple functions contained in that defining sequence $\left(\sum_{P \in \mathcal{P}_n} \varphi_P^{(n)} \mathbb{1}_P : n \in \mathbb{N}\right)$ fulfil the condition. This sequence can be chosen such that $\sum_{P \in \mathcal{P}_n} |\varphi_P^{(n)}| \mathbb{1}_P \leq |f(\theta')|$ $(n \in \mathbb{N}, \theta' \in \Theta', y^{\#}\mathbb{P}\text{-a.s.})$, pointwise in Ω as given in Lemma (B.4). Then $\sum_{P \in \mathcal{P}_n} \|\varphi_P^{(n)} g(\theta')\|_X \mathbb{1}_P(\theta') \leq \|f(\theta')g(\theta')\|_X$ and $\sum_{P \in \mathcal{P}_n} \|\varphi_P^{(n)} x(\theta)\|_X \mathbb{1}_P(z(\theta)) \leq \|f(z(\theta))x(\theta)\|_X$. Furthermore, $\mathbb{E}[\|f(z)x\|_X] < \infty$ and $\mathbb{E}[\|f(z)g(z)\|_X] < \infty$, since f, x, g are square integrable. Therefore, the dominated convergence theorem holds (dc, cf. [12, p. 45, Theorem 3]) and

$$\begin{aligned}
&\int \left(\int f(\theta', \omega) g(\theta', \omega)\mathrm{d}(z^{\#}\mathbb{P})(\theta') - \int f(z(\theta), \omega) x(\theta, \omega)\mathrm{d}\mathbb{P}(\theta) \right) \mathrm{d}\mu(\omega) \\
&= \int \Bigg(\int \lim_{n\to\infty} \sum_{P \in \mathcal{P}_n} \varphi_P^{(n)}(\omega) \mathbb{1}_P(\theta') g(\theta', \omega)\mathrm{d}(z^{\#}\mathbb{P})(\theta') \\
&- \int \lim_{n\to\infty} \sum_{P \in \mathcal{P}_n} \varphi_P^{(n)}(\omega) \mathbb{1}_P(z(\theta)) x(\theta, \omega)\mathrm{d}\mathbb{P}(\theta) \Bigg) \mathrm{d}\mu(\omega) \\
&\overset{\text{dc}}{=} \int \left(\lim_{n\to\infty} \sum_{P \in \mathcal{P}_n} \varphi_P^{(n)}(\omega) \left(\int_P g(\theta', \omega)\mathrm{d}(z^{\#}\mathbb{P})(\theta') - \int_{\{z \in P\}} x(\theta, \omega)\mathrm{d}\mathbb{P}(\theta) \right) \right) \mathrm{d}\mu(\omega) \\
&= \int \left(\lim_{n\to\infty} \sum_{P \in \mathcal{P}_n} \varphi_P^{(n)}(\omega) \right) \cdot 0 \mathrm{d}\mu(\omega) = 0.
\end{aligned}$$

$(2 \Rightarrow 1)$: Without loss of generality, assume $x, g \geq 0$ ($\mathbb{P}$-a.s. and μ-a.e.). For any $f \in \mathcal{J}$: $\langle g(z) - x, f(z)\rangle_{\mathcal{X}} = 0$. Let $F' \in \mathcal{F}'$, $B \in \mathcal{A}$ and $f(\theta', \omega) := \mathbb{1}_{F'}(\theta')\mathbb{1}_B(\omega)$ $(\theta' \in \Theta', \omega \in \Omega)$.

Thus,

$$\begin{aligned}
0 &= \langle g(z) - x, \mathbb{1}_{F'}(z)\mathbb{1}_B\rangle_{\mathcal{X}} \\
&= \int \langle g(\theta'), \mathbb{1}_{F'}(\theta')\mathbb{1}_B\rangle_X \mathrm{d}(z^{\#}\mathbb{P})(\theta') - \int \langle x(\theta), \mathbb{1}_{F'}(z(\theta))\mathbb{1}_B\rangle_X \mathrm{d}\mathbb{P}(\theta) \\
&= \int_{F'} \langle g(\theta'), \mathbb{1}_B\rangle_X \mathrm{d}(z^{\#}\mathbb{P})(\theta') - \int_{\{z\in F'\}} \langle x(\theta), \mathbb{1}_B\rangle_X \mathrm{d}\mathbb{P}(\theta) \\
&= \int_{F'}\int_B g(\theta', \omega)\mathrm{d}\mu(\omega)\mathrm{d}(z^{\#}\mathbb{P})(\theta') - \int_{\{z\in F'\}}\int_B x(\theta, \omega)\mathrm{d}\mu(\omega)\mathrm{d}\mathbb{P}(\theta) \\
&= \int_B \int_{F'} g(\theta', \omega)\mathrm{d}(z^{\#}\mathbb{P})(\theta')\mathrm{d}\mu(\omega) - \int_B \int_{\{z\in F'\}} x(\theta, \omega)\mathrm{d}\mathbb{P}(\theta)\mathrm{d}\mu(\omega),
\end{aligned}$$

by the theorem of Fubini. Since $B \in \mathcal{A}$ is chosen arbitrarily,

$$\begin{aligned}
\mathcal{A} \ni B \mapsto \int_B \int_{F'} g(\theta', \omega)\mathrm{d}(z^{\#}\mathbb{P})(\theta')\mathrm{d}\mu(\omega) \in [0, \infty), \\
\mathcal{A} \ni B \mapsto \int_B \int_{\{z\in F'\}} x(\theta, \omega)\mathrm{d}\mathbb{P}(\theta)\mathrm{d}\mu(\omega) \in [0, \infty)
\end{aligned}$$

are measures, which are identical on the given domain $\mathcal{A}$ and absolutely μ-continuous. By Radon-Nikodym, the densities of both measures are μ-a.e. identical, so

$$\int_{F'} g(\theta', \omega)\mathrm{d}(z^{\#}\mathbb{P})(\theta') = \int_{\{z\in F'\}} x(\theta, \omega)\mathrm{d}\mathbb{P}(\theta) \quad (\omega \in \Omega, \mu\text{-a.e.})$$

which is equivalent to the Definition of the conditional expectation (2.3).

The theorem of Fubini holds each time it is applied, since all functions are contained in $\mathcal{X} = \mathcal{L}^2(\Theta, \mathcal{F}, \mathbb{P}; X)$ and therefore absolutely integrable. □

(2.12) Remark. The expectation of x is indeed the best constant approximation of x. The case *given no additional information* is equivalent to the case *given some $\mathbb{P}$-a.s. constant random variable z*. In that case, $\mathbb{E}[x|z = \cdot] = \mathbb{E}[x]$ ($z^{\#}\mathbb{P}$-a.s.; [2, p. 224, Theorem 5.5.7 (a')]) minimises the expected square error loss function, by Theorem (2.11). •

Before concluding this section and moving on to the definition of the Bayes Linear estimator, some illustrative examples of conditional expectations are considered. The first example discusses a setting in which x is finite dimensional and z is discretely distributed. Hence, the conditional expectation can be derived using the formula:

$$\mathbb{E}[x|z = \tilde{z}] := \frac{1}{\mathbb{P}(z = \tilde{z})}\mathbb{E}[\mathbb{1}_{\{z=\tilde{z}\}}x] \quad (\text{whenever } \tilde{z} \in \Theta' : \mathbb{P}(z = \tilde{z}) > 0), \tag{2.13}$$

which coincides with the Definition (2.3) of the conditional expectation: [2, p. 210-211,

5.3.5].

(2.14) Example. Let Assumptions (2.1) hold and assume furthermore, that $X := \mathbb{R}, \Theta' := \{0,1\}$, that $x \sim \text{Unif}(0,1)$ is uniformly distributed, and

$$z := \begin{cases} 0, & \text{if } x \in (0, \frac{1}{2}), \\ 1, & \text{otherwise.} \end{cases}$$

i.e. $z \sim (\frac{1}{2}\delta_0 + \frac{1}{2}\delta_1)$, respectively $\mathbb{P}(z=0) = \frac{1}{2} = \mathbb{P}(z=1)$. So x is randomly picked from $(0,1)$ and z contains the information, whether x is in $(0, \frac{1}{2})$ or in $[\frac{1}{2}, 1)$. The expected value of x is $\mathbb{E}[x] = \frac{1}{2}$. One fairly obvious approach to adjust this expectation of x is to derive $\mathbb{E}[x|z=\cdot]$ respectively $\mathbb{E}[x|z]$:

$$\mathbb{E}[x|z=1] \overset{(2.13)}{=} \frac{\mathbb{E}[\mathbb{1}_{\{z=1\}}x]}{\mathbb{P}(z=1)} = \frac{\mathbb{E}[\mathbb{1}_{\{x \geq \frac{1}{2}\}}x]}{\mathbb{P}(x \geq \frac{1}{2})} = \frac{3}{4},$$
$$\mathbb{E}[x|z=0] \overset{(2.13)}{=} \frac{\mathbb{E}[\mathbb{1}_{\{z=0\}}x]}{\mathbb{P}(z=0)} = \frac{\mathbb{E}[\mathbb{1}_{\{x < \frac{1}{2}\}}x]}{\mathbb{P}(x < \frac{1}{2})} = \frac{1}{4}.$$

Thus, the conditional expectation of x given z is a random variable distributed according to $(\frac{1}{2}\delta_{\frac{1}{4}} + \frac{1}{2}\delta_{\frac{3}{4}})$. •

The next example illustrates the way in which the conditional expectation approximates a random variable, in contrast to a Taylor approximation.

(2.15) Example. Let Assumptions (2.1) hold and $n \in \mathbb{N}$, $\Theta := [0,1]$, $\Theta' := \{0, 1, ..., 2^n - 1\}$, $X := [1, \mathrm{e}]$ and $\mathbb{P} = \text{Unif}[0,1]$. Let $x = \exp$ and $z = \sum_{l=0}^{2^n-1} l \mathbb{1}_{(\frac{l}{2^n}, \frac{l+1}{2^n}]}$. Let $\widehat{z} \in \{0, ..., 2^n - 1\}$. Then:

$$\mathbb{E}[x|z=\widehat{z}] \overset{(2.13)}{=} \frac{\mathbb{E}[\mathbb{1}_{\{z=\widehat{z}\}}x]}{\mathbb{P}(z=\widehat{z})} = 2^n \int_{\widehat{z}/2^n}^{(\widehat{z}+1)/2^n} \exp(\theta) \mathrm{d}\theta = 2^n [\exp(\tfrac{\widehat{z}+1}{2^n}) - \exp(\tfrac{\widehat{z}}{2^n})],$$

so $\mathbb{E}[x|z] = \sum_{l=0}^{2^n-1} 2^n [\exp(\frac{l+1}{2^n}) - \exp(\frac{l}{2^n})] \mathbb{1}_{(\frac{l}{2^n}, \frac{l+1}{2^n}]}$ ($\mathbb{P}$-a.s.) is the approximation of x using the information provided by z. Another way to approximate x is a Taylor approximation (cf. e.g. [38, Theorem 2.1]), which uses and preserves the smoothness of $x = \exp$, rather than information provided by other random variables. This example is continued in Example (2.33). Figure 3 (a),(b), which is given within this later example, illustrates the conditional expectations for $n \in \{1,2,3\}$ and Taylor polynomials. •

The *conditional probability distribution given* $z = \cdot$ is a *Markov kernel* and defined by $\mathbb{P}(F|z=\cdot) := \mathbb{E}[\mathbb{1}_F|z=\cdot]$ ($z^{\#}\mathbb{P}$-a.s.), for any $F \in \mathcal{F}$. Given that it is possible to

integrate analytically with respect to this Markov Kernel, the conditional expectation can be derived in the following way:

$$\mathbb{E}[x|z=\cdot] := \int x(\theta)\mathrm{d}\mathbb{P}(\theta|z=\cdot). \tag{2.16}$$

This standard result is stated in [2, Section 5.3]. The example below presents a situation where this method is applied.

(2.17) Example. Let $X := Y := \Theta' := \mathbb{R}$, $\widehat{z} \in \Theta$, $z^{\#}\mathbb{P} = \mathsf{N}(0, \gamma^2)$ and $x^{\#}\mathbb{P}(\cdot|z=\widehat{z}) = \mathsf{N}(\widehat{z}, \sigma^2)$. Then: $\mathbb{E}[x|z=\widehat{z}] = \widehat{z}$. •

The last example of this section illustrates the Linear Inverse problem in the Bayesian setting (cf. [17]). The conditional expectation of the unknown parameter x can be derived analytically in this case.

(2.18) Example. Let $x \sim \mu_0 = \mathsf{N}(m_0, \mathcal{C}_0)$ be an X-valued random variable and data $\tilde{y} \in Y$. The operator $\mathcal{G} : X \to Y$ is linear and continuous. Furthermore, $\eta \in \mathcal{Y}$, $\eta \sim \mathsf{N}(0, \Gamma)$ is Gaussian distributed noise with positive definite covariance matrix $\Gamma \in \mathbb{R}^{k\times k}$. The conditional expectation of x given $\mathcal{G}x + \eta = \tilde{y}$ is a point estimate of the unknown x. In this case,

$$\mathbb{E}[x|\mathcal{G}x+\eta=\tilde{y}] = m_0 + \mathcal{C}_0\mathcal{G}^*(\mathcal{G}\mathcal{C}_0\mathcal{G}^* + \Gamma)^{-1}(\tilde{y} - \mathcal{G}m_0) \quad \text{(a.s.)},$$

where $(\mathcal{G}\mathcal{C}_0\mathcal{G}^* + \Gamma)$ is indeed invertible (cf. Remark (2.42)). This result is derived in [34] and [36], in a more general case. Furthermore, the conditional expectation is affine with respect to $\tilde{y}$. •

2.2. Construction of the Bayes Linear Estimator

The preceding Examples (2.14) and (2.15) contain a simple way to derive the conditional expectation of one random variable given another. This approach however is limited to discretely distributed z and therefore infeasible in the Bayesian inverse problem setting. Example (2.17) is based on the fact, that the conditional distribution is known and that it is possible to integrate analytically with respect to this conditional distribution. In the Bayesian inverse problem setting, the conditional distribution is implicitly given. However, since $\mathcal{G}$ is typically complex, it is generally non-trivial to integrate with respect to this distribution. Integration in that case is often based on Monte Carlo (MC) or Monte Carlo Markov Chains (MCMC). These methods are discussed in [1] and [10, Section 5.2].

Bayes Linear is an approach which uses some affine function to approximate the conditional expectation of some random variable x given that $y = \cdot$. Since it only depends on the expected values of x and y, the Covariance operator of y and the Covariance operator of the joint distribution of x, y, it is generally easier to determine than the precise conditional expectation.

(2.19) Definition (BLE)**.** Let Assumptions (2.1) hold and $\widehat{y} \in \mathbb{R}^M$ be some realisation of y. The *Bayes Linear estimator* (BLE) of x given $y = \widehat{y}$ is given by:

$$x^{\mathrm{BLE}}_{y=\widehat{y}} := \mathbb{E}[x] + \mathrm{Cov}(x,y)\mathrm{Cov}(y)^{\dagger}(\widehat{y} - \mathbb{E}[y]), \tag{2.20}$$

where $\mathrm{Cov}(x,y) := \mathbb{E}[(x - \mathbb{E}[x]) \otimes (y - \mathbb{E}[y])]$, $\mathrm{Cov}(y) := \mathbb{E}[(y - \mathbb{E}[y]) \otimes (y - \mathbb{E}[y])]$ and $A^{\dagger}$ is the Moore-Penrose pseudo inverse of some quadratic matrix $A \in \mathbb{R}^{M \times M}$. The *empirical Bayes Linear estimator* ($\widehat{\mathrm{BLE}}$) of x given $y = \widehat{y}$ given $J \in \mathbb{N}$ samples $(x^{(j)} : j \in \{1, ..., J\}) \in X^J, (y^{(j)} : j \in \{1, ..., J\}) \in Y^J$ is given by

$$x^{\widehat{\mathrm{BLE}},J}_{y=\widehat{y}} := \bar{x}_J + \widehat{\mathrm{Cov}}_J(x,y)(\widehat{\mathrm{Cov}}_J(y))^{\dagger}(\widehat{y} - \bar{y}_J), \tag{2.21}$$

where $\bar{x}_J := \frac{1}{J}\sum_{j=1}^{J} x^{(j)}$, $\bar{y}_J := \frac{1}{J}\sum_{j=1}^{J} y^{(j)}$, $\widehat{\mathrm{Cov}}_J(x,y) := \frac{1}{J}\sum_{j=1}^{J}(x^{(j)} - \bar{x}_J) \otimes (y^{(j)} - \bar{y}_J)$ and $\widehat{\mathrm{Cov}}_J(y) := \frac{1}{J}\sum_{j=1}^{J}(y^{(j)} - \bar{y}_J) \otimes (y^{(j)} - \bar{y}_J)$.

Analogously to Definition (2.3), one denotes $x^{\mathrm{BLE}}_y := x^{\mathrm{BLE}}_{y=\cdot} \circ y$ and $x^{\widehat{\mathrm{BLE}},J}_y := x^{\widehat{\mathrm{BLE}},J}_{y=\cdot} \circ y$. •

(2.22) Remark. 1. In the following, one only considers the case where $\mathrm{Cov}(y)$ (respectively $\widehat{\mathrm{Cov}}_J(y)$) is a regular matrix and the Moore-Penrose pseudo inverse is substituted by the 'normal' inverse. The following theorems also hold in the more general case (cf. Penrose [39]). Furthermore, the equivalent of $\mathrm{Cov}(y)$ (respectively $\widehat{\mathrm{Cov}}_J(y)$) in the inverse problem setting is always regular (cf. Remark (2.42)), due to the assumption that the noise covariance matrix Γ is strictly positive definite.

2. The empirical Bayes Linear estimator is discussed in Subsection 2.4.3. •

The Lemma below proves that $x^{\mathrm{BLE}}_{y=\widehat{y}}$ is indeed an element of X.

(2.23) Lemma. Let Assumptions (2.1) hold and $\widehat{y} \in Y$. Then $x^{\mathrm{BLE}}_{y=\widehat{y}} \in X$. •

Proof. The statement is true, if $X = \mathbb{R}^k$, since x, y are square integrable. Hence, assume that $X = \mathcal{L}^2(D, \mathcal{B}D, \mathrm{Leb}(k))$. The proof of $(f : D \to \mathbb{R}, \omega \mapsto \mathbb{E}[x(\omega)]), (g : D \to \mathbb{R}, \omega \mapsto \mathrm{Cov}(x(\omega), y_m)) \in \mathcal{L}^2(D, \mathcal{B}D, \mathrm{Leb}(k))$ for any $m \in \{1, ..., M\}$ is sufficient to show that the Lemma holds. Both mappings are each a composition of a linear and a measurable

mapping, which implies measurability. The Jensen inequality and Tonelli (x is an $\mathcal{L}^2$-function and $x^2 \geq 0$) implies:

$$\int \mathbb{E}[x(\omega)]^2 \mathrm{d}\omega \leq \int \mathbb{E}[x(\omega)^2] \mathrm{d}\omega = \mathbb{E}\left[\int x(\omega)^2 \mathrm{d}\omega\right],$$

which is bounded, since $x \in \mathcal{X}$. Let $m \in \{1, ..., M\}$, then

$$\begin{aligned}
\int \mathrm{Cov}(x(\omega), y_m)^2 \mathrm{d}\omega &\leq \int \mathrm{Var}(x(\omega))\mathrm{Var}(y_m)\mathrm{d}\omega \\
&= \mathrm{Var}(y_m) \cdot \int \mathrm{Var}(x(\omega))\mathrm{d}\omega \\
&= \mathrm{Var}(y_m) \cdot \int \mathbb{E}[x(\omega)^2] - (\mathbb{E}[x(\omega)])^2 \mathrm{d}\omega \\
&\leq \mathrm{Var}(y_m) \cdot \int \mathbb{E}[x(\omega)^2]\mathrm{d}\omega \\
&= \mathrm{Var}(y_m) \cdot \mathbb{E}\left[\int x(\omega)^2 \mathrm{d}\omega\right],
\end{aligned}$$

which again is bounded, since $x \in \mathcal{X}$. □

The conditional expectation is characterised as an $\mathcal{L}^2$-function that approximates some random variable optimally in terms of expected square error loss in Theorem (2.11). The Bayes Linear estimator is the equivalent optimal affine function.

(2.24) Theorem. Let Assumptions (2.1) hold and $\mathcal{K} := \{\widehat{k} : Y \to X : \exists k_0 \in X, k \in X^M : \widehat{k} = k_0 + \langle k, \cdot \rangle_Y\} \subseteq \mathcal{L}^2(Y, \mathcal{B}Y, y^{\#}\mathbb{P}; X)$. Then

$$x^{\mathrm{BLE}}_{y=\cdot} \in \mathrm{argmin}_{\widehat{k} \in \mathcal{K}} \mathbb{E}[\|\widehat{k}(y) - x\|^2_X] \tag{2.25}$$ •

Proof. 1. Prove the subset relation $\mathcal{K} \subseteq \mathcal{L}^2(Y, \mathcal{B}Y, y^{\#}\mathbb{P}; X)$. If X is finite dimensional, it is sufficient that y is square integrable and that $\mathcal{L}^2(Y, \mathcal{B}Y, y^{\#}\mathbb{P}; X)$ is a vector space. Let X be infinite dimensional and $\widehat{k} \in \mathcal{K}$. Then,

$$\begin{aligned}
\mathbb{E}[\|\widehat{k}(y)\|^2_X] &= \iint (k_0(\omega) + k(\omega)^T y(\theta))^2 \mathrm{d}\omega \mathrm{d}\mathbb{P}(\theta) \\
&\leq 2\int k_0(\omega)^2 \mathrm{d}\omega + 2\iint \|k(\omega)\|^2_Y \|y(\theta)\|^2_Y \mathrm{d}\omega \mathrm{d}\mathbb{P}(\theta) \\
&= 2\int k_0(\omega)^2 \mathrm{d}\omega + 2\int \|y(\theta)\|^2_Y \mathrm{d}\mathbb{P}(\theta) \int \|k(\omega)\|^2_Y \mathrm{d}\omega < \infty,
\end{aligned}$$

by the Cauchy-Schwarz-inequality. Thus, $\widehat{k} \in \mathcal{L}^2(Y, \mathcal{B}Y, y^{\#}\mathbb{P}; X)$.

