Bayesian modeling and uncertainty quantification with applications in computed tomography and hemodialysis

Bangsgaard, Katrine Ottesen

Publication date:
2023

Document Version
Publisher's PDF, also known as Version of record

Ph.D. Thesis
Doctor of Philosophy

DTU Compute
Department of Applied Mathematics and Computer Science

Bayesian modeling and uncertainty quantification with applications in computed tomography and hemodialysis

Katrine Ottesen Bangsgaard

Kongens Lyngby 2023

DTU Compute
Department of Applied Mathematics and Computer Science
Technical University of Denmark

Matematiktorvet
Building 303B
2800 Kongens Lyngby, Denmark
Phone +45 4525 3031
compute@compute.dtu.dk
www.compute.dtu.dk

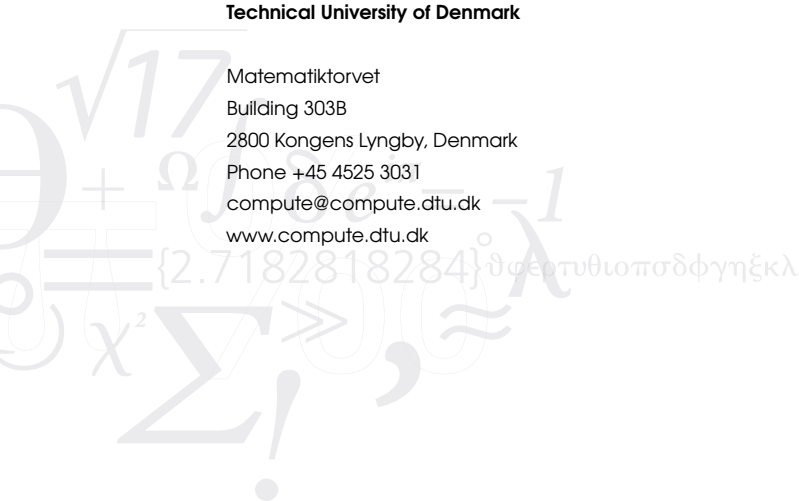

Summary (English)

In the healthcare system, a goal is to provide effective diagnosis and treatment to patients. However, gaining a comprehensive understanding of the underlying mechanisms within the human body is often necessary to achieve this objective. Often, we cannot not directly observe these mechanisms, since it is not feasible and/or practical to gain access directly. Instead, we can describe the mechanisms using mathematical models that are characterized by a set of parameters. The mathematical models relate the parameters of interest to available non-invasive measurements, and we can use the parameters to, e.g., predict the response to treatment or assess the severity of diseases. In this thesis we explore the two applications:

- Computed tomography: A non-invasive imaging technique.
- Hemodialysis: A treatment for patients with kidney failure.

The aim of this thesis is to obtain a better understanding of the underlying mechanisms for the two applications such that we for instance might enable tailor patient-specific treatments in hemodialysis and improve the quality of low dose experiments in computed tomography. We employ Bayesian modeling and uncertainty quantification to formulate mathematical models for computed tomography and hemodialysis that allow us to estimate the parameters of interest along with their statistical properties such as mean, covariance and credible intervals.

In computed tomography, radiation is emitted from a source which is partially absorbed as it travels through the object of interest, e.g., the human body. Based on the absorption, images of the interior can be reconstructed. To obtain these images, it is necessary to know the intensity of the source which is typically estimated by flat-field measurements, i.e., measurements without an object in the scanner. However, for low dose and/or time-limited experiments, discrepancies between the true and estimated source can introduce systematic model errors which appear as concentric rings in the image, also known as ring artifacts. We employ modeling techniques for mitigating ring artifacts caused by model errors, and the key finding is that, by using explicit modeling of the source, it is possible to eliminate ring artifacts from the images.

We have also considered ring reduction for spectral computed tomography. In spectral computed tomography, the attenuation is measured at multiple energies, which enables the reconstruction of images at individual energy levels, which is referred to as a spectral reconstruction. The spectral reconstruction suffers from ring artifacts, and we propose an extended flat-field model for mitigating ring artifacts which exploits high correlation across energy channels in the spectral flat-fields.

Patients with kidney failure often suffer from hyperphosphatemia which is associated with increased vascular calcification and mortality. Thus, it is of great importance to regulate the phosphate level of the kidney-failure patients. Hemodialysis is a method for removing phosphate from the blood of the patient. The patient is connected to a dialyzer where blood and dialysate fluid are separated by a semipermeable membrane allowing phosphate to move from the blood of the patient to the dialysate fluid. We consider the phosphate removal during hemodialysis in a Bayesian framework. We compare two types of hemodialysis treatments, i.e., single-pass and multiple-pass, and find that the uncertainty of the parameters estimated based on the single-pass model is greater than those estimated based on the multiple-pass model. Moreover, a key finding is that the uncertainty of the parameter estimates is greatly reduced by measuring the patients for consecutive treatments whereas measurements in the relapse phase has limited effect on the precision of the parameter estimates.

This thesis contributes to an enhanced understanding of artifacts in computed tomography and phosphate removal during hemodialysis using Bayesian modeling and uncertainty quantification.

Summary (Danish)

Et mål i sundhedsvæsnet er at tilbyde effektiv diagnosering og behandling til patienter. For at tilbyde dette er det ofte nødvendigt at opbygge en omfattende forståelse af de underliggende mekanismer i menneskekroppen. Ofte er det ikke muligt at tilgå disse mekanismer direkte og i stedet må vi beskrive mekanismerne ved matematiske modeller som er karakteriserede ved et sæt af parametre. De matematiske modeller relaterer parametrene til tilgængelige ikke-invasive målinger, og vi kan bruge parametrene til at f.eks. forudsige hvordan patienter vil respondere på behandling eller vurdere sygdomsprogressionen. Vi vil I denne afhandling undersøge de to applikationer:

- Computer tomografi: En ikke-invasiv scanningsteknik.
- Hæmodialyse: En behandling for patienter med nyresvigt.

Målet for denne afhandling er at opnå en bedre forståelse af de underliggende mekanismer, så vi f.eks. kan muliggøre patient specifik behandling i hæmodialyse og forbedre kvaliteten af scanninger ved lav stråling i computer tomografi. Vi benytter Bayesiansk modellering og uncertainty quantification til at formulere matematiske modeller for computer tomografi og hæmodialyse, hvorved vi kan estimere de relevante parametre sammen med deres statistiske egenskaber som middelværdi, covarians og credible intervaller.

I computer tomografi, udsendes stråling fra en kilde som bliver delvist absorberet når det bevæger sig gennem objekter, f.eks. en menneskekrop. Baseret på absorptionen af strålingen, kan man rekonstruere billeder af det indre af objektet. For at kunne beregne en rekonstruktion, er det nødvendigt at kende intensiteten af kildens stråling. Typisk estimeres kildens intensitet ud fra flat-field målinger, som er målinger uden et objekt i scanneren. For eksperimenter men lav stråling og/eller tidsbegrænsning er signal-støjforholdet lavt og der kan forekomme store forskelle mellem den faktiske kilde intensitet og den estimerede. Denne afvigelse kan give systematiske modelfejl og resultere i ring artefakter som er koncentriske ringe i rekonstruktionen. Vi anvender modelleringsteknikker til at dæmpe ring artefakter og hovedresultatet er at vi, ved eksplicit modellering af kilden, kan eliminere ring artefakter i rekonstruktionerne.

Vi har også betragtet ring reduktion for spektral computer tomografi. I spektral computer tomografi måles absorptionen ved forskellige energier,

og man kan rekonstruere billeder for hver energi hvilket også kaldes for en spektral rekonstruktion. Den spektrale rekonstruktion lider ofte af et lavt signal-støjforhold og ring artefakter. Vi har formuleret en udvidet flat-field metode som udnytter høj korrelation på tværs af de målte energier til at bekæmpe ring artefakter i spektral computer tomografi.

Patienter med nyresvigt lider ofte af hyperfosfatæmi, som er relateret til øget åreforkalkning og dødelighed. Det er derfor vigtigt at hjælpe nyresvigt patienter med at regulere deres fosfatniveauet i blodet. Hæmodialyse er en behandling hvorved fosfat fjernes fra blodet. Patienten kobles til en dialysator hvor blod og dialysevæske er adskilt af en semipermeable membran som tillader diffusion af fosfat fra blodet til dialysevæsken. Vi benytter en Bayesiansk tilgang til at undersøge fjernelse af fosfat ved hæmodialyse. Vi betragter og sammenligner to typer af behandling, single-pass og multiple-pass, og vores undersøgelser viser, at usikkerhederne knyttet til de estimerede parametre er større for single-pass end for multiple-pass. Et hovedresultat er at vores undersøgelser indikerer at vi kan sænke usikkerhederne forbundet med parametrene ved at tage målinger af patienten ved på hinanden følgende behandlinger mens at målinger i tilbagefaldsfasen, dvs. efter afsluttet behandling, har en begrænset effekt på præcisionen af parameter estimaterne.

Denne afhandling bidrager til en øget forståelse af ring artefakter i computer tomografi og fjernelse af fosfat under hæmodialyse ved hjælp af Bayesiansk modellering og uncertainty quantification.

Preface

This PhD thesis was prepared at the department of Applied Mathematics and Computer Science at the Technical University of Denmark in fulfillment of the requirements for acquiring a PhD degree. The research covered in this thesis was performed between September 1st 2019 and July 7th 2023 in the Section for Scientific Computing under principal supervisor Associate Professor Martin S Andersen and co-supervisors Senior Researcher Jakob S. Jørgensen, and Professor Per Christian Hansen. Part of the research was carried out during an external stay visiting Professor Johnny T. Ottesen and Associate Professor Morten Andersen at Roskilde University, Denmark. This work was supported by The Villum Foundation (grant no. 25893).

This thesis covers research performed during the PhD program related to modeling, optimization, and uncertainty quantification for inverse problems. The aim of this thesis is to summarize and present the work of three papers written during the PhD program and to provide an overview of the field.

Kongens Lyngby, July 7, 2023

Katrine Ottesen Bangsgaard

List of Papers

The following lists of papers are the scientific outcome of the research carried out during the PhD program. The thesis summarizes the work of Papers A, B and C with the initial chapters providing background knowledge and an overview of the field. Papers D and E are not directly included but briefly discussed and listed for completeness.

Papers included in thesis

A) **Katrine O. Bangsgaard** and Martin S. Andersen, "A statistical reconstruction model for absorption CT with source uncertainty", *Inverse Problems*, 2021, [3].

B) **Katrine O. Bangsgaard**, Genoveva Burca, Evelina Ametova, Martin S. Andersen and Jakob S. Jørgensen, "Low-rank flat-field correction for artifact reduction in spectral computed tomography", *Applied Mathematics in Science and Engineering*, 2023, [6].

C) **Katrine O. Bangsgaard**, Morten Andersen, James G. Heaf and Johnny T. Ottesen, "Bayesian parameter estimation for phosphate dynamics during hemodialysis", *Mathematical Biosciences and Engineering*, 2022, [4].

Papers not included in thesis

D) **Katrine O. Bangsgaard**, Morten Andersen, Vibe Skov, Lasse Kjær, Hans C Hasselbalch and Johnny T. Ottesen, "Dynamics of competing heterogeneous clones in blood cancers explains multiple observations-a mathematical modeling approach" *Mathematical Biosciences and Engineering*, 2020, [5].

E) Morten Andersen, **Katrine O. Bangsgaard**, James G. Heaf and Johnny T. Ottesen "Analytical solution of phosphate kinetics for hemodialysis", *Journal of Mathematical Biology*, 2023, [1].

Acknowledgments

I would like to thank my supervisors Martin S. Andersen, Jakob S. Jørgensen and Per Christian Hansen for their support and valuable guidance. I thank Johnny T. Ottesen and Morten Andersen for a fruitful and inspiring external stay.

My PhD was heavily affected by the world-wide Coronavirus pandemic, and I would like to thank my fellow PhD colleagues, and especially my office mate Hjørdis Amanda Schlüter, who I have shared both highs and lows of this journey with. Lastly, I would like to thank my friends and family for support during the completion of this project.

Contents

Summary (English) i

Summary (Danish) iii

Preface v

List of Papers vii

Acknowledgments ix

Contents xi

1 Introduction 1
- 1.1 Motivation and applications . . . 1
- 1.2 Inverse problems . . . 3
- 1.3 Uncertainty quantification . . . 3
- 1.4 Contribution . . . 4
- 1.5 Structure of the thesis . . . 4

2 Computed tomography 5
- 2.1 Experimental set-up . . . 5
- 2.2 The physical model . . . 6
- 2.3 Filtered Back Projection . . . 7
- 2.4 Discretized model . . . 8
- 2.5 Ring artifacts . . . 10
- 2.6 Spectral model . . . 12

3 Phosphate kinetics 15
- 3.1 Phosphate's role in the body . . . 15
- 3.2 Hemodialysis . . . 16
- 3.3 Compartment modeling . . . 17
- 3.4 Parameter estimation . . . 20

4 Discrete inverse problems 23
- 4.1 Classification of inverse problems . . . 23
- 4.2 Linear inverse problems . . . 24

4.3 Regularization 26

5 Uncertainty quantification **31**
5.1 Bayesian approach 31
5.2 Sampling 35
5.3 Metropolis Hastings 37
5.4 Gibbs sampler 39
5.5 Hamiltonian Monte Carlo 42
5.6 Convergence diagnostics 46

6 Computed tomography with uncertain source **49**
6.1 Motivation and goals 49
6.2 Contributions 50
6.3 Methods 50
6.4 Results 51
6.5 Summary 56

7 Ring artifacts in spectral computed tomography **57**
7.1 Motivation and goals 57
7.2 Contributions 58
7.3 Methods 58
7.4 Results 59
7.5 Summary 60

8 Hemodialysis modeling and estimation **61**
8.1 Motivation and goals 61
8.2 Contributions 62
8.3 Methods 62
8.4 Results 63
8.5 Summary 64

9 Discussion and Conclusion **67**
9.1 Discussion 67
9.2 Concluding remarks 68

A A statistical reconstruction model for absorption CT with source uncertainty **71**

B Low-rank flat-field correction for artifact reduction in spectral computed tomography **95**

C Bayesian parameter estimation and uncertainty quantification for hemodialysis **113**

Bibliography **153**

CHAPTER 1

Introduction

Mathematical models can describe complex systems and are widely used in natural sciences such as physics, biology, chemistry and engineering. Mathematical models can take many forms such as differential equations, algebraic equations, statistical models, etc. and play a key role in the quest to understand the world around us by providing quantitative insight into complex systems.

Real world phenomena are often extremely complicated and formulating a mathematical model requires a trade-off between simplicity and accuracy of the model. If the model is very complicated, it may accurately describe the measured phenomenon, but the trade-off is that it may not be tractable from a computational point of view. However, if the model is too simple then it might not characterize the system sufficiently well and conclusions drawn from the model may be misleading or even wrong. Thus, scientists aim to formulate a mathematical model based on the principle of Occam's razor which states that we should choose the simplest model that adequately describes the underlying system. Consequently, the models describing real world systems are rarely exact, but good models capture the essence of the underlying system while disregarding the unimportant aspects.

1.1 Motivation and applications

In this thesis, we employ mathematical modeling to improve the understanding of artifact reduction in computed tomography and phosphate removal during hemodialysis. Conceptual illustrations of the two applications considered in this thesis are shown in Figure 1.1.

1.1.1 Computed tomography

In computed tomography (CT), we want to image the interior of an object by exposing the object to radiation. The radiation is attenuated as it travels through the object, and an image of the interior can be reconstructed by measuring how the radiation is attenuated. As an example, CT can be used by medical doctors to detect bone fractures and other abnormalities without cutting or damaging the patient [17], i.e., it is a non-invasive scanning technique. The CT image can be represented using pixels, and a mathematical

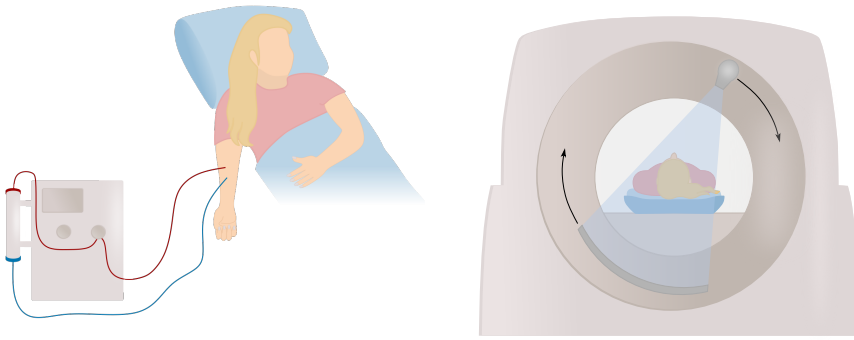

(a) Hemodialysis. **(b)** Computed tomography (CT).

Figure 1.1: Hemodialysis (a) and computed tomography (b) are the two applications considered in this thesis. Hemodialysis is a treatment where toxic substances such as phosphate are removed from the blood of the patient. Computed tomography is an imaging technique that uses X-ray radiation to image the interior of a patient or object.

model that relates the pixels in the image to the attenuation of the radiation can be formulated. In most cases, the formulated model is accurate enough to produce a high quality image of the interior. However, the model may fall short and introduce artifacts, i.e., a feature which appears in the image but is not present in the object, if the radiation dose is low and/or the exposure time is low. In Chapter 2, we give a general introduction to CT, and we develop methods for mitigating artifacts in CT in Chapter 6.

Spectral CT is a technique where the attenuation is measured at multiple energies, and spectral CT has received considerable interest in recent years because it contains much richer information about the object of interest. However, the radiation dose for a single energy is often low, and spectral CT is challenged by artifacts as well. In Chapter 7 we develop a method for mitigating artifacts for spectral CT.

1.1.2 Hemodialysis

Patients with renal-failure, i.e., kidney-failure, suffer from hyperphosphatemia which denotes excessive amount of phosphate in the blood. Hyperphosphatemia is associated with increased vascular calcification and mortality. Thus, it is essential to remove the excessive phosphate from the blood of renal-failure patients. Hemodia-lysis is a conventional treatment for patients with hyperphosphatemia that filters phosphate from the blood. We can formulate a mathematical model that describes the phosphate removal to obtain an improved understanding of the phosphate kinetics. We introduce modeling of hemodialysis in Chapter 3 and develop Bayesian models for two

types of hemodialysis treatments, the conventional single-pass and the novel multiple-pass treatments in Chapter 8.

1.2 Inverse problems

Both problems considered, i.e., hemodialysis and CT, can be characterized by a set of unknown parameters, a mathematical model, and a set of measured data. Moreover, both problems share the challenge that we cannot observe the parameters directly, i.e., the pixels in CT and phosphate kinetics in hemodialysis, and therefore we must rely on indirect measurements. Thus, both problems can be classified as inverse problems.

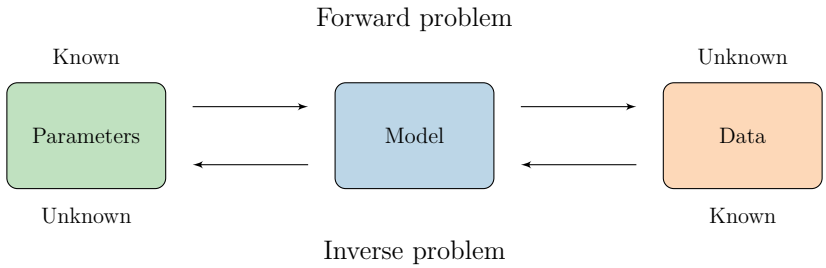

Figure 1.2: Conceptual diagram of a forward problem and the corresponding inverse problem. For the forward problem, both the parameters and model are known whereas for the inverse problem, we know the data (output) and the model, and we want to recover the unknown parameters.

The term inverse refers to the fact that the problem we consider is of a contrary nature, i.e., the inverse problem is the reverse problem of the forward problem. The forward and corresponding inverse problem are conceptualized in Figure 1.2. In the forward problem, we know the parameters and the model, and we aim to find the output (data). Solving the forward problem corresponds to simulating the system and is usually straightforward.

However, if we do not know the parameters and cannot access them directly through measurements, then we need to solve an inverse problem. Inverse problems can be difficult to solve as they are often highly sensitive to measurement noise and model errors. We give a brief introduction to inverse problems in Chapter 4.

1.3 Uncertainty quantification

The classical approach to solve inverse problems is to compute point estimates of the parameters. Obvious questions in this context are: *how much can we trust the estimated parameters? How sensitive is the model to the*

choice of parameters? A methodology aiming to address these statistical questions is uncertainty quantification (UQ). In UQ, we consider the parameters as random variables and characterize the statistical properties of the parameters such as mean, covariance and credible intervals. In Chapter 5, we introduce a Bayesian approach for inverse problems and Markov Chain Monte Carlo sampling for retrieving statistical information about the parameters.

1.4 Contribution

The thesis gives an overview of modeling of CT and hemodialysis and an introduction to a Bayesian approach for inverse problems. The aim of the thesis is two-fold. In paper A and B, we use mathematical modeling and uncertainty quantification to mitigate model errors in CT arising from too simplistic mathematical models. In paper C, we use mathematical modeling and uncertainty quantification to obtain an improved understanding of the dynamics of hemodialysis and to compare two types of hemodialysis treatments. The detailed contributions can be summarized as follows:

- New mathematical model for CT with source uncertainty (paper A).
- New method for removing ring artifacts in spectral CT (paper B).
- Bayesian model for phosphate removal during hemodialysis (paper C).

1.5 Structure of the thesis

We start out by introducing the two applications, CT and hemodialysis, in Chapter 2 and Chapter 3, respectively. We give a brief introduction to discrete inverse problems in Chapter 4, and this is followed by an introduction to the Bayesian approach for inverse problems in Chapter 5. Chapters 6, 7 and 8 summarize the findings in paper A, B, and C, respectively. Lastly, Chapter 9 discusses and comments on the work and concludes the thesis. The papers are attached in the appendix.

CHAPTER 2

Computed tomography

We introduce CT which is the inverse problem considered in both paper A and B. We present the physical model and derive a discretized measurement model for conventional and spectral CT. Moreover, we discuss how artifacts such as ring artifacts arise from model errors and how these degrade the image quality.

2.1 Experimental set-up

Tomography originates from the Greek words *Tomos* which is the Greek word for *slice* or *section* and *graphos* which means *to write*. Thus, CT is a non-invasive imaging technique that enable us to obtain cross-sectional images (slices) of the interior of an object from a set of projection images.

The projection images are acquired by illuminating the object by radiation from a source, e.g., X-rays or neutrons, and measuring the attenuated radiation. A panel of detectors is placed opposite of the source, and the detectors measure the attenuation of the radiation by counting the number of particles hitting the panel. The source and the detector panel are then rotated around the object such that projection images are collected from multiple angles, see Figure 2.1.

The experimental setup shown in Figure 2.1 for 2D imaging is called parallel-beam geometry since the beams travel in straight parallel lines. The parallel-beam geometry was used in early scanners and is used today in large-scale synchrotron facilities. There exist other configurations such as the 2D fan-beam and 3D cone-beam geometry, but from a mathematical point of view, all configurations result in a similar mathematical model structure. Thus, we focus on the parallel-beam geometry since it is the simplest configuration both mathematically and conceptually.

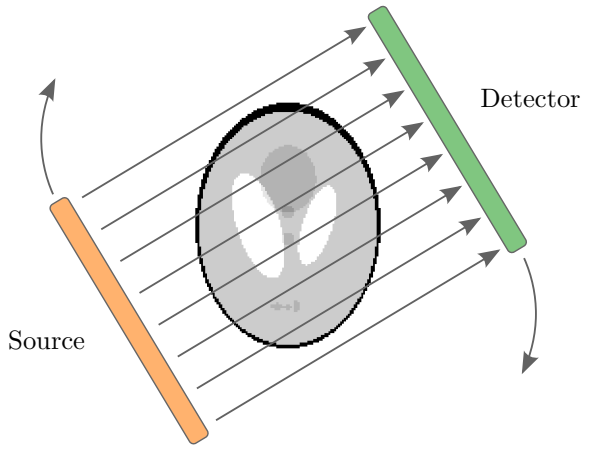

Figure 2.1: Illustration of the experimental set-up for CT. The Shepp-Logan phantom simulates a cross-section of a human head where darker shades correspond to higher absorption. A source emits radiation, e.g., X-rays or neutrons, which travels through the object and is attenuated. The attenuation of the rays is measured by a detector panel placed opposite of the source. The source and the detector are then rotated such that the attenuation is measured from different projection angles.

2.1.1 Types of radiation

We consider two types of radiation, i.e., photons (X-ray) and neutrons. Neutron and X-ray imaging provide complementary scanning techniques since the neutrons interact differently than X-rays with materials due to their zero charge. However, from a mathematical point-of-view, the imaging techniques are similar, and we can use the same model for both neutrons and X-rays, i.e., the models derived in this chapter apply to both types of radiation, and we will not distinguish between the two but use the terminology for X-ray CT.

2.2 The physical model

We can describe the attenuation of a single ray by Lambert-Beer's law. Lambert-Beer's law is a model of the physics of CT which describes the mean photon count at a specific energy. Let $I \in \mathbb{R}_+$ and $I_0 \in \mathbb{R}_+$ denote the intensity incident on the detector element and on the object, respectively. We can model I by the Lambert-Beer law,

$$I = I_0 \exp\left(- \int_\ell \mu(\mathbf{x}) \mathrm{d}\mathbf{x} \right), \tag{2.1}$$

where $\mathbf{x} \in \mathbb{R}^d$ is the spatial position with dimension $d \in \{2, 3\}$, ℓ denotes the line segment between the source and the detector element, and $\mu : \mathbb{R}^d \to \mathbb{R}_+$ is the energy-dependent spatial attenuation function that we aim to recover. We will focus on the two-dimensional case, i.e., we assume $d = 2$ henceforth.

Commonly, it is assumed that we know I_0 and that the source is static throughout the data acquisition. The aim is to reconstruct the attenuation function μ from the line integral equation,

$$-\log\left(\frac{I}{I_0}\right) = \int_\ell \mu(\mathbf{x}) \mathrm{d}\mathbf{x}. \tag{2.2}$$

In practice, statistical fluctuations will influence the photon count. It can be shown that the statistical fluctuations are well described by a Poisson distribution with mean and variance equal to I [14]. Consequently, the CT data in (2.2) is log-Poisson distributed. We can analyze the statistical properties of the log-Poisson CT data by making a first order approximation. Let X represent the measured data such that X is considered a Poisson distributed variable with mean and variance $\mathrm{E}[X] = \mathrm{V}[X] = I$, and let Y be the corresponding log-Poisson distributed variable such that $Y = \log(X)$. We can make a first order approximation around I for the Poisson distributed variable X by,

$$\log(X) \approx \log(I) + \frac{1}{I}(X - I),$$

and it follows that $\mathrm{E}[Y] \approx \log(I)$ and $\mathrm{V}[Y] \approx \frac{1}{I}$.

Thus, if I is high, then the variance is low and our measurements are close to the true photon count I. However, if I is low, then so is the variance and the measured data may fluctuate significantly from I which might introduce substantial noise.

2.3 Filtered Back Projection

We can reconstruct the attenuation function μ in (2.2) analytically using the inverse Radon transform. The forward model is called the Radon transform after mathematician Johan Radon who proved in 1917 that under certain regularity assumptions, the attenuation function may be reconstructed perfectly from a full set of line integrals over all angles [41].

The most commonly used reconstruction method for CT is the Filtered Back Projection (FBP) which inverts the Radon transform to reconstruct the attenuation function μ, see, e.g., [36] for a detailed description of FBP. The FBP is widely used in practice since it is fast and reliable when the model assumptions are valid, and the signal-to-noise ratio is high.

However, the FBP method is not robust with respect to noise, and FBP gives a low quality reconstruction when the signal-to-noise ratio is low. Moreover, FBP struggles if the underlying assumptions fail and it is not easy to

incorporate adjustments. Hence, we need a more robust and flexible method for computing reconstructions if we have CT data with a low signal-to-noise ratio or if the model assumptions are violated.

2.4 Discretized model

A method for obtaining a more flexible reconstruction method is to consider the discretized CT problem. Let n denote the number of pixels in the discretized domain and let π_l denote the lth pixel as illustrated in Figure 2.2.

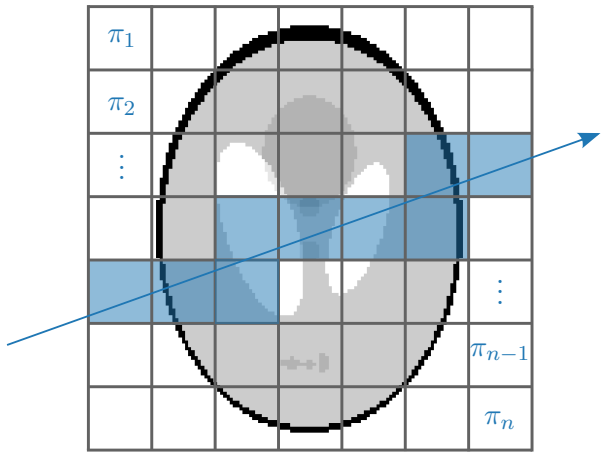

Figure 2.2: Discretization of the domain into n pixels. π_l denotes the lth pixel and the arrow symbolizes the traveled distance of the ith ray through the grid for the jth projection. The length of which the ray travels through each pixel is stored in the ith row of the matrix A_j such that the element $(A_j)_{i,l}$ corresponds to the traveled distance of the ith ray through pixel l for projection j.

We assume that we have measured p projection images which have been acquired by a detector panel with r detector elements. We can discretize the line integral in (2.2) by introducing the following parameterization of μ,

$$\mu(\mathbf{x}) \approx \sum_{l=1}^{n} u_l \chi_l(\mathbf{x}), \tag{2.3}$$

where u_l is the attenuation coefficient in pixel π_l and $\chi_l(\mathbf{x})$ is an indicator function defined by,

$$\chi_l(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in \pi_l, \\ 0 & \text{otherwise.} \end{cases}$$

The parameterization of μ allows us to express the line integral by,

$$\int_{\ell_{i,j}} \mu(\mathbf{x}) \mathrm{d}\mathbf{x} = e_i^T A_j u,$$

where $A_j \in \mathbb{R}^{r \times n}$ is the matrix with entries,

$$(A_j)_{i,l} = \int_{\ell_{i,j}} \chi_l(\mathbf{x}) \mathrm{d}\mathbf{x}$$

such that $(A_j)_{i,l}$ denotes the length of the traveled distance of ray i through pixel π_l at projection j. Thus, most entries in A_j are zero since each ray only passes through a small subset of the pixels as illustrated in Figure 2.2.