2. Let $k_0 := \mathbb{E}[x] - \mathrm{Cov}(x, y)\mathrm{Cov}(y)^{-1}\mathbb{E}[y]$ and $k := \mathrm{Cov}(x, y)\mathrm{Cov}(y)^{-1}$. By Lemma

(A.5), the vector $(k_0, k)^T \in \mathcal{K}$ is a minimiser, if and only if $\mathbb{E}[\langle k_0 + \langle k, y\rangle_Y - x, \widehat{v}(y)\rangle_X] = 0$, for any $\widehat{v} \in \mathcal{K}$. Furthermore,

$$\begin{aligned}\mathbb{E}[\langle \widehat{k}(y) - x, \widehat{v}(y)\rangle_X] &= \mathbb{E}[\langle k_0 + \langle k, y\rangle_Y - x, \widehat{v}(y)\rangle_X] \\ &= \mathbb{E}[\langle \mathbb{E}[x] + \mathrm{Cov}(x,y)\mathrm{Cov}(y)^{-1}(y - \mathbb{E}[y]) - x, \widehat{v}(y)\rangle_X] \\ &= \mathbb{E}[\langle \mathbb{E}[x] + \mathrm{Cov}(x,y)\mathrm{Cov}(y)^{-1}(y - \mathbb{E}[y]) - x, v_0 + \langle v, y\rangle_Y\rangle_X],\end{aligned}$$

which is 0, if $\mathbb{E}[(\mathbb{E}[x] + \mathrm{Cov}(x,y)\mathrm{Cov}(y)^{-1}(y - \mathbb{E}[y]) - x)y_m] = 0$, for any $m \in \{1, ..., M\}$, and $\mathbb{E}[\mathbb{E}[x] + \mathrm{Cov}(x,y)\mathrm{Cov}(y)^{-1}(y - \mathbb{E}[y]) - x] = 0$, which is indeed true:

$$\begin{aligned}&\mathbb{E}[(\mathbb{E}[x] + \mathrm{Cov}(x,y)\mathrm{Cov}(y)^{-1}(y - \mathbb{E}[y]) - x)y_m] \\ &\quad= \mathbb{E}[y_m\mathbb{E}[x] + \mathrm{Cov}(x,y)\mathrm{Cov}(y)^{-1}(y_m y - y_m\mathbb{E}[y]) - y_m x] \\ &\quad= \mathbb{E}[y_m]\mathbb{E}[x] + \mathrm{Cov}(x,y)\mathrm{Cov}(y)^{-1}(\mathbb{E}[y_m y] - \mathbb{E}[y_m]\mathbb{E}[y]) - \mathbb{E}[y_m x] \\ &\quad= \mathbb{E}[y_m]\mathbb{E}[x] - \mathbb{E}[y_m x] + \mathrm{Cov}(x,y)e_m \\ &\quad= \mathrm{Cov}(x,y)e_m - \mathrm{Cov}(x,y_m) = 0, \\ &\mathbb{E}[\mathbb{E}[x] + \mathrm{Cov}(x,y)\mathrm{Cov}(y)^{-1}(y - \mathbb{E}[y]) - x] \\ &\quad= \mathbb{E}[x] + \mathrm{Cov}(x,y)\mathrm{Cov}(y)^{-1}(\mathbb{E}[y] - \mathbb{E}[y]) - \mathbb{E}[x] = 0.\end{aligned}$$ □

The minimiser of the objective function in Theorem (2.24) is $\mathbb{P}$-a.s. (resp. $\mathbb{P}$-a.s. Leb(k)-a.e.) unique, i.e. the property above can be used to define the BLE $y^{\#}\mathbb{P}$-a.s. (resp. $y^{\#}\mathbb{P}$-a.s. Leb(k)-a.e.). However, this report always considers Definition (2.19) for practical reasons. Some corollaries of Theorem (2.24) and properties of the BLE are considered in the next Section 2.3.

2.3. Properties and Examples of Bayes Linear Estimations

2.3.1. Properties

The fundamental characterisation of the Bayes Linear Estimator in Section 2.2 is its optimality in square error loss sense in Theorem (2.24). The BLE also approximates $\mathbb{E}[x|y = \cdot]$ optimally in the sense of expected square error loss. This is formulated below and also discussed in [8] and [21].

(2.26) Corollary. Let Assumptions (2.1) hold and $\mathcal{K} := \{\widehat{k} : Y \to X : \exists k_0 \in X, k \in X^M :$

$\widehat{k} = k_0 + \langle k, \cdot \rangle_Y\} \subseteq \mathcal{L}^2(Y, \mathcal{B}Y, y^{\#}\mathbb{P}; X)$. Then

$$x^{\mathrm{BLE}}_{y=\cdot} \in \operatorname{argmin}_{\widehat{k}\in\mathcal{K}} \mathbb{E}[\|\widehat{k}(y) - \mathbb{E}[x|y]\|^2_X] \qquad \bullet$$

Proof. Consider the given target functional:

$$\begin{aligned}\mathbb{E}[\|\widehat{k}(y) - \mathbb{E}[x|y]\|^2_X] &= \mathbb{E}[\|\widehat{k}(y) - x + x - \mathbb{E}[x|y]\|^2_X] \\ &= \mathbb{E}[\|\widehat{k}(y) - x\|^2_X] + \mathbb{E}[\|x - \mathbb{E}[x|y]\|^2_X]\end{aligned}$$

It differs from the target functional that is given in Theorem (2.24) only by the constant summand $\mathbb{E}[\|x - \mathbb{E}[x|y]\|^2_X]$. Therefore, the solutions of the minimisation problems are equivalent. □

Moreover, [14] classifies the Bayes Linear estimator as a *best linear unbiased estimator (BLUE)*, i.e. the estimator is linear with respect to the data and unbiased. Furthermore, its variance is minimal compared with all other linear estimators. (cf. [44, p. 189])

(2.27) Corollary. Let Assumptions (2.1) hold. The Bayes Linear estimator $x^{\mathrm{BLE}}_{y=\cdot}$ of x given $y = \cdot$ is BLUE. •

Proof. The BLE is unbiased, since

$$\begin{aligned}\mathbb{E}[x^{\mathrm{BLE}}_y] &= \mathbb{E}\left[\mathbb{E}[x] + \mathrm{Cov}(x,y)\mathrm{Cov}(y)^{-1}(y - \mathbb{E}[y])\right] \\ &= \mathbb{E}[x] + \mathrm{Cov}(x,y)\mathrm{Cov}(y)^{-1}(\mathbb{E}[y] - \mathbb{E}[y]) = \mathbb{E}[x].\end{aligned}$$

It is linear by definition and its variance is minimal by Theorem (2.24). □

Part of the statement of Theorem (2.24) is that the elements of the set of affine functions $\mathcal{K}$ are square integrable with respect to $y^{\#}\mathbb{P}$. In particular, this means that the BLE equals the conditional expectation in the best case, but cannot be a better approximation (cf. [8, Corollary 3.5]):

(2.28) Corollary. Let Assumptions (2.1) hold. Then,

$$\mathbb{E}[\|\mathbb{E}[x|y] - x\|^2_X] \leq \mathbb{E}[\|x^{\mathrm{BLE}}_y - x\|^2_X]. \qquad \bullet$$

This corollary is applied in Example (2.32) and implicitly in Theorem (2.31). Before moving on to that, a further Bayes Linear based estimator is introduced.

In addition to the conditional expectation of x given $y = \cdot$, the *conditional covariance operator* of x given $y = \cdot$ can be a useful object to describe the conditional distribution $x^{\#}\mathbb{P}(\cdot|y = \cdot)$. Let $d_1, d_2 \in D$, respectively $\{1, ..., k\}$. The conditional covariance of $x(d_1)$ and $x(d_2)$ given $y = \cdot$ is defined by

(2.29)

$$\mathrm{Cov}(x(d_1), x(d_2)|y = \cdot) := \mathbb{E}\left[(x(d_1) - \mathbb{E}[x(d_1)|y = \cdot])(x(d_2) - \mathbb{E}[x(d_2)|y = \cdot])|y = \cdot\right],$$

and a covariance operator in higher or infinite dimensions can be defined analogously. Goldstein and Wooff [21, Definition 3.8] propose a straightforward Bayes Linear method to approximate the conditional covariance. It is the result of a Bayes Linear approximation of the conditional covariance (2.29):

(2.30) Definition. Let Assumptions (2.1) hold, $\widehat{y} \in Y$ and $d_1, d_2 \in D$ respectively $\{1, ..., k\}$. The *Bayes Linear covariance* of $x(d_1)$ and $x(d_2)$ given $y = \widehat{y}$ is defined by

$$\mathrm{Cov}(x(d_1), x(d_2))^{\mathrm{BLE}}_{y=\widehat{y}} := \mathrm{Cov}(x(d_1), x(d_2)) - \mathrm{Cov}(x(d_1), y)\mathrm{Cov}(y)^{-1}\mathrm{Cov}(y, x(d_2)). \quad \bullet$$

Bayes Linear estimator and Bayes Linear covariance are specifically useful in the case where (x, y) is Gaussian. Conditional and Bayes Linear mean and covariance estimates are equivalent in this setting:

(2.31) Theorem. Let Assumptions (2.1) hold. Furthermore, let $X := \mathbb{R}^k$, $m_x \in X$, $m_y \in Y$, $m = (m_x, m_y)^T$ and $C \in \mathbb{R}^{(M+k)\times(M+k)}$ be a strictly positive definite and symmetric matrix, given by

$$C := \begin{pmatrix} C_x & C_{xy} \\ C_{xy}^T & C_y \end{pmatrix},$$

where $C_x \in \mathbb{R}^{k\times k}, C_{xy} \in \mathbb{R}^{k\times M}$ and $C_y \in \mathbb{R}^{M\times M}$. Moreover, assume that the random variables x, y are jointly Gaussian distributed $(x, y) \sim \mathrm{N}(m, C)$. Then, $x^{\mathrm{BLE}}_{y=\cdot} = \mathbb{E}[x|y = \cdot]$ and $\mathrm{Cov}(x)^{\mathrm{BLE}}_{y=\cdot} = \mathrm{Cov}(x|y = \cdot)$ ($y^{\#}\mathbb{P}$-a.s.). $\bullet$

Proof. The conditional expectation and conditional covariance operator are well-known in this given setting and discussed in several textbooks (cf. e.g. [48, p. 289] or [24, p. 171]). Therefore they are not rigorously derived here but stated below.

$$\mathbb{E}[x|y = \widehat{y}] = m_x + C_{xy}C_y^{-1}(\widehat{y} - m_y) = \mathbb{E}[x] + \mathrm{Cov}(x, y)\mathrm{Cov}(y)^{-1}(\widehat{y} - \mathbb{E}[y]) = x^{\mathrm{BLE}}_{y=\widehat{y}},$$
$$\mathrm{Cov}(x|y = \widehat{y}) = C_x - C_{xy}C_y^{-1}C_{xy}^T = \mathrm{Cov}(x) - \mathrm{Cov}(x, y)\mathrm{Cov}(y)^{-1}\mathrm{Cov}(y, x) = \mathrm{Cov}(x)^{\mathrm{BLE}}_{y=\widehat{y}},$$

where the equations hold for $y^{\#}\mathbb{P}$-almost all $\widehat{y} \in Y$ $\square$

Hence, another way to interpret the Bayes Linear based mean and covariance estimators, as given in Definitions (2.19) and (2.30), is given in Theorem (2.31) above, by showing that the Bayes Linear approximation is almost surely equivalent to the conditional expectation in a Gaussian distributed case. Bayes Linear in a more general class of problems is equivalent to using a Gaussian-type formula in a non-Gaussian problem. This interpretation is, for instance, given by Schillings and Stuart [43, p. 4].

2.3.2. Examples

Examples (2.14) and (2.15) are considered again here. Bayes Linear in the setting of Example (2.17) is implicitly already studied in Theorem (2.31).

(2.32) Example (Example (2.14) revisited)**.** The conditional expectation is given by

$$\mathbb{E}[x|z=\widehat{z}] = \tfrac{1}{4}\mathbb{1}_{\{\widehat{z}=0\}} + \tfrac{3}{4}\mathbb{1}_{\{\widehat{z}=1\}} = \tfrac{1}{4} + \tfrac{1}{2}\widehat{z},$$

where $\widehat{z} \in \{0,1\}$ which is clearly an affine function. Corollary (2.28) implies that $\mathbb{E}[x|z=\cdot] = x^{\mathsf{BLE}}_{z=\cdot}$ ($z^{\#}\mathbb{P}$-a.s.), whenever $\mathbb{E}[x|z=\cdot]$ is $z^{\#}\mathbb{P}$-a.s. affine. Deriving the BLE without referring to Example (2.14) shows the convenience gained from its application:

$$\begin{aligned}\mathbb{E}[x] = \mathbb{E}[z] = \tfrac{1}{2}, \quad \mathrm{Var}(z) = \tfrac{1}{2} - \tfrac{1}{4} = \tfrac{1}{4},\\ \mathrm{Cov}(x,z) = \mathbb{E}[xz] - \mathbb{E}[x]\mathbb{E}[z] = \mathbb{E}[x\mathbb{1}_{x\geq\frac{1}{2}}] - \tfrac{1}{4} = \tfrac{3}{8} - \tfrac{1}{4} = \tfrac{1}{8},\end{aligned}$$

so that

$$x^{\mathsf{BLE}}_{z=\widehat{z}} = \mathbb{E}[x] + \mathrm{Cov}(x,z)\mathrm{Var}(z)^{-1}(\widehat{z} - \mathbb{E}[z]) = \tfrac{1}{2} + 4\cdot\tfrac{1}{8}(\widehat{z} - \tfrac{1}{2}) = \tfrac{1}{4} + \tfrac{1}{2}\widehat{z}. \qquad \bullet$$

(2.33) Example (Example (2.15) revisited)**.** The conditional expectation is given by

$$\mathbb{E}[x|z=\widehat{z}] = \frac{\mathbb{E}[\mathbb{1}_{\{z=\widehat{z}\}}x]}{\mathbb{P}(z=\widehat{z})} = 2^n\int_{\widehat{z}/2^n}^{(\widehat{z}+1)/2^n}\exp(\theta)\mathrm{d}\theta = 2^n[\exp(\tfrac{\widehat{z}+1}{2^n}) - \exp(\tfrac{\widehat{z}}{2^n})],$$

which is clearly a non-linear function whenever $n > 1$. Hence, the parameters of the BLE

have to be derived separately:

$$\begin{aligned}
\mathbb{E}[x] &= \int_{[0,1]} \exp(\theta)\mathrm{d}\theta = \exp(1) - \exp(0) = \mathrm{e} - 1,\\
\mathbb{E}[z] &= \mathbb{E}\left[\sum_{l=0}^{2^n-1} l\mathbb{1}_{(\frac{l}{2^n}, \frac{l+1}{2^n}]}\right] = \sum_{l=0}^{2^n-1} l\mathbb{P}((\tfrac{l}{2^n}, \tfrac{l+1}{2^n}]) = \frac{1}{2^n}\sum_{l=0}^{2^n-1} l = \frac{2^n(2^n-1)}{2^{n+1}} = 2^{n-1} - 2^{-1},\\
\mathbb{E}[z^2] &= \mathbb{E}\left[\sum_{l=0}^{2^n-1} l^2\mathbb{1}_{(\frac{l}{2^n}, \frac{l+1}{2^n}]}\right] = \sum_{l=0}^{2^n-1} l^2\mathbb{P}((\tfrac{l}{2^n}, \tfrac{l+1}{2^n}]) = \frac{1}{2^n}\sum_{l=0}^{2^n-1} l^2 = \tfrac{1}{3}2^{2n} - 2^{n-1} + \tfrac{1}{6},\\
\mathrm{Var}(z) &= \tfrac{1}{3}2^{2n} - 2^{n-1} + \tfrac{1}{6} - (2^{n-1} - 2^{-1})^2 = \tfrac{1}{12}2^{2n} - \tfrac{1}{12},\\
\mathbb{E}[xz] &= \int \left(\sum_{l=0}^{2^n-1} l\mathbb{1}_{(\frac{l}{2^n}, \frac{l+1}{2^n}]}(\theta)\right)\exp(\theta)\mathrm{d}\mathbb{P}(\theta) = \sum_{l=0}^{2^n-1} l\int_{(\frac{l}{2^n}, \frac{l+1}{2^n}]}\exp(\theta)\mathrm{d}\mathbb{P}(\theta)\\
&= \sum_{l=0}^{2^n-1} l[\exp(\tfrac{l+1}{2^n}) - \exp(\tfrac{l}{2^n})],\\
\mathrm{Cov}(x,z) &= \sum_{l=0}^{2^n-1} l[\exp(\tfrac{l+1}{2^n}) - \exp(\tfrac{l}{2^n})] - (\mathrm{e}-1)(2^{n-1} - 2^{-1}).
\end{aligned}$$

Thus, the Bayes Linear Estimator of x given $z = \widehat{z}$ is given by

$$x^{\mathrm{BLE}}_{z=\widehat{z}} = (\mathrm{e}-1) + \frac{\left(\sum_{l=0}^{2^n-1} l[\exp(\frac{l+1}{2^n}) - \exp(\frac{l}{2^n})]\right) - (\mathrm{e}-1)(2^{n-1} - 2^{-1})}{\frac{1}{12}2^{2n} - \frac{1}{12}}(\widehat{z} - 2^{n-1} + 2^{-1}).$$

In Figure 3, Bayes Linear estimations are compared with Taylor approximations of x and conditional expectations of x. A first order Taylor approximation and the Bayes Linear estimations are two different ways to linearise $x = \exp$. The Taylor approximation retains the smoothness of x and is a good (local) approximation in the neighbourhood of some specific point (in this case 0). The Bayes linearisation does not necessarily retain the smoothness. The domain $\Theta' = \mathrm{im}(z)$ of $x^{\mathrm{BLE}}_{z=\cdot}$ is finite in the given setting and the BLE is not constant, but piecewise constant. However, x^{BLE}_{z} is globally optimal and allows approximate integration on any $F \in \sigma(z)$. These subsets fulfil the Radon-Nikodym-equation in Definition (2.5) of the conditional expectation of x given z, so integration on these subsets of the conditional expectation is exact, and the BLE is the best linear approximation of the conditional expectation of x given z. (cf. Corollary (2.26)) •

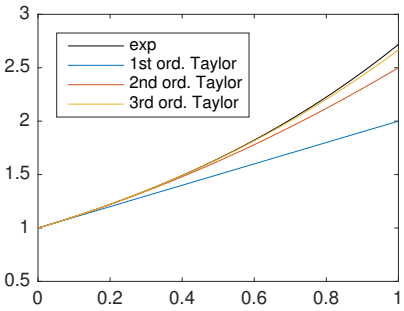

(a) Taylor approximations

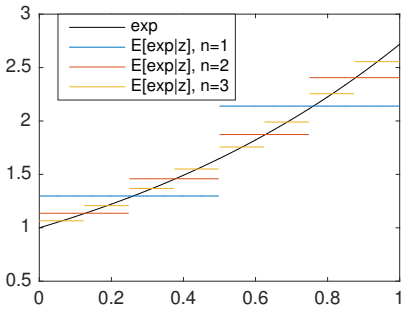

(b) Conditional expectations

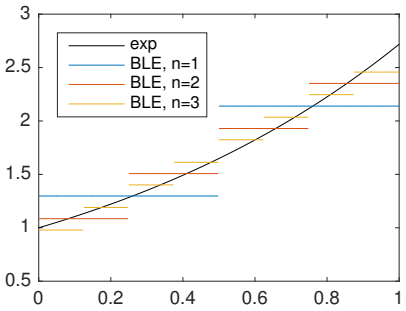

(c) Bayes Linear estimations

Figure 3 – Conditional expectations, Bayes Linear estimations and Taylor approximations of exp as considered in Examples (2.15) and (2.33).

2.4. Bayes Linear Applied to Inverse Problems

2.4.1. Goldstein's Approach

Michael Goldstein and various collaborators [9], [20], [49] have been successfully applying Bayes Linear methods to solve both finite dimensional Bayesian inverse problems and *Bayesian history matching problems*. In the latter, one is interested in intervals containing the true parameters. This can be seen as a reduction of the finite dimensional parameter space, rather than an estimation of the parameters themselves. This report is mostly concerned with applications of Bayes Linear methods that avoid or prevent a (typically expensive) fully Bayesian analysis of the given inverse problems. In contrast to Goldstein who uses Bayes Linear techniques to enable a fully Bayesian analysis. His method is briefly introduced here to provide a more general picture of Bayes Linear approaches to inverse problems, before proceeding to the main focus of this report. A thorough presentation of Goldstein's approach is given by him and Rougier in [20].

Given the inverse problem setting

$$y = \mathcal{O} \circ G(u) + \eta,$$

as first introduced in Subsection 1.1.2, and evaluations $\widehat{G} = (G(u^{(1)}), ..., G(u^{(J)}))$ of the model G with respect to an ensemble of input parameters $\widehat{u} = (u^{(1)}, ..., u^{(J)})$. In [9], not only do the parameters vary in the set of evaluations, but also the precision, i.e. the spatial mesh width, of the evaluation of the model G. The evaluations $(\widehat{u}, \widehat{G})$ are then used to construct an *emulator*. An emulator is a function f, less complex than G, with domain X depending on a vector of parameters v chosen such that $f \approx G$ or $f \approx \mathcal{G} = \mathcal{O} \circ G$. This parameter estimation concerning v is just another (noise free) inverse problem

$$\left(G(u^{(1)}), ..., G(u^{(J)})\right) = \left(f(u^{(1)}, v), ..., f(u^{(J)}, v)\right),$$

respectively

$$\left(\mathcal{G}(u^{(1)}), ..., \mathcal{G}(u^{(J)})\right) = \left(f(u^{(1)}, v), ..., f(u^{(J)}, v)\right),$$

which is approached using *Bayes Linear* (or a fully Bayesian method). The emulator f is then evaluated instead of G resp. $\mathcal{G}$ to solve the inverse or history matching problem in a fully Bayesian manner.

2.4.2. Analytical Bayes Linear Approach

Consider the Bayesian inverse problem setting as described in Subsection 1.1.2. The Bayes Linear estimator is now applied to approximate the conditional expectation of u given that $\mathcal{G}(u) + \eta = y$. The considered assumptions are repeated below:

(2.34) Assumptions. Let X be the *parameter space*, Y the *data space* and $\mathcal{G} : X \to Y$ be a continuous operator. The *unknown parameter* $u \in \mathcal{X}$ is a random variable distributed according to μ_0. The *noise* $\eta \in \mathcal{Y}$ is also a random variable, with $\eta^{\#}\mathbb{P} = \mathrm{N}(0, \Gamma) =: R$ and $\Gamma \in \mathbb{R}^{M \times M}$ is a symmetric and strictly positive definite matrix. η and u are independent. Furthermore, the operator $\mathcal{G}$ is chosen such that $\mathcal{G}(u) + \eta \in \mathcal{Y}$ is square integrable. The given *data* is $y \in Y$. •

The inverse problems considered within this subsection are contrived in the respect that the BLE can be derived analytically. Example (2.18), which focuses on the linear inverse problem, is now reconsidered.

(2.35) Example (Example (2.18) revisited)**.** Let $u \sim \mu_0 = \mathrm{N}(m_0, \mathcal{C}_0)$ be an X-valued random variable. The operator $\mathcal{G} : X \to Y$ is linear and continuous. The conditional expectation of u given $\mathcal{G}u + \eta = y$ is

$$\mathbb{E}[u|\mathcal{G}u + \eta = y] = m_0 + \mathcal{C}_0\mathcal{G}^*(\mathcal{G}\mathcal{C}_0\mathcal{G}^* + \Gamma)^{-1}(y - \mathcal{G}m_0),$$

which is linear with respect to y. In this case, Corollary (2.28) implies that

$$\mathbb{E}[u|\mathcal{G}u+\eta=\cdot] = u^{\mathsf{BLE}}_{\mathcal{G}(u)+\eta=\cdot} \quad \left((\mathcal{G}u+\eta)^{\#}\mathbb{P}\text{-a.s.}\right).$$

Furthermore, the BLE is derived below without considering Corollary (2.28). (cf. Example (2.32))

$$\begin{aligned}
\mathbb{E}[u] &= m_0, \\
\mathbb{E}[\mathcal{G}u+\eta] &= \mathbb{E}[\mathcal{G}u] = \mathcal{G}m_0, \\
\mathrm{Cov}(u, \mathcal{G}u+\eta) &= \mathrm{Cov}(u,\mathcal{G}u) + \mathrm{Cov}(u,\eta) = \mathrm{Cov}(u,\mathcal{G}u) \\
&= \mathbb{E}[(u-m_0)\otimes(\mathcal{G}(u-m_0))] = \mathcal{C}_0\mathcal{G}^*, \\
\mathrm{Cov}(\mathcal{G}u+\eta) &= \mathrm{Cov}(\mathcal{G}u) + \Gamma = \mathbb{E}[\mathcal{G}(u-m_0)\otimes\mathcal{G}(u-m_0)] + \Gamma \\
&= \mathcal{G}\mathcal{C}_0\mathcal{G}^* + \Gamma
\end{aligned}$$

Numerical experiments in Chapter 5 compare the analytical Bayes Linear solution of the linear inverse problem with a numerical solution given by the Ensemble Kalman Filter both with and without Bayes Linear line search. These methods are introduced in the Chapters 3 and 4. •

The second inverse problem considered here is a finite dimensional nonlinear inverse problem, which is similarly given in [46, Example 2.2]. Before considering the problem in the more general high-dimensional case, the problem is presented in one dimension.

(2.36) Example (Cubic inverse problem in 1D)**.** Let Assumptions (2.34) hold, $X := Y := \mathbb{R}$ and $\mathcal{G} : X \to Y$ be given by

$$v \mapsto \mathcal{G}(v) = v^3. \tag{2.37}$$

The prior distribution of u is given by $\mu_0 = \mathsf{N}(m_0, c_0^2)$, for some $c_0^2 > 0$. The parameters of the BLE can be derived using standard results concerning moments of the univariate Gaussian distribution (or cf. e.g. [29]):

$$\begin{aligned}
\mathbb{E}[u] &= m_0, \\
\mathbb{E}[\mathcal{G}(u)] &= \mathbb{E}[u^3] = m_0^3 + 3m_0c_0^2, \\
\mathrm{Cov}(u,\mathcal{G}(u)) &= \mathbb{E}[u^4] - \mathbb{E}[u]\mathbb{E}[u^3] = m_0^4 + 6m_0^2c_0^2 + 3c_0^4 - m_0^4 - 3m_0^2c_0^2 = 3m_0^2c_0^2 + 3c_0^4, \\
\mathrm{Var}(\mathcal{G}(u)) &= \mathbb{E}[u^6] - \mathbb{E}[u^3]^2 = m_0^6 + 15m_0^4c_0^2 + 45m_0^2c_0^4 + 15c_0^6 - (m_0^3 + 3m_0c_0^2)^2 \\
&= 9m_0^4c_0^2 + 45m_0^2c_0^4 + 15c_0^6.
\end{aligned}$$

Then,

$$u^{\mathrm{BLE}}_{\mathcal{G}(u)+\eta=y} = m_0 + \frac{m_0^2 c_0^2 + c_0^4}{3m_0^4 c_0^2 + 15 m_0^2 c_0^4 + 5c_0^6 + \frac{1}{3}\Gamma}(y - 3c_0^2 m_0 - m_0^3),$$

where $3m_0^4 c_0^2 + 15 m_0^2 c_0^4 + 5c_0^6 + \frac{1}{3}\Gamma > 0$, since $c_0 \geq 0$ and $\Gamma > 0$. (cf. Remark (2.42)). Particularly interesting is the case where $m_0 = 0$. In this case:

$$u^{\mathrm{BLE}}_{\mathcal{G}(u)+\eta=y} = \frac{c_0^4}{5c_0^6 + \frac{1}{3}\Gamma} y. \qquad \bullet$$

Proceeding to the multidimensional analogue of example (2.36), one is particularly interested in the situation where the components of u are not independent. Including this case requires the results of Isserlis' theorem [29], which addresses the derivation of (product) moments of a multivariate normal distribution. Although Isserlis' theorem provides a method to derive any (product) moment of the multivariate normal distribution, the Lemma below only presents the required ones.

(2.38) Lemma (Isserlis 1918 [29])**.** Let Assumptions (2.34) hold, $X := \mathbb{R}^2$ and $(x_1, x_2) \in \mathcal{X}$, $(x_1, x_2) \sim \mathrm{N}(0, C)$ be a Gaussian random vector, where $C \in \mathbb{R}^{2\times 2}$ is a symmetric and strictly positive definite matrix. Then,

$$\begin{aligned} \mathbb{E}[x_1 x_2^3] &= 3C_{1,2}C_{2,2} \\ \mathbb{E}[x_1^3 x_2^3] &= 9C_{1,2}C_{1,1}C_{2,2} + 6C_{1,2}^3. \end{aligned} \qquad \bullet$$

Proof. [29, p. 139]. □

Based on this result, the multivariate generalisation of Example (2.36) is studied next.

(2.39) Example (Cubic inverse problem)**.** Let Assumptions (2.34) hold, $X := Y := \mathbb{R}^k$ and $\mathcal{G} : X \to Y$ be given by

$$\begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_k \end{pmatrix} \mapsto \mathcal{G}(v_1, v_2, ..., v_k) = \begin{pmatrix} v_1^3 \\ v_2^3 \\ \vdots \\ v_k^3 \end{pmatrix}.$$

The prior distribution of u is given by $\mu_0 = \mathrm{N}(0, \mathcal{C}_0)$, for some symmetric and strictly positive definite matrix $\mathcal{C}_0 \in \mathbb{R}^{k\times k}$. The parameters of the BLE are derived using Lemma

(2.38). Let $i, j \in \{1, ..., k\}$. Then,

$$\begin{aligned}
\mathbb{E}[u_i] &= 0, \\
\mathbb{E}[\mathcal{G}(u)_i] &= \mathbb{E}[u_i^3] = 0, \\
\mathrm{Cov}(u_i, \mathcal{G}(u)_j) &= \mathbb{E}[u_i u_j^3] = 3(\mathcal{C}_0)_{i,j}(\mathcal{C}_0)_{j,j}, \\
\mathrm{Cov}(\mathcal{G}(u)_i, \mathcal{G}(u)_j) &= \mathbb{E}[u_i^3 u_j^3] = 9(\mathcal{C}_0)_{i,j}(\mathcal{C}_0)_{i,i}(\mathcal{C}_0)_{j,j} + 6(\mathcal{C}_0)_{i,j}^3.
\end{aligned}$$

The BLE is then derived using the standard formula, which is given in Definition (2.19):

$$u^{\mathrm{BLE}}_{\mathcal{G}(u)+\eta=y} = [(\mathcal{C}_0)_{i,j}(\mathcal{C}_0)_{j,j}]_{1\le i,j\le k} \cdot \left[3(\mathcal{C}_0)_{i,j}(\mathcal{C}_0)_{i,i}(\mathcal{C}_0)_{j,j} + 2(\mathcal{C}_0)_{i,j}^3 + \tfrac{1}{3}\Gamma_{i,j}\right]^{-1}_{1\le i,j\le k} \cdot y,$$

where $\left[3(\mathcal{C}_0)_{i,j}(\mathcal{C}_0)_{i,i}(\mathcal{C}_0)_{j,j} + 2(\mathcal{C}_0)_{i,j}^3 + \frac{1}{3}\Gamma_{i,j}\right]_{1\le i,j\le k}$ is indeed invertible (cf. Remark (2.42)).

The analytical Bayes Linear solution of the cubic inverse problem is also compared with a numerical solution given by the Ensemble Kalman Filter with and without Bayes Linear line search and a Monte Carlo approximation of $\mathbb{E}[u|\mathcal{G}(u) + \eta = y]$ using autonormalised importance sampling as given in [1, p. 7] in Chapter 5. •

2.4.3. Bayes Linear Approach with Simulated Parameters

The forward response operator $\mathcal{G}$ is typically not given in a closed form and it is rarely possible to actually derive expected values, variances and covariances of $\mathcal{G}(u)$ analytically ($u \sim \mu_0$). However, the parameters of the BLE can be estimated using a Monte-Carlo-Simulation.

Let an ensemble of $J \in \mathbb{N}$ samples $u^J := \left(u^{(1)}, ..., u^{(J)}\right) \in \mathcal{X}^J$, $u^J \sim \mu_0^{\otimes J}$ and evaluations $\mathcal{G}^J := \left(\mathcal{G}(u^{(1)}), ..., \mathcal{G}(u^{(J)})\right)$ be given. These evaluations $\mathcal{G}^J$ are then independent and identically distributed according to $\mathcal{G}(u)^{\#}\mathbb{P}$. Under Assumptions (2.34), the strong law of

large numbers (SLLN) implies the $\mathbb{P}$-a.s. convergence of the following estimators:

$$\begin{aligned}
\overline{u}_J &= \frac{1}{J}\sum_{j=1}^{J} u^{(j)} \overset{\mathbb{P}\text{-a.s.}}{\longrightarrow} \mathbb{E}[u], \\
\overline{\mathcal{G}(u)}_J &= \frac{1}{J}\sum_{j=1}^{J} \mathcal{G}(u^{(j)}) \overset{\mathbb{P}\text{-a.s.}}{\longrightarrow} \mathbb{E}[\mathcal{G}(u)], \\
\widetilde{\mathrm{Cov}}_J(u,\mathcal{G}(u)) &= \frac{1}{J-1}\sum_{j=1}^{J}(u^{(j)}-\overline{u}_J)\otimes(\mathcal{G}(u^{(j)})-\overline{\mathcal{G}(u^{(j)})}_J) \overset{\mathbb{P}\text{-a.s.}}{\longrightarrow} \mathrm{Cov}(u,\mathcal{G}(u)), \\
\widetilde{\mathrm{Cov}}_J(\mathcal{G}(u)) &= \frac{1}{J-1}\sum_{j=1}^{J}(\mathcal{G}(u^{(j)})-\overline{\mathcal{G}(u^{(\cdot)})}_J)\otimes(\mathcal{G}(u^{(j)})-\overline{\mathcal{G}(u^{(\cdot)})}_J) \overset{\mathbb{P}\text{-a.s.}}{\longrightarrow} \mathrm{Cov}(\mathcal{G}(u)),
\end{aligned}$$

as $J \to \infty$. The SLLN in a finite dimensional setting is discussed in [2, 6.2.5] and the infinite dimensional case is given in [7, Theorem 2.4]. Moreover, the convergence of the covariance operators is discussed in [44, p. 373].

(2.40) Remark. In Definition (2.19), the covariances are given by the (non-corrected) empirical covariance operator $\widehat{\mathrm{Cov}}_J(x) = \left(\frac{J-1}{J}\right)\widetilde{\mathrm{Cov}}_J(x)$ $(x \sim \mu_0)$, which reflects the covariance of the empirical distribution $\mu_0^J = \frac{1}{J}\sum_{j=1}^{J}\delta_{x^{(j)}}$, (where $(x^{(1)},\ldots,x^{(J)}) \sim \mu_0^{\otimes J}$):

$$\begin{aligned}
\overline{x}_J &= \frac{1}{J}\sum_{j=1}^{J} x^{(j)} = \int \xi \mathrm{d}\mu_0^J(\xi), \\
\widehat{\mathrm{Cov}}_J(x) &= \frac{1}{J}\sum_{j=1}^{J}(x^{(j)}-\overline{x}_J)\otimes(x^{(j)}-\overline{x}_J) = \int(\xi-\overline{x}_J)\otimes(\xi-\overline{x}_J)\mathrm{d}\mu_0^J(\xi).
\end{aligned}$$

Although $\widehat{\mathrm{Cov}}_J(x)$ is a biased estimator, it is preferred throughout this report over the unbiased estimator $\widetilde{\mathrm{Cov}}_J(x) = \frac{1}{J-1}\sum_{j=1}^{J}(x^{(j)}-\overline{x}_J)\otimes(x^{(j)}-\overline{x}_J)$. This is because the number of particles that are used to estimate the parameters is typically relatively small. Thus, a correct covariance estimator of the empirical distribution appears to be more important than the correct convergence. Furthermore, the sample covariance $\widehat{\mathrm{Cov}}_J(x)$ is equal to the maximum likelihood estimate of the covariance operator of a Gaussian random variable (cf. [24, p. 184ff]). •

An introduction to Monte Carlo methods is given in the textbook of Robert and Casella [42]. Other relevant Monte Carlo techniques are discussed in [1], [10] and [48].

The Bayes Linear estimator for inverse problems with simulated parameters is then given by

$$u^{\widehat{\mathrm{BLE}},J}_{\mathcal{G}(u)+\eta=y} := \overline{u}_J + \widehat{\mathrm{Cov}}_J(u,\mathcal{G}(u))(\widehat{\mathrm{Cov}}_J(\mathcal{G}(u))+\Gamma)^{-1}(y-\overline{\mathcal{G}(u)}_J). \tag{2.41}$$

This definition is slightly inconsistent with the notation given in Definition (2.19). However, since the covariance operator Γ of the noise is known in the context of Bayesian inverse problems and u, η are independent by assumption, an estimation is not necessary.

(2.42) Remark. Assumptions (2.34) propose a strictly positive definite noise covariance operator Γ. Furthermore, covariance matrices are generally positive semidefinite, implying that the sample covariances are also positive semidefinite (cf. Remark (2.40)). Hence, $[\widehat{\mathrm{Cov}}_J(\mathcal{G}(u)) + \Gamma]$ resp. $[\mathrm{Cov}(\mathcal{G}(u)) + \Gamma]$ is strictly positive definite and invertible. Thus, one can in fact substitute the Moore-Penrose pseudo inverse with the standard inverse. (cf. Remark (2.22)) •

3. The Ensemble Kalman Filter

The underlying assumption of the Bayes Linear estimator is that the conditional expectation of u given $\mathcal{G}(u) + \eta = y$ can be well approximated by a linear function of y. This assumption is for the most part not realistic. In fact, numerical results in Chapter 5 show that pure Bayes Linear estimations of u are quite poor in any of the considered non-linear inverse problems.

Let the forward response operator $\mathcal{G}$ be non-linear and assume that the Bayes Linear estimator and the conditional expectation of u given $\mathcal{G}(u) + \eta = y$ are not a.s. equivalent. Since these estimators are not equivalent, the BLE is not optimal in $\mathcal{L}^2$ (unlike the conditional expectation) and could be improved upon by simply incorporating an (in some way) adjusted version of the data set again using a further Bayes Linear estimation. The subsequent estimate is then possibly also not optimal and this idea can be repeated sequentially. However, the BLE is just a point estimate in X, rather than a random variable in $\mathcal{X}$ or a probability distribution on X, and data incorporation using (Bayes' rule) or similar approaches is not possible. Thus, one has to consider an approximate posterior distribution which has the BLE as posterior mean. These posterior distributions can be approximated given an ensemble of particles and their empirical probability distribution. The probability distributions can then be updated using a Bayes Linear type formula to update the particles (cf. Theorem (3.8)). This proposed update strategy is well-known, without necessarily involving a Bayes Linear interpretation, and referred to as the Ensemble Kalman Filter for inverse problems. This is proposed in [28] and its connection to Bayes Linear is discussed specifically in [14] and [37].

The Ensemble Kalman Filter and variations of it have been successfully applied to various inverse problems. (cf. [14], [26], [27], [43], [28]). Even though the ensemble does typically

not provide a good approximation of the true posterior distribution, the posterior mean approximations are generally good (cf. [14], [43]). Furthermore, the Ensemble Kalman Filter works well with small ensemble sizes (cf. [43]) and is thus a computationally cheap alternative to MCMC, albeit only if one is interested in an estimate of the posterior mean, rather than a posterior distribution.

In the following, the Ensemble Kalman Filter is motivated in the discrete data assimilation setting of [33, Section 2.1] as it is originally proposed by Evensen [15], [16] and its Bayes Linear interpretation given in [14], [37]. Thereafter, in section 3.2, the Ensemble Kalman Filter for inverse problems is derived as it is done (historically) in [28], based on the discrete data assimilation method. The report then proceeds to another version of the Ensemble Kalman Filter that uses the Ensemble Kalman/BLE update to approximate a Sequential Monte Carlo method. This Sequential Monte Carlo version is presented in [43] and considered throughout the remaining chapters of the report.

3.1. Filtering in Data Assimilation and the Ensemble Kalman Filter

(3.1) Definition (Filtering in discrete data assimilation)**.** Let Assumptions (2.34) hold, $\Psi : X \to X$ be a continuous operator and $H : X \to Y$ be a linear operator.

Consider the sequence of random variables $(v_n : n \in \mathbb{N}) \in \mathcal{X}^{\mathbb{N}}$ given by the dynamical system

$$v_n = \Psi(v_{n-1}) + \xi_{n-1} \quad (n \in \mathbb{N}), \tag{3.2}$$

where $v_0 \sim \mathrm{N}(m_0, \mathcal{C}_0)$, $(\xi_n : n \in \mathbb{N} \cup \{0\}) \sim \mathrm{N}(0, \Sigma)^{\otimes \mathbb{N} \cup \{0\}}$ and $\Sigma : X \to X$ is a covariance operator, and the data $(y_n : n \in \mathbb{N}) \in Y^{\mathbb{N}}$ given by realisations of the data stream

$$y_n = Hv_n + \eta_n \quad (n \in \mathbb{N}), \tag{3.3}$$

where $(\eta_n : n \in \mathbb{N}) \sim \mathrm{N}(0, \Gamma)^{\otimes \mathbb{N}}$,

Filtering problem: Let $Y_n := (y_1, ..., y_n)$ $(n \in \mathbb{N})$, $Y_0 = ()$, $H_n = (Hv_1 + \eta_1, ..., Hv_n + \eta_n)$ $(n \in \mathbb{N})$, $H_0 = ()$ and $P(X) := \{\mu : \mu \text{ is a probability measure on } X\}$.

A *filter* is a map

$$F : P(X) \to P(X),\ \mathbb{P}(v_{n-1} \in \cdot | H_{n-1} = Y_{n-1}) \mapsto \mathbb{P}(v_n \in \cdot | H_n = Y_n) \quad (n \in \mathbb{N}) \tag{3.4}$$

This filtering procedure is divided into two consecutive steps: *Prediction* and *analysis*. These steps are given by

$$\text{(Prediction)} \qquad \mathbb{P}(v_{n-1} \in \cdot | H_{n-1} = Y_{n-1}) \mapsto \mathbb{P}(v_n \in \cdot | H_{n-1} = Y_{n-1}),$$
$$\text{(Analysis)} \qquad \mathbb{P}(v_n \in \cdot | H_{n-1} = Y_{n-1}) \mapsto \mathbb{P}(v_n \in \cdot | H_n = Y_n).$$

•

The filtering problem can be solved explicitly given a linear operator $\Psi : X \to X$. This filter is called *Kalman Filter* and sketchily presented in the example below. A rigorous derivation is given in [33, Section 4.1].

(3.5) Example. Let $X := \mathbb{R}^k$ and $\Psi : X \to X$ be a linear operator given by some matrix $A \in \mathbb{R}^{k \times k}$. Furthermore, assume that Γ, the covariance operator of the noise η_n ($n \in \mathbb{N}$) is strictly positive definite.

Since Ψ is a linear operator and v_0 is Gaussian, all elements of $(v_n : n \in \mathbb{N})$ are also Gaussian distributed. Hence, the conditional distributions are uniquely defined by the conditional means and conditional covariances. Set $\mathbb{P}(v_n \in \cdot | H_{n-1} = Y_{n-1}) =: \mathrm{N}(\widehat{m}_n, \widehat{\mathcal{C}}_n)$ and $\mathbb{P}(v_n \in \cdot | H_n = Y_n) =: \mathrm{N}(m_n, \mathcal{C}_n)$ ($n \in \mathbb{N}$).

The distribution of the initial value v_0 is defined to be $\mathbb{P}(v_0 \in \cdot) = \mathrm{N}(m_0, \mathcal{C}_0)$. Since Ψ is linear and v_n, ξ_n, η_n are Gaussian ($n \in \mathbb{N}$), the conditional expectation and covariance of v_n given $H_n = Y_n$ can be derived using Theorem (2.31). As an example, the derivation of the prediction and analysis step from $n-1$ to n is sketched below. The steps are presented as maps $(m_{n-1}, \mathcal{C}_{n-1}) \mapsto (\widehat{m}_n, \widehat{\mathcal{C}}_n)$ and $(\widehat{m}_n, \widehat{\mathcal{C}}_n) \mapsto (m_n, \mathcal{C}_n)$.