Let $Y \in \mathbb{R}^{r \times p}$, often called the sinogram, denote the matrix with the jth projection image in the jth column as illustrated in Figure 2.3, and let $\nu \in \mathbb{R}^r$ denote the effective measured intensity incident on the object.

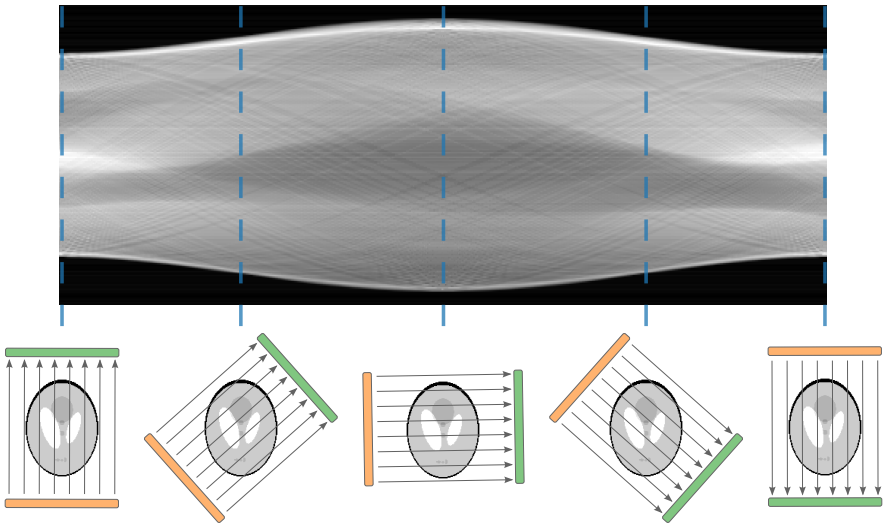

Figure 2.3: Illustration of the measurement matrix Y, i.e., sinogram. Darker shades correspond to higher intensity, i.e., measurements that are not attenuated are black and the measurements gradually become lighter as the attenuation increases. Each column of Y corresponds to a projection image at a specific rotation angle.

We can then formulate the discretized CT model for projection j by,

$$Y_j = \mathrm{diag}(\nu) \exp(-A_j u) \tag{2.4}$$

for $j = 1, 2 \ldots, p$ with $\mathrm{diag}(\nu) \in \mathbb{R}^{r \times r}$ defined as the diagonal matrix with the elements of ν on the diagonal. Equivalently, we can formulate the discretized

CT model for all projections by,

$$y = \operatorname{diag}(1_p \otimes \nu) \exp(-Au)\,, \tag{2.5}$$

where $y = \operatorname{vector}(Y) \in \mathbb{R}^{rp}$ is the vector obtained by stacking the columns of Y vertically, $\otimes$ is the Kronecker product and

$$A = \begin{bmatrix} A_1 \\ A_2 \\ \vdots \\ A_p \end{bmatrix}.$$

If we assume that ν is known, we arrive at the linear system of equations,

$$b = Au, \tag{2.6}$$

where $b = -(\log(y) - \log(1_p \otimes \nu))$. Thus, reconstructing the attenuation coefficients boils down to solving a linear system of equations.

2.5 Ring artifacts

Commonly, most reconstruction methods assume that we know the effective intensity incident on the object ν in (2.6). In practice, ν is estimated from so-called flat-field measurements (also known as white-fields) which are simply projection images acquired without an object in the scanner.

The flat-field measurements contain noise due to the statistical nature of the radiation, just as regular CT measurements. Often a couple of flat-fields are measured, and we denote the estimate of ν based on the sampled mean of the recorded flat-fields by $\hat{\nu}$. Correcting the measurements by the estimate $\hat{\nu}$ is often called flat-field correction.

Flat-field correction reduces the amount of noise introduced in (2.6), but often $\hat{\nu}$ is not sufficiently close to the true intensity ν when the acquisition time is limited or the dose is low.

Errors in $\hat{\nu}$, also called flat-field errors, will enter the model through b in a highly systematic manner which can be seen by considering the discretized model in (2.6). The discretized model reveals that $\hat{\nu}$ is replicated for each projection, i.e., the flat-field errors introduce stripes through the sinogram as illustrated in Figure 2.4. The stripes in the sinogram can give rise to concentric rings in the reconstruction known as ring artifacts. Ring artifacts degrade the quality of the reconstruction and might conceal features of the object as illustrated in Figure 2.4. Paper A proposes a mathematical model for mitigating ring artifacts for monochromatic CT.

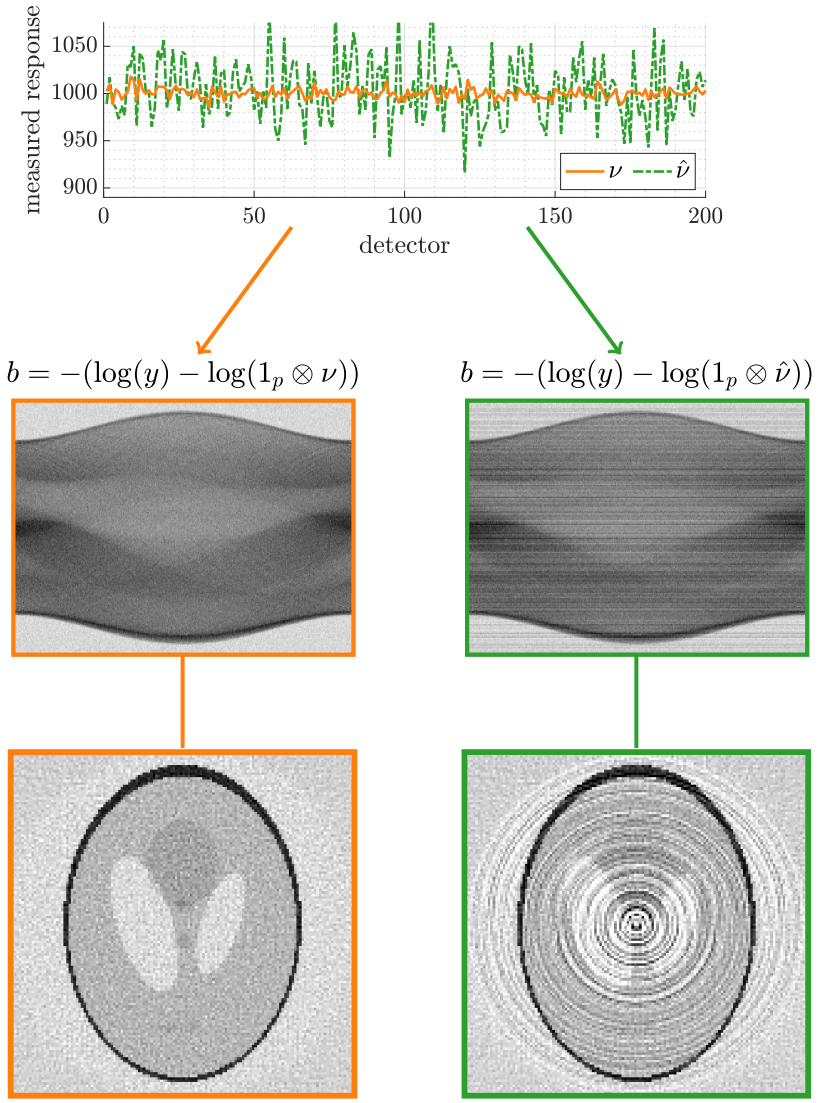

Figure 2.4: Illustration of the effect of ring artifacts. Darker shades correspond to higher values. If we know ν, then we normalize the data correctly and do not introduce systematic errors (left). However, if we use the estimate $\hat{\nu}$ based on flat-field measurements, we might introduce systematic errors in the normalized sinogram b which appear as horizontal stripes. Consequently, the horizontal stripes in b result in concentric rings in the corresponding reconstruction (right).

2.6 Spectral model

The attenuation is both material- and energy-specific as illustrated in Figure 2.5 where the attenuation of gold, lead, iodine, iron and water are depicted as a function of energy. The figure illustrates that different materials may look similar for some energies while distinct for other energies. Thus, spectral CT can be used to discern different materials in an object and thereby obtain a quantitative material decomposition as illustrated in Figure 2.6.

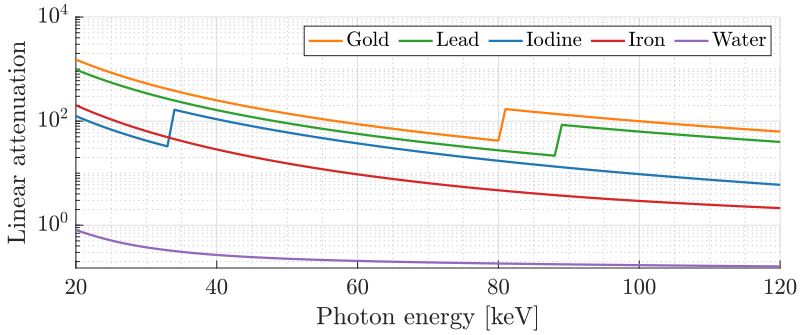

Figure 2.5: Linear attenuation for five materials (gold, lead, iodine, iron and water) as a function of energy for X-ray CT.

We can obtain a material decomposition by computing a reconstruction for each energy, also referred to as energy-wise reconstructions or a spectral reconstruction. Let m denote the number of energies. For each energy $k = 1, 2, \ldots, m$, we can use Lambert-Beer's law and arrive at the linear system of equations,

$$b_k = Au_k, \tag{2.7}$$

where $b_k = -(\log(y_k) - \log(1_p \otimes v_k))$. Thus, we have the same type of model for spectral CT as for monochromatic CT with the only difference that we have m times as many reconstruction problems. For a single energy, the photon count is often low, and consequently spectral CT is challenged by ring artifacts. We develop a method for mitigating ring artifacts for spectral CT in Paper B.

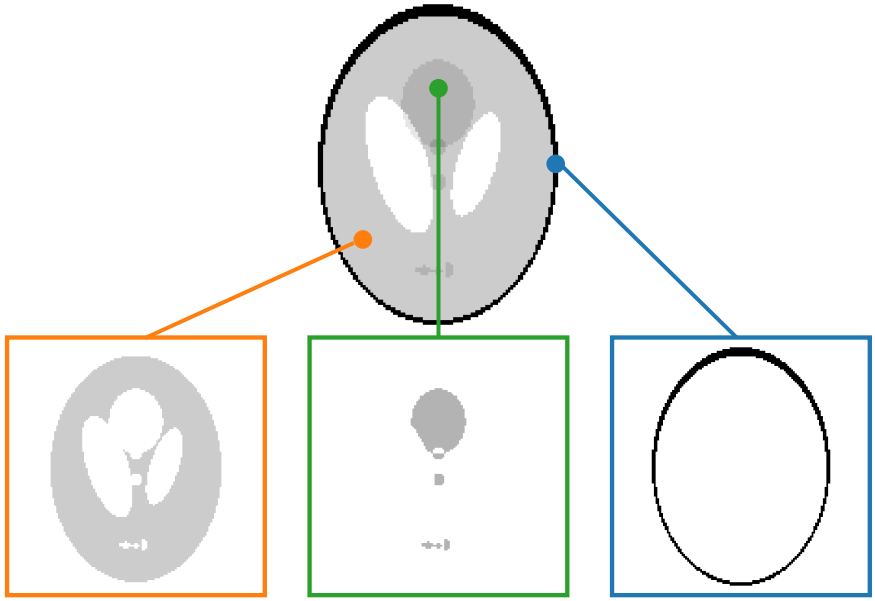

Figure 2.6: Illustration of a material decomposition. The phantom has three distinct materials illustrated by varying attenuation coefficients.

CHAPTER 3

Phosphate kinetics

In this chapter, we give a brief introduction to the regulation mechanisms of phosphate in the human body and introduce the treatment hemodialysis for patients with hyperphosphatemia. We introduce the concept of compartment modeling and derive differential equations describing the phosphate kinetics during hemodialysis, which are the foundation of the models analyzed in papers C and E. We refer the reader to [29, 19] for an introduction to modeling in life sciences.

3.1 Phosphate's role in the body

Phosphate plays a crucial role in the human body since it performs vital processes such as construction of nucleic acids, energy transport and bone tissue formation. [9] About 80-85% of the phosphate in an adult is stored in the bones and 15–20% is present in body fluids and soft tissues. Only 1% of the phosphate is found in the plasma and extracellular fluid, which is accessible through the blood and can be measured in clinical practice. [30]

The phosphate level is tightly regulated, and both hyperphosphatemia and hypophosphatemia can have fatal consequences [46]. The regulation of phosphate concentration in the blood is mainly maintained by the phosphate storage in bones and soft tissues, absorption from the intestines, and excretion and reabsorption in the kidneys. A simplified overview of the physiological mechanisms involved in the phosphate regulation is given in Figure 3.1.

The kidneys are crucial in the maintenance of the phosphate level since phosphate can only be renally cleared. For renally impaired patients, the ability to excrete phosphate is reduced leading to accumulation of phosphate in the blood. Hyperphosphatemia is a considerable clinical problem since it is associated with serious adverse outcomes for the patients and imposes a significant burden on both the patients and the healthcare system. [34]

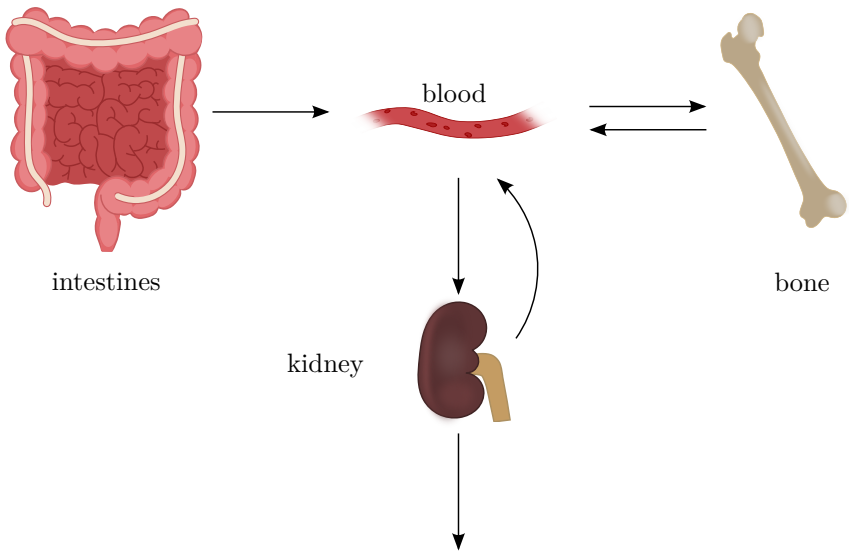

Figure 3.1: Conceptual diagram of the regulation mechanisms of phosphate in the human body. Phosphate enters the body through food intake and is absorbed into the systemic blood circulation from the intestines. The bones provide a phosphate storage which can both excrete and absorb phosphate to and from the systemic blood circulation. The kidney filters the blood and removes phosphate and other toxins through urinary excretion.

3.2 Hemodialysis

About half of all renal-failure patients suffer from hyperphosphatemia. Strategies to control phosphate levels include phosphate binders, low-phosphate diet and removal of phosphate by hemodialysis. Hemodialysis is a conventional treatment for hyperphosphatemia where the patient is connected to a dialysis machine for four to eight hours. [24] *Hemo* is the Greek word for blood, and *dialysis* originates from dialuete, which means separation. The idea of hemodialysis is to separate the excessive phosphate from the blood. The patient is connected to the dialysis machine through access to the blood vessels as illustrated in Figure 3.2. The blood of the patient is then circulated through the dialysis machine where waste products, e.g., phosphate and excess fluid is moved from the blood to the dialysate fluid by use of a semipermeable membrane called a dialyzer. Thus, the dialyzer works as an artificial kidney that cleanses and returns the blood to the body.

Hemodialysis cannot cure the patient since it only removes phosphate from the blood temporarily. After the treatment ends, the phosphate concentration starts to increase and the patient experiences a relapse. Thus, the patient needs recurrent treatments to keep the phosphate concentration

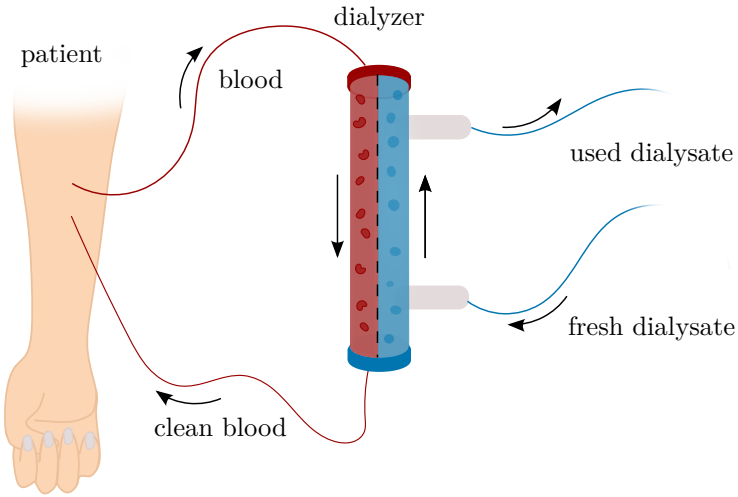

Figure 3.2: Illustration of the hemodialysis treatment. The blood of the patient is passed through the dialyzer where a semipermeable membrane allows phosphate and other toxins to flow to the dialysate fluid by diffusion.

at a non-fatal level. Moreover, the rebound in phosphate concentration is relatively fast, and consequently the patient often needs treatment trice a week to maintain a tolerable phosphate concentration.

3.3 Compartment modeling

We use a compartment model to describe the process of phosphate removal during hemodialysis. Compartment modeling is a general modeling technique that can model a broad class of biological systems such as chemical reactions, infectious diseases and population dynamics. The assumption is that we can describe a system as a set of interconnected, well-mixed compartments that exchange substances by simple linear kinetics. [19]

Each compartment represents a biological quantity, i.e., phosphate concentration in hemodialysis. We consider a three-compartment model consisting of a bone, blood (including plasma and extracellular fluid) and dialysate fluid compartment as illustrated in Figure 3.3. We can describe the change in concentration of phosphate in each compartment by the conservation of mass. The law of mass conservation states that the change of mass is equal to the substance flowing into the compartment subtracted that flowing out of the compartment.

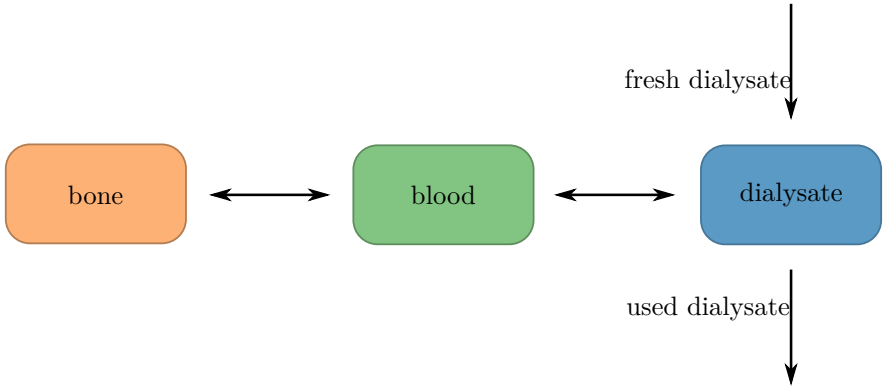

Figure 3.3: Compartment model for the conventional hemodialysis treatment denoted single-pass. We consider a three-compartment model consisting of bone, blood and dialyzer compartments. The arrows indicate that phosphate flows from the bones to the blood and from the blood to the dialysate fluid.

3.3.1 Single-pass modeling

First, we consider the conventional hemodialysis treatment denoted single-pass where the dialysis fluid is constantly replenished as illustrated in Figure 3.3. We consider the bone compartment as being an infinite storage which is motivated by the fact that most of the phosphate in the human body is stored in the bones and only a small fraction of this is removed during dialysis. Thus, we assume that the concentration of phosphate in the bones is constant and can be modeled as a source that continuously excretes phosphate to the blood.

The constant inflow of phosphate from the bones to the blood compartment can be modeled as a diffusion process. Diffusion is a metabolically cheap transport mechanism where a substance, e.g., phosphate, moves from a compartment of high concentration to a compartment of lower concentration. We assume that the compartments are separated by a thin membrane such that we may model the diffusion process as proportional to the concentration gradient between the two compartments. When the patient is connected to the dialyzer, phosphate can move from the blood to the dialysate fluid by diffusion as well. The dialysate fluid has an inflow of phosphate from the blood through the semipermeable membrane. However, the dialysate fluid is constantly replenished such that the concentration in the dialysis compartment is constantly low and can be assumed constant.

We can formalize the transport of phosphate in the blood by an ordinary differential equation (ODE). Let C_s, $C_b(t)$ and C_d denote the concentration of phosphate in the bones (source), blood and dialysate fluid, respectively.

We model the flow of phosphate during hemodialysis by,

$$V_b \frac{\mathrm{d}C_b(t)}{\mathrm{d}t} = K_s(C_s - C_b(t)) - K_b(C_b(t) - C_d), \tag{3.1}$$

where K_s is the diffusion coefficient, V_b is the volume of the blood compartment which is assumed to be constant during the hemodialysis, and K_b is the diffusion coefficient from blood to dialysate. For the system to have a unique solution, we equip the ODE with the initial condition $C_b(0)$ which is the concentration of phosphate in the blood at time $t = 0$.

Note that if the patient is not receiving treatment, then there is no outflow of phosphate, and we simply have $K_b = 0$. Thus, we can use the model to predict the relapse of the patient when $K_b = 0$.

3.3.2 Multiple-pass modeling

The quality of life for a hemodialysis patient can be compromised by the frequent hospital treatments. This may be eased by having a dialysis unit at home. However, home treatment requires significant training and logistics and approximately a 100 liters of dialysate. [27] As an alternative to the single-pass treatment, we can consider the novel hemodialysis treatment called multiple-pass. The multiple-pass treatment recirculates the dialysis fluid and requires less than 20% of the dialysate fluid compared to the single-pass treatment. [20] Thus, the multiple-pass treatment enables a transportable dialysis unit which can ease home treatment and enable treatment during travel. An illustration of the multiple-pass compartment model is depicted in Figure 3.4.

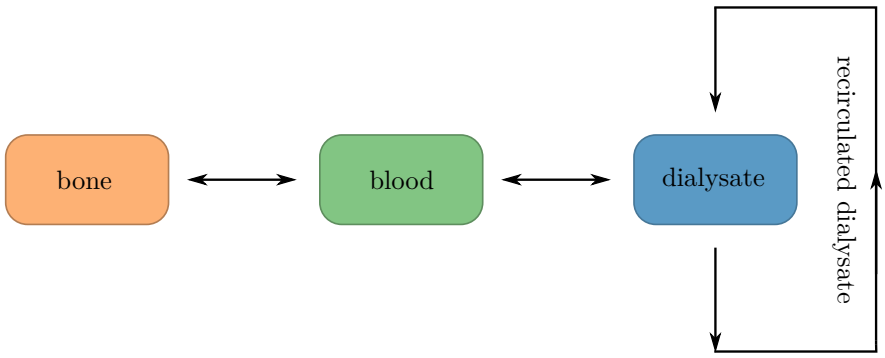

Figure 3.4: Compartment model for the novel hemodialysis treatment denoted multiple-pass. We consider a three-compartment model consisting of a bone, blood and dialyzer compartment. The arrows indicate that phosphate flow from the bones to the blood, from the blood to the dialysate fluid and that the dialysate fluid is recirculated.

We can formulate an ODE governing the flow of mass for the multiple-pass model. Note that the flow between the inner compartments, i.e., bone and blood, is unaffected by the change of the dialyzer mechanism, and thus we may model the in- and outflow of phosphate for the blood compartment by

$$V_b \frac{\mathrm{d}C_b(t)}{\mathrm{d}t} = K_s(C_s - C_b(t)) - K_b(C_b(t) - C_d(t)), \tag{3.2}$$

where the only change from (3.1) is that the concentration of phosphate in the dialysate fluid is now time dependent since it accumulates over time. For the multiple-pass treatment, the dialysate fluid is recirculated, and we can model the change of mass in the dialysate fluid by the following ODE,

$$V_d \frac{\mathrm{d}C_d(t)}{\mathrm{d}t} = K_b(C_b(t) - C_d(t)), \tag{3.3}$$

where V_d is the volume of the dialysate fluid. Thus, we can model the multiple-pass treatment by the coupled system of ODEs,

$$V_b \frac{\mathrm{d}C_b(t)}{\mathrm{d}t} = K_s(C_s - C_b(t)) - K_b(C_b(t) - C_d(t)), \tag{3.4a}$$

$$V_d \frac{\mathrm{d}C_d(t)}{\mathrm{d}t} = K_b(C_b(t) - C_d(t)), \tag{3.4b}$$

and we equip the ODEs with the initial conditions $C_b(0)$ and $C_d(0)$ which are the concentrations of phosphate in the blood and dialysate fluid at time $t = 0$, respectively. Note however, that at time $t = 0$, the concentration of phosphate in the dialysate fluid is zero, i.e., we have $C_d(0) = 0$.

3.4 Parameter estimation

We can measure V_d directly and we have estimates of V_b from measurements, but we cannot measure the remaining parameters, i.e., C_s, K_s, K_b, directly. Thus, the goal is to recover the set of parameters (C_s, K_s, K_b) that gave rise to the observed phosphate concentration in the blood during hemodialysis.

We can formalize the system of ODEs in a general setting. We assume that we have m measurements at time $t_1, t_2, \ldots, t_m$. Let $b_i \in \mathbb{R}^d$ denote the measurements of phosphate concentrations at time t_i (with $d = 1$ for single-pass and $d = 2$ for multiple-pass), let $u = \begin{bmatrix} C_s, K_s, K_b \end{bmatrix}^T$ denote the vector of unknown parameters, and let $F_i(t_i, u) \in \mathbb{R}^d_+$ be the solution to the system of ODEs at time t_i with respect to u, see e.g. [35] for methods for solving ODEs numerically or [39] for analytical solutions to ODEs. We can then formulate the parameter estimation problem as a non-linear system of equations,

$$b = F(u), \tag{3.5}$$

where $F(u) \in \mathbb{R}^{dm}$ and $b \in \mathbb{R}^{dm}$ are the vectors obtained by stacking $F_i(t_i, u)$ and b_i vertically, respectively.

There are great similarities in parameter estimation problem for hemodialysis in (3.5) and the reconstruction problem for CT in (2.6), and we now turn to the more general theory of how to deal with problems where we need indirect measurements to estimate the parameters.

CHAPTER 4

Discrete inverse problems

In this chapter, we introduce the reader to discrete inverse problems. The chapter begins with a general introduction to discrete inverse problems and the characteristics of inverse problems such as ill-posedness. We focus on the solution to the linear discrete inverse problem in the context of least-squares solutions and introduce the concept of regularization. We refer to [2, 36] for a thorough introduction to inverse problems.

4.1 Classification of inverse problems

We want to recover the parameters that characterize the system,

$$b = F(u), \tag{4.1}$$

where b is the vector of measurements, u is the vector of unknown parameters and F denotes the mathematical model that relates the parameters u to the measurements b. For inverse problems, we cannot measure u directly, and we can only obtain information about u indirectly through b.

The measurements b may be a function of time or a set of discrete observations, and the model F can take many forms and may arise from, e.g., ordinary differential equations, partial differential equations, or a system of algebraic equations. We will focus on the case where u and b are finite-dimensional vectors, which are often called discrete inverse problems or parameter estimation problems.

We use CT as an example to illustrate the characteristics of discrete inverse problems throughout the chapter. Figure 4.1 depicts the inverse problem for CT. In CT, the forward problem corresponds to computing projection images b from a known object u whereas the inverse problem is to reconstruct the object u from the set of acquired projection images b.

Discrete inverse problems can, despite their simple expression, be very challenging to solve since they are often *ill-posed*. To formally characterize ill-posed problems, we first introduce the complementary notion of a *well-*

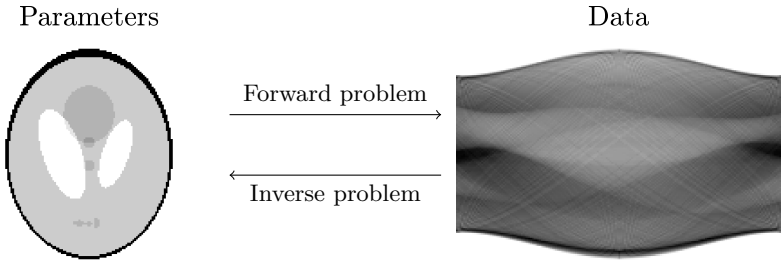

Figure 4.1: Illustration of the forward and inverse problem of CT. The parameters are the material-specific attenuation coefficients represented as pixels, and the data is the log-transformed sinogram defined in (2.6).

posed problem which was first introduced by Hadamard [25]. A problem is said to be well-posed if it satisfies the following three conditions:

- Existence: There should be at least one solution.
- Uniqueness: There should be at most one solution.
- Stability: The solution must depend continuously on data.

On the other hand, if a problem fails to satisfy one or more conditions, then the problem is said to be an ill-posed problem.

4.2 Linear inverse problems

To solve (4.1), we often linearize $F(u) \approx Au$ to obtain a linear system of equations, i.e.,

$$b = Au, \tag{4.2}$$

where $A \in \mathbb{R}^{m \times n}$ is the system matrix or model, $u \in R^n$ is the vector of unknowns and $b \in \mathbb{R}^m$ is the vector of measurements.

We can consider the linear discrete inverse problem in the context of the three conditions, existence, uniqueness and stability. If the linear system is consistent and A has fewer rows than columns, then the null space of A is non-trivial and there are infinitely many solutions, and consequently the uniqueness condition fails. On the other hand, if A has more rows than columns and full column rank then the solution exists if and only if b lies in the range of A. It is often the case that the observed data is imperfect due to measurement noise which moves b out of the range of A, and consequently there does not exist any solution.