Prediction:

$$\begin{aligned}
\widehat{m}_n &= \mathbb{E}[v_n | H_{n-1} = Y_{n-1}] = \mathbb{E}[Av_{n-1} + \xi_{n-1} | H_{n-1} = Y_{n-1}] \\
&= \mathbb{E}[Av_{n-1} | H_{n-1} = Y_{n-1}] + \mathbb{E}[\xi_{n-1}] = Am_{n-1}, \\
\widehat{\mathcal{C}}_n &= \mathrm{Cov}(v_n | H_{n-1} = Y_{n-1}) = A\mathrm{Cov}(v_{n-1} | H_{n-1} = Y_{n-1})A^T + \Sigma = A\mathcal{C}_{n-1}A^T + \Sigma,
\end{aligned}$$

Analysis:

$$
\begin{aligned}
m_n &= \mathbb{E}[v_n | H_n = Y_n] = \mathbb{E}[v_n | H_{n-1} = Y_{n-1}, Hv_n + \eta_n = y_n] \\
&= \mathbb{E}[v_n | H_{n-1} = Y_{n-1}] + \mathrm{Cov}(v_n, Hv_n | H_{n-1} = Y_{n-1}) \\
&\quad \times (\mathrm{Cov}(Hv_n | H_{n-1} = Y_{n-1}) + \Gamma)^{-1} (Y_n - H\mathbb{E}[v_n | H_{n-1} = Y_{n-1}]) \\
&= \widehat{m}_n + \widehat{\mathcal{C}}_n H^T (H\widehat{\mathcal{C}}_n H^T + \Gamma)^{-1} (Y_n - H\widehat{m}_n) \\
&= \widehat{\mathcal{C}}_n H^T (H\widehat{\mathcal{C}}_n H^T + \Gamma)^{-1} Y_n + (I - \widehat{\mathcal{C}}_n H^T (H\widehat{\mathcal{C}}_n H^T + \Gamma)^{-1} H)\widehat{m}_n \\
\mathcal{C}_n &= \mathrm{Cov}(v_n | H_n = Y_n) = \mathrm{Cov}(v_n | H_{n-1} = Y_{n-1}, Hv_n + \eta_n = y_n) \\
&= \mathrm{Cov}(v_n | H_{n-1} = Y_{n-1}) - \mathrm{Cov}(v_n, Hv_n | H_{n-1} = Y_{n-1}) \\
&\quad \times (\mathrm{Cov}(Hv_n | H_{n-1} = Y_{n-1}) + \Gamma)^{-1} \mathrm{Cov}(Hv_n, v_n | H_{n-1} = Y_{n-1}) \\
&= \widehat{\mathcal{C}}_n - \widehat{\mathcal{C}}_n H^T (H\widehat{\mathcal{C}}_n H^T + \Gamma)^{-1} H\widehat{\mathcal{C}}_n \\
&= \left(I - \widehat{\mathcal{C}}_n H^T (H\widehat{\mathcal{C}}_n H^T + \Gamma)^{-1} H\right) \widehat{\mathcal{C}}_n. \qquad \bullet
\end{aligned}
$$

The *Ensemble Kalman Filter* is a generalisation of the Kalman Filter. This generalisation is equipped to work with non-linear operators $\Psi : X \to X$. Given such a non-linear Ψ, the elements of $(v_n : n \in \mathbb{N})$ are, in contrast to the linear case, no longer Gaussian distributed. This means, that updating the mean and covariance operator is insufficient. The Ensemble Kalman Filter addresses this issue by approximating the distribution of v_n through an ensemble of particles. Updates are then applied to this ensemble as opposed to a mean and covariance operator. Apart from that difference, the derivation of the Ensemble Kalman Filter is similar to the derivation of the Kalman Filter, which is given above in Example (3.5).

(3.6) Definition (Ensemble Kalman Filter)**.** Let $J \in \mathbb{N}$ be the number of elements contained in the ensemble that is used to approximate the (conditional) probability distributions that are considered in a filtering problem and $n \in \mathbb{N} \cup \{0\}$ the current step. The ensemble of particles is given by $(v_n^{(j)} : j \in \{1, ..., J\}) \in \mathcal{X}^J$. Moreover, the initial ensemble $(v_0^{(j)} : j \in \{1, ..., J\})$ is distributed according to the distribution of the initial element of the sequence, $(v_0^{(j)} : j \in \{1, ..., J\}) \sim \mathrm{N}(m_0, \mathcal{C}_0)^{\otimes J}$. The *Ensemble Kalman Filter* is given by the following prediction and analysis step.

Prediction:

$$\begin{aligned}
\widehat{v}_{n+1}^{(j)} &:= \Psi(v_n^{(j)}) + \xi_n^{(j)} \quad (j \in \{1,...,J\}, (\xi_n^{(j)} : j \in \{1,...,J\}) \sim \mathrm{N}(0,\Sigma)^{\otimes J}),\\
\widehat{m}_{n+1} &:= \tfrac{1}{J}\sum_{j=1}^{J} \widehat{v}_{n+1}^{(j)},\\
\widehat{C}_{n+1} &:= \tfrac{1}{J}\sum_{j=1}^{J} (\widehat{v}_{n+1}^{(j)} - \widehat{m}_{n+1}) \otimes (\widehat{v}_{n+1}^{(j)} - \widehat{m}_{n+1}).
\end{aligned}$$

Analysis:

$$\begin{aligned}
K_{n+1} &:= \widehat{C}_{n+1} H^T (H \widehat{C}_{n+1} H^T + \Gamma)^{-1},\\
y_{n+1}^{(j)} &:= y_{n+1} + \eta_{n+1}^{(j)} \quad (j \in \{1,...,J\}, (\eta_{n+1}^{(j)} : j \in \{1,...,J\}) \sim \mathrm{N}(0,\Gamma)^{\otimes J}),\\
v_{n+1}^{(j)} &:= (I - K_{n+1}H)\widehat{v}_{n+1}^{(j)} + K_{n+1} y_{n+1}^{(j)} \quad (j \in \{1,...,J\}).
\end{aligned} \tag{3.7}$$

The operator K_{n+1} is called *Kalman gain*. A filter, as it is defined in (3.1), is actually a map defined on the set of probability distributions. As motivated above, these distributions are approximated by

$$\begin{aligned}
\mathbb{P}(v_n \in \cdot | H_n = Y_n) &\approx \mu_n^J = \tfrac{1}{J}\sum_{j=1}^{J} \delta_{v_n^{(j)}},\\
\mathbb{P}(v_{n+1} \in \cdot | H_n = Y_n) &\approx \widehat{\mu}_n^J = \tfrac{1}{J}\sum_{j=1}^{J} \delta_{\widehat{v}_n^{(j)}}.
\end{aligned} \qquad \bullet$$

The Ensemble Kalman Filter is closely related to the Bayes Linear estimator with simulated parameters, which is discussed in Subsection 2.4.3. This fact is discussed in [14, 4.2] and [37, p. 271], and also in the following theorem.

(3.8) Theorem. Consider the setting given in Definition (3.6). Let $j \in \{1,...,J\}$, $n \in \mathbb{N}$ and K_{n+1} be the Kalman gain as it is defined in Equation (3.7). Then,

$$K_{n+1}(y_{n+1}^{(j)} - H\widehat{v}_{n+1}^{(j)}) = -(v_{n+1})^{\widehat{\mathrm{BLE}},J}_{Hv_{n+1}+\eta_{n+1}=H\widehat{v}_{n+1}^{(j)}} + (v_{n+1})^{\widehat{\mathrm{BLE}},J}_{Hv_{n+1}+\eta_{n+1}=y_{n+1}^{(j)}}. \qquad \bullet$$

Proof. Let $\widehat{y} \in Y$. Then,

$$\begin{aligned}
(v_{n+1})^{\widehat{\mathrm{BLE}},J}_{Hv_{n+1}+\eta_{n+1}=\widehat{y}} &= \overline{[v_{n+1}^{(\cdot)}]_J} + \widehat{\mathrm{Cov}}_J(v_{n+1}, Hv_{n+1}+\eta_{n+1})(\widehat{\mathrm{Cov}}_J(Hv_{n+1}) + \Gamma)^{-1}(\widehat{y} - H\overline{[v_{n+1}^{(\cdot)}]_J})\\
&= \overline{[v_{n+1}^{(\cdot)}]_J} + \widehat{\mathrm{Cov}}_J(v_{n+1}) H^* (H\widehat{\mathrm{Cov}}_J(v_{n+1})H^* + \Gamma)^{-1}(\widehat{y} - H\overline{[v_{n+1}^{(\cdot)}]_J})\\
&= \overline{[v_{n+1}^{(\cdot)}]_J} + \widehat{C}_{n+1} H^* (H\widehat{C}_{n+1}H^* + \Gamma)^{-1}(\widehat{y} - H\overline{[v_{n+1}^{(\cdot)}]_J}).
\end{aligned}$$

Thus,

$$\begin{aligned}
-(v_{n+1})^{\widehat{\mathrm{BLE}},J}_{Hv_{n+1}+\eta_{n+1}=H\widehat{v}^{(j)}_{n+1}} &+ (v_{n+1})^{\widehat{\mathrm{BLE}},J}_{Hv_{n+1}+\eta_{n+1}=y^{(j)}_{n+1}} \\
&= -\overline{[v^{(\cdot)}_{n+1}]}_J - \widehat{C}_{n+1}H^*(H\widehat{C}_{n+1}H^*+\Gamma)^{-1}(H\widehat{v}^{(j)}_{n+1} - H\overline{[v^{(\cdot)}_{n+1}]}_J) \\
&\quad + \overline{[v^{(\cdot)}_{n+1}]}_J + \widehat{C}_{n+1}H^*(H\widehat{C}_{n+1}H^*+\Gamma)^{-1}(y^{(j)}_{n+1} - H\overline{[v^{(\cdot)}_{n+1}]}_J) \\
&= -\widehat{C}_{n+1}H^*(H\widehat{C}_{n+1}H^*+\Gamma)^{-1}H\widehat{v}^{(j)}_{n+1} + \widehat{C}_{n+1}H^*(H\widehat{C}_{n+1}H^*+\Gamma)^{-1}y^{(j)}_{n+1} \\
&= \widehat{C}_{n+1}H^*(H\widehat{C}_{n+1}H^*+\Gamma)^{-1}(y^{(j)}_{n+1} - H\widehat{v}^{(j)}_{n+1}) = K_{n+1}(y^{(j)}_{n+1} - H\widehat{v}^{(j)}_{n+1}). \qquad \square
\end{aligned}$$

Hence, the analysis step can be interpreted as a correction of the predicted $\widehat{v}^{(j)}_{n+1}$. The $\widehat{\mathrm{BLE}}$ reflects the conditional expectation on v_{n+1} given the prediction $H\widehat{v}^{(j)}_{n+1}$, respectively the data $y^{(j)}_{n+1}$. This correction can then be seen as a substitution of the information provided by the possibly wrong prediction from the information given by the indeed correct data.

3.2. The Ensemble Kalman Filter for Inverse Problems

This section considers two different but similar ways to apply the Ensemble Kalman Filter to inverse problems. The version from [28], which is based on an auxiliary data assimilation problem, is introduced before moving on to the more relevant method that is proposed by [43] as an approximate Sequential Monte Carlo (SMC) method.

3.2.1. From Data Assimilation to Inverse Problems

Consider the data assimilation setting given in Definition (3.1) and Assumptions (2.34). Let the operator $\Psi : X \oplus Y \to X \oplus Y$ be given by

$$\begin{pmatrix} u \\ z \end{pmatrix} \mapsto \Psi(u,z) = \begin{pmatrix} u \\ \mathcal{G}(u) \end{pmatrix}. \tag{3.9}$$

Furthermore, let $H := (0, I) : X \oplus Y \to Y$ be the projection mapping

$$\begin{pmatrix} u \\ \mathcal{G}(u) \end{pmatrix} \mapsto (0, I) \begin{pmatrix} u \\ \mathcal{G}(u) \end{pmatrix} = \mathcal{G}(u).$$

The initial ensemble is given by $(u^{(j)}_0 : j \in \{1, ..., J\}) \sim \mu_0^{\otimes J}$. The dynamical system induced by Ψ is deterministic, i.e. $\Sigma = 0$.

(3.10) Remark. X is either given by $\mathcal{L}^2(D, \mathcal{B}D, \mathrm{Leb}(k))$ or $\mathbb{R}^k$. The main assump-

tions and therefore the techniques of Section 3.1 do not cover the case that $X = \mathcal{L}^2(D, \mathcal{B}D, \mathrm{Leb}(k)) \oplus Y$. However, a generalisation of the filtering methods from Section 3.1 taking $X = \mathcal{L}^2(D, \mathcal{B}D, \mathrm{Leb}(k)) \oplus Y$ into account is straightforward and therefore not discussed. •

Apply the Ensemble Kalman filter to solve the discrete data assimilation problem presented above. Let $n \in \mathbb{N} \cup \{0\}$. The prediction step is given by

$$\begin{aligned}
\widehat{v}_{n+1}^{(j)} &:= \Psi(u_n^{(j)}, z_n^{(j)}) = \begin{pmatrix} u_n^{(j)} \\ \mathcal{G}(u_n^{(j)}) \end{pmatrix} \quad (j \in \{1, ..., J\}), \\
\widehat{m}_{n+1} &:= \tfrac{1}{J} \sum_{j=1}^{J} \widehat{v}_{n+1}^{(j)} := \tfrac{1}{J} \sum_{j=1}^{J} \begin{pmatrix} u_n^{(j)} \\ \mathcal{G}(u_n^{(j)}) \end{pmatrix}, \\
\widehat{C}_{n+1} &:= \tfrac{1}{J} \sum_{j=1}^{J} (\widehat{v}_{n+1}^{(j)} - \widehat{m}_{n+1}) \otimes (\widehat{v}_{n+1}^{(j)} - \widehat{m}_{n+1}) \\
&= \tfrac{1}{J} \sum_{j=1}^{J} \left(\begin{pmatrix} u_n^{(j)} \\ \mathcal{G}(u_n^{(j)}) \end{pmatrix} - \widehat{m}_{n+1} \right) \otimes \left(\begin{pmatrix} u_n^{(j)} \\ \mathcal{G}(u_n^{(j)}) \end{pmatrix} - \widehat{m}_{n+1} \right). \\
&= \widehat{\mathrm{Cov}}_J \left(\begin{pmatrix} u_n \\ \mathcal{G}(u_n) \end{pmatrix} \right)
\end{aligned}$$

The analysis step based on this prediction is then:

$$\begin{aligned}
K_{n+1} &:= \widehat{C}_{n+1} H^T (H \widehat{C}_{n+1} H^T + \Gamma)^{-1}, \\
&:= \widehat{\mathrm{Cov}}_J \left(\begin{pmatrix} u_n \\ \mathcal{G}(u_n) \end{pmatrix}, (0, I) \begin{pmatrix} u_n \\ \mathcal{G}(u_n) \end{pmatrix} \right) \left[\widehat{\mathrm{Cov}}_J \left((0, I) \begin{pmatrix} u_n \\ \mathcal{G}(u_n) \end{pmatrix} \right) + \Gamma \right]^{-1} \\
&= \widehat{\mathrm{Cov}}_J \left(\begin{pmatrix} u_n \\ \mathcal{G}(u_n) \end{pmatrix}, \mathcal{G}(u_n) \right) \left[\widehat{\mathrm{Cov}}_J (\mathcal{G}(u_n)) + \Gamma \right]^{-1}, \\
y_{n+1}^{(j)} &:= y + \eta_{n+1}^{(j)} \quad (j \in \{1, ..., J\}, (\eta_{n+1}^{(j)} : j \in \{1, ..., J\}) \sim \mathrm{N}(0, \Gamma)^{\otimes J}), \\
v_{n+1}^{(j)} &= \begin{pmatrix} u_{n+1}^{(j)} \\ z_{n+1}^{(j)} \end{pmatrix}, \\
u_{n+1}^{(j)} &:= ((I, 0) - (I, 0) K_{n+1} H) \widehat{v}_{n+1}^{(j)} + K_{n+1} y_{n+1}^{(j)} \\
&= u_n^{(j)} + \widehat{\mathrm{Cov}}_J(u_n, \mathcal{G}(u_n)) \left[\widehat{\mathrm{Cov}}_J (\mathcal{G}(u_n)) + \Gamma \right]^{-1} (y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})) \quad (j \in \{1, ..., J\}), \\
z_{n+1}^{(j)} &:= ((0, I) - (0, I) K_{n+1} H) \widehat{v}_{n+1}^{(j)} + K_{n+1} y_{n+1}^{(j)} \\
&= \mathcal{G}(u_n^{(j)}) + \widehat{\mathrm{Cov}}_J(\mathcal{G}(u_n)) \left[\widehat{\mathrm{Cov}}_J (\mathcal{G}(u_n)) + \Gamma \right]^{-1} (y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})) \quad (j \in \{1, ..., J\}).
\end{aligned}$$

Updating the second component $z \in Y$ of the vector describing the dynamical system is

not necessary, since Ψ is constant with respect to $z \in Y$. Hence, the Ensemble Kalman Filter for Inverse Problems can be defined in the following way.

(3.11) Definition (Ensemble Kalman Filter for Inverse Problems: Version 1)**.** Let Assumptions (2.34) hold and $J \in \mathbb{N}$. The *Ensemble Kalman Filter for Inverse Problems* is given by the initial ensemble $(u_0^{(j)} : j \in \{1, ..., J\}) \in \mathcal{X}^J$ with $(u_0^{(j)} : j \in \{1, ..., J\}) \sim \mu_0^{\otimes J}$ and the ensemble update for all $n \in \mathbb{N} \cup \{0\}$:

$$
\begin{aligned}
(3.12) \qquad u_{n+1}^{(j)} &:= u_n^{(j)} + C_n^{up}(C_n^{pp} + \Gamma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})) \quad (j \in \{1, ..., J\}), \\
y_{n+1}^{(j)} &:= y + \eta_n^{(j)} \quad (j \in \{1, ..., J\}, (\eta_n^{(j)} : j \in \{1, ..., J\}) \sim \mathrm{N}(0, \Gamma)^{\otimes J}) \\
C_n^{up} &:= \widehat{\mathrm{Cov}}_J(u_n, \mathcal{G}(u_n)) = \tfrac{1}{J} \sum_{j=1}^{J} (u_n^{(j)} - \overline{[u_n^{(\cdot)}]}_J) \otimes (\mathcal{G}(u_n^{(j)}) - \overline{[\mathcal{G}(u_n^{(\cdot)})]}_J), \\
C_n^{pp} &:= \widehat{\mathrm{Cov}}_J(\mathcal{G}(u_n)) = \tfrac{1}{J} \sum_{j=1}^{J} (\mathcal{G}(u_n^{(j)}) - \overline{[\mathcal{G}(u_n^{(\cdot)})]}_J) \otimes (\mathcal{G}(u_n^{(j)}) - \overline{[\mathcal{G}(u_n^{(\cdot)})]}_J). \qquad \bullet
\end{aligned}
$$

3.2.2. From Sequential Monte Carlo to a Sequential Bayes Linear Strategy

Consider again the general posterior distribution in the Bayesian inverse problem setting. It is motivated in Subsection 1.1.2 and presented in Theorem (1.8):

$$\mu^y = \frac{1}{Z(y)} \exp\left(-\Phi(\cdot; y)\right) \mu_0,$$

given that $Z(y) := \int \exp(-\Phi(u; y)) \mathrm{d}\mu_0(u) > 0$ ($y \in Y, R$-a.s.). Let $N \in \mathbb{N}$ and $h := \frac{1}{N}$. Sequential Monte Carlo (SMC) does not sample directly from this posterior distribution, but uses a finite sequence of probability measures $(\mu_n : n \in \{0, 1, ..., N\})$, to approach the posterior distribution stepwise. This sequence members ($\neq 0, N$) act as intermediate posterior distributions, where the nth one is reached using the $(n-1)$th distribution ($n \in \{1, ..., N\}$) as a prior distribution. I.e. the steps are given recursively by

$$\mu_n := \frac{1}{Z_n(y)} \exp\left(-h\Phi(\cdot; y)\right) \mu_{n-1} \quad (n \in \{1, ..., N\}),$$

given that $Z_n(y) := \int \exp\left(-h\Phi(u; y)\right) \mathrm{d}\mu_{n-1}(u) > 0$ ($y \in Y, R$-a.s.). Then, $\mu_N = \mu^y$.

A general introduction to SMC is given in [13]. Moreover, SMC methods in the Bayesian inverse problem setting are discussed in [10, 5.3], [30] and briefly also in [43, Chapter 2].

In Section 2.2, the Ensemble Kalman Filter's update is used to approximately transform μ_0 distributed samples into μ^y distributed samples. This Bayes Linear type update formula,

also discussed in Theorem (3.8), is here used to successively update the particles' distribution within the sequence $(\mu_n : n \in \{1, ..., N\})$, i.e. used to approximately transform μ_n distributed samples into μ_{n+1} $(n \in \{0, ..., N-1\})$ distributed samples.

The μ_n-density of the update step is μ_n-a.s. given by

$$\frac{\mathrm{d}\mu_{n+1}}{\mathrm{d}\mu_n}(u) \propto \exp\left(-h\Phi(u;y)\right) = \exp\left(-\frac{h}{2}\|y - \mathcal{G}(u)\|^2_{\Gamma^{-1}}\right) = \exp\left(-\frac{1}{2}\|y - \mathcal{G}(u)\|^2_{(h^{-1}\Gamma)^{-1}}\right).$$

So the only difference between this μ_n-density and the μ_0-density given in Theorem (1.8) (Bayes' rule for inverse problems) is the noise covariance. Therefore, the Bayes Linear type update from an μ_n distributed ensemble to an μ_{n+1} distributed ensemble can be defined as follows below, where just Γ is substituted by $h^{-1}\Gamma$.

(3.13) Definition (Ensemble Kalman Filter for Inverse Problems: Version 2)**.** Let Assumptions (2.34) hold, $N \in \mathbb{N}$ and $J \in \mathbb{N}$. The *Ensemble Kalman Filter for Inverse Problems (EnKF)* is given by the initial ensemble $(u_0^{(j)} : j \in \{1, ..., J\}) \in \mathcal{X}^J$ with $(u_0^{(j)} : j \in \{1, ..., J\}) \sim \mu_0^{\otimes J}$ and the ensemble update for all $n \in \{0, ..., N-1\}$

$$\begin{aligned}
(3.14) \qquad u_{n+1}^{(j)} &:= u_n^{(j)} + C_n^{up}(C_n^{pp} + h^{-1}\Gamma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})) \quad (j \in \{1, ..., J\}),\\
y_{n+1}^{(j)} &:= y + \eta_{n+1}^{(j)} \quad (j \in \{1, ..., J\}, (\eta_{n+1}^{(j)} : j \in \{1, ..., J\}) \sim \mathrm{N}(0, h^{-1}\Gamma)^{\otimes J})\\
C_n^{up} &:= \widehat{\mathrm{Cov}}_J(u_n, \mathcal{G}(u_n)) = \tfrac{1}{J}\sum_{j=1}^{J}(u_n^{(j)} - \overline{[u_n^{(\cdot)}]}_J) \otimes (\mathcal{G}(u_n^{(j)}) - \overline{[\mathcal{G}(u_n^{(\cdot)})]}_J),\\
C_n^{pp} &:= \widehat{\mathrm{Cov}}_J(\mathcal{G}(u_n)) = \tfrac{1}{J}\sum_{j=1}^{J}(\mathcal{G}(u_n^{(j)}) - \overline{[\mathcal{G}(u_n^{(\cdot)})]}_J) \otimes (\mathcal{G}(u_n^{(j)}) - \overline{[\mathcal{G}(u_n^{(\cdot)})]}_J).
\end{aligned}$$

The point estimate of the Ensemble Kalman Filter of u given $\mathcal{G}(u) + \eta = y$ is given by

$$u^{\mathrm{EnKF},J,N}_{\mathcal{G}(u)+\eta=y} := \frac{1}{J}\sum_{j=1}^{J} u_N^{(j)}. \qquad \bullet$$

(3.15) Remark. 1. By the way of construction, this version of the Ensemble Kalman Filter can be easily interpreted as a sequential version of the Bayes Linear estimator with simulated parameters, which is discussed in Section 2.4.3.

2. Applying Bayes Linear/EnKF updates rather than doing normal SMC diminishes the quality of the approximation of the posterior distribution. In fact, the numerical experiments of [14, Chapter 5] show that the ensemble tends to approximate the distribution badly. However, the sample mean is still a good estimate of the conditional expectation of u given $\mathcal{G}(u) + \eta = y$. Thus, the EnKF is treated as a point estimator. (cf. [43, p. 2])

3. Moreover, a major difference between the EnKF and SMC is that the EnKF particles traverse the sample space with fixed weight $\frac{1}{J}$, whereas the SMC particles are spatially fixed but reweighted in every update step. [43, p. 2] states that this second approach tends to result in an approximate posterior which is only concentrated on very few particles since the majority of particles have negligible weights (close to 0). In this case, the posterior distribution is approximated quite badly, which also motivates applying the EnKF in this situation. •

4. Line Search Strategies for the Ensemble Kalman Filter

The Ensemble Kalman Filter as given in Section 3.2 uses a constant step length $h := 1/N$ (cf. Definition (3.13)) respectively $h := 1$ (cf. Definition (3.11)). In numerical non-linear optimisation, constant step lengths are often problematic: an algorithm that uses a step length which is too short has to take many steps to reach the optimal point, each of which is computationally expensive. In contrast, a step length which is too long can result in an algorithm that cannot converge to a stationary point and might consequently not even terminate. This problem motivates a dynamical way to choose step lengths. A brief introduction to line search strategies in numerical non-linear optimisation is given in Section 4.1, which particularly introduces the *Wolfe conditions* [50], [51] in a Gradient descent method. This subsection is mostly based on [38].

Considering the misfit functional of the Ensemble Kalman Filter applied to the linear inverse problem in Figure 5 (a) raises the following questions.

- Is it necessary to choose a shorter step length for the first steps, where the impact of the EnKF on the estimate is relatively large?
- Is it possible to choose a longer step length as the optimum is approached? Here, the impact of the EnKF on the estimate is barely observable. The small steps appear to overfit the model and longer step lengths would reduce the amount of $\mathcal{G}$-evaluations.

Iglesias [26], [27] proposes a way to regularise the Ensemble Kalman Filter, given Hanke's Levenberg-Marquardt condition [23], which is interpreted in this report as a dynamical choice of step lengths. This approach is summarised in Section 4.2.

Bayes Linear methods are used in Section 4.3 to linearise the target functional in the classical inverse problem setting to construct a gradient-free Wolfe condition for the Ensemble Kalman Filter. The Bayes Linear estimator that is used to construct this gradient-free

Wolfe condition implicitly approximates the conditional expectation of the target functional given that the inverse problem holds. This conditional expectation is also examined below.

4.1. Line Search in Numerical Non-linear Optimisation

Let $X = \mathbb{R}^k$ and $f : X \to \mathbb{R}$ be a continuously differentiable and convex function that is bounded below. This section considers the problem of minimising f numerically, i.e. proposing a finite sequence $(x_0, ..., x_N) \in X^N$ approaching the minimum of f, given some initial point $x_0 \in X$.

The method discussed here is the *steepest descent* method, which is given by the following recursion

$$x_{n+1} = x_n - \alpha_n p_n \quad (n \in \{1, ..., N\}),$$

where $\alpha_n > 0$ is the length of the step in direction $p_n = \nabla f(x_n)$. This proposed direction $-\nabla f(x_n)$ is indeed the direction of the steepest descent in point x_n. However, the efficiency and convergence of this algorithm is crucially dependent on the step length α_n, as motivated in the introductory text of Chapter 4. Three different ways to approach this line search problem are briefly stated in the definition below.

(4.1) Definition. 1. Choose $\alpha_n := \alpha \in \mathbb{R}$ constant,

2. Choose $\alpha_n :\in \operatorname{argmin}_{\alpha>0} f(x_n - \alpha \nabla f(x_n))$, which is the optimal way to choose a step size, this is referred to as *exact line search*,

3. Choose some α_n that fulfils the *Wolfe conditions*:

$$f(x_n - \alpha_n p_n) \leq f(x_n) - c_1 \alpha_n \nabla f(x_n)^T p_n, \tag{4.2}$$

$$\nabla f(x_n - \alpha_n p_n)^T p_n \geq c_2 \nabla f(x_n)^T p_n, \tag{4.3}$$

where $0 < c_1 < c_2 < 1$ are predetermined tuning parameters. •

Applying a constant step size is cheap, but often inefficient. Deriving a perfect step length is expensive, but often very efficient. Wolfe conditions attempt to find an optimal compromise between these two competing goals. Condition (4.2) ensures that the update provides good descent results by comparing the value of the target function f in the proposed next step $x_n - \alpha_n p_n$ with a first order Taylor approximation of the target function in this point. However, very small step sizes $\alpha_n \approx 0$ always fulfil this equation. To account for

this, Condition (4.3) is employed to ensure step sizes are not allowed to shrink undeterred. If the Wolfe conditions hold in every update step and f fulfils some further smoothness conditions, the steepest descent method (or more general: any gradient descent method) does indeed converge to a local minimum (cf. [38, Theorem 3.2]).

The Wolfe conditions can be implemented in a backtracking manner. I.e. a maximal step length $\alpha > 0$ is proposed and decreased until the first Wolfe condition (4.2) is fulfilled. The second Wolfe condition (4.3) is negligible, since the decreasing of the step length itself ensures sufficiently large steps. A drawback of any backtracking method is a relatively large amount of f evaluations that are necessary to determine the step length of a single step.

4.2. Regularisation of the Ensemble Kalman Filter

Consider again the classical approach to inverse problems, which is given by the minimisation problem:

$$\min_{u \in X} \|\mathcal{G}(u) - y\|_Y^2.$$

Solving this minimisation problem numerically is problematic. A form of regularisation is suggested in Subsection 1.1.1 (and denoted in Equation (1.5)) to overcome problems given by the non-convexity of the target functional and complexity of $\mathcal{G}$. Assume that $\mathcal{G}$ is continuously differentiable. Then, another approach to solve this minimisation problem is the *Levenberg-Marquardt* (LM) algorithm, a *trust-region* method for non-linear least squares problems. Instead of first proposing a descent direction and then deriving a step size as given in Section 4.1, trust-region algorithms first propose a radius r_n and then derive an (approximately) optimal update $x_{n+1} = x_n + h_n$ within the sphere of this radius r_n around the current state, i.e. $\|h_n\|_X \leq r_n$. This update is given as the optimal solution of a linearisation of the (classical) inverse problem, which is regularised by the Lagrange multipliers of the constraint $\|h_n\|_X \leq r_n$, i.e.

$$h_n \in \operatorname{argmin}_{h \in X} \|y - \mathcal{G}(x_n) - \nabla\mathcal{G}(x_n)h\|_Y + \alpha_n \|h\|_X^2, \tag{4.4}$$

where $\alpha_n \in \mathbb{R}$ is the Lagrange multiplier. Hanke [23, Chapter 2] proposes a way to set this α_n, which ensures stable parameter estimations x_N. There, the Lagrange multiplier α_n is chosen such that

$$\|y - \mathcal{G}(x_n) - \nabla\mathcal{G}(x_n)h_{n,\alpha_n}\|_Y = c_1 \|y - \mathcal{G}(x_n)\|_Y,$$

where $h_{n,\beta} \in \operatorname{argmin}_{h\in X} \|y-\mathcal{G}(x_n)-\nabla\mathcal{G}(x_n)h\|_Y+\beta\|h\|_X^2$ is the solution of the optimisation problem in (4.4) based on the Lagrange multiplier β and $c_1 \in (0,1)$ is some predetermined parameter.

Algorithm 1: Regularised Ensemble Kalman Filter [27, p. 4f]

$N \in \mathbb{N}$ is the amount of steps, J the ensemble size, $\alpha \in (0,\infty)^N$ the minimal Lagrange multiplier in each step and $(u_0^{(j)} : j \in \{1,...,J\}) \sim \mu_0^{\otimes J}$ the inital ensemble.

for $n \in \{0,...,N-1\}$ **do**

  while $c_1\|\Gamma^{-\frac{1}{2}}(y - \overline{[\mathcal{G}(u_n^{(\cdot)})]_J})\|_Y > \alpha_n\|\Gamma^{\frac{1}{2}}(C_{n+1}^{pp} + \alpha_n\Gamma)^{-1}(y - \overline{[\mathcal{G}(u_n^{(\cdot)})]_J})\|_Y$ **do**

    $\alpha_n := 2 \cdot \alpha_n$

  end

  for $i \in \{1,...,J\}$ **do**

    $u_{n+1}^{(j)} := u_n^{(j)} + C_n^{up}(C_n^{pp} + \alpha_n\Gamma)^{-1}(y - \mathcal{G}(u_n^{(j)}))$

  end

end

return $u_{\mathcal{G}(u)+\eta=y}^{\mathrm{EnKFreg},J,N} := \frac{1}{J}\sum_{j=1}^{J} u_N^{(j)}$.

This approach is then applied in a gradient free manner in [26] and [27] to stabilise the estimate of the Ensemble Kalman Filter. Here, α_n is the weight of the regularisation rather than a Lagrange multiplier. Given is Algorithm 1, which uses a backtracking method to derive a minimal α_n fulfilling the gradient free Hanke condition

$$c_1\|\Gamma^{-\frac{1}{2}}(y - \overline{[\mathcal{G}(u_n^{(\cdot)})]_J})\|_Y \leq \alpha_n\|\Gamma^{\frac{1}{2}}(C_{n+1}^{pp} + \alpha_n\Gamma)^{-1}(y - \overline{[\mathcal{G}(u_n^{(\cdot)})]_J})\|_Y.$$

It can be interpreted as a backtracking line search method, since it determines a step length α_n^{-1}. However, the main objective of this algorithm is not a reduction of the amount of $\mathcal{G}$-evaluations, which is motivated in the introductory text of Chapter 4. Its focus is more concerned with regularisation and stabilisation of the EnKF.

Results of numerical experiments comparing this algorithm with the non-regularised Ensemble Kalman Filter are presented in [27, p. 8].

4.3. A Bayes Linear Backtracking Line Search Approach

A backtracking line search method is presented in Section 4.1 in order to determine a step length that is sufficiently large and fulfils the first Wolfe condition (4.2). The first Wolfe condition can be interpreted as a comparison between the target after the next update step and a first order Taylor approximation of the target in the next step. This Taylor approximation is constructed about the current state of the optimisation method.

Motivated by Wolfe, a line search for the EnKF is constructed here that also compares the evaluation of the target functional in the next update step with its linearisation. Since the EnKF is gradient free and its updates are based on Bayes Linear approximations of conditional expectations, a Bayes Linear approximation should form the right hand side of the condition, rather than a first order Taylor approximation. Particular, the Bayes Linear approximation of the conditional expectation of $\|\mathcal{G}(u) - y\|^2_{\Gamma^{-1}}$ given $\mathcal{G}(u) + \eta = y$.

Before this Bayes Linear approximation is discussed, the actual conditional expectation $\mathbb{E}[\|\mathcal{G}(u) - y\|^2_{\Gamma^{-1}} | \mathcal{G}(u) + \eta = y]$ is examined further, which generally states how small the misfit functional is on average, given the inverse problem and potentially provides insight into the influence of noise. Given a linear forward response operator $\mathcal{G} : X \to Y$, this conditional expectation can be derived analytically:

(4.5) Theorem. Let Assumptions (2.34) hold, $X := \mathbb{R}^k$, $\mathcal{G} : X \to Y$ be a linear operator and $u \sim \mu_0 = \mathrm{N}(m_0, \mathcal{C}_0)$. Then,

$$\mathbb{E}\left[\|\mathcal{G}u - y\|^2_{\Gamma^{-1}} | \mathcal{G}u + \eta = y\right] = \sum_{m=1}^{M} [(\Gamma^{-\frac{1}{2}}\mathcal{G}(\mathcal{C}_0 - \mathcal{C}_0\mathcal{G}^T(\mathcal{G}\mathcal{C}_0\mathcal{G}^T + \Gamma)^{-1}\mathcal{G}\mathcal{C}_0)\mathcal{G}^T\Gamma^{-\frac{1}{2}T})_{m,m} + (\Gamma^{-\frac{1}{2}}(\mathcal{G}(m_0 + C_0\mathcal{G}^T(\mathcal{G}C_0\mathcal{G}^T + \Gamma)^{-1}(y - \mathcal{G}m_0)) - y))^2_{m,m}]. \bullet$$

Proof. By the independence of u and η and since $\mathcal{G}$ is linear, $[\mathcal{G}u+\eta]^{\#}\mathbb{P} = \mathrm{N}(\mathcal{G}m_0, \mathcal{G}\mathcal{C}_0\mathcal{G}^T + \Gamma)$. Theorem (2.31) states, that $\mathbb{P}(u \in \cdot | \mathcal{G}u + \eta = y) = \mathrm{N}(m_1, \mathcal{C}_1)$, where

$$\begin{aligned} m_1 &= m_0 + \mathcal{C}_0\mathcal{G}^T(\mathcal{G}\mathcal{C}_0\mathcal{G}^T + \Gamma)^{-1}(y - \mathcal{G}m_0), \\ \mathcal{C}_1 &= \mathcal{C}_0 - \mathcal{C}_0\mathcal{G}^T(\mathcal{G}\mathcal{C}_0\mathcal{G}^T + \Gamma)^{-1}\mathcal{G}\mathcal{C}_0. \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbb{E}\left[\|\mathcal{G}u - y\|^2_{\Gamma^{-1}} | \mathcal{G}u + \eta = y\right] &= \int \|\Gamma^{-\frac{1}{2}}(\mathcal{G}u - y)\|^2_Y \mathrm{dN}(m_1, \mathcal{C}_1)(u) \\ &= \int \|v\|^2_Y \mathrm{dN}(\Gamma^{-\frac{1}{2}}(\mathcal{G}m_1 - y), \Gamma^{-\frac{1}{2}}\mathcal{G}\mathcal{C}_1\mathcal{G}^T\Gamma^{-\frac{1}{2}T})(v) \\ &= \sum_{m=1}^{M} \int v_m^2 \mathrm{dN}(\Gamma^{-\frac{1}{2}}(\mathcal{G}m_1 - y), \Gamma^{-\frac{1}{2}}\mathcal{G}\mathcal{C}_1\mathcal{G}^T\Gamma^{-\frac{1}{2}T})(v) \\ &= \sum_{m=1}^{M} \left[(\Gamma^{-\frac{1}{2}}\mathcal{G}\mathcal{C}_1\mathcal{G}^T\Gamma^{-\frac{1}{2}T})_{m,m} + (\Gamma^{-\frac{1}{2}}(\mathcal{G}m_1 - y))^2_m\right]. \end{aligned}$$

□

The content of this theorem is illustrated in a one dimensional example below.

(4.6) Example. Let $\mathcal{G}v := 5v$ ($v \in \mathbb{R}$), $\Gamma = 1$ and $\mu_0 = \mathrm{N}(0, 2)$. Then, $\mathcal{G}u+\eta \sim \mathrm{N}(0, 51)$,

$\mathbb{E}[u|\mathcal{G}u+\eta=y]=\frac{10}{51}y$, $\mathrm{Var}(u|\mathcal{G}u+\eta=y)=\frac{2}{51}$ and $\mathbb{E}\left[\|\mathcal{G}u-y\|^2_{\Gamma^{-1}}|\mathcal{G}u+\eta=y\right]=\frac{50}{51}+\frac{y^2}{2601}$. •

Since the norm is weighted by Γ^{-1}, one would normally assume that this conditional expectation is approximately 1 - at least given that the noise is not specifically high. Given the example above, one intuitively expects that $\mathcal{G}(u)+\eta\in(-\sqrt{51},\sqrt{51})$ with high probability and in fact $\mathbb{P}(\mathcal{G}(u)+\eta\in(-\sqrt{51},\sqrt{51}))\approx 0.6827$. If y is outside that interval, the conditional expectation is >1, which implies that the noise might have been higher than expected. When y is close to zero, this could imply a small noise or (seemingly more likely) that the noise η and u have different signs. Therefore, the conditional expectation is still approximately given by 1 in that case.

Wolfe and the connection between Bayes Linear and the Ensemble Kalman Filter appears to suggest the Bayes Linear approximation of the conditional expectation given above to be superior to the conditional expectation itself, which is in general hard to determine anyway. This Bayes linearisation is

$$\begin{aligned}&\left[\|\mathcal{G}(u)-y\|^2_{\Gamma^{-1}}\right]^{\mathrm{BLE}}_{\mathcal{G}(u)+\eta=y}\\&=\mathbb{E}\left[\|\mathcal{G}(u)-y\|^2_{\Gamma^{-1}}\right]+\mathrm{Cov}\left(\|\mathcal{G}(u)-y\|^2_{\Gamma^{-1}},\mathcal{G}(u)\right)(\mathrm{Cov}(\mathcal{G}(u))+\Gamma)^{-1}(y-\mathbb{E}[\mathcal{G}(u)]),\end{aligned}$$

and the empirical version that would be used in the n-th ($n\in\mathbb{N}\cup\{0\}$) step of the Ensemble Kalman Filter is then

(4.7)
$$\left[\|\mathcal{G}(u_n)-y\|^2_{\Gamma^{-1}}\right]^{\widehat{\mathrm{BLE}},J}_{\mathcal{G}(u_n)+\eta=y}=\overline{\left[\|\mathcal{G}(u_n^{(\cdot)})-y\|^2_{\Gamma^{-1}}\right]}_J+C_n^{zp}(C_n^{pp}+h_n^{-1}\Gamma)^{-1}(y-\overline{[\mathcal{G}(u_n^{(\cdot)})]}_J),$$

where $C_n^{zp}:=\widehat{\mathrm{Cov}}_J(\|\mathcal{G}(u_n^{(\cdot)})-y\|^2_{\Gamma^{-1}},\mathcal{G}(u_n^{(\cdot)}))$. This Bayes Linear estimation is then compared to the expected value of the target functional in the next step, i.e.

$$\overline{\left[\|\mathcal{G}(u_{n+1}^{(\cdot)})-y\|^2_{\Gamma^{-1}}\right]}_J=\overline{\left[\|\mathcal{G}(u_n^{(\cdot)}+C_n^{up}(C_n^{pp}+h_n^{-1}\Gamma)^{-1}(y_{n+1}^{(\cdot)}-\mathcal{G}(u_n^{(\cdot)})))-y\|^2_{\Gamma^{-1}}\right]}_J, \tag{4.8}$$

which should be smaller than the estimator in (4.7). A backtracking method would now, equivalently to Algorithm 1, decrease h_n by a factor until

$$\begin{aligned}&\overline{\left[\|\mathcal{G}(u_n^{(\cdot)}+C_n^{up}(C_n^{pp}+h_n^{-1}\Gamma)^{-1}(y_{n+1}^{(\cdot)}-\mathcal{G}(u_n^{(\cdot)})))-y\|^2_{\Gamma^{-1}}\right]}_J\\&\leq\overline{\left[\|\mathcal{G}(u_n^{(\cdot)})-y\|^2_{\Gamma^{-1}}\right]}_J+c_1C_n^{zp}(C_n^{pp}+h_n^{-1}\Gamma)^{-1}(\overline{[y_{n+1}^{(\cdot)}]}_J-\overline{[\mathcal{G}(u_n^{(\cdot)})]}_J),\end{aligned} \tag{4.9}$$

where $c_1\in(0,2)$ is a predetermined tuning parameter, which allows an adjustment of the covariance of u and $\mathcal{G}(u)$): $c_1\in(0,1)$ implies a weaker dependence, whereas $c_1\in(1,2)$

implies a stronger dependence.

A basic requirement for any backtracking condition is that the condition holds for step sizes close to 0 or at least for 0 itself. Inequality (4.9) is typically fulfilled (cf. Sections 5.1, 5.2, and 5.3) for small h_n, and certainly as $h_n \downarrow 0$:

(4.10) Lemma. Let Assumptions (2.34) hold and consider Definition (3.13). Then,

$$\overline{\left[\|\mathcal{G}(u_n^{(\cdot)}) - y\|^2_{\Gamma^{-1}}\right]_J} + C_n^{zp}(C_n^{pp} + h^{-1}\Gamma)^{-1}(\overline{[y_{n+1}^{(\cdot)}]_J} - \overline{[\mathcal{G}(u_n^{(\cdot)})]_J})) \xrightarrow{h\downarrow 0} \overline{\left[\|\mathcal{G}(u_n^{(\cdot)}) - y\|^2_{\Gamma^{-1}}\right]_J},$$
$$\overline{\left[\|\mathcal{G}(u_n^{(\cdot)} + C_n^{up}(C_n^{pp} + h^{-1}\Gamma)^{-1}(y_{n+1}^{(\cdot)} - \mathcal{G}(u_n^{(\cdot)}))) - y\|^2_{\Gamma^{-1}}\right]_J} \xrightarrow{h\downarrow 0} \overline{\left[\|\mathcal{G}(u_n^{(\cdot)}) - y\|^2_{\Gamma^{-1}}\right]_J}. \bullet$$

Proof. Let $(\lambda_m^h : m \in \{1, ..., M\})$ be the eigenvalues of $(C_n^{pp} + h^{-1}\Gamma)$, $(\lambda_m^C : m \in \{1, ..., M\})$ be the eigenvalues of C_n^{pp} and $(\lambda_m^\Gamma : m \in \{1, ..., M\})$ be the eigenvalues of Γ and let each of these vectors be in descending order. Since each of these three matrices is symmetric, [45, Theorem 1 (Wielandt)] can be applied given $A := C_n^{pp} + h^{-1}\Gamma$, $B := -C_n^{pp}$ and $A + B = h^{-1}\Gamma$ (A, B are given in [45]'s notation). It states that

$$\lambda_m^h \geq \lambda_M^C + h^{-1}\lambda_m^\Gamma \quad (m \in \{1, ..., M\}).$$

Since Γ is strictly positive definite, $\lambda_m^\Gamma > 0$ $(m \in \{1, ..., M\})$. Therefore, $\lambda_m^h \to \infty$ $(m \in \{1, ..., M\})$, as $h \downarrow 0$. Which implies that all eigenvalues of $(C_n^{pp} + h^{-1}\Gamma)^{-1}$, which are $((\lambda_m^h)^{-1} : m \in \{1, ..., M\})$, converge to 0 as $h \downarrow 0$. But then, $\|(C_n^{pp} + h^{-1}\Gamma)^{-1}\|_2 \to 0$ $(h \downarrow 0)$ and thus $(C_n^{pp} + h^{-1}\Gamma)^{-1} \to 0$ $(h \downarrow 0)$. (cf. [25, p. 151, p. 346])

Eventually, the continuity of $\mathcal{G}$ implies that the statements are true. □

The Ensemble Kalman Filter with Bayes Linear line search is given in Algorithm 2. Some further topics concerning this algorithm are discussed in Remark (4.11)

(4.11) Remark. 1. Even though the backtracking condition is fulfilled in the limit as $h \downarrow 0$, solving a system of linear equations given by $(C_n^{pp} + h^{-1}\Gamma)$ and some vector or inverting that matrix might be numerically infeasible, if $h \approx$ EPS. Hence, the algorithm should have a predefined minimal step length, which is chosen if backtracking reaches that point.