Even in the case where the existence and uniqueness conditions are satisfied and A is invertible, then the stability condition may still fail. The

stability condition can be interpreted as that a small change in the data should only result in a small change in the solution. The stability of a linear discrete inverse problem is closely connected to the condition number of the system matrix A. The condition number of an invertible matrix A is given by,

$$\kappa(A) = \|A\|_2 \|A^{-1}\|_2, \tag{4.3}$$

where $\|\cdot\|_2$ is the 2-norm.

We can relate the stability to the condition number of the matrix A by considering the observed data b as a vector consisting of two components: the clean data $\bar{b}$ and measurement noise ε, i.e.,

$$b = \bar{b} + \varepsilon. \tag{4.4}$$

Let $\bar{u} = A^{-1}\bar{b}$ denote the exact solution and let $u = A^{-1}b$ denote the computed solution affected by the noise component in b. The difference between the exact solution and the computed solution satisfies the following bound [48],

$$\frac{\|u - \bar{u}\|_2}{\|\bar{u}\|_2} \leq \kappa(A) \frac{\|\varepsilon\|_2}{\|\bar{b}\|_2}. \tag{4.5}$$

Thus, even small perturbations of the clean data $\bar{b}$ might lead to large errors in the computed solution u if the condition number is large.

4.2.1 Least-squares solution

We can address the issue with non-existence of the solution by considering the least-squares solution. The least-squares solution is defined as the solution to the following optimization problem,

$$\underset{u}{\text{minimize}} \ \|Au - b\|_2^2, \tag{4.6}$$

which is often referred to as a data fitting problem.

The least-squares problem is a convex optimization problem, which means that a local solution is also a global solution. We refer the reader to [12] for a thorough introduction to convex optimization. Consequently, the first order optimality condition is both sufficient and necessary for the solution to be optimal. The first order optimality condition requires that the gradient of the objective function is zero at the optimal solution. The optimal solution is therefore easily derived from the first order optimality condition and has the solution,

$$\left(A^T A\right) u = A^T b. \tag{4.7}$$

As a remark, we will briefly mention that even though there exists an analytical solution to the optimization problem in (4.6), it may not be practical nor feasible to compute the exact solution. For large-scale problems such as CT, it is often infeasible and/or impractical to compute the matrix product A^TA and even more troublesome to compute the inverse as needed to solve (4.7). Thus, instead we use iterative methods to approximate the solution of the optimization problem. For a general introduction to iterative methods, we refer the reader to [2, 26].

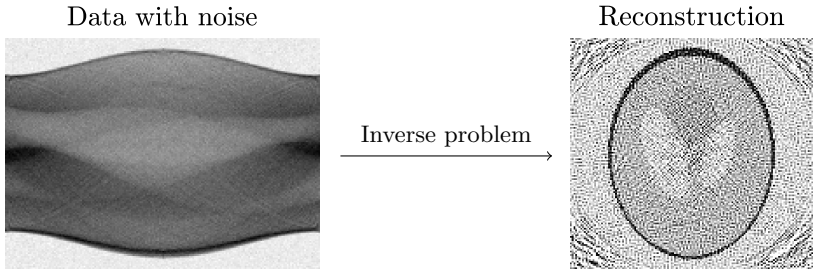

Figure 4.2: Illustration of the ill-posedness of the CT problem. We have simulated X-ray data, i.e., sinogram with noise (left). The resulting reconstruction which is the solution to the least-squares problem in (4.6) is severely affected by noise (right).

Even though the least-squares solution in (4.7) guarantees existence it does not guarantee uniqueness (only if A has rank n) nor stability. Figure 4.2 shows the complications that may arise when the stability condition fails. The reconstruction is the solution to (4.6) computed using a Matlab implementation of the iterative method FISTA, see [10] for details. The sinogram (data) has been perturbed by adding small amount of noise. The figure illustrates that even a small amount of noise in the sinogram can result in a reconstruction heavily affected by noise where most of the details of the original image are lost.

4.3 Regularization

To address uniqueness and stability, we can regularize the least-squares problem to obtain a modified problem that has a unique and stable solution. The simplest form for regularization is perhaps the 2-norm regularization also known as Tikhonov regularization. The Tikhonov solution is defined as the solution to the following optimization problem,

$$\underset{u}{\text{minimize}} \ \|Au - b\|_2^2 + \lambda \|u\|_2^2, \tag{4.8}$$

where $\lambda > 0$ is the regularization parameter, also referred to as a hyperparameter, that balances the data fitting term and the regularization term.

Figure 4.3 illustrates how Tikhonov regularization can be used to improve the image quality in CT. The naive solution is the least-squares solution in (4.6) and the regularized solution is the solution to (4.8) with $\lambda = 100$. Both solutions are computed using a Matlab implementation of the iterative method FISTA [10]. The figure illustrates how Tikhonov regularization suppresses noisy components and favors smooth solutions for CT.

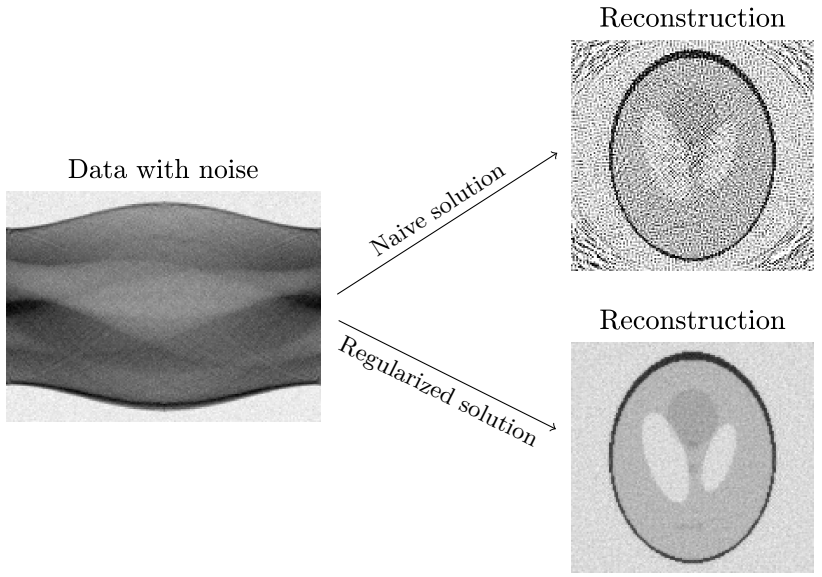

Figure 4.3: Illustration of the effect of Tikhonov regularization. The naive reconstruction is the solution to the least-squares optimization problem in (4.6), which is severely affected by measurement noise. The regularized reconstruction is the solution to the regularized optimization problem in (4.8) with $\lambda = 100$.

The regularization term represents our prior knowledge about the solution. However, for CT imaging, we are often interested in images with sharp edges, e.g., fracture detection in medical imaging and thus Tikhonov may not be the optimal form of regularization since Tikhonov favors smooth solutions. There exists a variety of different regularization techniques, e.g., the edge preserving total variation technique, non-negativity, and sparsity regularization. We refer the reader to [2] for details regarding advanced regularization techniques.

For a fixed λ, the objective function in (4.8) is differentiable and convex with the solution,

$$\left(A^T A + \lambda I\right) u = A^T b. \tag{4.9}$$

If λ is small, then we give more weight to the data fitting term whereas increasing λ gives more weight to the regularization term. We define the

optimal λ as the value that brings us closest to the true solution in terms of minimizing the relative error, i.e.,

$$\frac{\|\bar{u} - u\|_2}{\|\bar{u}\|_2}. \tag{4.10}$$

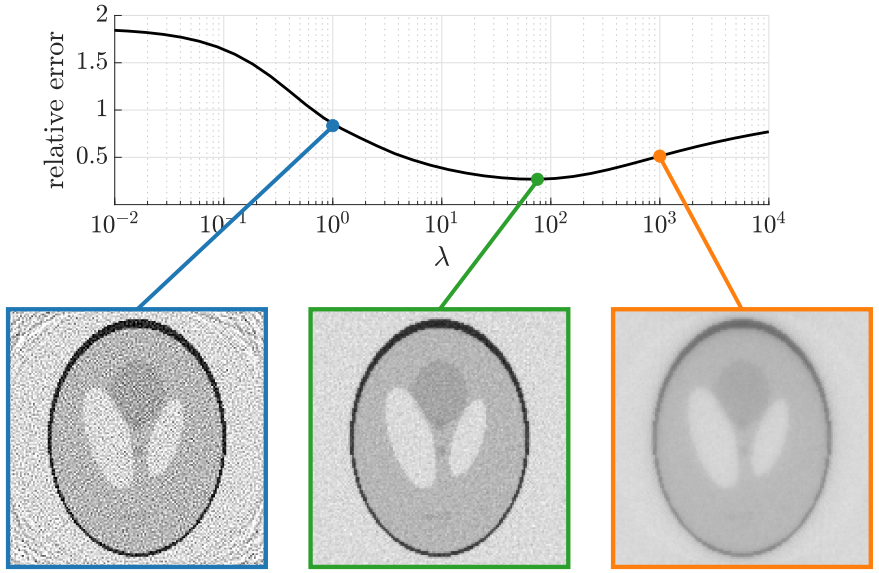

Figure 4.4: Relative error as a function of the regularization parameter λ. For small values of λ, we get a reconstruction dominated by noise whereas we obtain a smooth and blurred reconstruction for large values of λ. Our definition of the optimal λ is a value that brings us closest to the true solution in terms of relative error and balances the trade-off between data fitting and regularization.

Figure 4.4 shows the relative error for the CT example as a function of λ to illustrate the effect of regularizing the least-squares solution. If we choose λ too small, then we obtain a noisy reconstruction. If we choose λ too large, then we introduce too much bias and end up with an over-smoothed reconstruction.

4.3.1 Bias-variance trade-off

Regularization can bring us from a low-quality solution where important features are masked by noise to a reconstruction with distinct visible features, as illustrated in Figures 4.3 and 4.4. However, there is a price to pay. The statistical interpretation of regularization is that it introduces a trade-off between bias and variance. If we do not have regularization, then we have

a high variance in our estimate and we risk fitting our estimate to the noise component. If we have too much regularization, then we introduce a high bias, and the computed solution may be too far from the true solution.

4.3.2 Choice of regularization parameter

Choosing the regularization parameter can be a tedious and non-trivial task. If chosen too small, we get a reconstruction close to the noisy naive reconstruction like in Figure 4.2, but if chosen too large, we might smooth out important features of the image.

For small-scale simulated examples where we know the true solution $\bar{u}$, we can compute the relative error for different choices of λ to narrow down a suitable λ as in Figure 4.4. However, in practice, we do not know the true solution $\bar{u}$ and finding a suitable regularization parameter can, especially for large-scale problems, be a tedious task. There exist several strategies for selecting a suitable regularization parameter such as the L-curve, discrepancy principle, and generalized cross validation, and we refer the reader to [26] for a thorough overview of strategies for choosing a suitable regularization parameter.

CHAPTER 5

Uncertainty quantification

In this chapter, we introduce the reader to the Bayesian approach to inverse problems. In the Bayesian framework, we consider the inverse problem from a statistical point-of-view and consider the parameters to be random variables rather than deterministic quantities. We will focus on the case with Gaussian measurement noise and priors since these are commonly used and play a key role in paper A and C. We introduce Markov Chain Monte Carlo and consider the following sampling methods: Metropolis Hastings, the Gibbs sampler, and the No-U-Turn Sampler (NUTS) which are used in paper A and C. We refer the reader to [7, 15, 45] for a thorough introduction to Bayesian inference for inverse problems and to [15, 42, 45] for a rigorous introduction to sampling techniques for parameter estimation.

5.1 Bayesian approach

So far, we have only discussed how to find a solution, but how certain can we be of the computed solution? A way to address this question of uncertainty, is to consider the optimization problem from a Bayesian point-of-view.

From a Bayesian point-of-view, we consider the parameters and the measurement errors to be random variables and wish to infer their statistical properties such as mean, correlation and 95% credible intervals (CIs), i.e., the interval in which the parameter lies with a probability of 95%. Thus, we are not just interested in a point estimate, but the statistical properties which provide a useful, interpretable visual summary of the uncertainty of the model parameters.

In particular, we are interested in the posterior density of the parameters, and we will let $\pi(\cdot)$ denote probability density functions. We can write the posterior of the parameters u formally using Bayes' rule by

$$\pi(u|b) \propto \pi(b|u)\pi(u),$$

where $\pi(u|b)$ is the posterior of u given b, $\pi(b|u)$ is the density function of b given u, and $\pi(u)$ denotes the prior of u.

The exact expression for the posterior is,

$$\pi(u|b) = \frac{\pi(b|u)\pi(u)}{\pi(b)},$$

where $\pi(b) = \int \pi(b|u)\pi(u))du$ is called the evidence. The evidence can be extremely difficult to calculate since it in most cases cannot be solved analytically and numerical integration becomes expensive if there are more than just a few dimensions.

However, the evidence does not depend on the parameters and often simply plays the role of a normalizing constant. We will only consider cases where the posterior is only known up to a multiplicative constant, and we will later in this chapter introduce sampling techniques that allow us to avoid explicit calculation of the evidence.

5.1.1 Likelihood function

The likelihood function relates the observed data to the parameters, i.e., it constitutes a mechanism through which the data informs the posterior. The likelihood is specified by the mathematical model that maps the unknown parameters to the observations and by the assumptions made regarding the density of the measurement errors.

If we assume additive independent and identically distributed (iid) Gaussian noise with zero mean, then we can formulate the statistical linear forward problem,

$$b = Au + \varepsilon, \tag{5.1}$$

with $\varepsilon \sim \mathcal{N}\big(0, \sigma^2 I_{m\times m}\big)$, where b, u and ε are random variables representing measurements, parameters and measurement errors, respectively. The statistical model in (5.1) corresponds to b being Gaussian with mean Au and variance σ^2, i.e.,

$$\pi\big(b|u,\sigma^2\big) = \left(\frac{1}{2\pi\sigma^2}\right)^{\frac{m}{2}} \exp\left(-\frac{1}{2\sigma^2}\|Au-b\|_2^2\right). \tag{5.2}$$

The likelihood function is a function of u and σ^2 for a specific observation of b, b^{obs}, i.e.

$$L(u,\sigma^2|b^{\text{obs}}) = \left(\frac{1}{2\pi\sigma^2}\right)^{\frac{m}{2}} \exp\left(-\frac{1}{2\sigma^2}\|Au-b^{\text{obs}}\|_2^2\right). \tag{5.3}$$

If we fix σ^2 and simply seek the parameters u that maximize the likelihood function, then the solution that maximizes (5.3), also referred to as a maximum likelihood (ML) estimate, is equivalent to a least-squares solution in

(4.6). This can easily be shown by utilizing the fact that the Gaussian density is log-concave with respect to u, and thus maximizing the likelihood is equivalent to minimizing the negative log-likelihood, i.e.,

$$\underset{u}{\text{minimize}} \quad -\log\big(L(u, \sigma^2|b^{\text{obs}})\big)\,.$$

Inserting the expression from (5.3) we get

$$\underset{u}{\text{minimize}} \quad \frac{m}{2}\log\left(\frac{1}{2\pi\sigma^2}\right) + \frac{1}{2\sigma^2}\|Au - b^{\text{obs}}\|_2^2,$$

and by ignoring terms independent of u, we arrive at the least-squares problem

$$\underset{u}{\text{minimize}} \quad \|Au - b^{\text{obs}}\|_2^2.$$

5.1.2 Prior modeling

The prior density represents our prior beliefs about our parameters and allows us to formally include such knowledge into the model, e.g., clinical knowledge from previous studies or structural information about the image.

In the absence of high-quality data, the prior can help fill in missing information and stabilize the solution. The inclusion of the prior has a regularizing effect on the parameter estimation problem in the sense that the parameter estimates become less sensitive to measurement noise and can in some cases improve identifiability of the parameters [15, 45].

There are several different strategies for constructing a suitable prior, e.g., see [7, 15], but here we will focus on the Gaussian prior which is used in both paper A and C. If we assume that the prior of u is iid Gaussian with zero mean and covariance δI, then we can write the prior by,

$$\pi(u|\delta) = \left(\frac{1}{2\pi\delta}\right)^{\frac{n}{2}} \exp\left(-\frac{1}{2\delta}\|u\|_2^2\right). \tag{5.4}$$

The Gaussian prior has a strong connection to Tikhonov regularization since maximizing the Gaussian prior with respect to u is equivalent to minimizing the squared two-norm of u.

The Gaussian prior is a very convenient choice if the measurement noise is assumed Gaussian since the Gaussian prior is a conjugate prior for the likelihood function of a Gaussian. Conjugacy refers to the property that the prior and posterior have the same parametric form, i.e., the resulting posterior for u will also be Gaussian if both likelihood and prior are Gaussian. The use of conjugate priors is often very advantageous since the corresponding posterior will have a closed-form expression.

5.1.3 Posterior

For this specific choice of likelihood in (5.3) and prior in (5.4), we can write the posterior of u as,

$$\pi\big(u|b,\delta,\sigma^2\big) \propto \exp\!\left(-\frac{1}{2\sigma^2}\|Au-b\|_2^2 - \frac{1}{2\delta}\|u\|_2^2\right), \tag{5.5}$$

which corresponds to the Gaussian posterior,

$$u|b,\sigma^2,\delta \sim \mathcal{N}\!\left(\left(\frac{1}{\sigma^2}A^TA+\frac{1}{\delta}I\right)^{-1}\frac{1}{\sigma^2}A^Tb, \left(\frac{1}{\sigma^2}A^TA+\frac{1}{\delta}I\right)^{-1}\right). \tag{5.6}$$

However, for large-scale problems it is neither feasible nor practical to compute the exact statistical properties (mean and covariance) due to the computational expenses associated with the covariance matrix. Thus, just as for the least-squares solution in (4.7), we need to employ iterative methods to approximate the statistical properties of the posterior.

If we assume that both σ^2 and δ are known a priori, then the solution which maximizes the posterior of u, also referred to as a maximum a posteriori (MAP) estimate, is found by maximizing (5.5) or equivalently solving

$$\underset{u}{\text{minimize}} \quad \frac{1}{2\sigma^2}\|Au-b\|_2^2 + \frac{1}{2\delta}\|u\|_2^2,$$

which is equivalent to the optimization problem in (4.8) for the specific choice of $\lambda = \frac{\sigma^2}{\delta}$.

5.1.4 Hyperpriors

We often do not know σ^2 or δ and finding suitable hyperparameters can be a tedious task, just as for the regularization parameter λ in (4.8). In a Bayesian framework, the natural approach is to consider the hyperparameters as random variables and include them in the inference process by assigning prior densities. The prior densities are called hyperpriors, and we denote the hyperpriors for σ^2 and δ by $\pi(\sigma^2)$ and $\pi(\delta)$.

We have no information about the hyperparameters, except for the fact that they model variances and therefore cannot be negative. As commonly done when the prior information is of questionable accuracy, it is often better to choose a so called non-informative prior with minimal influence on the inference. [45]

Commonly, an inverse gamma hyperprior is chosen due to its non-negative property and its non-informative nature. Moreover, the inverse gamma hyperprior is often a convenient choice since it is a conjugate prior to the Gaussian, i.e., the posterior density of the hyperparameter will also be an inverse gamma. [7] The density of the inverse gamma is defined by,

$$\pi(\gamma) = \frac{\beta^\alpha}{\Gamma(\alpha)}\gamma^{-\alpha-1}\exp\big(-\beta\gamma^{-1}\big),$$

where α is the shape parameter and β is the scale parameter. Bardsley [7] proposes the choice $\alpha = 1$ and $\beta = 10^{-4}$, making the hyperpriors exponentially distributed with small decay parameters.

We can formulate the full posterior for u, σ^2 and δ, i.e.,

$$\pi\big(u, \delta, \sigma^2 | b\big) \propto \pi\big(b | u, \sigma^2\big)\, \pi(u|\delta)\, \pi(\sigma^2)\pi(\delta). \tag{5.7}$$

However, we only know the full posterior up to a multiplicative constant and thus we need to employ sampling techniques to characterize the full posterior.

5.2 Sampling

In many scientific problems, it is often infeasible or impractical to compute the statistical properties of the full posterior, e.g., when the full posterior is only known up to multiplicative constant, does not have a closed-form solution, and/or is high-dimensional. Instead, we can employ sampling techniques to generate samples that, as the number of samples increases, converge to the true full posterior in terms of statistical properties.

For the remainder of the chapter, we will let x denote the unknown parameters, i.e., for the problem in (5.7) $x = [u, \sigma^2, \delta]$. Moreover, we assume that we know the target density $\pi(x)$ up to a multiplicative constant $C > 0$, i.e.,

$$\pi(x) = \frac{f(x)}{C} \geq 0, \quad \forall\, x, \tag{5.8}$$

and that we can evaluate f but it is extremely impractical if not infeasible to compute C which is exactly the case for the full posterior in (5.7).

First, we consider a very simple method, namely the accept and reject method to sample from non-standard densities. We briefly consider this method since it illustrates elegantly the concept of sampling and serves as a prelude to the later introduced Metropolis Hastings (MH) algorithm.

Suppose f is a non-negative function that satisfies (5.8) on the interval $[a, b]$ and is bounded by a constant $M > 0$ such that $0 \leq f(x) < M$ for all $x \in [a, b]$. The fundamental theorem of simulation [42] states that if we generate uniform samples (x^i, u^i) on the set $[a, b] \times [0, M]$ and only accept the set of samples that satisfy $u^i < f(x^i)$, then the normalized histogram of the accepted samples gives an approximation of the density of π. Thus, we can approximate samples from π by generating uniform samples and evaluating f. The accept and reject method is illustrated in Figure 5.1.

The accept and reject method is a very simple and elegant method for computing samples from a possibly very complicated density. However, the performance of the accept and reject method is highly dependent on the choice of M and the interval $[a, b]$. For the one-dimensional case, these are

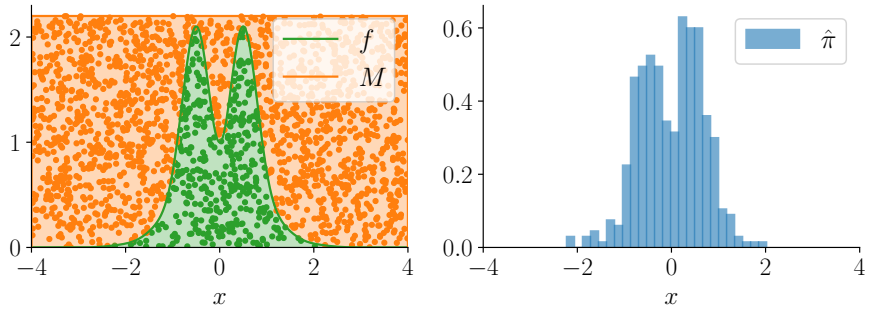

Figure 5.1: Illustration of the accept and reject method for the one-dimensional case with the function $f(x) = \exp(-x^2)\left(\sin\left(\frac{1}{2}x\right)^2 + 3\cos(x)^2\sin(2x)^2 + 1\right)$. (Left) Points (x^i, u^i) are sampled uniformly on the domain $[a, b] \times [0, M]$. The i'th sample is accepted if $u^i < f(x^i)$ (green dots) and rejected otherwise (red dots). (Right) Normalized histogram of the accepted samples.

easy to choose since we can visually inspect f and find an appropriate support. However, it can be very hard to find a suitable choice of support for high-dimensional cases. The interval $[a, b]$ that supports f is generally unknown, and M must be larger than the maximum of f, but if chosen too large we end up rejecting too many samples leading to computational inefficiency. Moreover, each sample is completely independent of previous samples, and uniform sampling in a high-dimensional space is highly inefficient since the support of f is often restricted to only a small subset. Thus, we risk ending up with a large number of rejected samples.

5.2.1 Markov Chain Monte Carlo

Markov Chain Monte Carlo (MCMC) is a sampling technique for approximating the statistical properties of the posterior. In MCMC, we generate a sequence of random samples $(x^1, x^2, \ldots, x^s)$ whose density asymptotically approaches the target density π. A Markov chain is a sequence of samples where the next sample x^{i+1} only depends on the current sample x^i, i.e.,

$$\pi(x^{i+1}|x^i, x^{i-1}, \ldots, x^0) = \pi(x^{i+1}|x^i),$$

and the Markov chain is completely characterized by the transition kernel K,

$$K(x^i, x^{i+1}) = \pi(x^{i+1}|x^i),$$

which describes the probability of moving from the current state x^i to the next state x^{i+1}. Moreover, we will consider only time-homogeneous Markov

chains satisfying,

$$K(x^i = x, x^{i+1} = y) = K(x^0 = x, x^1 = y), \quad \forall i,$$

which implies that the transition probability is independent of how the chain evolves, i.e., the probability of moving from a sample x to y is the same for all time steps.

A finite time-homogeneous Markov chain which is irreducible, aperiodic and satisfies the detailed balance equation is guaranteed to converge to the target density [44]. A Markov chain is irreducible if any state x^i can be reached from any other state x^j in a finite number of steps and aperiodic if $K(x, x)$ is non-zero for every state x. The detailed balance equation is given by

$$\pi(y)K(y, x) = \pi(x)K(x, y), \tag{5.9}$$

where π is the target density that we know up to a multiplicative constant and K is the transition kernel. The key challenge for MCMC methods is how to construct a transition kernel that satisfies (5.9). We refer the reader to [42, 44] for a thorough theoretical introduction to MCMC.

5.3 Metropolis Hastings

A commonly used MCMC method is the MH algorithm. We introduce the MH algorithm since the later introduced Gibbs sampler and NUTS can be interpreted as special cases of the MH algorithm. The introduction to MH is conceptual, and we refer the reader to [15] for an in depth introduction.

The goal of MH is to construct a transition kernel that satisfies (5.9). Let the density q be a proposal generating kernel such that if we are at the current state x^i then $q(x^i, y)$ generates a candidate y. Note that q should be chosen such that it is easy to obtain samples and preferably is close to the target density. If we sample a candidate y from q and it happens that y satisfies

$$\pi(y)q(y, x^i) = \pi(x^i)q(x^i, y), \tag{5.10}$$

then we accept y and set $x^{i+1} = y$. However, in most cases we will have a situation where y does not satisfy the detailed balance equation, e.g.,

$$\pi(y)q(y, x^i) < \pi(x^i)q(x^i, y). \tag{5.11}$$

To correct for this imbalance, we introduce a correcting function $\theta(x^i, y)$ such that

$$\pi(y)\theta(y, x^i)q(y, x^i) = \pi(x^i)\theta(x^i, y)q(x^i, y), \tag{5.12}$$

where $\theta(x^i, y)$ can be interpreted as a probability that we make the move from x^i to y.

The problem has now been reduced to finding a suitable correcting function $\theta(x^i, y)$. The correcting function does not need to be symmetric and we may choose $\theta(y, x^i) = 1$ for simplicity. An interpretation of θ is that the choice $\theta(y, x^i) = 1$ corresponds to maximizing the chance of moving from y to x^i which is desirable since without the correction we make the move from y to x^i too rarely. The choice of $\theta(x^i, y)$ which satisfies (5.12) is therefore

$$\theta(x^i, y) = \frac{\pi(y)q(y, x^i)}{\pi(x^i)q(x^i, y)} < 1. \tag{5.13}$$

Note that if the inequality was flipped in (5.11) then we simply interchange x^i and y and set $\theta(x^i, y) = 1$. Hence, to ensure that the detailed balance equation is satisfied for the proposal generating kernel q, we pick the transition kernel

$$K(x^i, y) = \theta(x^i, y)q(x^i, y)$$

with

$$\theta(x^i, y) = \min\left\{\frac{\pi(y)q(y, x^i)}{\pi(x^i)q(x^i, y)}, 1\right\}. \tag{5.14}$$

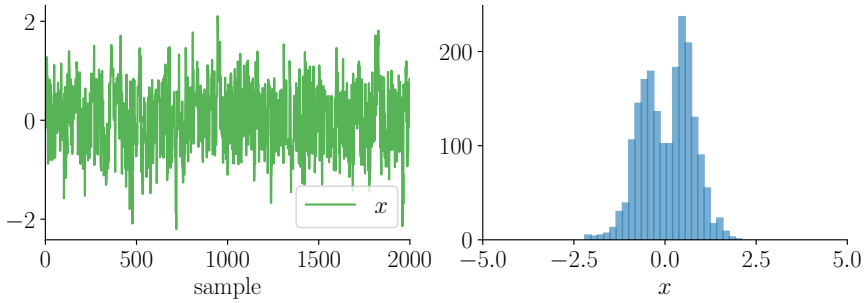

Figure 5.2: Illustration of the MH algorithm. The MCMC chain generated by the MH algorithm (left) and the histogram of the sample chain (right). The histogram closely resembles the target density depicted in Figure 5.1.

This choice of transition kernel K guarantees that the accepted samples converge asymptotically to the target density $\pi(x)$ since it satisfies the detailed balance equation. In practice, we generate from $q(x^i, y)$ a candidate y and accept it with probability $\theta(x^i, y)$. Implementation-wise, we introduce an accept or reject step such that for each proposed sample y, we generate a uniform random variable u on the interval $[0, 1]$ and if $u < \theta(x^i, y)$, we accept and set $x^{i+1} = y$, otherwise we reject and set $x^{i+1} = x^i$.

Note that the target density π appears as a ratio in the correcting function, and we can therefore write θ as,

$$\theta(x^i, y) = \min\left\{\frac{f(y)q(y, x^i)}{f(x^i)q(x^i, y)}, 1\right\}, \tag{5.15}$$

which means that we only need to know π up to a multiplicative constant.

One strength of the MH algorithm compared to the accept and reject method is that we use the information from previous sample to guide the evolution of the sample chain. Consequently, a weakness of the MH algorithm is that the samples are correlated, and it might require a lot of samples to approximate the target density. Figure 5.2 shows a one-dimensional example for the MH algorithm. The left figure shows the sample chain, and the right figure shows the corresponding histogram of the sample chain.

The main challenge of the MH algorithm is how to choose a suitable proposal generating density q, since if chosen naively the MH algorithm is essentially performing a random walk without taking information of the target density into account. However, choosing q in a non-naive way is especially challenging in the case of high-dimensional problems where the MH algorithm is prone to exhibit slow convergence.

Figure 5.3 illustrates the sensitivity of the MH algorithm to the choice of the proposal kernel and the correlation-rejection trade off. We have chosen the simple Gaussian density with zero mean and standard deviation σ as the proposal density. The value of the standard deviation σ can be interpreted as the step size for the MH algorithm. When chosen too small (top row, $\sigma = 0.5$), consecutive samples of the chain tend to be very close to each other which is a sign that the simulation is not moving quickly through its space. On the other hand, if we choose σ too large (bottom row, $\sigma = 50$), we end up rejecting too many proposals which might lead to inefficient exploration of the target density. Chosen adequately (middle row, $\sigma = 6$), we achieve satisfying exploration of the target density.

5.4 Gibbs sampler

The Gibbs sampler is a method that avoids the search for a suitable proposal generating density by using a coordinate-wise proposal. [16, 38] The idea is to reduce the sampling from a complex high-dimensional density to a sequence of simpler low-dimensional densities by sampling each parameter in turn from their conditional density while keeping the remaining parameters fixed. This is exactly the case for the Bayesian model in paper A.

We assume that we have n random variables $x_1, x_2, \ldots, x_n$, and that the target density is the joint density

$$\pi(x_1, x_2, \ldots, x_n),$$

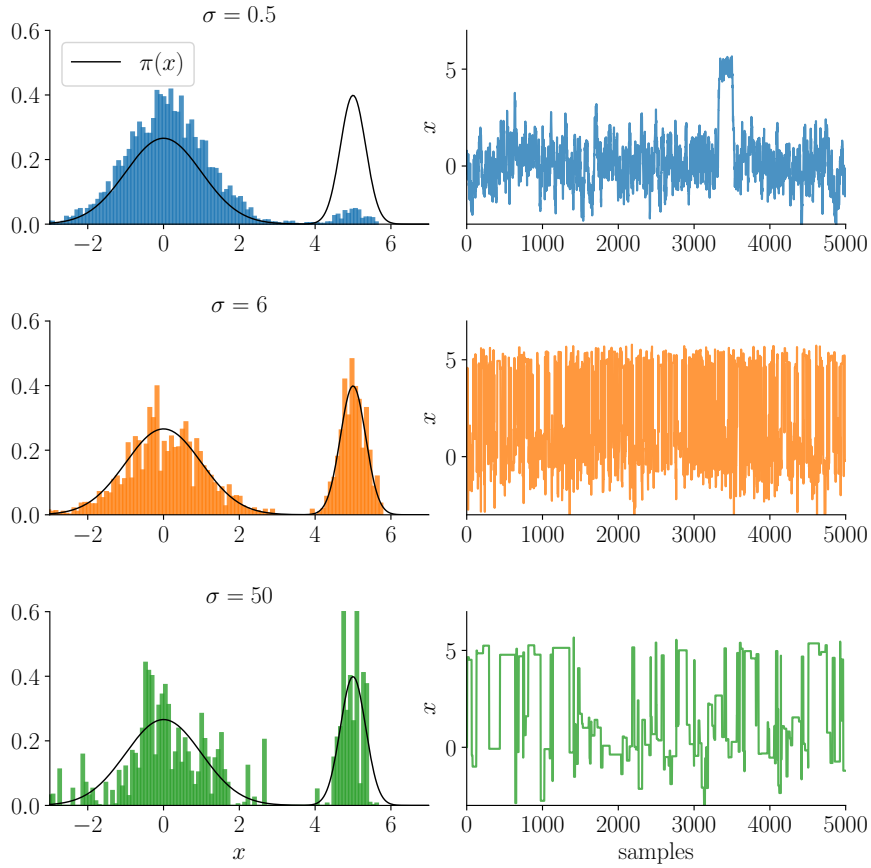

Figure 5.3: Illustration of the MH algorithm with Gaussian proposal density for different choices of variance (σ^2). (Left) depicts a histogram of the samples, the target density $\pi(x)$ chosen as $\pi(x) = \frac{4}{6}\pi_1(x) + \frac{2}{6}\pi_2(x)$ with π_1 being Gaussian distributed with $\mu_1 = 0$ and $\sigma_1 = 1$ and π_2 with $\mu_2 = 5$ and $\sigma = \frac{1}{3}$. (Right) shows the sample chains obtained by the MH algorithm.

and we have the conditional densities available

$$\pi(x_i|x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n),$$

for $1 \leq i \leq n$. The Gibbs sampler updates the elements sequentially by generating proposals from the conditional densities while keeping the remaining parameters fixed.

Suppose that we sample a candidate y to replace x_i from the conditional density of x_i, i.e.,

$$y \sim \pi(x_i|x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n).$$

To ease notation, we introduce $x \in \mathbb{R}^n$ as the vector with elements x_j for $1 \leq j \leq n$ and let $z \in \mathbb{R}^n$ denote the vector with elements $z_j = x_j$ for all $j \neq i$ and $z_i = y$. Moreover, we denote the vector with all the elements of x except x_i by x_{-i}.

We can rewrite the density for x by the rule for conditional densities [40], i.e.,

$$\pi(x) = \pi(x_i|x_{-i})\pi(x_{-i})$$

and similarly for the density of z we obtain

$$\pi(z) = \pi(y|z_{-i})\pi(z_{-i}) = \pi(y|x_{-i})\pi(x_{-i}),$$

where we have utilized that $z_{-i} = x_{-i}$. The proposal generating kernel for moving from x to z is simply

$$q(x, z) = \pi(y|x_{-i}),$$

and the probability of moving from z to x is

$$q(z, x) = \pi(x_i|z_{-i}) = \pi(x_i|x_{-i}).$$

The correcting function $\theta(x, z)$ for the MH algorithm can then be written as

$$\begin{aligned}\theta(x, z) &= \frac{\pi(z)q(z, x)}{\pi(x)q(x, z)} \\ &= \frac{\pi(y|x_{-i})\pi(x_{-i})\pi(x_i|x_{-i})}{\pi(x_i|x_{-i})\pi(x_{-i})\pi(y|x_{-i})} \\ &= 1.\end{aligned}$$

Hence, the Gibbs sampler is indeed a special case of the MH algorithm where the proposal y from the conditional density is always accepted since the acceptance probability is one.

So far, we have considered breaking the full set of parameters into scalar products. However, in some cases it may be advantageous to formulate the Gibbs sampler for blocks of parameters. Thus, the Gibbs sampler is ubiquitous in Bayesian inference since we often have simple conditionals for the parameters and hyperparameters which form two natural blocks or stages for the Gibbs sampler [47]. Thus, this method is applicable exactly for the full posterior in (5.7) when choosing the conjugate priors and hyperpriors. This specific type of Gibbs sampler is often referred to as the hierarchical Gibbs sampler.

To illustrate the strengths and weakness of the Gibbs sampler, we can consider the two-dimensional Gaussian characterized by

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix} \right), \tag{5.16}$$

where ρ is a measure of the correlation between the two components. The conditionals can be shown to be the following one-dimensional Gaussian densities [21]

$$x_1|x_2 \sim \mathcal{N}\big(\rho x_2, (1-\rho^2)\big)\,,$$
$$x_2|x_1 \sim \mathcal{N}\big(\rho x_1, (1-\rho^2)\big)\,.$$

Thus, we can reduce the sampling to two one-dimensional Gaussian samples.

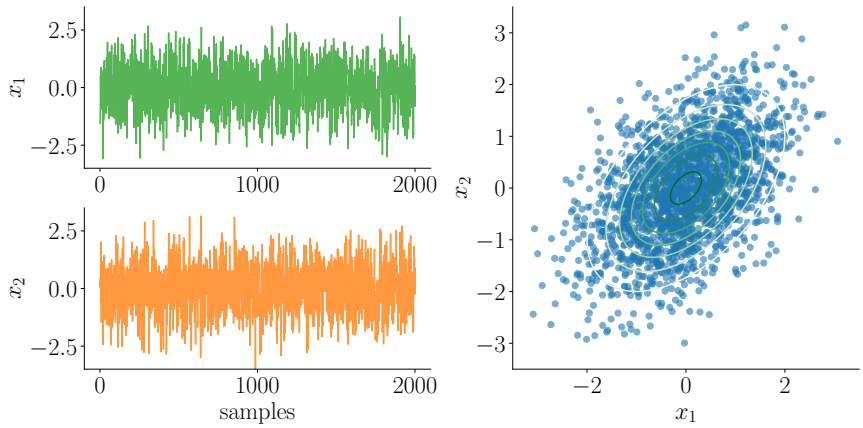

Figure 5.4: Illustration of the Gibbs sampler for the two-dimensional Gaussian density specified by (5.16) with $\rho = 0.5$. The sample chain moves rapidly through the space and the chains appear well-mixed (left). The green-shaded ellipsoids show the contour of the target density.

The main strength of the Gibbs sampler is that it is parameter-free in the sense that it does not require choosing a proposal as required in the MH algorithm. However, the Gibbs sampler is severely challenged if the different blocks of the parameters are highly correlated. High correlation between blocks or components may lead to painfully high correlation of the Markov chain since we only move a very small step in each conditional sample. [44] If we consider the two-dimensional Gaussian in (5.16) and increase the correlation between the two coordinates such that $\rho = 0.95$, the samples become highly correlated as illustrated in Figure 5.5.

5.5 Hamiltonian Monte Carlo

Hamiltonian Monte Carlo (HMC) is an MCMC technique that utilizes Hamiltonian dynamics to propose a new state by taking a series of first-order gradient informed steps that is distant from the current state with high probability of acceptance, even in the case of highly correlated parameters. Thus, the

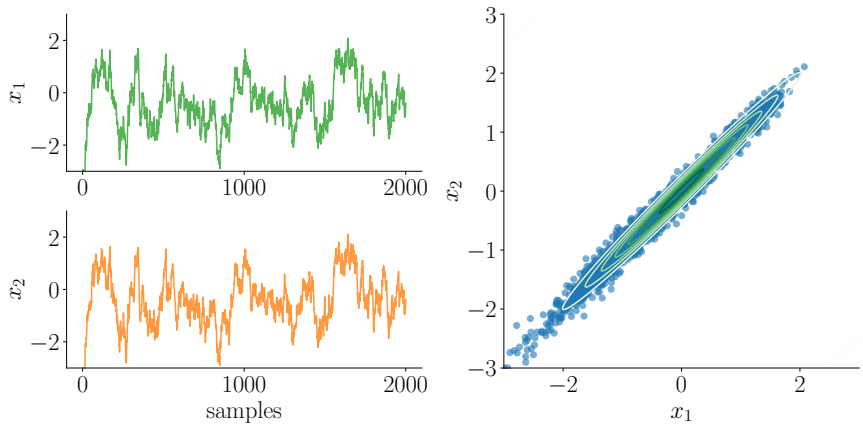

Figure 5.5: Illustration of the Gibbs sampler for the two-dimensional Gaussian density specified by (5.16) with $\rho = 0.95$. The correlation between the two coordinates is high and the Gibbs sampler only moves a very small step in each conditional sample (left). The green-shaded ellipsoids show the contour of the target density.

HMC can eliminate the slow convergence that may haunt the MH and Gibbs samplers in case of correlated or high-dimensional parameter spaces. The motivation for briefly explaining the concept of HMC is that a variant of the HMC sampler called NUTS is used in Paper C. Note that the explanation of the HMC method is only meant to provide a conceptual understanding, and we refer the reader to [13, 32] for an in dept introduction to HMC.

A intuitive interpretation of the HMC is that we can consider the sample chain as a particle moving through space where x^i is the current position of the particle. [28] The movement of the particle is driven by the potential and kinetic energy of the particle. We can interpret the negative logarithm of the target density as the potential of the particle such that, if no kinetic energy is applied to the particle, it will follow the gradient and move towards regions of higher probability and get stuck at a mode of the target density. However, we are not interested in the particle getting stuck at a mode since this would give a poor exploration of the target density. Thus, to ensure a sufficient exploration of the target density we introduce auxiliary momentum variables. [11]

Let x denote the parameters of interest and r denote the corresponding auxiliary momentum variables. We can write the joint density

$$\pi(r, x) = \pi(r|x)\pi(x),$$

which defines the Hamiltonian function

$$\begin{aligned} H(r,x) &= -\log(\pi(r,x)) \\ &= -\log(\pi(r|x)) - \log(\pi(x)) \\ &= E_K(r,x) + E_P(x), \end{aligned}$$

where $E_K(r,x) = -\log(\pi(r|x))$ is the kinetic energy and $E_P(x) = -\log(\pi(x))$ is the potential energy (our target density). As commonly done, we choose the kinetic energy to be a multivariate Gaussian, i.e.,

$$r \sim \mathcal{N}(0, I) \tag{5.17}$$

such that the kinetic energy is independent of the parameters, i.e., $E_K(x,r) = E_K(r)$.

The evolution of x and r is determined by Hamilton's equations,

$$\begin{aligned} \frac{\partial x}{\partial t} &= \frac{\partial H(x,r)}{\partial r}, \\ \frac{\partial r}{\partial t} &= -\frac{\partial H(x,r)}{\partial x}, \end{aligned}$$

and one important property of the Hamiltonian dynamics is that it is time reversible, i.e., if we reverse time then the trajectory will remain the same. Thus, there exists functions ϕ_x and ϕ_r such that $x_T = \phi_x(x_0, r_0)$ and $r_T = \phi_r(x_0, r_0)$ and if we reverse time then $x_0 = \phi_x(x_T, -r_T)$ and $r_0 = \phi_r(x_T, -r_T)$. Another important property is that the Hamiltonian dynamics conserve energy, i.e.,

$$\frac{\partial H(x,r)}{\partial t} = 0,$$

and lastly, the Hamiltonian dynamics preserve volume in the (x,r) space. This property ensures that we do not need to include a Jacobian for the mapping determined by the Hamiltonian dynamics.

We want to compute samples from $\pi(x,r)$ which we know only up to a multiplicative constant through $H(x,r)$, i.e.,

$$\pi(x,r) \propto \exp\{-H(x,r)\}.$$

The idea is that for a current state (x^i, r^i), we propose a candidate (y_T, r_T) and then use an MH acceptance step to evaluate whether or not the candidate should be accepted. First, we generate a random momentum vector r_0 from (5.17), and then we apply the Hamiltonian dynamics to the current position $y_0 = x^i$ and momentum r_0 and let the dynamics run for time T to obtain a proposal (y_T, r_T).

The Hamiltonian dynamics is deterministic and the probability of moving from (y_0, r_0) to (y_T, r_T) is therefore always 1. However, the probability of moving from (y_T, r_T) to (y_0, r_0) is zero. Thus, we modify our proposal by

negating the momentum variable, i.e., $(y_T, -r_T)$ to ensure that we have a valid proposal. In fact, this negation makes the proposal symmetric since the probability of moving from (y_0, r_0) to $(y_T, -r_T)$ and from $(y_T, -r_T)$ to (y_0, r_0) is 1 and therefore the proposal cancels out in the MH acceptance probability. Hence, we accept the proposed candidate with probability

$$\begin{aligned}\theta(y_0, r_0, y_T, -r_T) &= \min\left(1, \frac{\pi(y_T, -r_T)}{\pi(y_0, r_0)}\right) \\ &= \min\left(1, \frac{\exp(-H(y_T, -r_T))}{\exp(-H(y_0, r_0))}\right) \\ &= \min(1, \exp(-H(y_T, -r_T) + H(y_0, r_0))) .\end{aligned}$$

Note that since the Hamiltonian conserves energy then

$$H(y_0, r_0) = H(y_T, -r_T),$$

and thus, we always accept the proposal in theory. However, in practice the Hamiltonian equations are too complex to be solved analytically and numerical integration must be applied. The numerical methods introduce errors, and therefore we need the correction step to ensure that the detailed balance equation is satisfied in practice.

The most used numerical integrator for HMC is the leapfrog integrator [43] which depends on two hyperparameters, the number of steps and the step size. These two hyperparameters are crucial for the performance of the HMC and require non-trivial hand-tuning to obtain efficient sampling.

5.5.1 NUTS

To make HMC widely applicable, we need a method for tuning the hyperparameters of HMC automatically. The No-U-Turn Sampler (NUTS) provides exactly such a framework. NUTS is an extension to HMC that eliminates the need for hand-tuning the hyperparameters [28]. The idea is to stop the simulation when taking more steps, no longer increases the distance between the current state x^i and the proposed state y_T. Thus, we simulate the trajectory until we start moving back towards the current state whereof the name No-U-Turn originates. We refer the reader to [28] for details.

The great advantage of the NUTS (and HMC in general) is that it produces less correlated samples with a higher acceptance rate compared to MH and the Gibbs sampler which is illustrated by comparing the NUTS sample chain in Figure 5.6 to the Gibbs sample chain in Figure 5.5.

However, there is a price to pay. The price for the increased performance is that we need to compute the gradients of the Hamiltonian function, and thus NUTS requires the log-density of the parameters to be continuously differentiable. Computing the gradients can be troublesome for large-scale problems which might make the HMC impractical. However, for small to

moderate scale problems, NUTS provides an efficient method for sampling the full posterior for inverse problems.

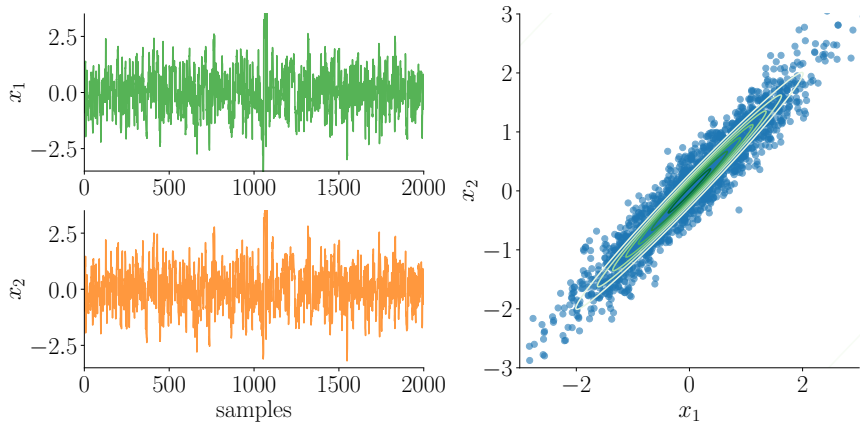

Figure 5.6: Illustration of the NUTS sampler for the two-dimensional Gaussian density specified by (5.16) with $\rho = 0.95$. The sample chain moves rapidly through the space and the chains appear well-mixed (left). The green-shaded ellipsoids show the contour of the target density.

5.6 Convergence diagnostics

In theory, with enough samples, our chain of samples will well-approximate the target density. However in practice, sampling can be expensive, especially for large-scale problems, and we only have a limited number of samples. Thus, we need convergence diagnostics to ensure that our sample chain has converged such that the statistical properties of the sample chain well-approximate our target density. We will not give a rigorous introduction to convergence diagnostics for the sampling algorithms, but briefly discuss methods used for assessing the convergence of the sample chain, see e.g. [42] for details.

The initial stage of a sample chain is called burn-in since the initial samples move from a region with low probability to a region of relative high probability. The samples in the burn-in stage should be removed since they can otherwise introduce a bias in the statistics computed based on the sample chain. Once the burn-in phase has been removed from the chain we say that the chain has converged, and if the chain has enough samples, we can compute statistics that well-approximate the statistics of the target density.

The simplest method to assess whether the sample chain has converged is by visual inspection of the one-dimensional sample chains. The visual inspec-

tion suggests that the sample chain has converged if it appears well-mixed. The chain appears well-mixed if the autocorrelation between neighboring samples is low and the samples fluctuate around a stationary mean. Figure 5.3 shows an example of a chain that appears well-mixed (middle row) and a chain with poor mixing properties (bottom row).

There exist quantitative methods for assessing convergence of the sample chain. Two widely used tests are the Geweke test [22] that compares the statistical properties of the first and last part of the sample chain to determine convergence, and the potential scale reduction statistic $\hat{R}$ [23] that based on several generated sample chains measures the ratio of the average variance of samples within each sample chain to the variance of the pooled samples across chains.

We have now introduced the foundation of the methods, i.e., Bayesian modeling and UQ, which are used in paper A, B and C, and we will now turn to the individual contributions of the papers.

CHAPTER 6

Computed tomography with uncertain source

Relevant paper

A) **Katrine O. Bangsgaard** and Martin S. Andersen, "A statistical reconstruction model for absorption CT with source uncertainty", *Inverse Problems*, 2021, [3].

6.1 Motivation and goals

In CT we often assume that we know the source intensity and that it is constant throughout the data acquisition process. However, for time demanding experiments, e.g., large-scale synchrotron neutron facilities, the source is prone to vary in intensity during the data acquisition, and we denote this phenomenon intensity drift. Intensity drift may occur because of source instabilities, vibrating beamline components, time-varying detector properties, etc. The result is a mismatch between the reconstruction model and the underlying physics which introduces systematic model errors that can compromise the image quality.

Commonly, the source intensity and detector response are estimated based on flat-field images prior to reconstruction. Existing methods for estimating the detector response can be divided into four main categories:

- Acquisition methods
- Preprocessing methods
- Post-processing methods
- Extended measurement models

Acquisition methods perturb the object between projections to smear out the systematic error in the sinogram. This can lead to notable improvements, but it is not always feasible nor practical. Preprocessing methods work by manipulating the projection images prior to the reconstruction, i.e., they modify the measured data to fit the model. Post-processing methods remove or reduce rings in the image domain after the reconstruction step. The extended reconstruction models either implicitly or explicitly include the detector response as an unknown quantity to be estimated along with the attenuation function, i.e., the model is modified to fit the measured data. However, these approaches are oblivious to intensity drift.

There exist methods that extract information about the intensity drift from flat-field images and use this information to manipulate the measurement data prior to the reconstruction. Commonly, linear interpolation based on flat-field images is used to correct for intensity drift, provided that flat-field images were acquired both before and after the projection images. Another heuristic method for estimating the intensity drift is by using air-pixel, i.e., detector pixels that are not affected by the object in some or all projections. However, this method is only applicable if the object does not fill the entire sinogram. Both methods can be viewed as preprocessing methods since they manipulate the sinogram prior to reconstruction of the image. Moreover, they depend on the noisy flat-fields and air-pixels which, just as for the detector response, may introduce systematic errors in the normalized sinogram and consequently in the image.

6.2 Contributions

The main contributions of paper A can be summarized as follows. We propose a CT reconstruction model that allows us to jointly estimate both reconstruction image, intensity drift, and the detector response. We formulate the model in a Bayesian framework and propose two models, an optimization model computing the MAP estimate denoted Approximated Maximum a Posteriori (AMAP) and a full Bayesian model denoted Approximated Posterior Mean (APM).

6.3 Methods

Consider a modified version of the discretized CT model in (2.5),

$$y = (\omega \otimes \nu) \exp(-Au)\,, \tag{6.1}$$

where $y \in \mathbb{R}^{rp}$ is a vector of measurements, $u \in \mathbb{R}^n$ is the vector of unknown attenuation coefficients, $\nu \in \mathbb{R}^r$ is the effective measured intensity incident on the object, $A \in \mathbb{R}^{rp \times n}$ is the system matrix, and $\omega \in \mathbb{R}^p$ models the time-varying intensity, i.e., ω_j is the intensity at the j'th projection.

For convenience, we introduce a change of variables

$$\nu = \operatorname{diag}(\hat{\nu}) \exp(-v)\,, \tag{6.2a}$$
$$\omega = \operatorname{diag}(\hat{\omega}) \exp(-w)\,, \tag{6.2b}$$

where v and w are the new variables and $\hat{\omega}$ and $\hat{\nu}$ are estimates presenting prior knowledge, e.g., obtained by flat-field measurements and air-pixels.

Inserting (6.2) into (6.1) gives the extended measurement model,

$$y = (\hat{\omega} \otimes \hat{\nu}) \exp(-Au - Gv - Hw)\,, \tag{6.3}$$

where $G = (1_p \otimes I_r)$ and $H = (I_p \otimes 1_r)$. We can rewrite (6.3) into a linear system of equations, i.e.,

$$b = Au + Gv + Hw,$$

where $b = -(\log(y) - \log(\hat{\omega} \otimes \hat{\nu}))$. To regularize the problem, we formulate the model in a Bayesian framework and model the measurement noise as a weighted Gaussian, i.e.,

$$b|u, v, w \sim \mathcal{N}(Au + Gv + Hw, \Sigma_b)\,, \tag{6.4}$$

where $\Sigma_b = \operatorname{diag}(y)^{-1}$ and assign priors to u, v and w, i.e.,

$$u|\delta \sim \mathcal{N}(0, \delta I_n) \tag{6.5a}$$
$$v|\alpha \sim \mathcal{N}(0, \alpha I_r) \tag{6.5b}$$
$$w|\beta \sim \mathcal{N}(0, \beta I_p)\,. \tag{6.5c}$$

For a fixed set of the hyperparameters, we can compute the AMAP estimate by solving

$$\underset{u,v,w}{\text{minimize}} \quad \|Au + Gv + Hw - b\|^2_{\Sigma_b^{-1}} + \frac{1}{2\delta}\|u\|_2^2 + \frac{1}{2\alpha}\|v\|_2^2 + \frac{1}{2\beta}\|w\|_2^2. \tag{6.6}$$

A full Bayesian model (AMP) is formulated by assigning the hyperparameters with hyperpriors and use a hierarchical Gibbs sampler to compute the posterior statistics such as posterior mean and 95% CIs.

6.4 Results

In this section we briefly summarize the main findings in paper A. We compare our method to a commonly used preprocessing method proposed by Münch et al. [37] which combines wavelet and Fourier filtering to mitigate ring artifacts in the reconstruction. We denote a preprocessed reconstruction with the prefix P. The preprocessing method is either combined with

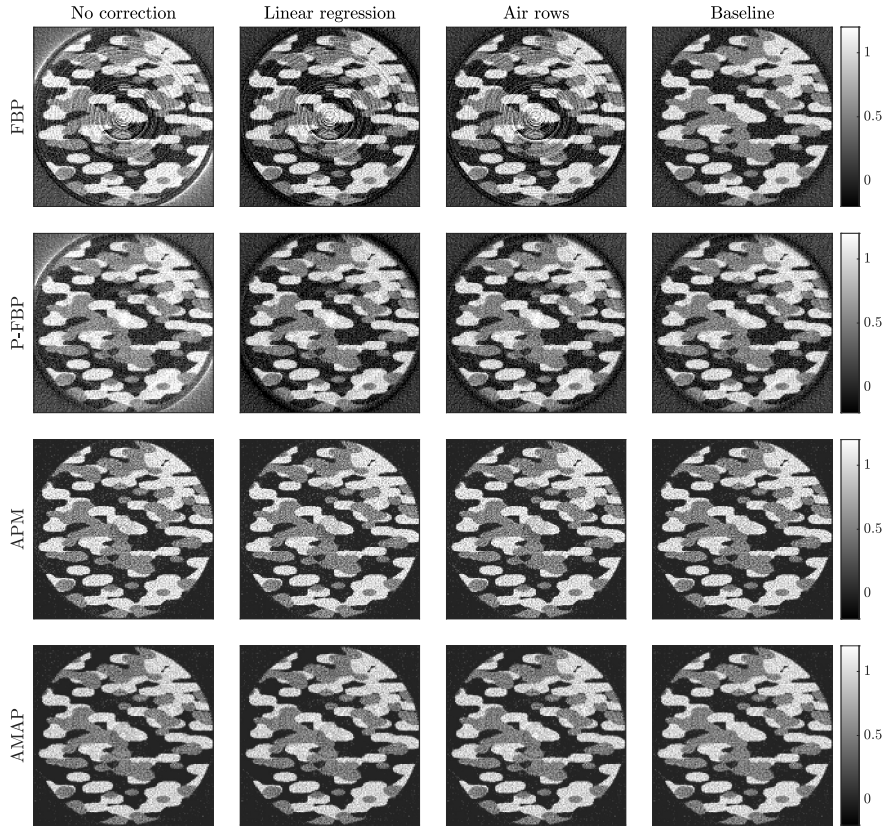

Figure 6.1: Reprint of main result figure from paper A.

the conventional FBP reconstruction method or the weighted least squares (WLS) reconstruction method which solves

$$\underset{u,v,w}{\text{minimize}} \quad \|Au + Gv + Hw - b\|^2_{\Sigma_b^{-1}},$$

with $\Sigma_b = \text{diag}(y)^{-1}$.

Figure 6.1 demonstrates that the AMAP and APM models can reduce artifacts and are competitive with existing methods. We note that although the computational cost associated with the new methods exceeds the cost of many existing methods, the extended model offers a clear advantage when working with datasets from challenging experimental setups that are expensive and/or difficult to improve.

6.4.1 Neutron Data Experiment

The findings of paper A have all been for simulated data, and we will now present results which have not been published elsewhere, where the proposed models are tested on a real neutron dataset. The neutron dataset [31] has been recorded at the IMAT Beamline at the ISIS Neutron and Muon Source based at the STFC Rutherford Appleton Laboratory, Harwell, UK [33]. During data acquisition, 186 projection images were acquired using the golden angle ratio for the projection angles with an acquisition time of 30 minutes per projection image. Six flat-fields have been recorded; a single flat-field before the experiment and five flat-fields after.

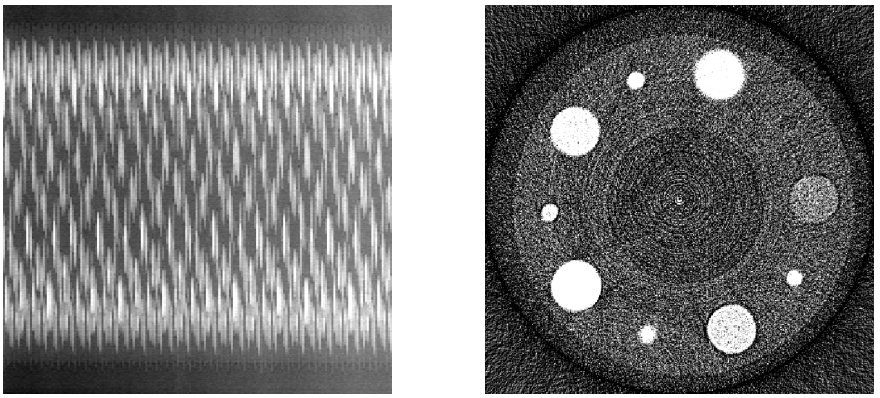

(a) Normalized sinogram. **(b)** FBP reconstruction.