2. All components of the estimator in (4.7) can be determined using the current ensemble and do not require further evaluations of $\mathcal{G}$. The estimator in (4.8) however requires J further evaluations of $\mathcal{G}$ per step length h that is tested. While doing that, $\mathcal{G}$ is also evaluated with respect to the correct update step, which is the update step based on the step length that is finally chosen by the backtracking. Consequently, the algorithm reuses these $\mathcal{G}$-evaluations to derive the next update step rather than evaluating the updated

ensemble again. If the maximal step size fulfils condition (4.9) in every step, the algorithm evaluates $\mathcal{G}$ a total of $JN + J$ times, rather than JN times in the normal EnKF. The additional J times are the evaluations checking whether the last step is indeed short enough. Since the algorithm terminates after that step, this ensemble evaluation cannot be reused.

3. Numerical experiments show that the RHS of condition (4.9) sometimes tends to be negative in the first few steps. Naturally, the condition fails to hold in this case, even

Algorithm 2: Ensemble Kalman Filter with Bayes Linear line search

Let J be the ensemble size, $h_{\max}$ be the maximal and $h_{\min}$ the be minimal step length, $h_\sum := 0$, $(u_0^{(j)} : j \in \{1, ..., J\}) \sim \mu_0^{\otimes J}$ be the inital ensemble, $n := 0$, $c_1 \in (0, 2)$ and $c_2 \in (0, 1)$

$\mathcal{G}_0^{(j)} := \mathcal{G}(u_0^{(j)})$ $(j \in \{1, .., J\})$

while $h_\sum < 1 - \text{EPS}$ **do**

- $h := \min\{1 - h_\sum, h_{\max}\}$ `/* (4.11) 4. */`
- $C_n^{pp} := \frac{1}{J}\sum_{j=1}^J (\mathcal{G}_n^{(j)} - \overline{[\mathcal{G}_n^{(\cdot)}]_J}) \otimes (\mathcal{G}_n^{(j)} - \overline{[\mathcal{G}_n^{(\cdot)}]_J})$
- $C_n^{up} := \frac{1}{J}\sum_{j=1}^J (u_n^{(j)} - \overline{[u_n^{(\cdot)}]_J}) \otimes (\mathcal{G}_n^{(j)} - \overline{[\mathcal{G}_n^{(\cdot)}]_J})$
- $C_n^{zp} := \frac{1}{J}\sum_{j=1}^J (\|\mathcal{G}_n^{(j)} - y\|^2_{\Gamma^{-1}} - \overline{[\|\mathcal{G}_n^{(\cdot)} - y\|^2_{\Gamma^{-1}}]_J}) \otimes (\mathcal{G}_n^{(j)} - \overline{[\mathcal{G}_n^{(\cdot)}]_J})$
- $y_{n+1}^{(j)} := y + \eta_n^{(j)}$ $(j \in \{1, ..., J\}, (\eta_n^{(j)} : j \in \{1, ..., J\}) \sim \mathrm{N}(0, h^{-1}\Gamma)^{\otimes J})$
- $u_{\text{prop}}^{(j)} := u_n^{(j)} + C_n^{up}(C_n^{pp} + h^{-1}\Gamma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}_n^{(j)})$ $(j \in \{1, ..., J\})$
- $\mathcal{G}_{\text{prop}}^{(j)} := \mathcal{G}(u_{\text{prop}}^{(j)})$ $(j \in \{1, .., J\})$
- $\text{RHS} := \overline{\left[\|\mathcal{G}_n^{(\cdot)} - y\|^2_{\Gamma^{-1}}\right]_J} + c_1 C_n^{zp}(C_n^{pp} + h^{-1}\Gamma)^{-1}(\overline{[y_{n+1}^{(\cdot)}]_J} - \overline{[\mathcal{G}_n^{(\cdot)}]_J})$
- **while** $\overline{\left[\|\mathcal{G}_{\text{prop}}^{(\cdot)} - y\|^2_{\Gamma^{-1}}\right]_J} > \text{RHS} \quad \&\& \quad h > h_{\min}$ `/* (4.11) 1. */`

  do
  - $h := c_2 h$
  - $y_{n+1}^{(j)} := y + \eta_n^{(j)}$ $(j \in \{1, ..., J\}, (\eta_n^{(j)} : j \in \{1, ..., J\}) \sim \mathrm{N}(0, h^{-1}\Gamma)^{\otimes J})$
  - $u_{\text{prop}}^{(j)} := u_n^{(j)} + C_n^{up}(C_n^{pp} + h^{-1}\Gamma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}_n^{(j)})$ $(j \in \{1, ..., J\})$
  - $\mathcal{G}_{\text{prop}}^{(j)} := \mathcal{G}(u_{\text{prop}}^{(j)})$ $(j \in \{1, .., J\})$
  - $\text{RHS} := \overline{\left[\|\mathcal{G}_n^{(\cdot)} - y\|^2_{\Gamma^{-1}}\right]_J} + c_1 C_n^{zp}(C_n^{pp} + h^{-1}\Gamma)^{-1}(\overline{[y_{n+1}^{(\cdot)}]_J} - \overline{[\mathcal{G}_n^{(\cdot)}]_J})$

  end
- $u_{n+1}^{(j)} := u_{\text{prop}}^{(j)}$ $(j \in \{1, .., J\})$
- $\mathcal{G}_{n+1}^{(j)} := \mathcal{G}_{\text{prop}}^{(j)}$ $(j \in \{1, .., J\})$ `/* (4.11) 2. */`
- $n := n + 1$
- $h_\sum := h_\sum + h$

end

return $u_{\mathcal{G}(u)+\eta=y}^{\text{EnKFLS},J} := \frac{1}{J}\sum_{j=1}^J u_n^{(j)}$.

if the update would have resulted in a descent in the target functional. A solution to this: substituting the RHS by $\overline{\left[\|\mathcal{G}(u_n^{(\cdot)}) - y\|^2_{\Gamma^{-1}}\right]_J}$, thus ensuring that descent still occurs while allowing large steps. However, initial numerical experiments show this strategy to

be inefficient, accepting that the condition cannot be fulfilled by large steps in this case is more profitable.

4. Iglesias' LM method refers to the Ensemble Kalman Filter that is studied in section 3.2.1 considering the data assimilation setting rather than the Sequential Monte Carlo setting. This report however focuses on the SMC setting, which is why the algorithm ensures that the step lengths add up to 1. •

5. Numerical Results and Discussion

The following methods are numerically investigated: The analytical Bayes Linear estimator, as given in Definition (2.19); the simulated Bayes Linear estimator, as discussed in Subsection 2.4.3; the Ensemble Kalman Filter with fixed step size (EnKF), which is motivated using sequential Monte Carlo and defined in (3.13), and the Ensemble Kalman Filter with Bayes Linear (backtracking) line search (EnKFLS), as stated in Algorithm 2.

These methods are applied to compare estimation results, the variances of the estimators, the data misfit and the accuracy of the parameters. The following test settings are especially noteworthy:

- EnKF and analytical BLE, considering linear (cf. Example (2.18)) and cubic inverse problem (cf. Example (2.39)),
- EnKF with ensemble size J and number of steps N and the BLE, where $J \cdot N$ Monte Carlo samples are used to simulate the parameters, considering the Groundwater flow inverse problem (cf. Example (1.2)),
- EnKF and the expected value of the true posterior, where the true posterior is simulated using self-normalised importance sampling (cf. [1]) considering the cubic inverse problem,
- EnKF and a preconditioned Monte Carlo Markov Chain (pCN MCMC; cf. [10, Algorithm 5.10]) considering a Bayesian inverse problem, where the posterior distribution is multi-modal and MCMC methods are generally problematic. This inverse problem concerns the estimation of parameters of the Glucose-Insulin model of Sturis et al. [47] (cf. also [31]).
- EnKF and EnKFLS, considering the linear, cubic and Groundwater flow inverse prob-

lem.

The following sections deal with the respective inverse problem and present results of numerical experiments.

5.1 the linear inverse problem,

5.2 the cubic inverse problem,

5.3 the Groundwater flow inverse problem,

5.4 Glucose modelling inverse problem.

5.1. Linear Inverse Problem

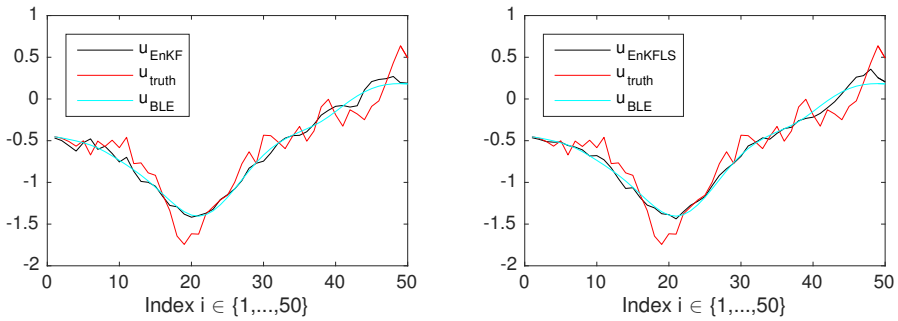

(a) True value, BLE/Conditional expectation estimate and EnKF estimate, given $J = 100$, $N = 50$.

(b) True value, BLE/Cond. exp. estimate and EnKF with Bayes Linear backtr. estimate, given $J = 100$, $h_{\max} = 0.02$.

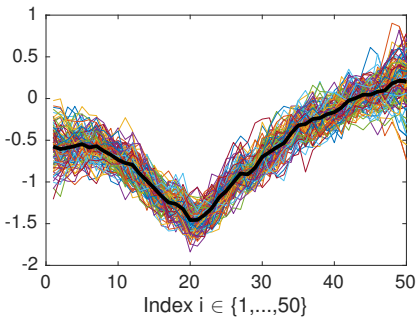

(c) EnKF Particles in step $n = 50$ and their mean (black).

Figure 4 – Linear inverse problem: Estimations.

Given is the linear inverse problem (cf. Example (2.18)) in the following setting: Let $X := Y := \mathbb{R}^{50}$, $\mathcal{G} \in \mathbb{R}^{50\times 50}$ be a 50-dimensional discretisation of the inverse Laplacian $(-\Delta)^{-1}$ on $[0, 1]$, which was derived using the series expansion (truncated after 1000

summands) of the Mercer kernel of the inverse Laplician given by Mercer's theorem (cf. [48, Theorem 11.3]). The prior is a centred Gaussian distribution with covariance matrix $\mathcal{C}_0 = (\frac{1}{2}\exp(-\frac{1}{500}|i-j|))_{1\leq i,j\leq 50}$. The given data $y \in Y$ was simulated using this prior distribution and noise $\eta \sim \mathrm{N}(0, 0.02I)$, which corresponds to a noise level of around 1%. The parameters c_2, which controls the descent of the step size, is set to be 0.25, and c_1, which controls the dependence of u and $\mathcal{G}u$, is 0.3. Furthermore, $h_{\min} := 0.01 h_{\max}$.

Method $\widehat{u} =$	$\frac{\|u^{\text{true}}-\widehat{u}\|_X}{\|u^{\text{true}}\|_X}$	(±SD)	$\frac{\|y-\mathcal{G}(\widehat{u})\|_{\Gamma^{-1}}}{\|y\|_{\Gamma^{-1}}}$	(±SD)	$\mathcal{G}_{ev}$	(±SD)
$\mathbb{E}[u\|\mathcal{G}(u)+\eta=y] = u^{\text{BLE}}_{\mathcal{G}(u)+\eta=y}$	2.257E-1		3.561E-2			
EnKF, $J=25$, $N=5$	2.798E-1	(3.68E-2)	3.688E-2	(9.68E-5)	125	
EnKF, $J=25$, $N=11$	2.857E-1	(4.04E-2)	3.688E-2	(1.02E-3)	275	
EnKFLS, $J=25$, $h_{\max}=0.2$, $c_1=0.3$	2.810E-1	(3.89E-2)	3.690E-2	(1.01E-3)	274	(11.15)
EnKF, $J=50$, $N=25$	2.568E-1	(2.51E-2)	3.626E-2	(5.99E-4)	1250	
EnKF, $J=50$, $N=31$	2.567E-1	(2.66E-2)	3.623E-2	(6.20E-4)	1550	
EnKFLS, $J=50$, $h_{\max}=0.04$, $c_1=0.3$	2.567E-1	(2.73E-2)	3.623E-2	(6.10E-4)	1550	(0)
EnKF, $J=100$, $N=50$	2.436E-1	(2.02E-2)	3.599E-2	(4.33E-4)	5000	
EnKF, $J=100$, $N=56$	2.422E-1	(1.90E-2)	3.561E-2	(4.08E-4)	5600	
EnKFLS, $J=100$, $h_{\max}=0.02$, $c_1=0.3$	2.429E-1	(1.93E-2)	3.595E-2	(4.04E-4)	5600	(0)

Table 1 – Linear inverse problem: Numerical results

Chosen are three different settings $(J, N) \in \{(25, 5), (50, 25), (100, 50)\}$ for numerical experiments given this inverse problem. One further simulation with the EnKF is performed each with an amount of steps $\widehat{N}$ such that $\widehat{N} \cdot J \approx \mathcal{G}_{ev}$, where $\mathcal{G}_{ev}$ is the number of $\mathcal{G}$-evaluations that were necessary to perform the EnKFLS in the settings above each with $h_{\max} := 1/N$. The results of the Ensemble Kalman Filters (EnKF, EnKFLS) that are presented in Table 1, are based on 500 test runs each.

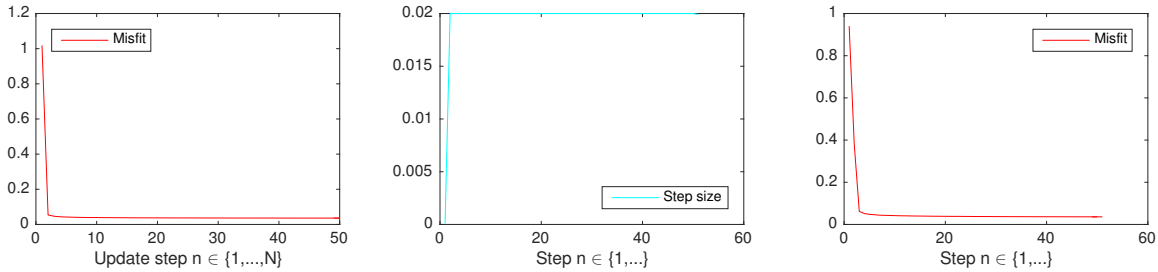

(a) Misfit functional EnKF (b) Step sizes that were chosen by the line search (c) Misfit functional EnKFLS

Figure 5 – Linear inverse problem: Misfit and step sizes of the EnKF as given in Figure 4 (b).

The Bayes Linear estimator and the conditional expectation are a.s. identical in this case, i.e. the Bayes Linear estimator is the optimal $\mathcal{L}^2$-function approximating u given the inverse problem. Therefore all the (non-linear) Ensemble Kalman Filter estimates, which were realised with a relatively small amount of particles, are worse than this Bayes Linear estimate. The line search does not appear to be significantly useful in this case either.

Further tuning of the parameters c_1, c_2 might improve its results or at least decrease the number of $\mathcal{G}$-evaluations that are necessary to determine the step length by backtracking. However, since the Bayes Linear estimate is optimal anyway, putting a lot of effort into further tuning might be rather unreasonable. The misfit functional and the line search behave as proposed in the introduction of Chapter 4: The first steps have a much bigger impact on the misfit than all the others and the line search algorithm takes this into account by choosing a specifically short step size for these initial steps. The impact of the first steps on the misfit in the line search case appears to be slightly smaller than in the standard case, i.e. the trajectory is steeper in Figure 5(a) than it is in Figure 5(c). This is expected, since the chosen step sizes and thus the impact of the EnKF updates are smaller here.

5.2. Cubic Inverse Problem

Method $\widehat{u} =$	$\frac{\|u^{\text{true}}-\widehat{u}\|_X}{\|u^{\text{true}}\|_X}$	(±SD)	$\frac{\|y-\mathcal{G}(\widehat{u})\|_{\Gamma^{-1}}}{\|y\|_{\Gamma^{-1}}}$	(±SD)	$\mathcal{G}_{ev}$	(±SD)
$u^{\text{BLE}}_{\mathcal{G}(u)+\eta=y}$	8.315E-1		9.842E-1			
$u^{\text{MC}}_{\mathcal{G}(u)+\eta=y}$	5.465E-1		7.229E-1			
EnKF, $J = 25, N = 5$	6.834E-1	(7.57E-2)	8.732E-1	(5.25E-2)	125	
EnKF, $J = 25, N = 12$	6.927E-1	(7.68E-2)	8.360E-1	(5.64E-2)	300	
EnKFLS, $J = 25, h_{\max} = 0.2, c_1 = 0.6$	6.819E-1	(8E-2)	8.565E-1	(5.67E-2)	294.3	(88.25)
EnKF, $J = 50, N = 25$	6.450E-1	(6.15E-2)	7.657E-1	(3.81E-2)	1250	
EnKF, $J = 50, N = 32$	6.464E-1	(5.91E-2)	7.614E-1	(3.94E-2)	1600	
EnKFLS, $J = 50, h_{\max} = 0.04, c_1 = 0.6$	6.445E-1	(6.03E-2)	7.594E-1	(3.81E-2)	1597.7	(245.85)
EnKF, $J = 100, N = 50$	6.218E-1	(4.59E-2)	7.261E-1	(2.12E-2)	5000	
EnKF, $J = 100, N = 57$	6.222E-1	(4.36E-2)	7.283E-1	(2.30E-2)	5700	
EnKFLS, $J = 100, h_{\max} = 0.02, c_1 = 0.6$	6.251E-1	(4.56E-2)	7.245E-1	(2.18E-2)	5630	(455.80)

Table 2 – Cubic inverse problem: Numerical results

Numerical experiments are done given the 50-dimensional version of the cubic inverse problem (cf. Example (2.39)), i.e. again $X := Y := \mathbb{R}^{50}$. The covariance operator of the centred Gaussian distribution that is used to simulate the data $y \in Y$ and also used as the prior distribution is proportional to the inverse Laplacian $(-\Delta)^{-1}$ on $[0, 1]$, particularly $\mathcal{C}_0 := 0.1(-\Delta)^{-1}$. It is again discretised using Mercer's theorem. The noise η is also centred Gaussian with covariance $\Gamma = 0.00001I$, which corresponds to a noise level of around 1%. The analytical Bayes Linear estimator is derived (cf. Example (2.39)) and the conditional expectation of u given $\mathcal{G}(u) + \eta = y$ is simulated using autonormalised importance sampling (cf. [1, p. 7]) with 5E6 samples. The settings of the Ensemble Kalman Filter with and without Bayes Linear (backtracking) line search are chosen as in the linear case, apart from the parameter c_1, which is set to be 0.6 here. The EnKF and EnKFLS results are again based on 500 testruns.

The analytical BLE is a relatively simple way to estimate the true parameter u and although its results are the worst presented, the rough shape of the estimate appears to be correct. The Ensemble Kalman Filter estimates however are better and even relatively close to the Monte Carlo simulated conditional expectation of u given $\mathcal{G}(u)+\eta = y$, which is a Monte Carlo approximation of the optimal estimator. Furthermore, the Ensemble Kalman Filter estimates are cheaper ($J \cdot N = 5\text{E}3$ $\mathcal{G}$-evaluations) than the Monte Carlo approximation (5E6 $\mathcal{G}$-evaluations), which proved unreliable when tested with samples. Therefore, the EnKF appears to be a legitimate alternative to a Monte Carlo approximation whenever one is only interested in an approximation of the conditional expectation.

The misfit functionals differed from expectations. Especially the impact of the first steps on the misfit is not as big as in the linear or the groundwater flow case. Despite this, the backtracking line search surprisingly chooses the step sizes more or less equivalently to the linear case.

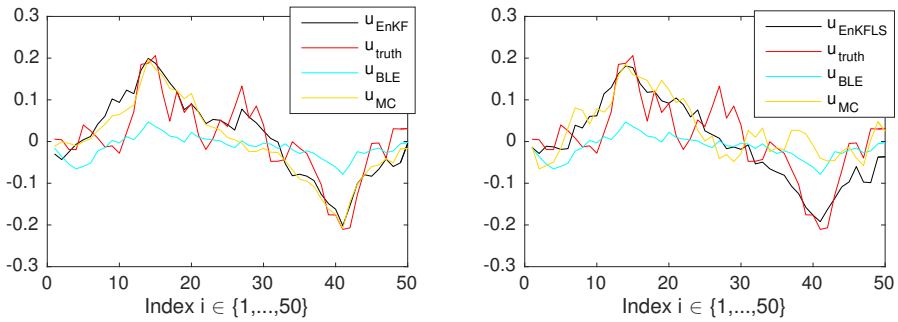

(a) True value, BLE/Conditional expectation estimate and EnKF estimate, given $J = 100$, $N = 50$.

(b) True value, BLE/Cond. exp. estimate and EnKF with Bayes Linear backtr. estimate, given $J = 100$, $h_{\max} = 0.02$.

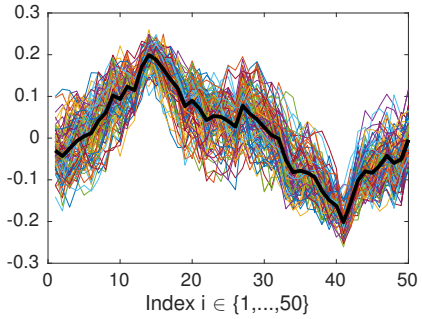

(c) EnKF Particles in step $n = 50$ and their mean (black).

Figure 6 – Cubic inverse problem: Estimations.

Even if the rough shape of the true parameter u^{true} is well approximated by any of these methods, none of the relative errors of the estimation result are overwhelmingly good. To prevent against explosion of the iterative methods one would have expected given the

cubic inverse problem, which would have deteriorated the estimation results, an appropriate prior and noise distribution should be chosen. In particular, they should ensure that all the particles are always in $(-1,1)^{50}$ with high probability. This (possibly overly informative) prior however might be the reason that this estimation problem appears to be much harder than the linear inverse problem, even if the cubic and the linear $\mathcal{G}$ are both bijections and hence relatively well-posed.

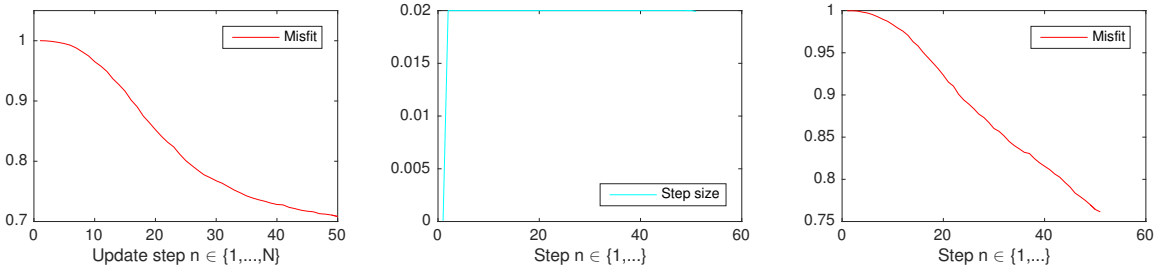

(a) Misfit functional EnKF (b) Step sizes (c) Misfit functional EnKFLS

Figure 7 – Cubic inverse problem: Misfit and step sizes of the EnKF as given in Figure 6(b).

5.3. Groundwater Flow Inverse Problem

Method $\widehat{u} =$	$\frac{\|u^{\text{true}}-\widehat{u}\|_X}{\|u^{\text{true}}\|_X}$	(±SD)	$\frac{\|y-\mathcal{G}(\widehat{u})\|_{\Gamma^{-1}}}{\|y\|_{\Gamma^{-1}}}$	(±SD)	$\mathcal{G}_{ev}$	(±SD)
BLE, $J = 125$	1.014	(3.20E-2)	1.445E-1	(2.15E-2)	125	
EnKF, $J = 25$, $N = 5$	9.342E-1	(9.35E-2)	3.941E-2	(6.73E-3)	125	
EnKF, $J = 25$, $N = 8$	9.282E-1	(8.92E-2)	3.591E-2	(7.362E-3)	200	
EnKFLS, $J = 25$, $h_{\max} = 0.2$, $c_1 = 0.8$	9.308E-1	(9.00E-2)	4.282E-2	(1.56E-2)	193.5	(123.2)
BLE, $J = 1250$	1.004	(1.27E-2)	1.347E-1	(3.53E-3)	1250	
EnKF, $J = 50$, $N = 25$	8.431E-1	(5.67E-2)	3.871E-2	(3.28E-3)	1250	
EnKF, $J = 50$, $N = 27$	8.487E-1	(5.75E-2)	3.756E-2	(3.48E-3)	1350	
EnKFLS, $J = 50$, $h_{\max} = 0.04$, $c_1 = 0.98$	8.381E-1	(5.07E-2)	3.89E-2	(3.62E-3)	1313	(62.2)
BLE, $J = 5000$	9.995E-1	(6.16E-3)	1.344E-1	(2.40E-3)	5000	
EnKF, $J = 100$, $N = 50$	7.996E-1	(3.75E-2)	4.415E-2	(2.27E-3)	5000	
EnKF, $J = 100$, $N = 52$	8.049E-1	(4.09E-2)	4.357E-2	(2.38E-3)	5200	
EnKFLS, $J = 100$, $h_{\max} = 0.02$, $c_1 = 1.05$	7.973E-1	(4.03E-2)	4.412E-2	(2.18E-3)	5161	(184.2)
BLE, $J = 5E4$	9.999E-1	(4.88E-1)	1.339E-1	(6.53E-2)	5E4	

Table 3 – Groundwater flow inverse problem: Numerical results

Consider the groundwater flow inverse problem, that is exemplified in (1.2) in the setting of the identical problem in [43, Section 5.2]. Particularly, $D := (-1,1)^2$, $X := \mathcal{L}^2(D, \mathcal{B}D, \text{Leb}(2))$ (rather than $\mathcal{L}^\infty(D, \mathcal{B}D, \text{Leb}(2))$, as proposed in Example (1.2), cf. [10, Theorem 2.18, Example 2.19]), $Y := \mathbb{R}^{49}$, $f(d) := 100$ $(d \in D)$, $\mu_0 := \text{N}(0, (-\Delta|_D)^{-2})$, $\Gamma := I$. The 49 observations are (deterministically) uniformly picked in D, i.e. $\mathcal{G} : X \to Y$,

$$v \mapsto (p_v(d) : d \in \tfrac{1}{4}\mathbb{Z}^2 \cap D),$$

where p_v is the solution of PDE (1.3) given log-permeability v, which is numerically evaluated by a linear finite elements method with spatial mesh width 2^{-4}. True parameter and data y is simulated using this model $\mathcal{G}$, prior μ_0 and noise distribution $R = \mathrm{N}(0, \Gamma)$. Ensemble size J, the number of steps N, and the step size boundaries $h_{\max}, h_{\min}$ are chosen exactly as they are given in the other examples above. The parameters c_2, which is 0.5 here, and c_1, the choice of which is discussed below. The simulated Bayes Linear Esti-

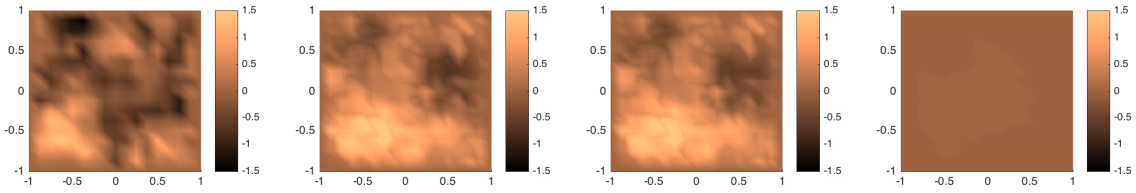

(a) True Parameter (b) EnKF (c) EnKFLS (d) BLE

Figure 8 – Groundwater flow inverse problem: Estimation results of EnKF ($J = 100, N = 50$), EnKF with Bayes Linear line search ($J = 100, h_{\max} = 0.02$) and BLE ($J = 5000$)

mator is based on $J^{\mathrm{BLE}} := J \cdot N$ samples/evaluations of $\mathcal{G}$ in each of the given settings. Furthermore, [43] provides evidence that the EnKF gives reasonable estimates, even if the ensemble size is rather small. However, simulating the parameters of the BLE with a small sample size does not appear adequate. Therefore, a further BLE is simulated using $J^{\mathrm{BLE}} :=$ 5E4 particles. Due to the high computational complexity, the numerical results are based on only 100 test runs per setting. The tuning of the parameter c_1 resulted in three different parameters for the three given settings, which increases as J increases. The line search can be interpreted here as a correction mechanism that ensures that small steps are chosen whenever sample means and covariances are badly approximated, which is especially critical within the first steps. Moreover, the approximations are generally better, given some greater ensemble size J. The resulting increased credibility of the simulated means and covariances allow for steeper descent to be encouraged, justifying how parameter c_1 is tuned here.

Given the results in Table 3, the line search appears to be generally more helpful given this inverse problem than in the linear and cubic case. However, the variation of the estimators is relatively high and considering Figure 8, fundamental differences depicted in the comparison of the Ensemble Kalman Filter and the regularised Ensemble Kalman Filter (cf. [27, Figure 1]) do not appear.

In Figure 9, the misfit functionals of both the EnKF with and without line search are similar to those of the linear inverse problem. Again, the impact of the first steps is greater than that of the later steps. The line search takes this into account and chooses the step sizes according to this fact.

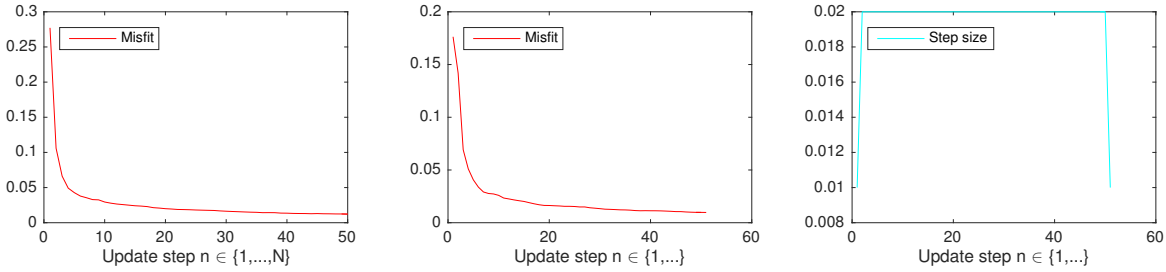

(a) Misfit EnKF (b) Misfit EnKFLS (c) BLE

Figure 9 – Groundwater flow inverse problem: Misfit functional and step sizes of EnKF and EnKFLS as given in Figure 8

Simulating the (one step) BLE in this setting is not useful at all. Comparing some EnKF estimation using J particles and $N \geq 2$ steps with some Bayes Linear estimation using J^{BLE} particles. The BLE solution of that inverse problem was always worse than the EnKF solution, even if $J^{\mathrm{BLE}} \gg J \cdot N$. Actually, even the poorest non-linear estimation (EnKF; $J = 25, N = 5$) is better than the best linear estimation (BLE; $J = 5\mathrm{E}4$). This implies that the linearity assumption the BLE is based on is generally indeed too strong, a fact which is used as a motivation of the Ensemble Kalman Filter in the introduction of Chapter 3.

5.4. Glucose Modelling Inverse Problem

A parameter estimation problem concerning the insulin and glucose model that is derived in [47] and also given in [31, Section 16.3] is studied. The biological background of glucose and insulin, their ultradian oscillations and the construction of $\mathcal{G}$ are also discussed in [47] and [31]. The notations below are borrowed from [31], even though the model is treated purely mathematically. It is given by the following system of ordinary differential equations

$$
\begin{aligned}
\frac{\mathrm{d}I_p}{\mathrm{d}t} &= f_1(G) - E\left(\frac{I_p}{V_p} - \frac{I_i}{V_i}\right) - \frac{I_p}{t_p}, \\
\frac{\mathrm{d}I_i}{\mathrm{d}t} &= E\left(\frac{I_p}{V_p} - \frac{I_i}{V_i}\right) - \frac{I_i}{t_i}, \\
\frac{\mathrm{d}G}{\mathrm{d}t} &= f_4(h_3) + I_G(t) - f_2(G) - f_3(I_i)G, \\
\frac{\mathrm{d}h_1}{\mathrm{d}t} &= \frac{I_p - h_1}{t_d}, \\
\frac{\mathrm{d}h_2}{\mathrm{d}t} &= \frac{h_1 - h_2}{t_d}, \\
\frac{\mathrm{d}h_3}{\mathrm{d}t} &= \frac{h_2 - h_3}{t_d},
\end{aligned}
\tag{5.1}
$$

where the functions f_1, f_2, f_3 and f_4 are given by

$$
\begin{aligned}
f_1(G) &= \frac{R_m}{1+\exp\left(\frac{-G}{V_g C_1}+a_1\right)},\\
f_2(G) &= U_b\left(1-\exp\left(\frac{-G}{V_g C_2}\right)\right),\\
f_3(I_i) &= \frac{1}{V_g C_3}\left(U_0+\frac{U_m-U_0}{1+(\frac{1}{C_4}(\frac{1}{V_i}-\frac{1}{E t_i})I_i)^{-\beta}}\right),\\
f_4(h_3) &= \frac{R_g}{1+\exp\left(\alpha\left(\frac{h_3}{V_p C_5}-1\right)\right)},
\end{aligned}
$$

$a_1 = 6.67, C_1 = 300, C_2 = 144, C_3 = 100, C_4 = 80, C_5 = 26, R_g = 180, R_m = 209, t_p = 6, t_d = 12, U_b = 72, U_0 = 4, U_m = 94, V_i = 11, V_g = 10, \alpha = 7.5, \beta = 1.77$ and $I_G : [0,\infty) \to \mathbb{R}$ is a continuous function. The vector $v := (E, V_p, t_i) \in (0,\infty)^3$ is a parameter and $F_v : [0,\infty) \to \mathbb{R}^6$ is the solution of this ode parametrised by v.

Method $\widehat{u} =$	$\frac{\|u^{\text{true}}-\widehat{u}\|_X}{\|u^{\text{true}}\|_X}$	(±SD)	$\frac{\|y-\mathcal{G}(\widehat{u})\|_{\Gamma^{-1}}}{\|y\|_{\Gamma^{-1}}}$	(±SD)	$\mathcal{G}_{ev}$
EnKF, $J = 12, N = 5$	7.583E-1	(4.96E-1)	6.975E-2	(1.73E-3)	60
EnKF, $J = 24, N = 10$	5.710E-1	(3.25E-1)	6.936E-2	(3.74E-4)	240
EnKF, $J = 36, N = 15$	5.939E-1	(2.57E-1)	7.001E-2	(2.13E-3)	540
EnKF, $J = 48, N = 20$	4.641E-1	(3.12E-1)	7.015E-2	(1.15E-3)	960
preconditioned MCMC	8.749E-1	(8.30E-1)	4.45E-2	(7.5E-3)	

Table 4 – Glucose modelling inverse problem: Numerical results

The *glucose modelling inverse problem* is given by $X := \mathbb{R}^3, Y := \mathbb{R}^M$, the forward response operator $\mathcal{G} : X \to Y$,

$$
\widehat{u} \mapsto \left(\pi_3 \circ F_{(\exp(\widehat{u}_1),\exp(\widehat{u}_2),\exp(\widehat{u}_3))}(t_m) : m \in 1,...,M\right),
$$

where $\pi_3 : \mathbb{R}^6 \to \mathbb{R}$ is the projection on the third coordinate and $(t_1,...,t_M) \in [0,\infty)^M$ are strictly ascending time steps. The given data $y \in Y$ is an evaluation of $\mathcal{G}(u) + \eta$, where

$$
\begin{aligned}
&u \sim \mu_0 := \mathrm{N}(\log(0.2), \log(0.2)^2) \otimes \mathrm{N}(\log(3), \log(3)^2) \otimes \mathrm{N}(\log(100), \log(100)^2),\\
&\eta \sim R := \mathrm{N}(0, 25I).
\end{aligned}
$$

It would have been possible to denote the prior as a log-normal distribution of the genuine parameters E, V_p and t_i. The notation using a Gaussian prior here is chosen for the sake of consistency. Moreover, the noise distribution R corresponds to a noise level of around 5%.

The glucose modelling inverse problem is numerically solved using the fixed step size En-

semble Kalman Filter and a preconditioned Monte Carlo Markov Chain (the MCMC results were provided by David Albers and Matthew Levine, Columbia University, New York). The numerical evaluation of the given model $\mathcal{G}$ is notably more expensive than any of the other examined models. This is mostly due to its stiffness, which was detected while doing numerical experiments with explicit solvers which did not use sufficiently small step sizes. To protect against further issues arising from the stiffness of the model, the model is evaluated using an implicit solver (Matlab ode15s). This is, in general a conservative and inefficient choice. (cf. [41, Section 11.10]) Due to this, and the low dimensional parameter space X, relatively small ensemble sizes are considered, more precisely $(J, N) \in \{(12, 5), (24, 10), (36, 15), (48, 20)\}$. The numerical results that are presented in Table 4 are based on 20 runs of the EnKF and 7 runs of the preconditioned Monte Carlo Markov Chain (MCMC). The MCMC method is based on 1.5E4 steps, respectively 1.3E4 steps and a burn-in period of 2E3 steps.

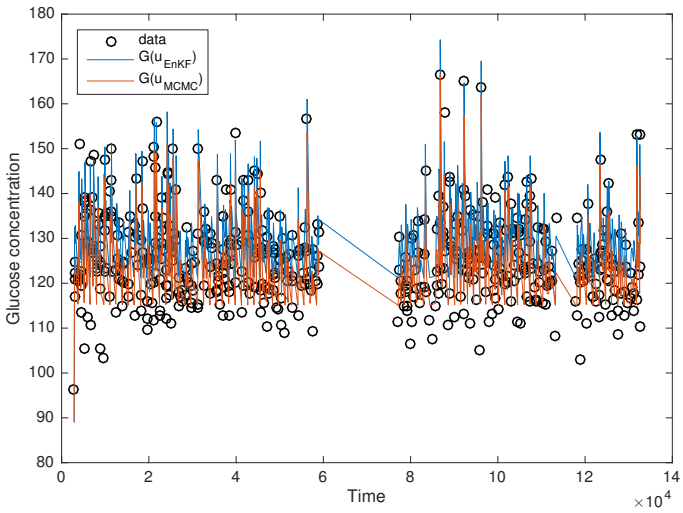

(a) EnKF estimate, MCMC estimate and data

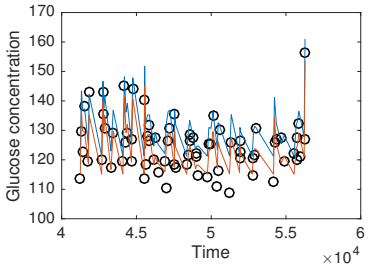

(b) Enlarged view on a section of (a)

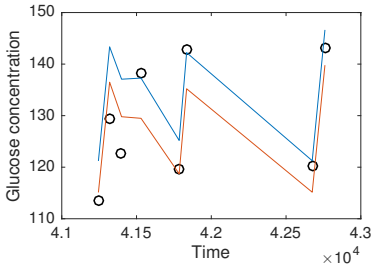

(c) More enlarged view on a section of (a)

Figure 10 – Glucose Modelling Inverse Problem: Data and evaluations of EnKF $\mathcal{G}(u^{\mathrm{EnKF},J=36,N=15})$ and MCMC $\mathcal{G}(u^{\mathrm{MCMC}})$estimates

Figure 10 presents the data and $\mathcal{G}$ evaluations of the best MCMC estimate and the best

EnKF estimate (best refers to the relative error of the estimated parameter u) of the simulations whose results are given in Table 4. Considering the noise, both evaluations restore the shape of the data well. However, the standard deviation of the relative error of the parameter estimated by both methods, is substantially higher than the standard deviation of the corresponding loss functionals. This implies, that the corresponding classical inverse problem might have several local minima, respectively the posterior distribution is multi-modal. Figure 11(a) shows the seven Monte Carlo Markov Chains, none of which seem to converge to their stationary distribution. The paths appear to be stuck in at least two different areas/modes - the yellow and dark(er) blue path are close to the true parameter. Comparing this with Figure 11(b) which presents (the updates of) the ensembles of seven independent EnKF simulations. The particles of each EnKF simulation cover the presented area of the posterior distribution's support much better than any of the paths of the MCMC. The combination of all the EnKF particles presented in this Figure would cover the area perfectly, which indicates that a higher amount of particles would improve the situation even more.

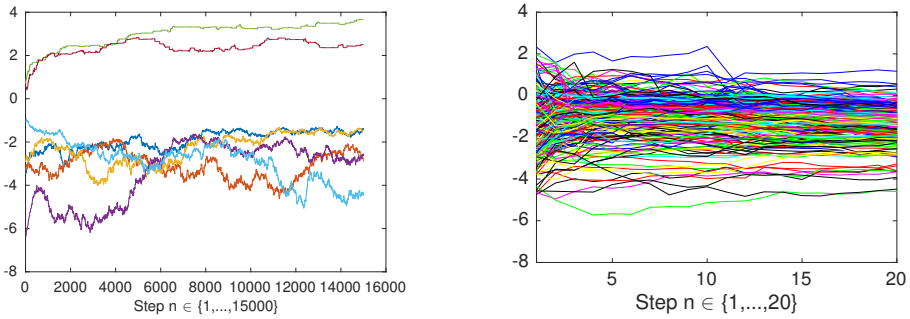

(a) Paths of the seven MCMC simulations presented in Table 4

(b) Particles of seven EnKF estimation given $J = 48, N = 20$

Figure 11 – Glucose Modelling Inverse Problem: EnKF and MCMC particles - each of them projected onto the first coordinate

Therefore the Ensemble Kalman Filter with sufficiently high J and N (cf. Table 4) is able to handle this multi-modal case better, than the MCMC method does.

6. Conclusions

Bayes Linear statistics is a huge class of methods that can be applied successfully in the framework of Bayesian inverse problems. The mathematical foundations for such applications are presented in this report (cf. Chapter 2).

The BLE is optimal, if $\mathcal{G}$ is linear, and thus, better than the EnKF in this case. However,

the EnKF provided better estimates than the BLE in the non-linear problems that are numerically examined in this report. The latter results are independent of the derivation of the BLE, which is equally inadequate in both analytically and simulation based variations. Allowing the same or a greater number of $\mathcal{G}$-evaluations as in a comparable EnKF estimation does not improve the estimation results of the BLE significantly. (cf. Subsections 5.2 and 5.3).

Numerical experiments have also shown that the EnKF is a legitimate alternative to a Monte Carlo or Monte Carlo Markov Chain simulation. The EnKF solution of the cubic inverse problem even appeared to converge to the mean estimate provided by (fully Bayesian) self-normalised importance sampling as $J, N \to \infty$. As mentioned above, it cannot be used to approximate the posterior distribution, but it is a computationally cheap alternative to estimate the posterior mean which can be applicable even in cases where a comparable MCMC method is problematic due to a multi modal posterior distribution. (cf. Subsections 5.2 and 5.4)

Once it is calibrated correctly the Bayes Linear (backtracking) line search can improve the estimation results of the Ensemble Kalman Filter. The correct calibration is crucial here and notably expensive. Furthermore, backtracking itself is a relatively expensive method to determine step lengths, since it requires a great number of $\mathcal{G}$-evaluations. The results of the numerical experiments given here neither confirm, nor discourage against using it in practice. (cf. Subsections 5.1, 5.2 and 5.3)

6.1. Future Research

This report examines only one Bayes Linear method in the inverse problems framework: the conditional mean estimate. Applying further Bayes Linear methods, such as Bayes Linear sufficiency, the Bayes Linear covariance estimator or some of the other methods presented in [21], could also improve upon the Ensemble Kalman Filter or other methods.