Figure 6.2: Normalized sinogram in (a) and FBP reconstruction with conventional flat-field correction in (b). The sinogram reveals an intensity drift during the time of acquisition, and the FBP reconstruction is affected by noise and source model errors. All images are displayed in display range $[-0.2, 0.6]$.

The normalized sinogram obtained by conventional flat-field correction and the corresponding FBP reconstruction are depicted in Figure 6.2. The sinogram in Figure 6.2a shows a gradual change in the intensity of the air-pixels which is an indication that an intensity drift has occurred during the data acquisition. The corresponding FBP reconstruction in Figure 6.2b is affected by the source model errors. Moreover, we see that the FBP reconstruction is of poor quality due to the low signal-to-noise ratio. Thus, we test the different ring reduction techniques in combination with the WLS reconstruction model instead of FBP. We extend the WLS model to include Tikhonov regularization and non-negativity constraints [8], i.e., the prior modeling for the WLS model is chosen as,

$$\varphi(u) = P_+(u) + \frac{1}{2}\|u\|_2^2,$$

where

$$P_+(u) = \begin{cases} u & u \geq 0, \\ 0 & u < 0. \end{cases}$$

APM and AMAP models are based on the Bayesian model,

$$b|u, v, w, \lambda \sim \mathcal{N}\big(Au + Gv + Hw, \lambda^{-1}\Sigma_b\big)\,,$$

priors defined in (6.5) with the addition of a non-negativity constraint for u (see [3] for details) and Gamma hyperpriors with $\alpha_\gamma = 1$ and $\beta_\gamma = 10^{-4}$ for the hyperparameters.

	Geweke p-value	95% CI
λ	0.998	[1688, 1723]
δ	0.998	[16.07, 16.36]
α	0.971	[589.2, 780.6]
β	0.952	[0.07431, 0.10986]

Table 6.1: Chain statistics for the hyperparameter chain depicted in Figure 6.4a.

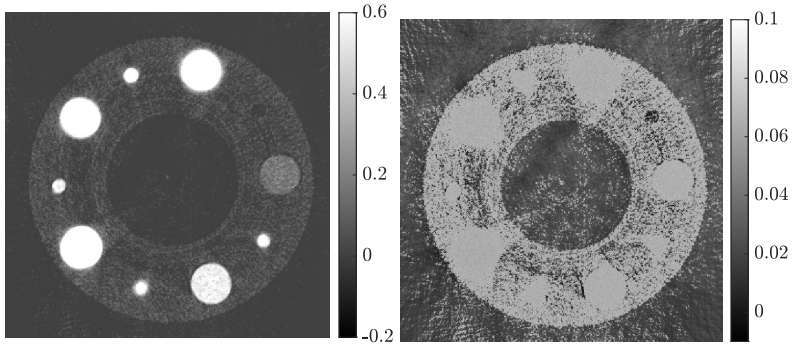

(a) APM. **(b)** Width of 95% CI.

Figure 6.3: APM reconstruction for the neutron dataset. The posterior mean is shown in (a) and (b) shows the width of the 95% CI.

The reconstruction for APM is shown in Figure 6.3, and chain statistics are shown in Figure 6.4 and listed Table 6.1. Figure 6.4 shows that the hyperparameter chains appear well-mixed which indicates that the sampler has converged. This observation is supported by Table 6.1 where the Geweke values are all close to one. The intensity of the APM reconstruction seems a bit too low and some remaining ring artifacts are still visible which indicates that, as for the simulated experiments in paper A, the Gibbs sampler has likely converged to a too low value for the hyperparameter for the detector response (α). Considering the width of the 95% CI in Figure 6.3b, we see a

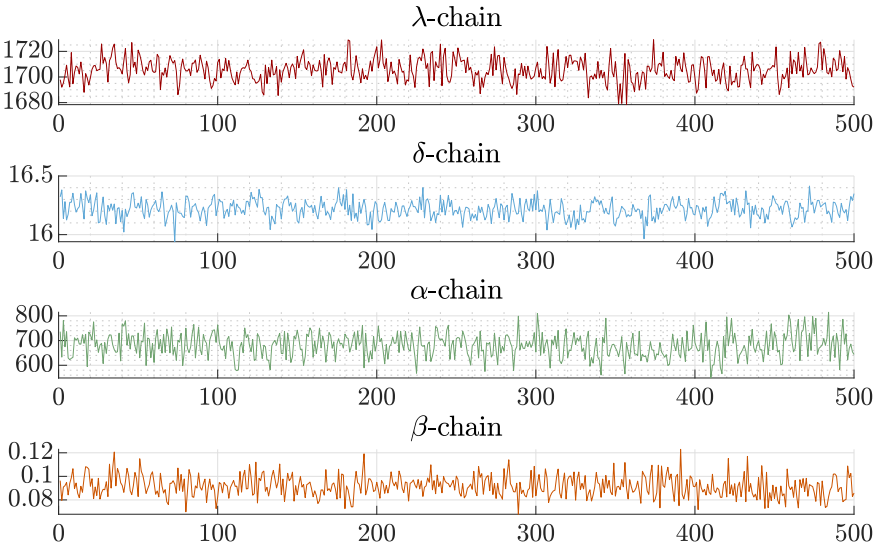

(a) Chains of the sampled hyperparameters.

Figure 6.4: Hyperparameter chains for the hierarchical Gibbs sampler.

high uncertainty in the pixels where the remaining ring artifacts are present and also where the intensity is high.

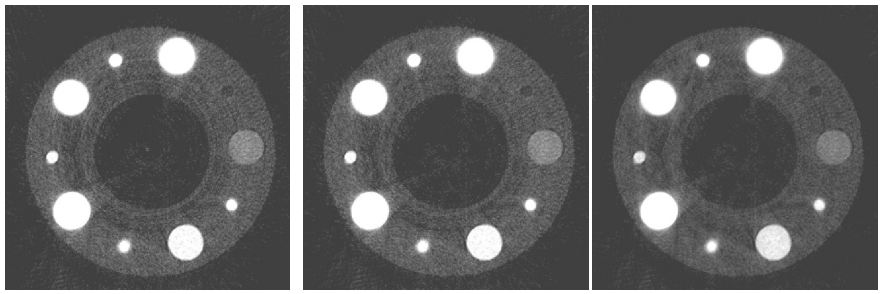

(a) WLS. **(b)** P-WLS. **(c)** AMAP.

Figure 6.5: WLS (a), P-WLS (b) and AMAP (c) reconstructions for the neutron data with air-pixel correction and different ring reduction techniques. All images are displayed in the intensity range $[-0.2, 0.6]$.

The reconstructions for WLS, P-WLS and AMAP are shown in Figure 6.5. We choose the hyperparameters for the AMAP model based on the hyperparameter estimates obtained from the APM model shown in Table 6.1.

The figure shows that the quality of the reconstruction is greatly in-

creased by using WLS with non-negativity compared to the FBP reconstruction in Figure 6.2b, but that there are still source model error present for the WLS reconstruction. Considering the ring reduction methods in Figures 6.5b and 6.5c, it seems that both methods (P-WLS and AMAP) succeed in reducing the ring artifacts and produce reconstructions that are similar in quality. However, the AMAP model seems to produce a slightly more homogeneous reconstruction compared to the P-WLS reconstruction where low-frequency ring artifacts are observed similar to the simulation experiments in paper A. Thus, the new results on real data support the findings of paper A, namely that the AMAP model is applicable for CT reconstruction with intensity drift, and moreover that the method is competitive with existing preprocessing methods.

6.5 Summary

In this chapter, we summarized the findings of paper A. We proposed a new reconstruction model using a Bayesian framework that jointly estimate both reconstruction image, intensity drift, and the detector response for CT. In addition, we presented previously unpublished results on real data that support the findings of paper A.

CHAPTER 7

Ring artifacts in spectral computed tomography

Relevant paper

B) **Katrine O. Bangsgaard**, Genoveva Burca, Evelina Ametova, Martin S. Andersen and Jakob S. Jørgensen, "Low-rank flat-field correction for artifact reduction in spectral computed tomography", *Applied Mathematics in Science and Engineering*, 2023, [6].

7.1 Motivation and goals

In Chapter 6 we proposed a new reconstruction model tailored to mitigate ring artifacts for monochromatic CT. Spectral CT is also challenged by ring artifacts since we compute a spectral reconstruction, i.e., a reconstruction for each energy, and the intensity for a single energy is low. Consequently, we have a low signal-to-noise ratio which may result in severe ring artifacts in the reconstructions. Thus, we need to employ ring reduction techniques to compute reliable reconstructions.

Extended models for mitigating ring artifacts in spectral CT have been proposed but they all rely on computational expensive algorithms. Conventional preprocessing methods for ring reduction in monochromatic CT, as described in Chapter 6, can also be applied to spectral CT. However, these methods struggle when the ring artifacts are severe and may unintentionally introduce new artifacts in the reconstructions [3].

To compute the spectral reconstructions, we need spectral flat-fields. A spectral flat-field is a collection of flat-fields measured for each energy. The ring artifacts arise because the spectral flat-fields are very noisy. Thus, we need several spectral flat-fields to avoid introducing systematic errors if we simply compute the spectral detector response as the mean of the spectral flat-fields.

However, for spectral CT, there is a correlation of the energy-wise measurements which also appears in the spectral flat-fields. Thus, a single spectral flat-field carries significant redundant information which can be exploited. Thus, the goal is to formulate a method for ring artifact reduction in spectral CT by utilizing the redundant information in the spectral flat-fields.

7.2 Contributions

The main contribution of paper B can be summarized as follows. We propose an extended flat-field model that exploits high correlation across energy-channels in the spectral flat-fields to mitigate ring artifacts in the spectral reconstruction.

7.3 Methods

We assume that we have m energies, r detector elements, p projection images, s flat-field images and n pixels. The spectral CT measurement model can be written,

$$y_k = \operatorname{diag}(1_p \otimes Z_k) \exp(-Au_k)\,, \tag{7.1}$$

for $k = 1, 2, \ldots, m$, where $u_k \in \mathbb{R}^n$ and $Z_k \in \mathbb{R}^r$ are the attenuation coefficients and detector response for energy k, respectively.

Conventionally, an estimate of the spectral detector response, $\hat{Z} \in \mathbb{R}^{r\times m}$ is computed by,

$$\hat{Z} = \frac{1}{s}\sum_{j=1}^{s} F_j = \frac{1}{s}\big(1_s^T \otimes I_r\big)\, F, \tag{7.2}$$

where

$$F = \begin{bmatrix} F_1 \\ \vdots \\ F_s \end{bmatrix} \in \mathbb{R}^{rs\times m},$$

is the matrix with the spectral flat-fields stacked vertically. The matrix F is contaminated by noise and carries redundant information. Thus, we aim to substitute F in (7.2) by a low-rank approximation which only carries the significant information about the spectral flat-fields.

We can formulate a low-rank matrix using a singular value decomposition (SVD). SVD is a technique that allows us to express F, and in fact any matrix of rank q, in terms of a sum of rank-one matrices [48], i.e.,

$$F = \sum_{i=1}^{q} \sigma_i U_i V_i^T, \tag{7.3}$$

where σ_i is the ith singular value that satisfies $\sigma_1 \geq \sigma_2 \geq \ldots \geq \sigma_q > 0$, and U_i and V_i are the ith left and right singular vectors, respectively. The Eckart–Young–Mirsky theorem [18] states that the best rank-l approximation of F that solves

$$F^l \in \underset{z}{\operatorname{argmin}} \left\{ \|F - Z\|_F \,|\, \operatorname{Rank}(Z) \leq l \right\},$$

can be computed by truncating (7.3) after the first l terms, i.e.,

$$F^l = \sum_{i=1}^{l} \sigma_i U_i V_i^T. \tag{7.4}$$

Note that l is the only parameter of the method, and that this parameter can be easily chosen by visual inspection of the singular values of the spectral flat-fields.

7.4 Results

In this section we briefly summarize the findings of paper B. The extended spectral flat-field model which we denote LR (Low-Rank) is tested on a real neutron dataset. The neutron data was acquired at the Imaging and Materials Science and Engineering (IMAT) beamline operating at the ISIS neutron spallation source (Rutherford Appleton Laboratory, UK). [33]

We compare the LR method to two existing ring reduction techniques: a preprocessing method proposed by Münch et al. [37] as described in Chapter 6 and a preprocessing method proposed by Vo et al. [49] that uses a combination of sorting and smoothing (non-local means) to reduce ring artifacts by mitigating stripes in the sinogram. We denote the methods of Münch et al. [37] and Vo et al. [49] by WF (Wavelet Fourier) and NLM (non-local means), respectively.

All three methods are preprocessing methods and need to be combined with a reconstruction model. We have chosen the conventional FBP and WLS combined with TV regularization.

The results demonstrate that our proposed LR method can successfully mitigate ring artifacts in spectral CT reconstruction. Moreover, our method is shown to be robust, i.e., even in case of severe ring artifacts where the conventional ring reduction methods NLM and WF struggle, our method prevails in mitigating the ring artifacts, and in particular, our method needs only a single spectral flat-field for ring artifact reduction whereas existing methods need multiple spectral flat-field images to reach a similar level of ring reduction as illustrated in Figure 7.1.

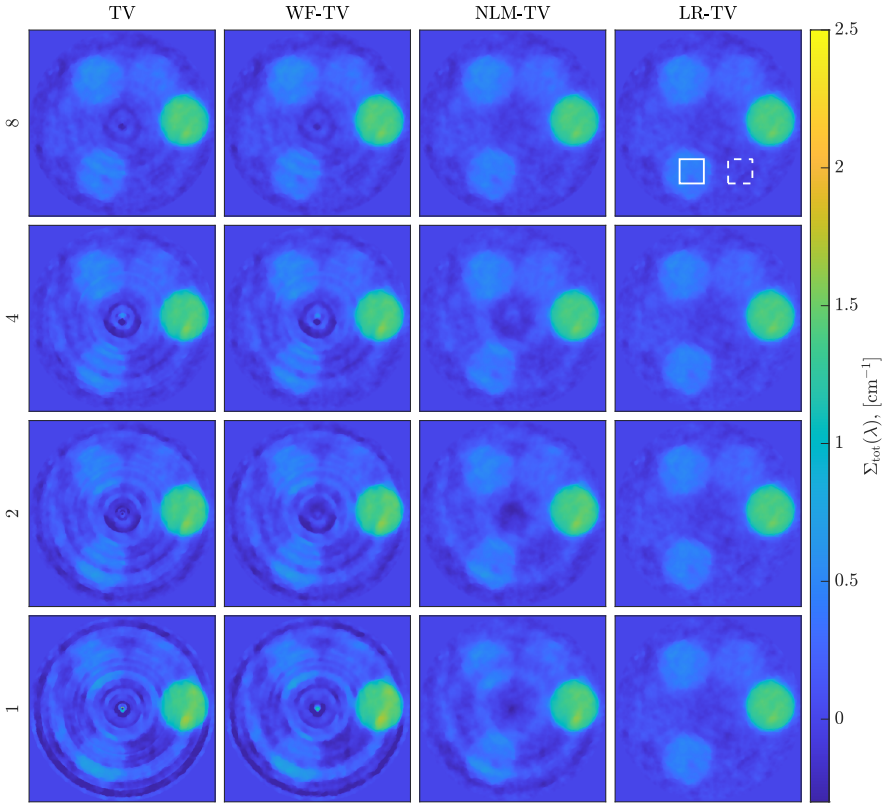

Figure 7.1: Reprint of main result from paper B. The numbers in each row indicate the number of flat-fields used for flat-field correction.

7.5 Summary

In this chapter, we summarized the findings of paper B. We proposed a new extended flat-field model that utilizes redundancy in the spectral flat-fields to mitigate ring artifacts for spectral CT, and we briefly summarized the results of paper B.

CHAPTER 8

Hemodialysis modeling and estimation

Relevant paper

C) **Katrine O. Bangsgaard**, Morten Andersen, James G. Heaf and Johnny T. Ottesen, "Bayesian parameter estimation for phosphate dynamics during hemodialysis", *Mathematical Biosciences and Engineering*, 2022, [4].

8.1 Motivation and goals

As described in Chapter 3, phosphate kinetics during hemodialysis may be modeled by a diffusion process and ODEs. The parameters are patient-specific and can potentially help clinicians tailor individual hemodialysis treatment in the future by providing improved insight into the physiological mechanisms and individual responses.

However, especially for single-pass, we have very few data points from which we can estimate the physiological parameters and to our knowledge, there has not been investigations of the uncertainty associated with the computed estimates. Moreover, since the parameters can be interpreted as physiological quantities, we have qualified prior knowledge such as a physiological meaningful parameter range and mean value from previous clinical trials. However, current models do not utilize this additional information to inform the model.

The goal of paper C is to investigate the single- and multiple-pass models using sampling techniques and extend them into a Bayesian framework such that we can incorporate prior clinical knowledge into the model.

8.2 Contributions

The main contributions of paper C can be summarized as follows. We propose a Bayesian approach for estimating patient-specific parameters for phosphate dynamics during hemodialysis and use UQ to assess the reliability of our parameter estimates. We address the identifiability of the parameters for the single- and multiple-pass and the combination of the two, denoted combined-pass. In addition, we conduct experiments on real and synthetic data to investigate how the parameter estimation can be improved by including relapse measurements and / or measure consecutive sessions using UQ.

8.3 Methods

We solve the parameter estimation problem using a Bayesian approach such that we can incorporate prior clinical knowledge into the model and regularize the parameter space. We assume that we have measurements of the blood and dialysate volumes, i.e., V_b and V_d, such that the unknown parameters are the phosphate concentration in the bones and the diffusion coefficients between bone and blood, and blood and dialysate, i.e., C_s, K_s and K_b. In addition, we also include the initial measurement of phosphate concentration in the blood, $C_b(0)$ as a parameter since $C_b(0)$ is subject to measurement noise.

We assume that the measurement noise of the measured phosphate concentrations b is Gaussian, i.e.,

$$b|u, \sigma_{\mathrm{d}}^2 \sim \mathcal{N}\big(F(u), \sigma_{\mathrm{d}}^2 I_{dm}\big)$$

where $u \in \mathbb{R}^4$ is the unknown parameters and initial condition for the phosphate concentration in the blood, $b \in \mathbb{R}^{dm}$ is the measured phosphate concentrations and $F(u) \in \mathbb{R}^{dm}$ is the solution to the system of ODEs in (3.1) with $d = 1$ for single-pass or (3.4) with $d = 2$ for multiple-pass, respectively.

We equip the parameters and initial condition with Gaussian priors. For the initial condition, we choose the Gaussian prior with mean equal to the measurement at time zero with variance modeled by σ_{d}, i.e., we assume that the measurement noise of the first data point is equal to the later measured concentrations. We choose the Gaussian priors for the parameters to incorporate clinical knowledge about the range and mean but choose to fix the variance to ensure stability of the sampling process. We use the NUTS sampler to sample the posterior since preliminary results showed that the parameters had high correlation.

8.4 Results

We have tested the Bayesian model on both real and synthetic data. The results for the synthetic data is found in the supplementary of paper C and show that if we use the single-pass model with fixed C_s and without incorporating the prior clinical knowledge as commonly done, then we get highly correlated parameter estimates for K_s and K_b with very large CIs, see Figure 8.1. Thus, the model used for conventional parameter estimation in hemodialysis gives estimates with high uncertainty. The corresponding estimates for multiple-pass without clinical knowledge produced estimates with much lower uncertainties which indicates that the multiple-pass model is more informative about the parameters than the single-pass model.

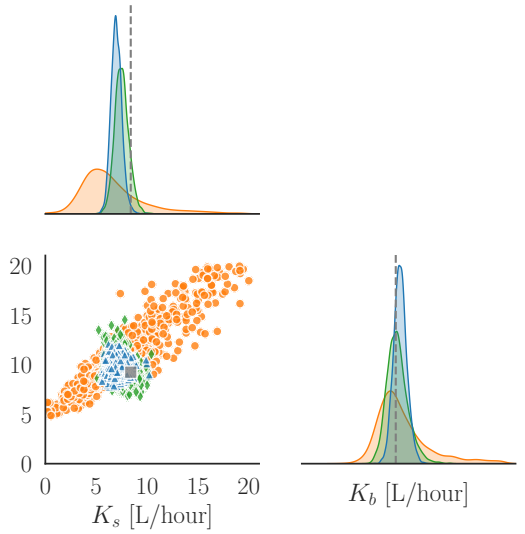

Figure 8.1: Reprint from the supplementary of paper C. Estimated posterior for the parameters using synthetic data with fixed C_s for both single- and multiple-pass. The orange, green and blue dots represent the samples for single-pass, multiple-pass and combined-pass, respectively. The gray lines and square represent the true parameter values.

Including the prior clinical knowledge through a Bayesian framework showed to greatly decrease the uncertainty associated with the parameter estimates and the correlation between parameters as seen in Figure 8.2. The results show that considering multiple-pass data for estimation greatly reduces the uncertainty of the parameters C_s and K_b compared to single-pass data whereas the uncertainty associated with the parameter K_s remains unaffected by the additional knowledge provided by multiple-pass.

Lastly, investigations of synthetic and real data revealed that we can reduce the uncertainty of the parameter estimates greatly by measuring the

same patient for consecutive sessions whereas measurement in the relapse phase (after ended treatment) had little impact on the parameter estimates as illustrated in Figure 8.3.

8.5 Summary

In this chapter, we summarized the findings of paper C. We proposed a Bayesian framework for estimating patient-specific parameters for phosphate kinetics during hemodialysis treatment. We used UQ to compare two existing hemodialysis treatments, the conventional single-pass and the novel multiple-pass, and we briefly summarized the results of paper C.

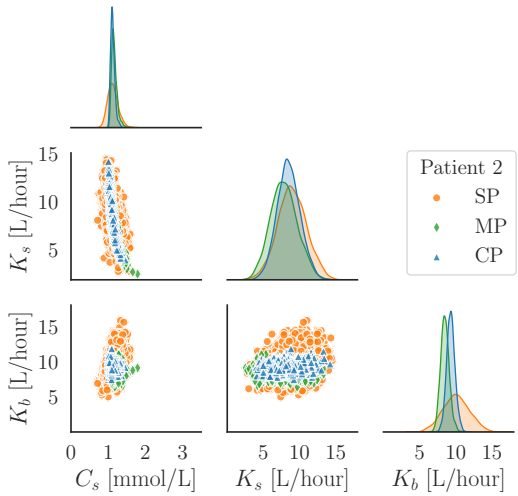

Figure 8.2: Reprint from paper C. Estimated posterior for the parameters using real data for both single- and multiple-pass. The orange, green and blue dots represent the samples for single-pass (SP), multiple-pass (MP) and combined-pass (CP), respectively.

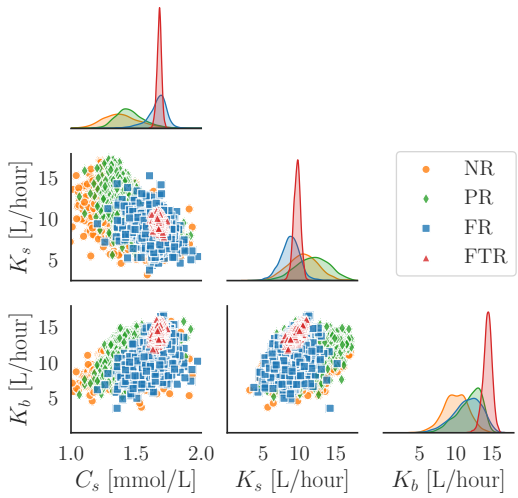

Figure 8.3: Reprint from paper C. Estimated posterior for the parameters using real data for a patient with consecutive single-pass measurements. The orange, green, blue and ret dots represent the samples for single-pass data with no relapse (NR), partial relapse (PR), full relapse (FR) and lastly full three relapse (FTR) corresponding to three consecutive sessions.

CHAPTER 9

Discussion and Conclusion

In this Chapter, we give a brief discussion bridging the contributions and potential future work. Lastly, we end the thesis by providing a summary and concluding remarks.

9.1 Discussion

The general aim of the PhD project was to study Bayesian modeling for inverse problems and apply the Bayesian framework to real life applications, i.e., CT and hemodialysis. The two CT related papers (paper A and paper B) were motivated by spectral neutron CT data.

In paper A, we extended the CT reconstruction model to include additional parameters, i.e., source intensity and detector response. However, introducing these additional parameters decreased the regularity of the problem and made the reconstruction problem significantly harder. Thus, prior modeling was of crucial importance to obtain a high-quality solution to the extended reconstruction problem. In particular, the non-negativity assumption was needed for both the AMAP and APM models to converge to a physical meaningful reconstruction.

In paper A, we used very simplistic priors but incorporating edge-preserving priors for the reconstruction might have a regularizing effect of the reconstruction problem since it would enhance the piece-wise-constant behavior which might stabilize the remaining parameters as well. We emphasize at this point that our models are computationally expensive, thus they are only justified for data with a low signal-to-noise ratio where the conventional preprocessing methods struggle and start to introduce new artifacts.

The reconstruction model proposed in paper A could readily be applied to the spectral CT case. The spectral neutron data was acquired at large synchroton scanners where both intensity drift and ring artifacts are sources to artifacts in the spectral reconstruction. The intensity drift and detector response correlates across the energy-channels, thus a possible improvement

of the extended reconstruction model for spectral CT could incorporate this correlation through the model and/or the priors.

The extended flat-field model proposed in paper B could be viewed as an improved prior mean and be used in combination with the extended reconstruction model in paper A (or potentially the modified version for spectral CT). Our Bayesian model heavily depends on the quality of the prior model, especially when the data quality is low. Thus, we would most likely ease the reconstruction problem significantly by improving the prior mean.

The general challenge for Bayesian CT is the curse of dimensionality. We have many parameters and the sampling algorithms are computational expensive in practice. However, for biological applications such as hemodialysis and blood cancer, we have fewer parameters such that the Bayesian approach does not impose an unreasonable computational cost compared to the standard parameter estimation methods.

Including the prior clinical knowledge in the hemodialysis model in paper C, showed to improve the stability of the estimation process, especially for the conventional treatment single-pass, compared to the parameter estimates found by the conventional optimization approach in paper E. The sampling process revealed and confirmed that the single-pass treatment without prior knowledge is not sufficient to estimate the patient-specific parameters reliable. Moreover, the Bayesian approach provided valuable insight into how the data experiments could be optimally designed.

We chose very simple Gaussian priors for the hemodialysis model making it readily applicable to other biological applications. Potential future work could be to apply the Bayesian approach to the blood cancer model in paper D for parameter estimation. Moreover, it would be interesting to use Bayesian inference to investigate if the experimental design can be optimized by measuring the patients at different time intervals, etc.

9.2 Concluding remarks

In this thesis, we have studied a Bayesian approach to solving inverse problems. The motivation has been driven by the two real-life applications CT and hemodialysis where conventional methods struggled to compute satisfying solutions possibly due to the lack of stability and prior modeling. Chapter 2 and 3 introduced the physics and biological mechanisms behind CT and hemodialysis and showed, that despite their very different applications, both applications can be narrowed down to the same problem structure described as a discrete inverse problem. A general introduction to discrete inverse problems and to uncertainty quantification for inverse problems were presented in Chapter 5 and 6 to provide the necessary intuition for the methods developed in papers A, B and C.

We have developed an extended reconstruction model in paper A for monochromatic CT that jointly estimate reconstructions and source to mitigate artifacts in the reconstruction. We employed a Bayesian approach and included prior modeling to the source uncertainty. The method showed that, in case of low signal-to-noise data, our model succeeded in mitigating ring artifacts even when other existing methods struggled.

In paper B, we analyzed the spectral flat-fields using singular value decomposition and the findings suggested that a low-rank approximation of the spectral flat-fields could well-approximate the measured flat-fields. Thus, inspired by these findings, we formulated an extended flat-field model to mitigate noise in the spectral flat-fields and consequently mitigate ring artifacts in the spectral reconstructions.

Lastly, we considered the parameter estimation problem for hemodialysis in paper C. Here we found that the prior modeling was crucial to stabilize the parameter estimates and that the Bayesian approach elegantly incorporated the clinical knowledge.

To summarize, we have applied the Bayesian approach for inverse problems to two very different applications and shown that the Bayesian approach is a powerful modeling tool for inverse problems.

APPENDIX A

A statistical reconstruction model for absorption CT with source uncertainty

Published in Inverse Problems, Volume 37, Number 8, 27 July 2021, DOI: 10.1088/1361-6420/ac11c7.

Katrine O. Bangsgaard and Martin S. Andersen.

IOP Publishing Inverse Problems
Inverse Problems **37** (2021) 085009 (23pp) https://doi.org/10.1088/1361-6420/ac11c7

A statistical reconstruction model for absorption CT with source uncertainty*

Katrine O Bangsgaard** and **Martin S Andersen**

Department of Applied Mathematics and Computer Science, Technical University of Denmark, Denmark

E-mail: kaott@dtu.dk and mskan@dtu.dk

Received 7 April 2021, revised 21 June 2021
Accepted for publication 6 July 2021
Published 27 July 2021

Abstract

Reconstruction methods for computed tomography are often based on the assumption that the source intensity and the detector response are known and static. In practice, however, both are unknown and must be estimated. An estimate of the combined source intensity and detector response is typically obtained by acquiring a number of so-called flat-field measurements, but this approach is oblivious to intensity drift, e.g. due to source instabilities, vibrating beamline components, etc. Discrepancies between the estimated response and true response can lead to severe artifacts in the reconstruction, especially for dose- and/or time-limited experiments. We propose a new, extended reconstruction model that jointly estimates the reconstruction, the detector response, and the (possibly time-varying) source intensity. We demonstrate through simulated experiments that the proposed reconstruction model leads to reconstructions with significantly reduced artifacts.

Keywords: computed tomography, model errors, reconstruction methods, intensity drift, ring artifacts

(Some figures may appear in colour only in the online journal)

1. Introduction

Absorption computed tomography (CT) is a non-invasive imaging technique that makes it possible to obtain cross-sectional images of the interior of an object from a set of projection images. The projection images are formed by illuminating an object in different orientations by radiation from a source (e.g. x-ray or neutron radiation) and recording the attenuated radiation using

*This work was supported by The Villum Foundation (Grant No. 25893) and The Novo Nordisk Foundation (Grant No. NNF20OC0061894).
**Author to whom any correspondence should be addressed.

a detector. The attenuation is material-specific and obeys Lambert–Beer's law which characterizes the attenuation of a beam of radiation as it propagates through a material or medium [8]. The problem of computing a spatial attenuation image from the projection images is an inverse problem that is commonly referred to as a reconstruction problem. Numerous reconstruction algorithms exist and, roughly speaking, each algorithm corresponds to a reconstruction model with different underlying assumptions.

Most reconstruction models rely on the assumption that the source intensity and the detector response are known. In practice, however, these are typically estimated from a set of flat-field measurements (also known as white-fields) which are simply projection images acquired without an object in the scanner [23]. The flat-field images contain noise due to the statistical nature of the radiation and the absorption process, and this noise can give rise to a type of reconstruction artifacts known as ring artifacts [16]. Such artifacts can mask important features in the reconstruction [21], and they may arise, e.g. if only a few flat-field images are available, if the exposure time is short, and/or if the source intensity is low [1].

The assumption that the source intensity is known is typically violated if the source intensity varies during the data acquisition process. For example, intensity drift may occur because of source instabilities, vibrating beamline components, time-varying detector properties, or other factors disturbing the stationarity of the experiment [25]. The resulting mismatch between the reconstruction model and the underlying physics introduces systematic model errors which, in turn, can lead to reconstruction artifacts that compromise diagnostic or quantitative measures of image quality.

Existing methods for ring artifact reduction can be divided into four main categories: (i) acquisition methods, (ii) preprocessing methods, (iii) postprocessing methods, and (iv) reconstruction methods based on extended models. Roughly speaking, acquisition methods mitigate the systematic errors that lead to ring artifact by perturbing the projection geometry between consecutive projections [10]. This can lead to notable improvements, but it is not always feasible nor practical. Preprocessing methods work by manipulating or filtering the projection images prior to the reconstruction [17, 26], and postprocessing methods are designed to remove or reduce rings in the image domain after the reconstruction step [20, 23, 27, 28]. The extended reconstruction models either implicitly or explicitly include the detector response as an unknown quantity to be estimated along with the attenuation function [1, 19, 22, 24]. A common trait of all the aforementioned approaches is that they are oblivious to intensity drift.

Existing methods for combating artifacts that arise because of intensity drift may be viewed as preprocessing methods that extract information about intensity drift from flat-field images [18, 25] and use this information to manipulate the measurement data prior to the reconstruction. Moreover, the parts of the projection images that are not affected by the object of interest are essentially flat-field measurements that can be used to estimate the time-varying intensity. We will refer to such measurements as 'air pixels' since they correspond to rays that only travel through air.

To address the limitations of existing methods, we propose a new reconstruction model that treats both the (possibly time-varying) source intensity and the detector response as unknown model parameters along with the unknown attenuation of the object of interest. The model generalizes the reconstruction methods of Aggrawal *et al* [1] and Salehjahromi *et al* [22], and it includes an array of other models as special cases.

Outline. Section 2 introduces our model assumptions, and we discuss how the implicit assumptions inherent in many existing models can lead to reconstruction artifacts. In section 3, we propose a new, extended reconstruction model, and we outline a Bayesian reconstruction framework that allows the inclusion of suitable priors. We also discuss different strategies for

selecting a set of hyperparameters. Section 4 contains some numerical experiments based on simulated data, and section 5 includes a discussion and concludes the paper.

Notation. The set $\mathbb{R}^n$ is the n-dimensional Euclidean space, $\mathbb{R}^n_+$ is the non-negative orthant of $\mathbb{R}^n$, and $\mathbb{R}^{m\times n}$ denotes the set of real-valued $m \times n$ matrices. Lower case letters denote vectors, and $x = vec(X)$ is the vector obtained by stacking the columns of the matrix X. If x is a vector in $\mathbb{R}^n$, then $\text{diag}(x) \in \mathbb{R}^{n\times n}$ denotes the diagonal matrix with the elements of x on its diagonal, and the functions $\log(x)$ and $\exp(x)$ are defined as an elementwise operation. The identity matrix is denoted by I, the vector $\mathbf{1}$ is a vector of ones, and e_i denotes the ith canonical basis vector; the dimension can be inferred from the context. Given two matrices A and B, $A \otimes B$ denotes the Kronecker product of A and B. Finally, $\|\cdot\|_2$ denotes the Euclidean norm and $\|\cdot\|_\text{F}$ denotes the Frobenius norm.

2. Extended reconstruction model

2.1. Measurement model

We start by introducing an extended measurement model that may be derived from the Lambert–Beer law by including the detector response and a time-varying source intensity. We will assume that the set of measurements consists of p images $y_1, \ldots, y_p \in \mathbb{R}^r$, each of which has been acquired using a detector with r detector elements. We associate with the jth image a tuple $(t_j, \theta_j, \delta_j)$ where t_j is the acquisition time, θ_j is the projection angle, and δ_j is a binary variable that indicates whether the jth image is a projection image ($\delta_j = 1$) or a flat-field image ($\delta_j = 0$). We will denote by s the total number of flat-field images, i.e.

$$s = \sum_{j=1}^{p} (1 - \delta_j).$$

Now, using the Lambert–Beer law, we express the intensity of the incident radiation on detector element i at time t_j as

$$I_i(t_j, \theta_j) = I_0 \nu_i \psi(t_j) \exp\left(-\delta_j \int_{\ell_i(\theta_j)} \mu(x) \mathrm{d}x \right), \quad i = 1, \ldots, r, \tag{1}$$

where $\mu : \mathbb{R}^d \to \mathbb{R}_+$ is a spatial attenuation function, I_0 denotes the nominal source intensity, the vector $\nu = (\nu_1, \ldots, \nu_r)$ models the detector response (including any effects due to the direction of incidence), the function $\psi : \mathbb{R} \to \mathbb{R}_+$ models a time-varying intensity, and $l_i(\theta_j)$ is the line segment between the source and the ith detector element for the jth projection angle.

The integral in (1) is discretized by introducing a parameterization of the attenuation function of the form

$$\mu(x) = \sum_{k=1}^{n} u_k \chi_k(x), \tag{2}$$

where $u \in \mathbb{R}^n$ is a vector of unknown parameters (e.g. pixel or voxel values) and $\chi_1(x), \ldots, \chi_n(x)$ are basis functions. For example, a parameterization based on a rectilinear grid with nearest-neighbor interpolation corresponds to $\chi_k(x)$ being an indicator function that takes the value 1 if x is inside the kth pixel or voxel and 0 otherwise. The parameterization (2) allows us to express the line integral in (1) as

$$\int_{\ell_i(\theta_j)} \mu(x) \mathrm{d}x = e_i^T A_j u,$$

where $A_j \in \mathbb{R}^{r\times n}$ is the matrix with entries

$$(A_j)_{ik} = \delta_j \int_{\ell_i(\theta_j)} \chi_k(x)\,\mathrm{d}x, \quad i = 1,\dots,r, \quad k = 1,\dots,n.$$

Note that $A_j = 0$ if the jth image is a flat-field measurement (i.e. $\delta_j = 0$). Our discretized forward model may now be expressed as

$$I_i(\theta_j) = I_0\nu_i\omega_j \exp(-e_i^T A_j u), \quad i = 1,\dots,r, \quad j = 1,\dots,p,$$

where $\omega_j = \psi(t_j)$ and $\omega = (\omega_1,\dots,\omega_p)$. Equivalently, if we define $\bar{y}_j = (I_1(\theta_j),\dots,I_r(\theta_j))$, then

$$\bar{y}_j = I_0\omega_j \operatorname{diag}(\nu)\exp(-A_j u), \quad j = 1,\dots,p.$$

To simplify our notation in subsequent derivations, it is convenient to make a change of variables. Specifically, we define

$$I_0\nu = \operatorname{diag}(\hat{\nu})\exp(-v) \tag{3}$$

$$\omega = \operatorname{diag}(\hat{\omega})\exp(-w), \tag{4}$$

where $v \in \mathbb{R}^r$ and $w \in \mathbb{R}^p$ are the new variables, and the vectors $\hat{\nu} \in \mathbb{R}^r_{++}$ and $\hat{\omega} \in \mathbb{R}^p_{++}$ are given. Our forward model may then be expressed as

$$\bar{y}_j(u,v,w;\hat{\nu},\hat{\omega}) = \hat{\omega}_j \operatorname{diag}(\hat{\nu})\exp(-A_j u - v - \mathbf{1}w_j), \quad j = 1,\dots,p,$$

or equivalently, in vectorized form,

$$\bar{y}(u,v,w;\hat{\nu},\hat{\omega}) = \operatorname{diag}(\hat{\omega}\otimes\hat{\nu})\exp(-Au - Gv - Hw), \tag{5}$$

where $G = \mathbf{1}\otimes I$ and $H = I\otimes\mathbf{1}$, and A is the vertical concatenation of $A_1,\dots,A_p$ such that $\bar{y} \in \mathbb{R}^{rp}$. We will simply write $\bar{y}(u,v,w)$ instead of $\bar{y}(u,v,w;\hat{\nu},\hat{\omega})$ whenever the vectors $\hat{\nu}$ and $\hat{\omega}$ are assumed to be given. We will refer to $\bar{y}$ as an *augmented*, *raw sinogram*, since the sinogram in conventional CT does not include the flat-field images. Figure 1(a) shows an example of an augmented, raw sinogram for a simulated experiment with a time-varying intensity and with projections acquired using the golden angle radial sampling [15] and periodic flat-field samples.

The forward model (5) does not take the statistical nature of the photon arrival process into account [8]. A common assumption in tomographic imaging is that the measurements obey Poisson distributions whose means are prescribed by the Lambert–Beer law. Adopting this assumption, we will model the jth image, y_j, as a sample of a random variable $\mathbf{y}_j$ with the conditional distribution

$$\mathbf{y}_j|u,v,w \sim \text{Poisson}(\bar{y}_j(u,v,w)). \tag{6}$$

Denoting the conditional probability of observing y_j by $\pi(y_j|u,v,w)$, the log-likelihood may be expressed as

$$\log(\pi(y_j|u,v,w)) = -\mathbf{1}^T\bar{y}_j(u,v,w) + y_j^T\,\log(\bar{y}_j(u,v,w)) - \mathbf{1}^T\,\log(y_j!), \tag{7}$$

where the logarithm and factorial functions are applied elementwise.

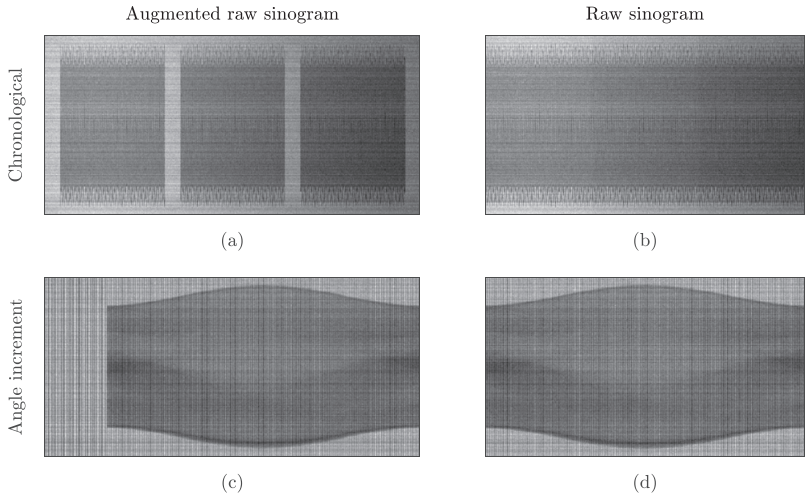

Figure 1. Illustration of measurements. (a) Shows an augmented raw sinogram for a simulated experiment with a time-varying intensity and golden angle radial sampling and periodic flat-field measurements (flat-fields are measured for every 200 projections). (b) Shows the raw sinogram without the flat-field projections, and (c) is the augmented raw sinogram with the flat-field projections in the beginning and the remaining projections ordered by angle. Finally, (d) is the raw sinogram ordered by angle but without flat-field projections.

2.2. Conventional reconstruction approach

Before we turn our attention to a Bayesian extension of the measurement model (6), it is instructive to see how a number of existing reconstruction models may be derived as special cases. This also allows us to illustrate when and how existing methods may fall short. The conventional approach is to estimate the flat-field, intensity drift, and the attenuation image separately. We now outline how each of these estimation problems is related to the model assumption (6).

Flat-field estimation. Perhaps the most common approach to flat-field estimation is to compute the empirical mean of the flat-field images. This estimate may be viewed as a maximum likelihood (ML) estimate based on the flat-field images. Specifically, in terms of our extended model (6), it corresponds to the assumptions that $v = 0$, $w = 0$, and $\hat{\omega} = \mathbf{1}$ such that

$$\begin{aligned}\hat{\nu} &= \operatorname*{argmin}_{\nu} \left\{ -\sum_{j=1}^{p} (1-\delta_j) \log \pi(y_j | u, v=0, w=0; \nu, \hat{\omega}=\mathbf{1}) \right\} \\ &= \frac{1}{s} \sum_{j=1}^{p} (1-\delta_j) y_j. \end{aligned} \tag{8}$$

Note that the vector u does not affect this estimate since only flat-field images are included in the estimation problem.

Intensity estimation based on flat-field images. Sometimes the flat-field images may also be used to estimate variations in the intensity. For example, a common heuristic is to use linear interpolation to correct for intensity drift, provided that flat-field images were acquired both before and after the projection images. Here we outline a somewhat more formal and flexible approach based on ML estimation and our statistical model (6). We will assume that a flat-field estimate $\hat{\nu}$ is available, and we will assume a parametric representation $\hat{\omega} = \omega(\alpha)$ where $\alpha \in \mathbb{R}^l$ is a vector of parameters. An ML estimate of α, based on only the flat-field images, may then be expressed as

$$\hat{\alpha} = \underset{\alpha \in \mathbb{R}^l}{\operatorname{argmin}} \left\{ -\sum_{j=1}^{p} (1 - \delta_j) \log \pi(y_j | u, v = 0, w = 0; \hat{\nu}, \omega(\alpha)) \right\}. \tag{9}$$

This problem is convex if $\omega(\alpha)$ is an affine function of α. For example, linear regression (LR) corresponds to the special case where $\omega(\alpha) = \begin{bmatrix} \mathbf{1} & t \end{bmatrix} \alpha$ with $\alpha \in \mathbb{R}^2$. In this case, the estimate (9) may be expressed as

$$\hat{\alpha} = \underset{\alpha \in \mathbb{R}^2}{\operatorname{argmin}} \left\{ \sum_{j=1}^{p} (1 - \delta_j) \mathbf{1}^T \left(\left(\alpha_1 + t_j \alpha_2 \right) \hat{\nu} - y_j \log \left(\alpha_1 + t_j \alpha_2 \right) \right) \right\}. \tag{10}$$

Only the flat-field samples contribute to the sum in both (9) and (10), so these estimates are independent of u.

Intensity estimation based on air pixels. Another heuristic technique for estimating intensity variations is applicable when the object of interest does not fill the sinogram, i.e. there are 'air rows (AR)' on either or both sides of the object in some or all projections. In this case, we can use the flat-field projections and additional measurements that are essentially flat-field measurements to estimate ω. Elements of the sinogram that correspond to flat-field measurements may be identified manually or by means of segmentation. To identify the flat-field measurements, we define a binary mask $C \in \mathbb{R}^{r \times p}$ such that $C_{ij} = 1$ if $\delta_j = 0$ (i.e. the jth projection is a flat-field projection) or if the ith pixel of the jth projection is essentially a flat-field measurement, and otherwise $C_{ij} = 0$. We then arrive at an ML estimate of the form $\omega(\hat{\alpha})$ where

$$\hat{\alpha} = \underset{\alpha \in \mathbb{R}^l}{\operatorname{argmin}} \left\{ \sum_{i=j}^{p} c_j^T \left(\hat{\nu} \omega_j(\alpha) - y_j \log(\omega_j(\alpha)) \right) \right\}, \tag{11}$$

and where c_j denotes the jth column of C. Note that (9) is obtained as a special case of (11) by letting $c_j = (1 - \delta_j)\mathbf{1}$.

Attenuation image estimation. Given some estimates $\hat{\nu}$ and $\hat{\omega}$, the distribution of the measurements y, conditioned on u, may be expressed as

$$\pi(y|u) = \prod_{j=1}^{p} \pi(y_j | u, v = 0, w = 0; \hat{\nu}, \hat{\omega}).$$

This allows us to estimate the attenuation image from the projection images, e.g. as an ML estimate

$$\hat{u} = \underset{u}{\operatorname{argmin}} \{ -\log \pi(y|u) \},$$

which is a convex problem. A common variation on this approach is to use a Gaussian approximation of the likelihood function which may also be viewed as constructing a quadratic approximation of the log-likelihood function. Specifically, if we define $b = \log(\hat{\omega} \otimes \hat{\nu}) - \log(y)$, the resulting approximation may be expressed as

$$\hat{u} \in \underset{u}{\operatorname{argmin}} \left\{ \frac{1}{2} \|Au - b\|^2_{\widehat{\Sigma}^{-1}} \right\}, \qquad \widehat{\Sigma}^{-1} = \operatorname{diag}(y), \tag{12}$$

which is a weighted least-squares (WLS) problem.

Effect of estimation errors. The conventional reconstruction approach that we have just outlined can easily be extended to include suitable priors on the attenuation function u, the flat-field ν, and/or the time-varying intensity ω. The inclusion of priors allows us to make use of maximum *a posteriori* (MAP) estimation or sampling methods such as Markov chain Monte Carlo (MCMC) instead of ML estimation. However, the outlined approach, with or without priors, is essentially a three-stage procedure: the estimation of $\hat{\nu}$, $\hat{\omega}$, and $\hat{u}$ is not performed jointly but separately. Roughly speaking, this approach can be expected to work well when the flat-field images are many and/or have a high signal-to-noise ratio, allowing us to obtain high-accuracy estimates of $\hat{\nu}$ and $\hat{\omega}$. However, in noisy imaging environments (e.g. imaging with low dose and/or short exposure), the sequential approach may introduce systematic errors in $\hat{u}$, as we will now demonstrate with an example based on the WLS estimator (12).

To illustrate the effect that inaccurate flat-field and intensity estimates have on b, we start by rewriting b. It follows from the variable transformations (3) and (4) that

$$v = \log(\hat{\nu}) - \log(I_0\nu), \qquad w = \log(\hat{\omega}) - \log(\omega),$$

and if we express y as $y = \operatorname{diag}(\bar{y}) \exp(-e)$, we may decompose b as

$$\begin{aligned} b &= \log(\hat{\omega} \otimes \hat{\nu}) - \log(y) \\ &= \log(\omega \otimes (I_0\nu)) - \log(\bar{y}) + e + Gv + Hw \\ &= \bar{b} + e + Gv + Hw. \end{aligned} \tag{13}$$

Here $\bar{b}$ represents noise-free measurements (i.e. $\bar{b} = Au$), e is a vector that represents measurement noise, and the terms Gv and Hw are systematic errors due to flat-field and intensity estimation errors, respectively. Figure 2 shows an example of what these terms may look like for a simulated set of measurements with a time-varying source intensity. As can be seen from the figure, the error terms that arise because of flat-field and intensity estimation errors are highly systematic in nature.

The decomposition of b in (13) allows us to analyze the effect of the different error terms when a linear reconstruction operator is applied to b. Examples of linear reconstructions operators include WLS estimator (12) and filtered backprojection (FBP). Figure 3 shows the FBP reconstructions of the individual terms in the decomposition of b. This example demonstrates that model errors propagate to the reconstruction when the underlying assumptions are violated, e.g. ignoring a time-varying intensity and/or flat-field errors.

2.3. Bayesian hierarchical model

We now turn our attention to a Bayesian hierarchical model that is based on the extended measurement model (6), i.e. we will assume that the detector response and intensity drift are independent and that the measurements are Poisson with the mean determined by Lambert–Beer's law.

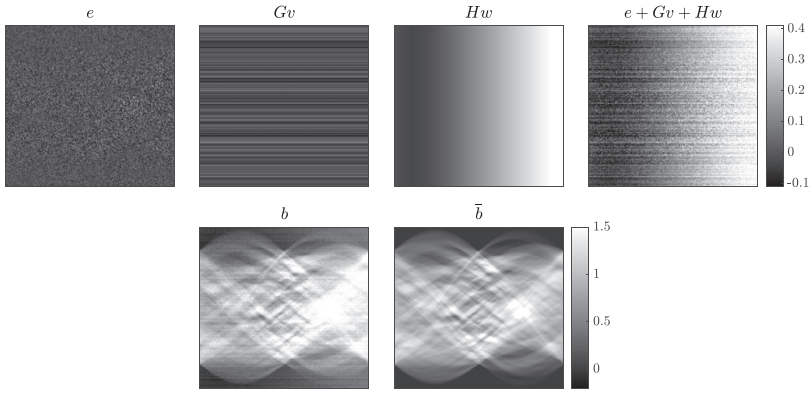

Figure 2. Illustration of the error terms in (13) introduced by conventional flat-field correction.

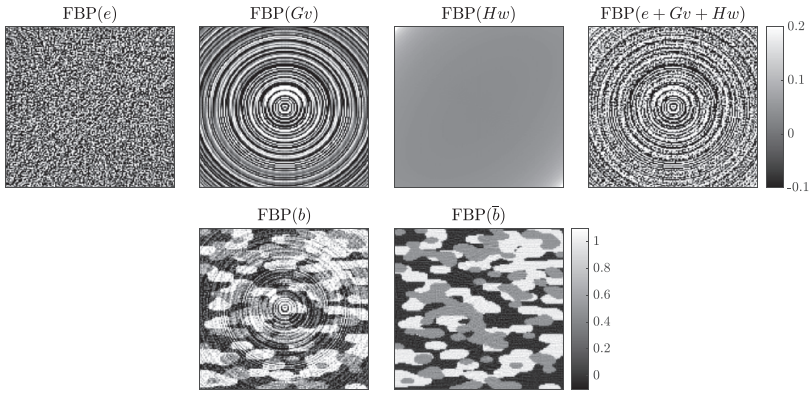

Figure 3. FBP reconstruction of the error terms in (13). The measurement noise, e appears as random noise as expected and the flat-field error, Gv appears as concentric rings in the reconstruction. The contribution from the intensity drift, Hw, is smeared out with varying intensity affecting all the pixels in the reconstruction, especially in the top left and bottom right corner of the reconstruction. The FBP reconstruction of the sum of the error components is depicted in the most right image, illustrating that they significantly degrade the quality of the reconstruction.

In order to formulate a Bayesian hierarchical model, we treat u, v and w as independent random variables with suitable prior distributions, say, $\pi(u|\delta)$, $\pi(v|\alpha)$, and $\pi(w|\beta)$ where δ, α, and β represent independent hyperparameters with hyperpriors $\pi(\delta)$, $\pi(\alpha)$, and $\pi(\beta)$. Let $x = (u, v, w)$ and $\eta = (\delta, \alpha, \beta)$ represent the model parameters and hyperparameters, respectively. Using Bayes' rule, the joint posterior distribution of our parameters and hyperparameters may be expressed as

$$\pi(x, \eta|y) = \frac{\pi(y|x)\pi(x|\eta)\pi(\eta)}{\pi(y)}. \tag{14}$$

This distribution may be used to compute point estimates (MAP, ML, posterior mean) and interval estimates (credible intervals) which, in turn, provide a useful, interpretable visual summary of the uncertainty of the model parameters. The MAP and ML point estimates generally require the solution of an optimization problem whereas the computation of the posterior mean and interval estimates typically involve multidimensional integration. In practice, sampling methods such as MCMC are often used to approximate multidimensional integrals, thus avoiding deterministic numerical integration methods that are subject to the curse of dimensionality.

The choice of priors plays a crucial role in that it provides regularization for an otherwise ill-posed inverse problem [13]. Several methods exist for selecting or estimating the hyperparameters when the hyperpriors are omitted, e.g. the GCV and L-curve criteria [13]. However, the use of the full posterior with weakly informative hyperpriors allows us to obtain useful statistical information about the hyperparameters through samples from the posterior distribution. This may come at the cost of a significant increase in computational cost, e.g. if a large optimization problem must be solved for each sample [3].

To address the computational cost associated with sampling in CT problems, we also propose an approximate posterior distribution that, with certain choices of priors and hyperpriors, may be cheaper to sample from than the posterior (14). Specifically, we will replace the Poisson likelihood $\pi(y|x)$ (see (6)) by a Gaussian approximation $\tilde{\pi}(b|x)$ where we define $b = \log(\hat{\omega} \otimes \hat{\nu}) - \log(y)$ as in (13) and let

$$b|x \sim \mathcal{N}(Mx, \Sigma_b), \qquad \Sigma_b = \operatorname{diag}(y)^{-1}, \qquad M = \begin{bmatrix} A & G & H \end{bmatrix}. \tag{15}$$

The derivation is included in appendix A. The resulting posterior then satisfies

$$\tilde{\pi}(x, \eta|b) \propto \tilde{\pi}(b|x)\pi(x|\eta)\pi(\eta). \tag{16}$$

We end this section by mentioning that both the model of Aggrawal *et al* [1] and that of Salehjahromi *et al* [22] may be viewed as special cases of the hierarchical model (14). Specifically, the model of Aggrawal *et al* [1] corresponds to the assumption that the prior on v is a gamma distribution and that $\hat{\omega} = \mathbf{1}$ and $w = 0$. Furthermore, the so-called SWLS reconstruction model may be viewed as a special case based on the approximation (16). In other words, the reconstruction models of Aggrawal *et al* [1] include the flat-field but ignore intensity drift. Similarly, the model of Salehjahromi *et al* [22] can be expressed as a special case of (16) with $\hat{\omega} = \mathbf{1}$ and $w = 0$ and specific priors on u and v. Specifically, it corresponds to a Laplace prior on the elements of v and a Gibbs prior on u that is derived from a particular directional total variation penalty that discourages ring artifacts. The fact that both models are special cases of our extended model implies that both models can readily be extended to include a time-varying intensity.

3. Implementation

3.1. Choice of priors and hyperpriors

The likelihood function associated with the Poisson measurement model (6) and its Gaussian approximation $\tilde{\pi}(b|x)$ are log-concave in x. Thus, the posterior will also be log-concave if we choose a log-concave prior. To simplify the exposition, we will limit our attention to the

Gaussian approximation (16) which is used for the numerical experiments in section 4. A convenient prior on x is a Gaussian prior of the form $x|\eta \sim \mathcal{N}(0, \Sigma_x(\eta))$ with hyperparameters $\eta = (\delta, \alpha, \beta)$ and covariance matrix

$$\Sigma_x(\eta) = \text{blkdiag}(\delta^{-1}\Sigma_u, \alpha^{-1}\Sigma_v, \beta^{-1}\Sigma_w), \tag{17}$$

where $\Sigma_u \in \mathbb{R}^{n\times n}$, $\Sigma_v \in \mathbb{R}^{r\times r}$, and $\Sigma_w \in \mathbb{R}^{p\times p}$ are symmetric and positive definite. This is a conjugate prior for the Gaussian likelihood, and it yields the conditional posterior

$$x|b, \eta \sim \mathcal{N}((M^T\Sigma_b^{-1}M + \Sigma_x(\eta)^{-1})^{-1}M^T\Sigma_b^{-1}b, (M^T\Sigma_b^{-1}M + \Sigma_x(\eta)^{-1})^{-1}), \tag{18}$$

which we will denote $\tilde{\pi}(x|b, \eta)$. This distribution coincides with the distribution of the solution to the stochastic quadratic problem

$$x(\xi) = \underset{x}{\text{argmin}} \left\{ \frac{1}{2}\|Mx - b\|^2_{\Sigma_b^{-1}} + \frac{1}{2}\|x\|^2_{\Sigma_x(\eta)^{-1}} - \xi^T x \right\}, \tag{19}$$

where $\xi \sim \mathcal{N}(0, M^T\Sigma_b^{-1}M + \Sigma_x(\eta)^{-1})$, which suggests that samples can be approximated by computing inexact solutions to instances of this problem. Indeed, the gradient scan Gibbs sampler (GSGS) [11] produces samples that are asymptotically distributed according to the target distribution (18) by applying a small, fixed number of iterations of the conjugate gradient method to instances of (19).