Deriving the simulation based Bayes Linear estimator using K samples is equivalent to some noise-free Ensemble Kalman Filter given $J = K$ ensemble members and $N = 1$ step. This is typically worse than most other EnKF estimations given combinations of J, N with $J \cdot N = K$. Similarly, one expects that an EnKF simulation with only one particle and K steps would lead to a poor approximation. These facts raise the question whether there is an optimal strategy to choose the proportion of N and J, given a maximal amount of $\mathcal{G}$-evaluations $K = N \cdot J$.

Even though the EnKF returns good posterior mean estimates, the posterior distribution approximately given by the empirical distribution of the particles is useless. However, having a posterior distribution is crucial in Bayesian inverse problems, which implies to consider strategies to improve this poor approximation in future. (cf. various approaches in [37])

The Bayes Linear backtracking condition is only used here to adjust the step length. Can it also be used to adjust the amount of particles J or the spatial mesh width of the $\mathcal{G}$-evaluations (in a Multi-level manner)? Are there alternative conditions that one could use for those other adjustments or for instance to increase the step length rather than decreasing it?

The step lengths suggested by the Bayes Linear line search method are roughly equivalent throughout all numerical experiments. Instead of calibrating a relatively expensive backtracking method, one could try to construct an EnKF with multiple step lengths that are chosen in a predetermined way, ascendingly and independently of the particles behaviour at runtime.

References

[1] S. Agapiou, O. Papaspiliopoulos, D. Sanz-Alonso, and A. M. Stuart. Importance sampling: computational complexity and intrinsic dimension. Submitted.

[2] R.B. Ash and C. Doléans-Dade. *Probability and Measure Theory*. Harcourt/Academic Press, 2000.

[3] P. Billingsley. *Probability and Measure*. Wiley Series in Probability and Statistics. Wiley, 2012.

[4] A. Bobrowski. *Functional Analysis for Probability and Stochastic Processes*. Cambridge University Press, 2005. Cambridge Books Online.

[5] S. Bochner. Integration von Funktionen, deren Werte die Elemente eines Vektorraumes sind. *Fundamenta Mathematicae*, 20(1):262–176, 1933.

[6] V.I. Bogachev. *Gaussian Measures*. Mathematical Surveys and Monographs. American Mathematical Soc., 1998.

[7] D. Bosq. *Linear Processes in Function Spaces: Theory and Applications*. Lecture Notes in Statistics. Springer New York, 2000.

[8] H. Bühlmann and A. Gisler. *A Course in Credibility Theory and its Applications*. Universitext (1979). Springer, 2005.

[9] J.A. Cumming and M. Goldstein. Bayes linear uncertainty analysis for oil reservoirs based on multiscale computer experiments. *The Oxford handbook of applied Bayesian analysis. Oxford University Press, Oxford*, pages 241–270, 2010.

[10] M. Dashti and A.M. Stuart. The Bayesian Approach to Inverse Problems. In R. Ghanem, D. Higdon, and H. Owhadi, editors, *Handbook of Uncertainty Quantification*. Springer.

[11] M. Dashti and A.M. Stuart. Uncertainty Quantification and Weak Approximation of an Elliptic Inverse Problem. *SIAM J. Numerical Analysis*, 49:2524–2542, 2011.

[12] J. Diestel and J.J. Uhl. *Vector Measures*. Mathematical surveys and monographs. American Mathematical Society, 1977.

[13] A. Doucet and A.M. Johansen. A tutorial on particle filtering and smoothing: Fifteen

years later. In *Handbook of Nonlinear Filtering*. Oxford, UK: Oxford University Press, 2011.

[14] O.G. Ernst, B. Sprungk, and H.-J. Starkloff. Analysis of the Ensemble and Polynomial Chaos Kalman Filters in Bayesian Inverse Problems. *ArXiv e-prints*, April 2015.

[15] G. Evensen. The Ensemble Kalman Filter: theoretical formulation and practical implementation. *Ocean Dynamics*, 53(4):343–367, 2003.

[16] G. Evensen. *Data Assimilation: The Ensemble Kalman Filter*. Earth and Environmental Science. Springer Berlin Heidelberg, 2009.

[17] J.N. Franklin. Well-posed stochastic extensions of ill-posed linear problems. *Journal of Mathematical Analysis and Applications*, 31(3):682 – 716, 1970.

[18] M. Goldstein. A Note on Some Bayesian Nonparametric Estimates. *Ann. Statist.*, 3(3):736–740, 05 1975.

[19] M. Goldstein. Approximate Bayes Solutions to Some Nonparametric Problems. *Ann. Statist.*, 3(2):512–517, 03 1975.

[20] M. Goldstein and J. Rougier. Bayes Linear Calibrated Prediction for Complex Systems. *Journal of the American Statistical Association*, 101(475):1132–1143, 2006.

[21] M. Goldstein and D. Wooff. *Bayes linear statistics, theory and methods*, volume 716. John Wiley & Sons, 2007.

[22] J. Hadamard. Sur les problèmes aux dérivés partielles et leur signification physique. *Princeton University Bulletin*, 13:49–52, 1902.

[23] M. Hanke. A regularizing Levenberg - Marquardt scheme, with applications to inverse groundwater filtration problems. *Inverse Problems*, 13(1):79, 1997.

[24] W.K. Härdle and L. Simar. *Applied Multivariate Statistical Analysis*. Springer Berlin Heidelberg, 2007.

[25] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, second edition, 2012. Cambridge Books Online.

[26] M.A. Iglesias. Iterative regularization for ensemble data assimilation in reservoir models. *Computational Geosciences*, 19(1):177–212, 2015.

[27] M.A. Iglesias. A Regularizing Iterative Ensemble Kalman Method for PDE-constrained Inverse Problems. *Inverse Problems*, 32(2):025002, 2016.

[28] M.A. Iglesias, K.J.H. Law, and A.M. Stuart. The Ensemble Kalman filter for Inverse Problems. *Inverse Problems*, 29(4):045001, 2013.

[29] L. Isserlis. On a Formula for the Product-Moment Coefficient of any Order of a Normal Frequency Distribution in any Number of Variables. *Biometrika*, 12(1/2):134–139, 1918.

[30] N. Kantas, A. Beskos, and A. Jasra. Sequential Monte Carlo methods for high-dimensional inverse problems: a case study for the Navier–Stokes equations. 2(1):464–489, 2014.

[31] J. Keener and J. Sneyd. *Mathematical Physiology*. Springer, 2009.

[32] P. Krug. The conditional expectation as estimator of normally distributed random variables with values in infinitely dimensional Banach spaces. *Journal of Multivariate Analysis*, 38(1):1 – 14, 1991.

[33] K.J.H. Law, A.M. Stuart, and K.C. Zygalakis. *Data Assimilation: A Mathematical Introduction*. Springer, 2015.

[34] M.S. Lehtinen, L. Paivarinta, and E. Somersalo. Linear inverse problems for generalised random variables. *Inverse Problems*, 5(4):599, 1989.

[35] G. J. Lord, C. E. Powell, and T. Shardlow. *An Introduction to Computational Stochastic PDEs*. Cambridge University Press, 2014. Cambridge Books Online.

[36] A. Mandelbaum. Linear estimators and measurable linear transformations on a Hilbert space. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 65(3):385–397, 1984.

[37] H.G. Matthies, E. Zander, B.V. Rosić, A. Litvinenko, and O. Pajonk. *Inverse Problems in a Bayesian Setting*, pages 245–286. Springer International Publishing, Cham, 2016.

[38] J. Nocedal and S. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer New York, 2006.

[39] R. Penrose. On best approximate solutions of linear matrix equations. *Mathematical Proceedings of the Cambridge Philosophical Society*, 52:17–19, 1 1956.

[40] J. Peypouquet. *Convex Optimization in Normed Spaces: Theory, Methods and Examples*. SpringerBriefs in Optimization. Springer International Publishing, 2015.

[41] A. Quarteroni, R. Sacco, and F. Saleri. *Numerical Mathematics*. Texts in Applied Mathematics. Springer, Paris, FR, second edition, 2007. Written for graduate students, researchers.

[42] C. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer New York, 2005.

[43] C. Schillings and A.M. Stuart. Analysis of the ensemble Kalman filter for inverse problems. *ArXiv e-prints*, February 2016.

[44] J. Shao. *Mathematical Statistics*. Springer Texts in Statistics. Springer, 2003.

[45] M. F. Smiley. Inequalities related to Lidskii's. *Proceedings of the American Mathematical Society*, 19(5):1029–1034, 1968.

[46] A.M. Stuart. Inverse problems: a Bayesian perspective. In *Acta Numerica 2010*, volume 19, pages 451–559.

[47] J. Sturis, K. S. Polonsky, E. Mosekilde, and E. Van Cauter. Computer model for mechanisms underlying ultradian oscillations of insulin and glucose. *American Journal of Physiology - Endocrinology and Metabolism*, 260(5):E801–E809, 1991.

[48] T.J. Sullivan. *Introduction to Uncertainty Quantification*, volume 63. Springer, 2015.

[49] I. Vernon, M. Goldstein, and R.G. Bower. Galaxy formation: a Bayesian uncertainty analysis. *Bayesian Anal.*, 5(4):619–669, 12 2010.

[50] P. Wolfe. Convergence Conditions for Ascent Methods. *SIAM Review*, 11(2):226–235, 1969.

[51] P. Wolfe. Convergence Conditions for Ascent Methods. II: Some Corrections. *SIAM Review*, 13(2):185–188, 1971.

A. Appendix: Optimisation

Some fundamental definitions and results from the theory of non-linear optimisation are summarised in this section. The stated results are based on [38], [40] and [48]. This appendix does not contain any proofs.

(A.1) Definition. Let $f : X \to \mathbb{R}$ be a function.

- x^* is a *global minimum* of f, if for all $x \in X \backslash \{x^*\}$: $f(x) \geq f(x^*)$ and *strict global minimum* if the inequality holds strictly,
- The *set of minimisers* is denoted by $\operatorname{argmin}_{x \in X} f(x) := \{x^* \in X : f(x^*) \leq f(x)\ (x \in X)\}$ and $\min_{x \in X} f(x) := f(x^*)$, if $x^* \in \operatorname{argmin}_{x \in X} f(x)$,
- $x^* \in X$ is a *(strict) local minimum* if a neighbourhood $B \subseteq X$ of x^* exists, such that x^* is a (strict) global minimum of $f|_B$,
- $x^* \in X$ is a *(strict) global maximum* of f, if x^* is a (strict) global minimum of $-f$. The other definitions apply for maximum analogously. •

The objective of optimisation is to find a minimum of a function. Finding a minimum numerically is well-posed, if that function has a local minimum and is convex.

(A.2) Definition. A subset V of some $\mathbb{R}$- or $\mathbb{C}$-vector space is *convex*, if for all $t \in (0, 1)$ and $u, v \in V$ their linear combination $tu + (1 - t)v \in V$.

A function $f : V \to \mathbb{R}$ is *convex*, if V is convex and for all $t \in (0, 1)$ and $u, v \in V$ the inequality

$$tf(u) + (1 - t)f(v) \geq f(tu + (1 - t)v)$$

holds. Furthermore, f is *strictly convex*, if f is convex and the inequality holds strictly. •

(A.3) Theorem. Let $f : V \to \mathbb{R}$ be a convex function. Then, statements 1-3 hold.

1. Let x^* be a local minimum. Then, x^* is also a global minimum.
2. Let f be strictly convex and x^* be a local minimum. Then, x^* is the unique global minimum. •

Proof. [48, Theorem 4.19] □

A.1. (Linear) Least Squares Problems

Let $\mathcal{X}, \mathcal{Y}$ be Hilbert spaces, $b \in \mathcal{Y}$ and $\mathcal{A} : \mathcal{X} \to \mathcal{Y}$ be a linear Operator. Furthermore, assume that $\mathrm{im}(\mathcal{A})$ is a closed subspace of $\mathcal{Y}$. A *linear least squares problem* is given by

$$\min_{x \in \mathcal{X}} \|\mathcal{A}x - b\|_{\mathcal{Y}}^2. \tag{A.4}$$

Let $V \subseteq \mathcal{X}$ be a closed subspace of $\mathcal{X}$ and $b \in \mathcal{X}$. Consider the following minimisation problem:

$$\min_{x \in V} \|x - b\|_{\mathcal{X}}^2. \tag{A.5}$$

(A.6) Lemma. $\hat{x} \in V$ is the unique minimiser of minimisation problem (A.5) if and only if $(\hat{x} - b) \perp V$. Thus, the map $\mathcal{X} \ni b \mapsto \hat{x} \in V$ $(\{\hat{x}\} = \mathrm{argmin}_{v \in V} \|v - b\|^2)$ is the *orthogonal projection* from $\mathcal{X}$ to V. •

Proof. [48, Lemma 4.25, Lemma 4.26] □

(A.7) Theorem. $\hat{x} \in \mathcal{X}$ is a minimiser of minimisation problem (A.4) if and only if $\mathcal{A}^*\mathcal{A}\hat{x} = \mathcal{A}^* b$. •

Proof. [48, Lemma 4.27] □

B. Appendix: Probability and Measure Theory

This section contains chosen standard theorems and notations from measure and probability theory the report refers to. An elementary introduction to measure and probability theory is given in several books, such as [2] or [3]. This appendix is based on [2], [3], [5], [6], [12], and [32] and does not contain any proofs.

B.1. Foundations

(B.1) Definition. 1. For a given family of sets $\mathcal{E}$, the *generated σ-algebra* is defined by $\sigma(\mathcal{E}) := \bigcap\{\mathcal{F} : \mathcal{F}\ \sigma\text{-algebra} : \mathcal{F} \supseteq \mathcal{E}\}$.

2. Let $(T, \mathcal{T})$ be a topological space. The *Borel-σ-algebra* on $(T, \mathcal{T})$ is defined by $\mathcal{B}(T, \mathcal{T}) := \sigma(\mathcal{T})$. If there is some standard topology $\mathcal{T}$ associated with T or

$\mathcal{T}$ induced by a given metric or norm, $\mathcal{B}T := \mathcal{B}(T) := \mathcal{B}(T, \mathcal{T})$. Furthermore, $\mathcal{B}^k := \mathcal{B}(\mathbb{R}^k)$ and $\mathcal{B} := \mathcal{B}^1$.

3. Let $(\Theta, \mathcal{F}), (\Theta', \mathcal{F}')$ be measurable spaces. A function $f : \Theta \to \Theta'$ is called $\mathcal{F}$-$\mathcal{F}'$-*measurable*, if $f^{-1}[F'] \in \mathcal{F}$ for any $F' \in \mathcal{F}'$ and denoted by $f : (\Theta, \mathcal{F}) \to (\Theta', \mathcal{F}')$. If $(\Theta', \mathcal{F}') = (\mathbb{R}, \mathcal{B})$, the function f is called *measurable numerical* or just *measurable*. Denote $\mathcal{M} := \{f : \Theta \to \mathbb{R} : f : (\Theta, \mathcal{F}) \to (\mathbb{R}, \mathcal{B})\}$ the set of all measurable numerical functions given some measurable space and $\mathcal{M}_+ := \{f : \Theta \to \mathbb{R} : f : (\Theta, \mathcal{F}) \to ([0, \infty), \mathcal{B}[0, \infty))\}$.

4. Let Θ be a set, $(\Theta', \mathcal{F}')$ be a measurable space and $f : \Theta \to \Theta'$ be a function. The *initial σ-algebra* of f is defined by

$$\sigma(f) := \{f^{-1}[F'] : F' \in \mathcal{F}'\}.$$

By definition, $f : (\Theta, \sigma(f)) \to (\Theta', \mathcal{F}')$. •

(B.2) Definition. 1. Let $(\Theta, \mathcal{F}, \mu)$ be a measure space, $(\Theta', \mathcal{F}')$ be a measurable space and $f : (\Theta, \mathcal{F}, \mu) \to (\Theta', \mathcal{F}')$ be a measurable function. The *pushforward measure* on $(\Theta', \mathcal{F}')$ is given by $f^{\#}\mu : \mathcal{F}' \to [0, \infty], F' \mapsto \mu(f \in F')$. It is welldefined as a measure. Furthermore, if μ is a probability measure, $f^{\#}\mu$ is also a probability measure. In the latter case, f is called *(Θ'-valued) random variable* and $f^{\#}\mu$ is the *probability distribution of f*.

2. Let $(\Theta, \mathcal{F}, \mu)$ be a measure space. It is called *finite*, if $\mu(\Theta) < \infty$, and *σ-finite*, if a countable partition $\mathcal{P} \subseteq \mathcal{F}$ of Θ exists, such that $\mu(P) < \infty$ $(P \in \mathcal{P})$. •

B.2. Integration

(B.3) Definition (Bochner-integral of simple functions)**.** Let $(\Theta, \mathcal{F}, \mu)$ be a finite measure space and X be a Banach space. A function $g : (\Theta, \mathcal{F}) \to (X, \mathcal{B}X)$ is *simple*, if a finite Partition $\mathcal{P} \subseteq \mathcal{F}$ of Θ and a vector $(x_P : P \in \mathcal{P}) \in X^{\mathcal{P}}$ exist, such that $g(\theta) = \sum_{P \in \mathcal{P}} x_P \mathbb{1}_P(\theta)$ $(\theta \in \Theta)$. The *(Bochner-)integral* of g is defined by

$$\int g \mathrm{d}\mu := \sum_{P \in \mathcal{P}} x_P \mu(P),$$

if $\sum_{P \in \mathcal{P}} \|x_P\|_X \mu(P) < \infty$.

A function $f : \Theta \to X$ is *μ-measurable* (and this is denoted by $f : (\Theta, \mathcal{F}) \to X$), if a

sequence of simple functions $(f_n : n \in \mathbb{N})$ exists, such that $\lim_{n\to\infty} \|f_n - f\|_X = 0$ (μ-almost everywhere). •

(B.4) Lemma. Let $(\Theta, \mathcal{F}, \mu)$ be a finite measure space, $f : (\Theta, \mathcal{F}) \to ([0,\infty), \mathcal{B}[0,\infty))$ be a measurable numerical function. Then, one can find an isotonic sequence of simple functions $f_n : (\Theta, \mathcal{F}) \to ([0,\infty), \mathcal{B}[0,\infty))$ $(n \in \mathbb{N})$, such that $f_n \uparrow f$ as $n \to \infty$. In particular, positive measurable numerical functions are μ-measurable. •

Proof. [3, p. 195, Theorem 13.5]. □

The definition of the Bochner-integral requires a condition containing the μ-integral of some positive measurable numerical function.

(B.5) Definition (μ-integral)**.** Let $(\Theta, \mathcal{F}, \mu)$ be a measure space, $f \in \mathcal{M}_+$ and $(f_n : n \in \mathbb{N}) \in \mathcal{M}_+^{\mathbb{N}}$ be an isotonic sequence of simple functions, such that $f_n \uparrow f$ (μ-a.e.) as $n \to \infty$ (well-defined, cf. Lemma (B.4)). The *μ-integral* of f is defined by

$$\lim_{n\to\infty} \int f_n \mathrm{d}\mu =: I(f, \mu). \qquad \bullet$$

(B.6) Definition (Bochner-integral)**.** Let $(\Theta, \mathcal{F}, \mu)$ be a finite measure space and X a Banach space, $f : (\Theta, \mathcal{F}) \to X$ be a μ-measurable function and $(f_n)_{n\in\mathbb{N}}$ be the sequence of simple functions, with $\lim_{n\to\infty} f_n = f$ μ-a.e.. f is *(Bochner-)integrable*, if

$$\lim_{n\to\infty} I(\|f - f_n\|_X, \mu) = 0.$$

Given that case, the *(Bochner-)integral* of f is

$$\lim_{n\to\infty} \int f_n \mathrm{d}\mu = \int f \mathrm{d}\mu.$$

If $(\Theta, \mathcal{F}, \mu) = (\mathbb{R}^k, \mathcal{B}^k, \mathrm{Leb}(k))$, the integral of f is called *Lebesgue-integral* and denoted by

$$\int f(\theta) \mathrm{d}\theta := \int f \mathrm{d}\mu. \qquad \bullet$$

(B.7) Definition. Let $(\Theta, \mathcal{F}, \mu)$ be a measure space and $f \in \mathcal{M}_+$.

1. The measure $\zeta : \mathcal{F} \to [0,\infty], A \mapsto \int \mathbb{1}_A f \mathrm{d}\mu$ has *μ-density f* (or *density f*, if $\mu = \mathrm{Leb}(k)$, $k \in \mathbb{N}$) and is denoted by $f\mu := \zeta$.

2. Let ζ be a further measure on $(\Theta, \mathcal{F})$. ζ is called *absolutely μ-continuous* (or *absolutely continuous*, if $\mu = \text{Leb}(k)$, $k \in \mathbb{N}$), if $\mu(F) = 0$ implies $\zeta(F) = 0$, for any $F \in \mathcal{F}$. This is denoted by $\zeta \ll \mu$. •

(B.8) Theorem (Radon-Nikodym)**.** Let $(\Theta, \mathcal{F})$ be a measurable space and μ, ζ measures on it. Furthermore, $(\Theta, \mathcal{F}, \mu)$ and $(\Theta, \mathcal{F}, \zeta)$ are σ-finite. 1-2 are equivalent.

1. $\zeta \ll \mu$,

2. a μ-a.e. unique function $f \in \mathcal{M}_+$ exists, such that $\zeta = f\mu$. Define $\frac{\mathrm{d}\zeta}{\mathrm{d}\mu} := f$. •

Proof. [3, p. 449, Theorem 32.2]. □

(B.9) Theorem. Let $(\Theta, \mathcal{F}, \mu)$ be a finite measure space, $(\Theta', \mathcal{F}')$ be another measurable space and X be a Banach space, $f : (\Theta', \mathcal{F}') \to X$ be an integrable function and $T : (\Theta, \mathcal{F}) \to (\Theta', \mathcal{F}')$ be a measurable function. f is integrable with respect to $T^{\#}\mu$, if and only if $f \circ T$ is integrable with respect to μ. In that case, the following equation holds for any $F' \in \mathcal{F}'$,

$$\int_{\{T \in F'\}} f \circ T \mathrm{d}\mu = \int_{F'} f \mathrm{d}(T^{\#}\mu) \quad \text{(B.10)} \quad \bullet$$

Proof. [3, p. 229, Theorem 16.13] for a version, where $X = \mathbb{R}$. [32, p. 4, Lemma 4] for the general theorem. □