In order to form the full approximate posterior $\tilde{\pi}(x, \eta|b)$, we need to define a prior on the hyperparameters η. The gamma distribution is a conjugate prior, and hence it is a convenient choice. We define $\pi(\eta) = \pi(\delta)\pi(\alpha)\pi(\beta)$ where

$$\delta \sim \text{Gamma}(\alpha_\gamma, \beta_\gamma), \quad \alpha \sim \text{Gamma}(\alpha_\gamma, \beta_\gamma), \quad \beta \sim \text{Gamma}(\alpha_\gamma, \beta_\gamma), \tag{20}$$

and we will assume that the shape parameter α_γ and the inverse scale parameter β_γ are known (e.g. $\alpha_\gamma = 1$ and $\beta_\gamma = 10^{-4}$, as suggested in [3], which yields an exponential distribution with rate β_γ). The posterior distribution of η, conditioned on x, may then be expressed as the product $\pi(\eta|x) = \pi(\delta|u)\pi(\alpha|v)\pi(\beta|w)$ where

$$\pi(\delta|u) \propto \delta^{n/2+\alpha_\gamma-1} \exp\left(-\delta\left(\frac{1}{2}u^T\Sigma_u^{-1}u + \beta_\gamma\right)\right) \tag{21a}$$

$$\pi(\alpha|v) \propto \alpha^{r/2+\alpha_\gamma-1} \exp\left(-\alpha\left(\frac{1}{2}v^T\Sigma_v^{-1}v + \beta_\gamma\right)\right) \tag{21b}$$

$$\pi(\beta|w) \propto \beta^{p/2+\alpha_\gamma-1} \exp\left(-\beta\left(\frac{1}{2}w^T\Sigma_w^{-1}w + \beta_\gamma\right)\right). \tag{21c}$$

As an alternative to the Bayesian hierarchical model (16) with the priors (18) and (20), we also consider an implicit prior that restricts u to be nonnegative. Following the approach by Bardsley and Fox [5] and Bardsley and Hansen [6], we then define the conditional posterior $\tilde{\pi}(x|b, \eta) \propto \tilde{\pi}(b|x)\pi(x|\eta)$ as the distribution of the solution to the constrained stochastic quadratic problem

$$x(\xi) = \underset{x\in C}{\text{argmin}} \left\{ \frac{1}{2}\|Mx - b\|^2_{\Sigma_b^{-1}} + \frac{1}{2}\|x\|^2_{\Sigma_x(\eta)^{-1}} - \xi^T x \right\}, \tag{22}$$

where $\xi \sim \mathcal{N}(0, M^T\Sigma_b^{-1}M + \Sigma_x(\eta)^{-1})$ and $C = \mathbb{R}^n_+ \times \mathbb{R}^r \times \mathbb{R}^p$. The resulting conditional posterior may be viewed as a mixture of two distributions: a truncated Gaussian, defined on the

interior of C, and a distribution on the boundary of C that is obtained from (18) by transporting the mass outside C to the boundary of C. A sample from $\tilde{\pi}(x|b,\eta)$ can be computed by solving (22) for a random realization of ξ. We note that without the constraint $x \in C$ in (22), the prior reduces to the explicit prior $x|\eta \sim \mathcal{N}(0, \Sigma_x(\eta))$, resulting in the conditional posterior (18). Finally, with the nonnegativity assumption, we replace the conditional posterior (21a) by

$$\pi(\delta|u) \propto \mathbf{1}_{\mathbb{R}^n_+}(u)\delta^{\bar{n}/2+\alpha_\gamma-1} \exp\left(-\delta\left(\frac{1}{2}u^T\Sigma_u^{-1}u + \beta_\gamma\right)\right), \tag{23}$$

where

$$\mathbf{1}_{\mathbb{R}^n_+}(u) = \begin{cases} 1, & u \in \mathbb{R}^n_+, \\ 0, & u \notin \mathbb{R}^n_+, \end{cases}$$

is the characteristic function associated with the nonnegative orthant and $\bar{n}$ is the number of nonzeros in u; see [5] for details. Both (21b) and (21c) remain unchanged.

3.2. Two-stage hierarchical Gibbs sampler

The hierarchical Gibbs sampler [9] is a natural choice for our hierarchical model due to the availability of the full conditionals, i.e. (18) and (21) or, with the nonnegativity assumption, (22), (23), (21b), and (21c). In both cases, the conditional posterior of x is a high-dimensional distribution, and to mitigate the high cost of sampling with direct methods, we will compute inexact solutions to (19) or (22) using an iterative method. Specifically, we will use FISTA [7] with a fixed number of iterations and using the previous sample, x^k, as initialization. On a conceptual level, the resulting sampler is very similar to the GSGS, but unlike the conjugate gradient method used in the GSGS, FISTA can readily be applied to both the unconstrained problem (19) and the constrained problem (22). We note that terminating FISTA after a fixed number of iterations incurs an approximation error, and hence the resulting Markov chain may not have the exact target distribution as its stationary distribution. The resulting two-stage hierarchical Gibbs sampler is shown in algorithm 1. We note that it is parameter-free in the sense that there is no need to tune or estimate the hyperparameters η. In fact, it is sufficient to initialize x (e.g. using FBP or some point estimate) since the hyperparameter vector η^1 only depends on x^0.

4. Numerical experiments

We now turn to a numerical investigation of our proposed extended reconstruction models. We will consider reconstructions based on four different methods. The first method is an approximate MAP (AMAP) reconstruction based on (22) with $\xi = 0$, i.e. it is the mode of the conditional posterior $x|b,\eta$ under the assumption that u is nonnegative[1]. The second reconstruction method that we will explore is based on the mean of $x,\eta|b$, also with the nonnegativity assumption on u. We will refer to this as an approximate posterior mean (APM) reconstruction. Note that the AMAP reconstruction method depends on η and involves the solution of an optimization problem. By contrast, the APM reconstruction method estimates x and η simultaneously using algorithm 1, and it requires an inexact solution of an optimization problem

[1] The variant of our reconstruction models that include a nonnegativity assumption on the attenuation image outperformed the unconstrained counterpart in all of our experiments, and hence we limit the presentation of our experimental results to those obtained with the nonnegativity assumption.

Algorithm 1. Two-stage hierarchical Gibbs sampler.

Input: Number of iterations K
Data: $y \in \mathbb{R}^{rp}$ and $(t_j, \theta_j, \delta_j)$, $j = 1, \ldots, p$
Compute $\hat{\nu} \in \mathbb{R}^r$ and $\hat{\omega} \in \mathbb{R}^p$, e.g., using (8) and (9), respectively
Compute $b = \log(\hat{\omega} \otimes \hat{\nu}) - \log(y)$
Compute initialization $x^0 \in \mathbb{R}^{n+r+p}$, e.g., $x^0 = (u^0, 0, 0)$ with $u^0 \in \mathbb{R}^n$ or $u^0 \in \mathbb{R}^n_+$
if *nonnegativity prior* **then**
  for $k = 0, 1, ..., K-1$ **do**
    1. Compute $\eta^{k+1} \sim \pi(\eta | x^k)$ by sampling from (21b), (21c), (23) with $x = x^k$
    2. Compute x^{k+1} by inexactly solving an instance of (22) with $\eta = \eta^{k+1}$
  end
else
  for $k = 0, 1, ..., K-1$ **do**
    1. Compute $\eta^{k+1} \sim \pi(\eta | x^k)$ by sampling from (21) with $x = x^k$
    2. Compute x^{k+1} by inexactly solving an instance of (19) with $\eta = \eta^{k+1}$
  end
end

Table 1. Reconstruction methods used in experiments.

Abbreviation	Reconstruction method
AMAP	Approximate MAP est.; FISTA applied to (22) with $\xi = 0$
APM	Approximate posterior mean est.; algorithm 1 with nonnegativity
FBP	Filtered backprojection
P-FBP	Filtered backprojection with preprocessing

in each iteration. In all experiments, we have used 20 iterations of FISTA to compute inexact solutions to (19) in algorithm 1. We will compare our reconstructions to those obtained using FBP with and without a sinogram preprocessing step based on the wavelet and FFT filtering approach proposed by Münch *et al* [17]. In all experiments with preprocessing, we used the damping factor 0.9 and the Daubechies 5 wavelet with a three-level decomposition; see [17] for details. We will use the shorthand notation P-FBP to refer to the FBP algorithm with preprocessing. Table 2 contains an overview of the four reconstruction methods and their abbreviations.

In all of the models, we define $b = \log(\hat{\omega} \otimes \hat{\nu}) - \log(y)$ with the implicit assumption that y, $\hat{\omega}$, and $\hat{\nu}$ are positive vectors. We use (8) to compute $\hat{\nu}$, and to address the intensity drift, we consider four different ways of estimating $\hat{\omega}$, as outlined in table 3. The first option is 'no correction (NC)' (i.e. $\hat{\omega} = \mathbf{1}$), and the second option is to use the linear parametric model (10). The third option uses the more general model (11) with a binary matrix $C \in \mathbb{R}^{r \times p}$ defined as a $C_{ij} = 1 - \tilde{\delta}_i \delta_j$ where $\tilde{\delta}_i \in \{0, 1\}$ indicates whether the ith detector pixel is influenced by the object of interest ($\tilde{\delta}_i = 1$) or provides direct flat-field measurements in all projections ($\tilde{\delta}_i = 0$). In other words, the estimation of ω relies on flat-field projections and a set of 'AR' of the augmented, raw sinogram. In all experiments with this parameterization, we include the first

Table 2. Intensity estimates based on (11).

Method	Abbreviation	Parameterization	Binary mask
No correction	NC	$\omega(\alpha) = \mathbf{1}$	N/A
Linear regression	LR	$\omega(\alpha) = \begin{bmatrix}\mathbf{1} & t\end{bmatrix}\alpha$	$C_{ij} = 1 - \delta_j$
Air rows	AR	$\omega(\alpha) = \alpha$	$C_{ij} = 1 - \tilde{\delta}_i\delta_j$
Baseline	B	$\widehat{\nu} = \nu, \widehat{\omega} = \omega$	N/A

and last five rows of the augmented, raw sinogram (i.e. we define $\tilde{\delta}_i = 1$ for $i = 6, \ldots, r-5$, and otherwise $\tilde{\delta}_i = 0$). As a fourth option, we use the ground truth $\hat{\nu} = \nu$ and $\hat{\omega} = \omega$ as a baseline (B) to illustrate the effect of perfect knowledge of the detector response and intensity drift.

We report our results from two simulated experiments. The first experiment compares reconstructions based on our extended models, AMAP and APM, to those obtained using FBP and P-FBP. The second experiment investigates the statistical properties of our reconstruction method by comparing the reconstruction quality for reconstructions obtained with FBP, P-FBP, and AMAP based on 100 realizations of the measurements.

Before presenting our results in sections 4.3 and 4.4, we briefly present some error measures and outline our simulation setup.

4.1. Quantitative error measures

In addition to our qualitative assessment based on a visual comparison of the reconstructions, we will report two quantitative error measures. The first one is the relative attenuation error (RAE), defined as

$$\mathrm{RAE}(\hat{u}) = \frac{\|\hat{u} - u\|_2}{\|u\|_2}, \tag{24}$$

where $\hat{u}$ is the estimated reconstruction and u is the ground truth. The second error measure is the relative source intensity error (RSE), defined as

$$\mathrm{RSE}(\hat{\nu}, \hat{\omega}) = \frac{\|\hat{\nu}\hat{\omega}^T - \nu\omega^T\|_{\mathrm{F}}}{\|\nu\omega^T\|_{\mathrm{F}}}, \tag{25}$$

where $\hat{\nu}$ and $\hat{\omega}$ are the estimated detector response and time-varying intensity, respectively. We will also visualize the elementwise absolute error for both reconstruction and sinogram, i.e. $\hat{u} - u$ and $\hat{\nu}\hat{\omega}^T - \nu\omega^T$, respectively.

In the experiments where several samples are available, we assess the statistical properties of the results by visualizing the sample mean and the empirical 95% credibility interval (CI). The latter is defined by the empirical 0.025 and 0.975 quantiles of the samples. For the reconstructions, we visualize this as an image of the pixelwise width of the interval (i.e. the difference between the 0.975 quantile and the 0.025 quantile).

To assess the convergence of our Gibbs sampler (algorithm 1), we use the Geweke diagnostic [12] which tests for equality of means based on two subsets of samples, one from the beginning of the chain and one from the end. We will also show the hyperparameter chains and histograms of these, allowing us to qualitatively assess the convergence of the hyperparameter chains [4].

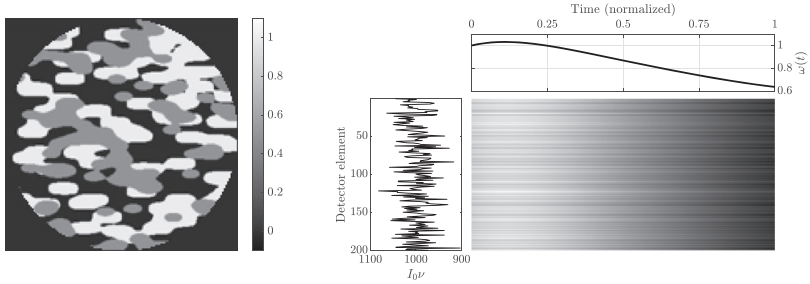

(a) Phantom. (b) Source intensity.

Figure 4. Phantom, the detector response and source intensity used to generate the simulated data for the numerical experiments.

4.2. Experimental setup

In order to assess the ability of our model to mitigate artifacts due to flat-field estimation errors and intensity drift, we simulate an experimental setup in which the source intensity is relatively low and subject to significant drift. Such an experimental setup often presents a significant challenge to existing reconstruction methods and may lead to severe reconstruction artifacts if the source uncertainty is ignored. We simulate a parallel beam projection geometry with 720 equidistant projection angles from 0 to 180 degrees, a detector with 200 detector elements, and a 2×2 cm reconstruction domain discretized into a 140×140 pixel grid (i.e. $n = 140^2$). To avoid committing an inverse crime, we generate the data using a finer grid (280×280 pixels).

The number of flat-field measurements vary between the experiments. In the first experiment, we simulate two flat-field projections prior to the scan and a single flat-field projection after the scan (i.e. $p = 723$). In the second experiment, we only have two flat-field projections prior to the scan (i.e. $p = 722$). We use the MATLAB package AIR Tools II [14] to generate the phantom and the augmented system matrix. In all simulations, we use the 'threephases' phantom, shown in figure 4(a). The true flat-field is generated by simulating $I_0\nu_i \sim \text{Poisson}(10^3)$, and the source intensity drift, $\omega(t)$, is given by the fifth degree polynomial[2]

$$\omega(t) = 0.798t^5 - 2.580t^4 + 3.441t^3 - 2.311t^2 + 0.396t + 1.2629, \quad t \in [0, 1].$$

Both $I_o\nu$ and $\omega(t)$ are shown in figure 4(b).

4.3. Intensity drift correction

We start by investigating how the different intensity drift correction methods, listed in table 3, impact the reconstructions obtained using the four methods in table 2. Figure 5 shows the sinogram for the four different intensity drift estimation methods. Without intensity drift correction, the absolute error grows toward the end of the scan due to the decreasing intensity.

The LR estimate reduces the overall magnitude of the error, but the horizontal gradient in the visualization of the absolute error reveals that the error has a significant systematic component. The 'air row' (AR) estimate uses additional information for the estimation of the intensity drift,

[2] The intensity drift was obtained from a real data set by fitting a fifth degree polynomial to the mean of a set of air rows in the sinogram. We were not able to include this data set due to license restrictions.

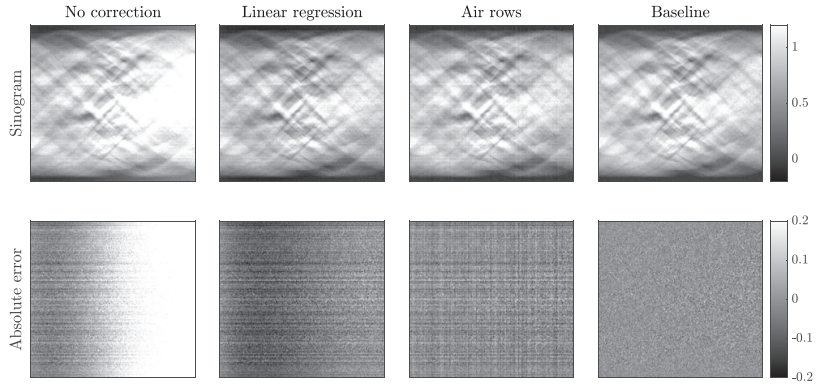

Figure 5. Visualization of the sinogram and the pixelwise absolute error for each of the four different source intensity correction methods listed in table 3.

Table 3. RAE and RSE for the intensity drift correction experiment depicted in figure 6.

	RAE				
	FBP	P-FBP	APM	AMAP	RSE
NC	0.459	0.409	0.235	0.193	0.304
LR	0.312	0.262	0.227	0.192	0.102
AR	0.313	0.257	0.225	0.188	0.080
B	0.246	0.259	0.216	0.181	0.068

and the resulting estimate yields a more accurate sinogram. However, as can be seen from the visualization of the error, it has introduced some vertical stripes in the sinogram. These are due to the noise in the measurements and can possible be mitigated by including a suitable smoothness prior in the intensity drift estimation problem (11). Figure 5 also includes the B correction where the effect of estimation errors is eliminated completely. As a result, the error does not have a systematic component, only independent measurement noise. The RSE associated with each of the four intensity drift estimates are included in table 4. These numbers verify that the RSE decreases from left to right in figure 5. Note that the most significant decrease in the RSE occurs when going from NC to the LR correction. Both the LR and AR correction are relatively close to the B in terms of their RSE.

Next, we consider the reconstructions which are shown in figure 6 with the corresponding RAEs listed in table 4. Recall that the AMAP reconstruction depends on the hyperparameters $\eta = (\delta, \alpha, \beta)$. We used the sample mean of the hyperparameter chains produced by algorithm 1 as an initial estimate of the hyperparameters for AMAP. These values generally resulted in slightly underregularized reconstructions that were similar to the APM reconstructions, and we were able to improve the AMAP reconstruction by increasing δ and α relative to the mean of the sampled hyperparameters. The hyperparameters used for the AMAP reconstructions are listed in table 5.

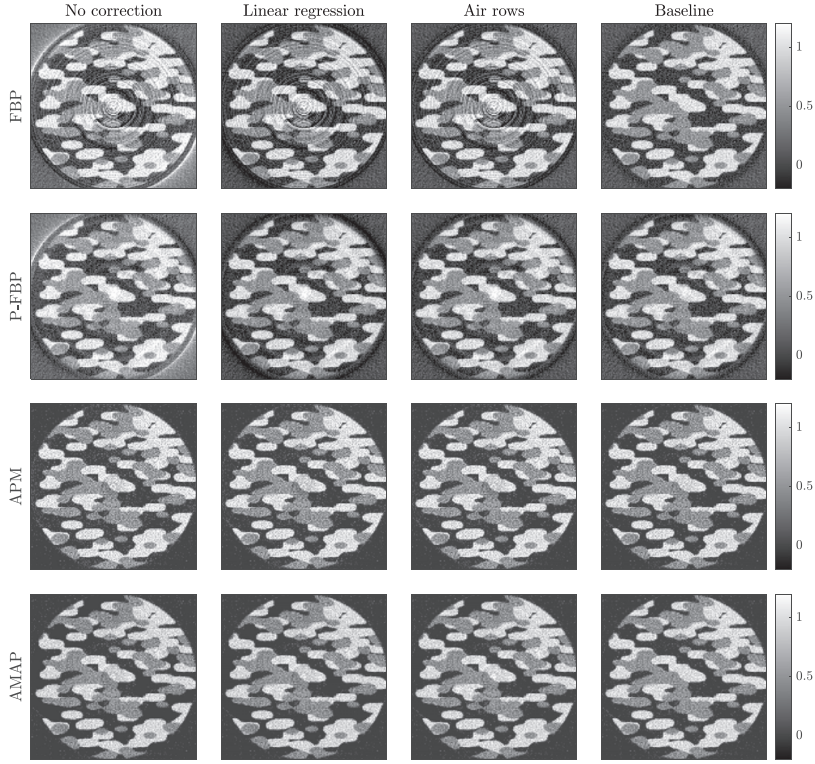

Figure 6. Visualization of the reconstruction obtained by FBP, P-FBP, APM and AMAP for the four different intensity drift correction methods.

Table 4. Hyperparameters used for AMAP reconstructions depicted in figure 6.

	δ	α	β
NC	11.48	1.292×10^4	2.830×10^1
LR	11.62	1.547×10^4	8.279×10^2
AR	11.63	2.672×10^4	1.726×10^3
B	11.52	7.603×10^4	2.221×10^5

It is clear from the reconstructions in figure 6 that, to a large extent, the P-FBP reconstruction method mitigates ring artifacts. However, both the FBP and P-FBP reconstructions contain noticeable drift artifacts, especially without intensity correction (i.e. $\hat{\omega} = \mathbf{1}$). The LR and AR estimates reduce the drift artifacts for both methods which can be verified both visually and in terms of the RAEs. Unsurprisingly, the best FBP and P-FBP reconstructions are obtained with the B intensity estimate, and in this case, the FBP reconstruction is marginally better that

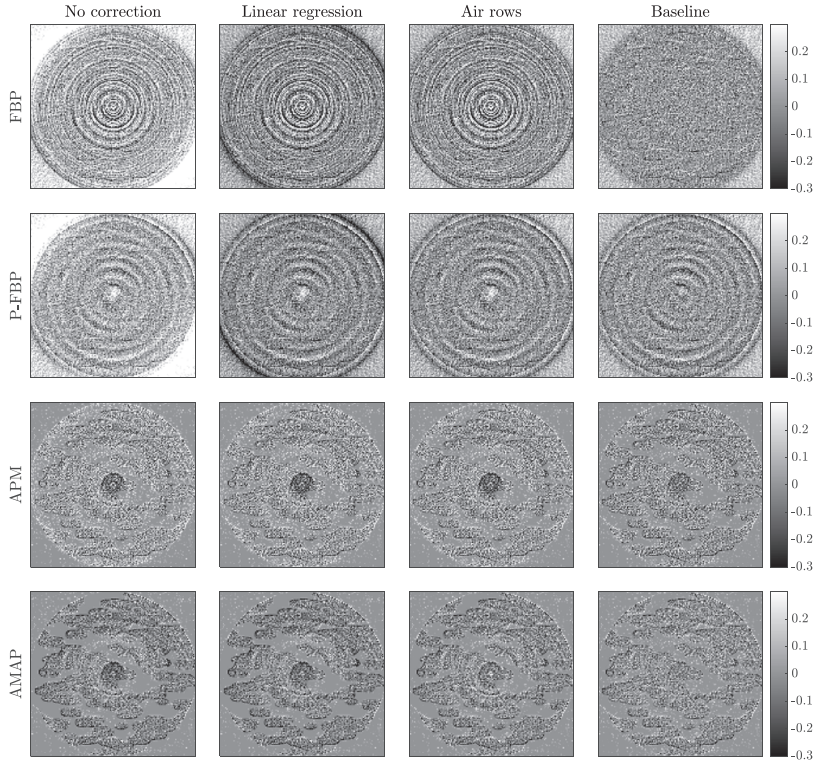

Figure 7. Visualization of the absolute errors for the reconstructions in figure 6.

the P-FBP one. The reason for this is that the preprocessing used in P-FBP is unnecessary and may introduce errors.

Unlike the FBP and P-FBP reconstructions, the APM and AMAP reconstructions are not very sensitive to the intensity drift estimate. This is not surprising since the two reconstruction models are based on the same extended model. This is a clear advantage as the LR and AR estimates rely on the availability of temporal information in the form of, e.g. flat-field images before and after the regular projections and/or entire rows of air pixels. The results demonstrate that the extended model can yield a significant reduction in ring and drift artifacts, although both the APM and AMAP reconstructions contain a faint ring or disc-like artifact in the center. This can be seen more clearly from the error images included in figure 7, and it is consistent with the analysis of the effect of flat-field estimation errors by Aggrawal *et al* [1]. The AMAP reconstructions appear to be somewhat less noisy that the corresponding APM reconstructions. This is partly due to the fact that the AMAP reconstructions are based on a larger value for δ relative to the posterior mean estimate obtained with APM.

Figure 8 shows the APM, the pixelwise 95% CIs, and the hyperparameter chains for the case without correction (i.e. $\hat{\omega} = \mathbf{1}$). Notice that pixelwise attenuation CIs reveal that the

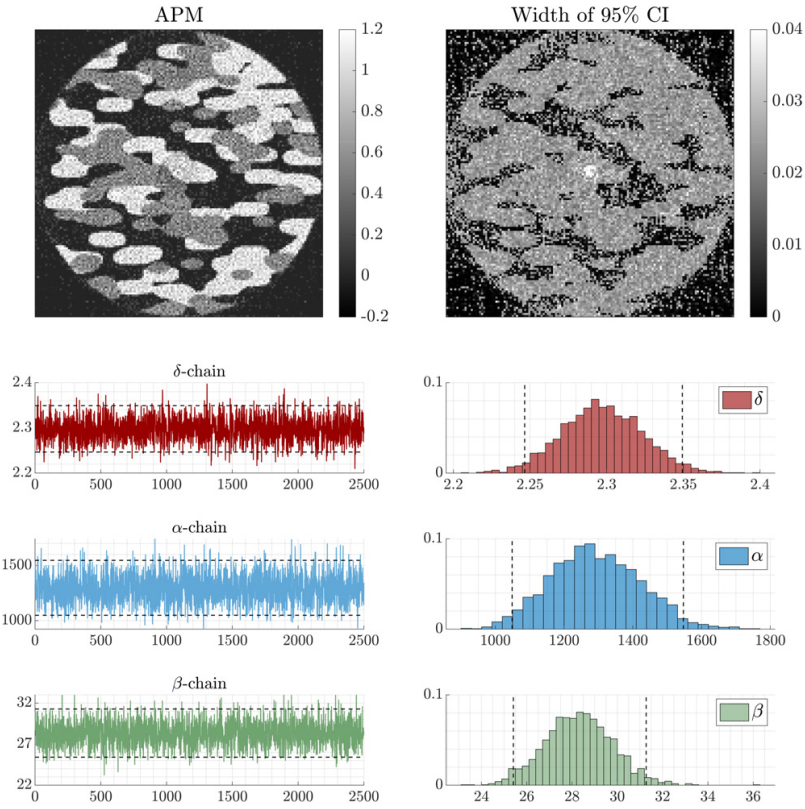

Figure 8. APM reconstruction obtained using algorithm 1 with NC, i.e. $\widehat{\omega} = \mathbf{1}$. The hyperparameter chains and histograms suggest that these have converged. The dashed lines represent the 95% CIs. This is further supported by the Geweke p-value [12], shown in table 5.

uncertainty is relatively large in the center of the reconstruction which is consistent with what we observe from the error images in figure 7. The chains suggest that the hierarchical Gibbs sampler has converged. Similar results were obtained for the three other intensity drift correction estimates which we omit to save space. Additional chain statistics are included in table 5 for all four intensity drift estimates.

4.4. Ensemble experiment

The aim of our next experiment is to empirically investigate pixelwise reconstruction bias and uncertainty based on an ensemble of random realizations of measurements. To this end, we generate 100 realizations of the measurements based on (7) with two flat-field images in the very beginning. For each realization, we compute a reconstruction with three methods: (i) FBP

Table 5. Chain statistics for the APM reconstruction method. The 'NC' results are shown in figure 8.

	Samples	Burn-in		Geweke p-value	95% CI quantiles	
					Lower	Upper
NC	5000	2500	δ	0.999	2.25×10^0	2.35×10^0
			α	0.996	1.05×10^3	1.55×10^3
			β	0.996	2.54×10^1	3.13×10^1
LR	2000	1000	δ	0.997	2.27×10^0	2.38×10^0
			α	0.978	1.25×10^2	1.85×10^2
			β	0.992	7.43×10^2	9.20×10^2
AR	2000	1000	δ	0.999	2.27×10^0	2.38×10^0
			α	0.996	4.36×10^2	6.41×10^2
			β	0.987	1.55×10^3	1.91×10^3
B	2000	1000	δ	0.997	2.25×10^0	2.36×10^0
			α	0.970	6.16×10^3	9.23×10^3
			β	0.999	1.73×10^5	2.79×10^5

with flat-field correction but without intensity drift correction, (ii) P-FBP with flat-field correction and the AR intensity drift correction (see table 3), and (iii) the extended reconstruction model, AMAP, defined in (23) and using the same $\hat{\nu}$ and $\hat{\omega}$ as for P-FBP. We used the hyperparameters $\delta = 12.84$, $\alpha = 3.52 \times 10^4$ and $\beta = 3.59 \times 10^3$ for AMAP; these were adjusted manually using the sample posterior mean estimate obtained by the hierarchical Gibbs sampler as an initial guess. Our results are summarized in figure 9, which shows reconstructions based on a single realization, the empirical mean, and the width of the 95% CI. Moreover, figure 10 includes histograms of the RAE for the 100 noise realizations.

The FBP reconstruction in figure 9 contains severe ring and drift artifacts. The effect of intensity drift is especially noticeable toward the lower right and upper left corners. The empirical mean, on the other hand, does not contain ring artifacts, but it is clear that the drift incurs a low-frequency bias. This is not surprising: the flat-field estimate is a stochastic variable whereas the intensity drift is fixed and has a deterministic effect on the sinogram. Finally, the 95% CI reveals that the uncertainty is largest in the center. Compared to the FBP reconstruction, the P-FBP reconstruction is much less affected by ring artifacts. This is consistent with the much narrower 95% CI. However, a closer inspection of the P-FBP reconstruction reveals that new artifacts have been introduced. Specifically, notice the embossed edge of the object. The effect is even more pronounced in the empirical mean which also reveals a wave-like pattern throughout the object. The AMAP reconstruction contains the least noise and artifacts of the three reconstructions, and although the empirical mean reveals a small bias in the center, the AMAP reconstruction is a clear improvement over the P-FBP reconstruction. This is confirmed by figure 10 which shows RAE histograms for the three methods. Considering the width of the 95% CI for AMAP, we see a clear structural difference compared to those obtained with the two other methods. Specifically, the AMAP reconstruction appears to have low uncertainty in areas of the image where the pixel values are zero. This is a consequence of the nonnegativity assumption that is included in the underlying reconstruction model. Neither FBP nor P-FBP makes such an assumption, and hence the support of the object is not visible in the CI images.

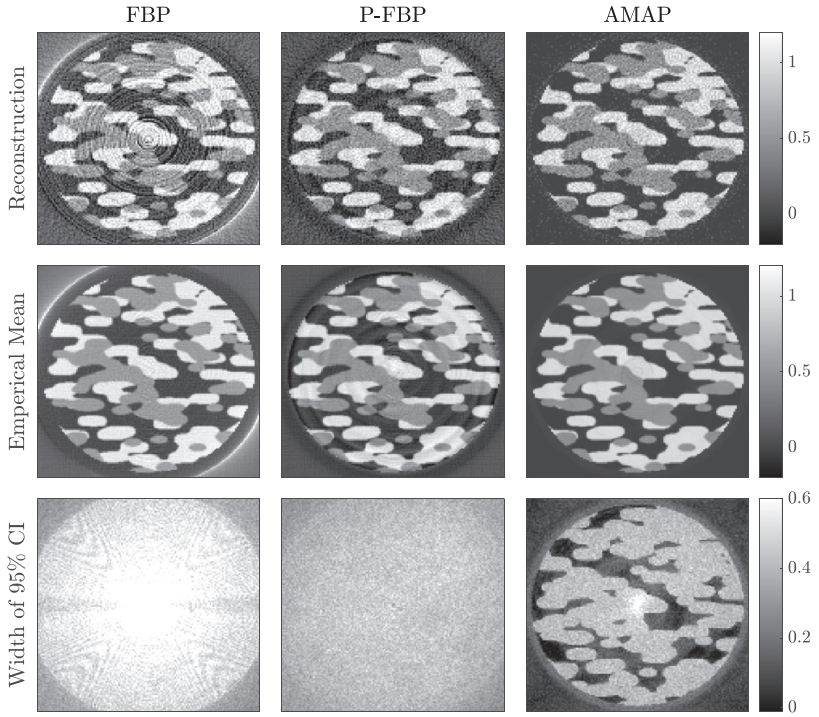

Figure 9. Statistical simulation study for 100 noise realizations, comparing FBP with NC and P-FBP and AMAP with air row correction. Top row depicts a reconstruction for a single noise realization whereas the middle row depicts the empirical mean of reconstructions for the 100 noise realizations. Lastly, the bottom row shows the width of the 95% CI for the 100 reconstructions.

5. Discussion and Conclusion

In CT, the source intensity is typically assumed to be constant and is estimated together with the detector response based on a set of flat-field images. However, intensity drift and noisy flat-field estimates give rise to systematic errors in the sinogram. These errors can lead to severe reconstruction artifacts, especially in experimental setups with a short exposure time or a low source intensity. By including the source intensity and the detector response as parameters in an extended reconstruction model, we have demonstrated that the attenuation image, the detector response and intensity drift can be estimated simultaneously, resulting in improved reconstructions under challenging conditions.

Our extended statistical measurement model (6) is readily embedded in a Bayesian framework, allowing the inclusion of suitable priors and the computation of quantities of interest such as point estimates and credible intervals. We have also derived a simplified model which is based on a Gaussian approximation of the likelihood function. This enables us to view the extended model as a generalization of many existing reconstruction models.

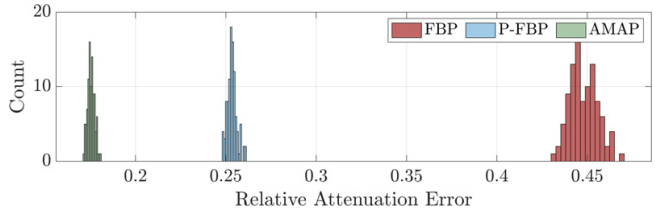

Figure 10. Histograms of RAE defined in (24) for figure 9. FBP has a mean RAE of 0.448 with a standard deviation of 7.7×10^{-3}, P-FBP has a mean RAE of 0.253 with a standard deviation of 2.7×10^{-3}, and finally, AMAP has a mean RAE of 0.175 with a standard deviation of 2.0×10^{-3}.

To overcome the computational burden associated with sampling from high-dimensional posterior distributions, we have investigated the use of a two-stage hierarchical Gibbs sampler. We propose to compute approximate solutions to a stochastic optimization problem within one of the two stages of the Gibbs sampler in order to allow the inclusion of an implicit nonnegativity prior on the attenuation image without incurring a high computational cost. The resulting method is parameter-free in the sense that the hyperparameters are estimated as part of the sampling process.

In the numerical experiments, we have compared two new reconstruction methods that are based on the extended model (AMAP and APM) to existing reconstruction methods with a heuristic ring reduction and intensity drift correction. The results, which are based on simulated data, demonstrate that our model is capable of reducing artifacts and is competitive with existing methods. We note that although the computational cost associated with the new methods exceeds the cost of many existing methods, the extended model offers a clear advantage when working with data sets from challenging experimental setups that are expensive and/or difficult to improve and redo.

The model that we have used in the experiments may be improved in several ways. We have chosen a Tikhonov prior because of its connection to the Gaussian density, but the reconstruction quality can possibly be improved by adopting more advanced priors. For example, one could impose an edge-preserving prior such as total variation on the attenuation image u and/or a smoothness prior on the intensity drift w, e.g. based on the spline kernel [2]. Moreover, the estimates $\hat{\nu}$ and $\hat{\omega}$ implicitly affect the prior on v and w, and hence it may be of interest to improve these estimates. For example, as an extension of the estimation problem (11), one could estimate ν and ω jointly or formulate a matrix completion problem.

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

Appendix A. Quadratic approximation of log-likelihood

The Gaussian approximation (15) may be obtained by constructing a quadratic approximation to the log-likelihood function (7), as we will now show. First, assume that $\mathbf{y}$ is Poisson distributed according to (6). An ML estimate of u, v, and w may then be expressed as

$$\begin{aligned}(\hat{u},\hat{v},\hat{w}) &= \underset{u,v,w}{\operatorname{argmin}} \left\{-\log(\pi(y|u,v,w))\right\} \\ &= \underset{u,v,w}{\operatorname{argmin}} \left\{y^T(Au+Gv+Hw)\right. \\ &\quad \left. + (\hat{\omega}\otimes\hat{\nu})^T \exp(-Au-Gv-Hw)\right\},\end{aligned}$$

or using a more compact notation,

$$\hat{x} = \underset{x}{\operatorname{argmin}} \left\{y^T Mx + (\hat{\omega}\otimes\hat{\nu})^T \exp(-Mx)\right\},$$

where $x = (u,v,w)$ and $M = \begin{bmatrix} A & G & H \end{bmatrix}$. Now, let $z = Mx$ and define

$$g(z) = y^T z + (\hat{\omega}\otimes\hat{\nu})^T \exp(-z),$$

such that $\hat{x} = \operatorname{argmin}_x\{g(Mx)\}$. We now construct a quadratic approximation of g from a second-order Taylor expansion of g around $b = \log(\hat{\omega}\otimes\hat{\nu}) - \log(y)$, i.e.

$$g(z) \approx g(b) + \nabla g(b)(z-b) + (z-b)^T \nabla^2 g(b)(z-b). \tag{26}$$

For this particular choice of b, we have that

$$\nabla g(b) = y - \operatorname{diag}(\hat{\omega}\otimes\hat{\nu})\exp(-b) = 0,$$

and

$$\nabla^2 g(b) = \operatorname{diag}(\hat{\omega}\otimes\hat{\nu})\operatorname{diag}(\exp(-b)) = \operatorname{diag}(y),$$

and hence (26) reduces to

$$g(z) \approx (z-b)^T \operatorname{diag}(y)(z-b).$$

Using the right-hand side, we arrive at the approximate ML estimation problem

$$\hat{x}_{\mathrm{AML}} = \underset{x}{\operatorname{argmin}} \left\{(Mx-b)^T \operatorname{diag}(y)(Mx-b)\right\},$$

or equivalently,

$$\hat{x}_{\mathrm{AML}} = \underset{x}{\operatorname{argmin}} \left\{\frac{1}{2}\|Mx-b\|^2_{\Sigma_b^{-1}}\right\},$$

with $\Sigma_b^{-1} = \operatorname{diag}(y)$. The approximate ML estimate $\hat{x}_{\mathrm{AML}}$ may be viewed as an ML estimate based on a Gaussian likelihood, corresponding to the model assumption $b|x \sim \mathcal{N}(Mx, \Sigma_b)$.

ORCID iDs

Katrine O Bangsgaard https://orcid.org/0000-0001-7491-0064
Martin S Andersen https://orcid.org/0000-0002-4654-3946

References

[1] Aggrawal H O, Andersen M S, Rose S D and Sidky E Y 2018 A convex reconstruction model for x-ray tomographic imaging with uncertain flat-fields *IEEE Trans. Comput. Imaging* **4** 17–31

[2] Andersen M S and Chen T 2020 Smoothing splines and rank structured matrices: revisiting the spline kernel *SIAM J. Matrix Anal. Appl.* **41** 389–412

[3] Bardsley J M 2018 Uncertainty *Computational Quantification for Inverse Problems* (Philadelphia, PA: SIAM)

[4] Bardsley J M 2012 MCMC-based image reconstruction with uncertainty quantification *SIAM J. Sci. Comput.* **34** A1316–32

[5] Bardsley J M and Fox C 2012 An MCMC method for uncertainty quantification in nonnegativity constrained inverse problems *Inverse Probl. Sci. Eng.* **20** 477–98

[6] Bardsley J M and Hansen P C 2020 MCMC algorithms for computational UQ of nonnegativity constrained linear inverse problems *SIAM J. Sci. Comput.* **42** A1269–88

[7] Beck A and Teboulle M 2009 A fast iterative shrinkage-thresholding algorithm for linear inverse problems *SIAM J. Imag. Sci.* **2** 183–202

[8] Buzug T M 2008 *Computed Tomography* (Berlin: Springer)

[9] Casella G and George E I 1992 Explaining the Gibbs sampler *Am. Statistician* **46** 167–74

[10] Doran S J, Koerkamp K K, Bero M A, Jenneson P, Morton E J and Gilboy W B 2001 A CCD-based optical CT scanner for high-resolution 3D imaging of radiation dose distributions: equipment specifications, optical simulations and preliminary results *Phys. Med. Biol.* **46** 3191–213

[11] Feron O, Orieux F and Giovannelli J-F 2016 Gradient scan Gibbs sampler: an efficient algorithm for high-dimensional Gaussian distributions *IEEE J. Sel. Top. Signal Process.* **10** 343–52

[12] Geweke J 1992 Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments *Bayesian Statistics 4* (Oxford: Clarendon) pp 169–93

[13] Hansen P C 2010 *Discrete Inverse Problems* (Philadelphia, PA: SIAM)

[14] Hansen P C and Jørgensen J S 2017 AIR Tools II: algebraic iterative reconstruction methods, improved implementation *Numer. Algorithms* **79** 107–37

[15] Kohler T 2004 A projection access scheme for iterative reconstruction based on the golden section *IEEE Symp. Conf. Record Nuclear Science* (IEEE)

[16] Kowalski G 1977 The influence of fixed errors of a detector array on the reconstruction of objects from their projections *IEEE Trans. Nucl. Sci.* **24** 2006–16

[17] Münch B, Trtik P, Marone F and Stampanoni M 2009 Stripe and ring artifact removal with combined wavelet-Fourier filtering *Opt. Express* **17** 8567

[18] Van Nieuwenhove V, De Beenhouwer J, De Carlo F, Mancini L, Marone F and Sijbers J 2015 Dynamic intensity normalization using eigen flat fields in x-ray imaging *Opt. Express* **23** 27975

[19] Paleo P and Mirone A 2015 Ring artifacts correction in compressed sensing tomographic reconstruction *J. Synchrotron Radiat.* **22** 1268–78

[20] Prell D, Kyriakou Y and Kalender W A 2009 Comparison of ring artifact correction methods for flat-detector CT *Phys. Med. Biol.* **54** 3881–95

[21] Rashid S, Lee S Y and Hasan M K 2012 An improved method for the removal of ring artifacts in high resolution CT imaging *EURASIP J. Adv. Signal Process.* **2012** 93

[22] Salehjahromi M, Wang Q, Gjesteby L A, Harrison D, Wang G and Yu H 2019 A directional TV based ring artifact reduction method *Medical Imaging 2019: Physics of Medical Imaging* ed H Bosmans, G-H Chen and T G Schmidt (SPIE)

[23] Sijbers J and Postnov A 2004 Reduction of ring artefacts in high resolution micro-CT reconstructions *Phys. Med. Biol.* **49** N247–53

[24] Six N, Beenhouwer J D, Nieuwenhove V V, Vanroose W and Sijbers J 2018 2018 Joint reconstruction and flat-field estimation using support estimation *IEEE Nuclear Science Symp. and Medical Imaging Conf. Proc. (NSS/MIC)* (IEEE)

[25] Titarenko V, Titarenko S, Withers P J, De Carlo F and Xiao X 2010 Improved tomographic reconstructions using adaptive time-dependent intensity normalization *J. Synchrotron Radiat.* **17** 689–99

[26] Vo N T, Atwood R C and Drakopoulos M 2018 Superior techniques for eliminating ring artifacts in x-ray micro-tomography *Opt. Express* **26** 28396

[27] Wei Z, Wiebe S and Chapman D 2013 Ring artifacts removal from synchrotron CT image slices *J. Instrum.* **8** C06006

[28] Yan L, Wu T, Zhong S and Zhang Q 2016 A variation-based ring artifact correction method with sparse constraint for flat-detector CT *Phys. Med. Biol.* **61** 1278–92

APPENDIX B

Low-rank flat-field correction for artifact reduction in spectral computed tomography

Published in Applied Mathematics in Science and Engineering, Volume 31, Number 1, 6 March 2023, DOI: 0.1080/27690911.2023.2176000.

Katrine O. Bangsgaard, Genoveva Burca, Evelina Ametova, Martin Skovgaard Andersen and Jakob Sauer Jørgensen.

APPLIED MATHEMATICS IN SCIENCE AND ENGINEERING
2023, VOL. 31, NO. 1, 2176000
https://doi.org/10.1080/27690911.2023.2176000

Taylor & Francis
Taylor & Francis Group

Low-rank flat-field correction for artifact reduction in spectral computed tomography

Katrine Ottesen Bangsgaard [a], Genoveva Burca [b,c,d], Evelina Ametova [c,e], Martin Skovgaard Andersen [a] and Jakob Sauer Jørgensen [a,c]

[a]Department of Applied Mathematics and Computer Science, Technical University of Denmark, Kongens Lyngby, Denmark; [b]ISIS Pulsed Neutron and Muon Source, STFC, UKRI, Rutherford Appleton Laboratory, Didcot, UK; [c]Department of Mathematics, The University of Manchester, Manchester, UK; [d]Diamond Light Source Ltd, Didcot, UK; [e]Laboratory for Applications of Synchrotron Radiation, Karlsruhe Institute of Technology, Karlsruhe, Germany

ABSTRACT

Spectral computed tomography has received considerable interest in recent years since spectral measurements contain much richer information about the object of interest. In spectral computed tomography, we are interested in the energy channel-wise reconstructions of the object. However, such reconstructions suffer from a low signal-to-noise ratio and share the challenges of conventional low-dose computed tomography such as ring artifacts. Ring artifacts arise from errors in the flat fields and can significantly degrade the quality of the reconstruction. We propose an extended flat-field model that exploits high correlation in the spectral flat fields to reduce ring artifacts in channel-wise reconstructions. The extended model relies on the assumption that the spectral flat fields can be well-approximated by a low-rank matrix. Our proposed model works directly on the spectral flat fields and can be combined with any existing reconstruction model, e.g. filtered back projection and iterative methods. The proposed model is validated on a neutron data set. The results show that our method successfully diminishes ring artifacts and improves the quality of the reconstructions. Moreover, the results indicate that our method is robust; it only needs a single spectral flat-field image, whereas existing methods need multiple spectral flat-field images to reach a similar level of ring reduction.

ARTICLE HISTORY
Received 8 June 2022
Accepted 17 January 2023

KEYWORDS
Ring artifacts; spectral computed tomography; neutron imaging; flat-field correction; low-rank approximation

MATHEMATICS SUBJECT CLASSIFICATIONS
68U10; 65R32; 65K10

1. Introduction

Computed Tomography (CT) is a non-invasive imaging technique that allows us to obtain structural knowledge about the interior of objects from a set of projection images. Projection images are acquired by illuminating the object from different angles with radiation from a source, e.g., an X-ray beam or beam of neutron radiation. The beam is attenuated as it travels through the object, and the attenuated beam is measured by a detector

CONTACT Jakob S. Jørgensen jakj@dtu.dk Department of Applied Mathematics and Computer Science, Technical University of Denmark, Richard Petersens Plads, Building 324, Kongens Lyngby 2800, Denmark; Department of Mathematics, The University of Manchester, Oxford Road, Alan Turing Building, Manchester M13 9PL, UK

placed opposite the source. The attenuation is governed by absorption in X-ray CT and by scattering and absorption in neutron CT – in both cases the attenuation is material- and energy-specific. In *spectral CT* [1, 2], attenuation is measured at multiple energies, either in turn or simultaneously, in other words at each energy level a full set of projection images is measured, which enables the reconstruction of object images at individual energy levels. Sometimes cases with few and many energy levels are distinguished by referring to the latter as *hyperspectral CT* [3, 4].

Spectral CT data can be reconstructed in different ways, the most straightforward being simply the reconstruction of each energy channel independently. From reconstructed energy images one may then, possibly with knowledge of the attenuation as a function of energy for the constituting materials seek to decompose into material images - this is known as material decomposition [5, 6]. A variety of material decomposition methods exist, sometimes with the decomposition carried out on the projection images before reconstruction and sometimes jointly. In this work, we focus solely on improving the step of energy-channel-wise reconstruction considered as a step toward improved material decomposition.

Most reconstruction methods rely on the assumption that the detector response is known. In practice, however, the detector response is subject to various errors and must be estimated from measurements acquired without an object in the scanner, i.e. from flat fields, also referred to as air scans [7], white fields [8] or open beams. The flat fields are noisy due to factors such as measurement noise, miscalibration, defective pixel elements with non-linear response, and may introduce concentric rings in the reconstruction, which are known as ring artifacts [9]. Ring artifacts are a great challenge for experimental CT set-ups with low-dose and/or short exposure time [10] and can significantly degrade the quality of the reconstruction. In spectral CT, we measure spectral flat fields, i.e. flat fields for each energy. However, the spectral measurements share the characteristics of low-dose CT since each energy channel has a low signal-to-noise ratio (SNR), and hence ring artifacts present a challenge in spectral CT [11, 12].

To illustrate the challenges of spectral CT, let us consider a neutron CT data set [13] which is described in detail in Section 3. Filtered back projection (FBP) reconstructions of the neutron data are shown for two energies in Figure 1, and the reconstructions reveal the presence of ring artifacts.

1.1. Existing methods for ring reduction

Several reconstruction methods have been proposed to combat ring artifacts as part of the spectral reconstruction step. Wu et al. [14] propose a reconstruction method that exploits the similarity across spectral images by computing a polychromatic reconstruction (average across the spectral dimension) as a reference image combined with total variation (TV). Lv et al. [15] and Fang et al. [16] both propose deep-learning approaches to suppress noise and remove ring artifacts for spectral CT. However, all methods rely on computationally expensive algorithms where the ring reduction is part of the reconstruction process.

Conventional preprocessing methods for ring reduction in monochromatic CT, i.e. single-energy CT, can also be applied to the spectral CT data, e.g. [17, 18]. The main drawback of these methods is that they are not designed for data with extremely low SNR. Moreover, they are sensitive to the choice of parameters.

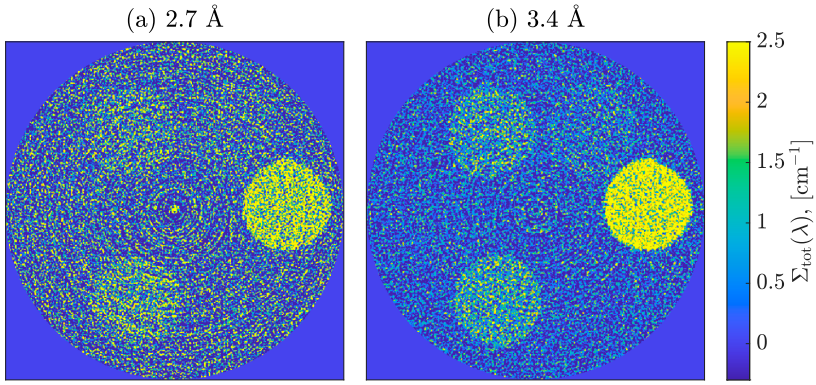

Figure 1. FBP reconstructions at 2.7 Å and 3.4 Å. Ring artifacts with varying severity are present in both reconstructions.

1.2. Contribution

Figure 2 shows the eight measured spectral flat fields stacked vertically and the corresponding singular values. A visual inspection suggests that the spectral flat fields carry significant redundant information and that we can improve the SNR in the spectral flat fields by approximating the spectral flat fields with a low-rank matrix. In particular, the singular values indicate that the spectral flat fields can be well-approximated by a rank-one matrix due to the large jump in magnitude between the first and second singular values. A similar idea where principal component analysis (PCA) is used to reduce ring artifacts in case of beam instability has been proposed by Hagemann et al. [19] and Nieuwenhove et al. [20]. The underlying assumptions in these studies are related, but the nature of the problems solved differs.

Inspired by Figure 2, we propose an extended flat-field model that exploits high correlation across channels in the spectral flat fields to reduce ring artifacts in the reconstructions. The extended model relies on the assumption that the spectral flat fields can be well-approximated by a low-rank matrix. Our method does not depend on a specific

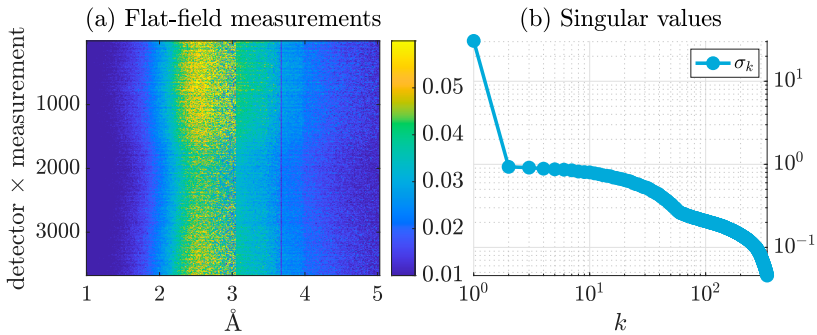

Figure 2. Visualization of the eight spectral flat fields (each 460 detectors) stacked vertically, with the neutron energy range on the first axis and the detector index from 1 to $8 \cdot 460 = 3680$ on the second axis (a) and a loglog plot of the singular values (b). The singular values indicate that the spectral flat fields are well-approximated by a rank-one matrix.

reconstruction method since it works directly on the spectral flat fields. Thus, our method can be combined with existing reconstruction models such as FBP and more advanced spectral methods, e.g. [14, 21–23]. Moreover, our method does not need to be applied channel-wise in the sense that the low-rank spectral flat field simply replaces the measured spectral flat field in the reconstruction step.

1.3. Outline

Section 2 introduces the spectral CT model, existing methods for ring reduction and the proposed methodology. Section 3 describes the experimental set-up for the neutron data set, and our numerical results are reported in Section 4. Section 5 discusses the results, and Section 6 concludes the paper.

1.4. Notation

The set $\mathbb{R}^n$ is the n-dimensional Euclidean space, $\mathbb{R}^n_+$ is the non-negative orthant, and $\mathbb{R}^{m\times n}$ denotes the set of real-valued $m \times n$ matrices. The vector $\mathbf{1}_n \in \mathbb{R}^n$ is a vector of ones, $I_n \in \mathbb{R}^{n\times n}$ denotes the identity matrix, and the transpose of A is denoted A^T. The exponential function $\exp(\cdot)$ applied to a vector or a matrix is interpreted element-wise. Given a vector x, $\operatorname{diag}(x)$ denotes the diagonal matrix with the elements of x on its diagonal. Lastly, the 2-norm of a vector x is denoted $\|x\|_2$.

2. Methods

Consider a spectral data set with m energy channels, and let E_k denote the energy associated with the kth energy channel. The incident intensity of a beam with energy E_k on a detector element is prescribed by the Beer–Lambert law [24],

$$I(E_k) = I_0(E_k) \exp\left(-\int_\ell \mu(\mathbf{x}, E_k)\, \mathrm{d}\mathbf{x}\right), \tag{1}$$

where $I(E_k)$ and $I_0(E_k)$ are the energy-dependent intensity incident on the detector element and on the object, respectively. Furthermore, ℓ is the line segment between the source and the detector, and $\mu : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}_+$ is the energy-dependent spatial attenuation function.

Let $Y_k \in \mathbb{R}^{rp}$ denote the measurements at energy E_k obtained from p projection images and a detector with r detector elements. We discretize the reconstruction domain into n pixels such that (see, e.g. [25])

$$Y_k = \operatorname{diag}\left(\mathbf{1}_p \otimes Z_k\right) \exp\left(-AX_k\right) + \Xi, \tag{2}$$

where $\otimes$ denotes the Kronecker product, $Z_k \in \mathbb{R}^r$ is the flat field and $X_k \in \mathbb{R}^n$ are the unknown attenuation coefficients associated with the kth channel, $A \in \mathbb{R}^{rp\times n}$ is the system matrix, and $\Xi \in \mathbb{R}^{rp}$ represents noise. In practice, the flat field Z_k is estimated by measuring the detector response. We assume that s flat-field measurements are available for each channel. An estimate $\hat{Z} \in \mathbb{R}^{r\times m}$ of the spectral flat field $Z = [\, Z_1 \cdots Z_m \,]$ is then given by

the sample mean of the s flat-field measurements $F_1, \ldots, F_s \in \mathbb{R}^{r\times m}$, i.e.

$$\hat{Z} = \frac{1}{s}\sum_{j=1}^{s} F_j = \frac{1}{s}\left(\mathbf{1}_s^T \otimes I_{r\times r}\right) F, \tag{3}$$

where $F = [F_1^T, F_2^T, \ldots, F_s^T]^T$. The flat-field estimate $\hat{Z}$ generally contains noise, which can give rise to ring artifacts in reconstructions.

2.1. Low-rank approximation

Each of the spectral flat-field measurements carries information about the detector response and the spectrum of the incident beam. Our aim is to exploit the high correlation in the spectral dimension, which is motivated by the observations in Figure 2. The figure suggests that F is well-approximated by a low-rank matrix. The best rank-l approximation of F (in the spectral norm) can be computed by means of a singular value decomposition (SVD) [26] of F, i.e.

$$F^l = \sum_{i=1}^{l} \sigma_i U_i V_i^T, \tag{4}$$

where $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_l \geq 0$ are the l largest singular values of F, and U_i and V_i are left and right singular vectors associated with σ_i. It follows from the Eckart–Young–Mirsky theorem that the relative approximation error is given by

$$\frac{\|F^l - F\|_2}{\|F\|_2} = \frac{\sigma_{l+1}}{\sigma_1}. \tag{5}$$

Figure 3 shows such approximations of rank 1 and rank 5 for the neutron data. The relative approximation error for the rank-1 and rank-5 approximations are 0.030 and 0.028, respectively, i.e. we only observe a minor reduction in the relative approximation error when including four additional singular vectors. As shown in Figure 3, the rank-1 approximation yields a substantial reduction of the noise, and the same is true for the rank-5 approximation. However, the rank-5 approximation appears to be noisier than the rank-1 approximation as can be seen in Figure 3(e,f), which show the difference between the spectral flat fields and the low-rank approximations. Thus, in our numerical experiments, we will only consider rank-1 approximations of the spectral flat fields, i.e. we compute the estimate for $\hat{Z}$ by replacing F by F^1 in (3), i.e.

$$\hat{Z} = \frac{1}{s}\left(\mathbf{1}_s^T \otimes I_{r\times r}\right) F^1. \tag{6}$$

The cost grows as $\mathcal{O}(\min(rs, m)^2 \max(rs, m))$ if the Golub–Reinsch algorithm is used to compute a thin SVD. We note that by employing a randomized SVD [27], this can be reduced to $\mathcal{O}(rsm)$, which is linear in size of the flat-field data, and hence it is fast. Moreover, it can be combined with any reconstruction method, e.g. FBP, iterative methods, statistical models, etc. since it is applied to the spectral flat field measurements as a form of preprocessing.

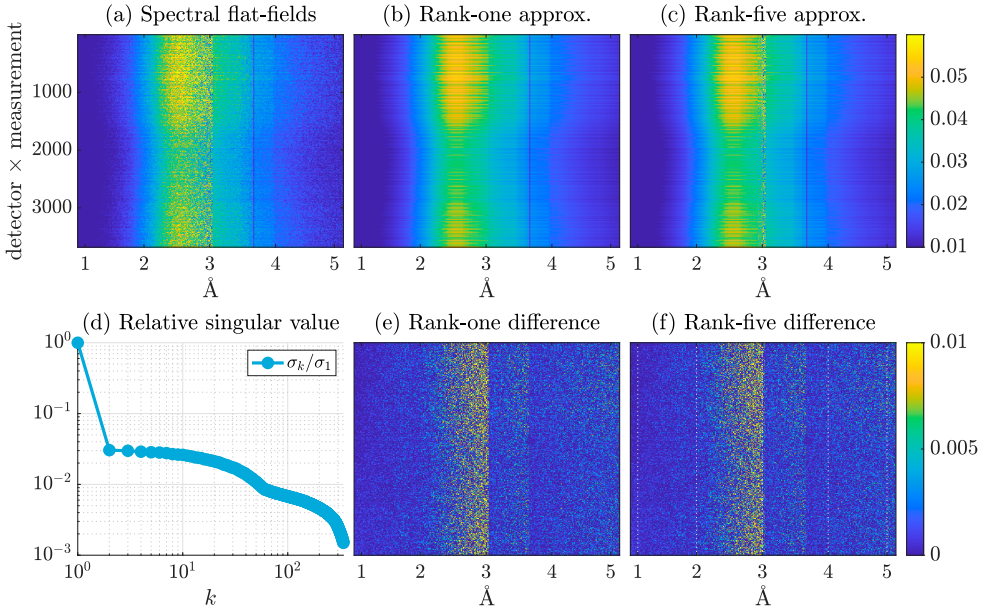

Figure 3. Rank-one and five approximations of the spectral flat field are shown in (b) and (c). The corresponding difference images between the spectral flat fields and low-rank matrices are shown in (e) and (f) for rank-1 and rank-5, respectively. In all plots except (d), the first axis is the neutron energy and the second axis is the detector index over the eight stacked flat fields.

2.2. Existing ring reduction method

Our method works solely on the spectral flat fields, and hence it separates the ring reduction from the reconstruction step. We compare our method to two existing ring reduction techniques for monochromatic CT. The first method is the preprocessing method proposed by Münch et al. [17], which combines wavelet and Fourier filtering to mitigate ring artifacts in the reconstruction. The method is computationally inexpensive and does not increase the overall computational cost significantly. The method depends on three parameters; we use a damping factor of 0.9 and the Daubechies 5 wavelet with a three-level decomposition for all numerical experiments; see [17] for details. The second preprocessing method is proposed by Vo et al. [18] and uses a combination of sorting and smoothing (non-local means) to mitigate stripes in the sinogram and thereby reducing ring artifacts. The method depends on a parameter related to the smoothing filter; see [18] for further details. We used the value 31 in our experiments. We denote the methods of Münch et al. [17] and Vo et al. [18] by WF (Wavelet Fourier) and NLM (non-local means), respectively.

3. Neutron data

We validate the proposed methodology on a neutron CT data set [13]. The neutron data were acquired at the Imaging and Materials Science and Engineering (IMAT) beamline operating at the ISIS neutron spallation source (Rutherford Appleton Laboratory, UK).

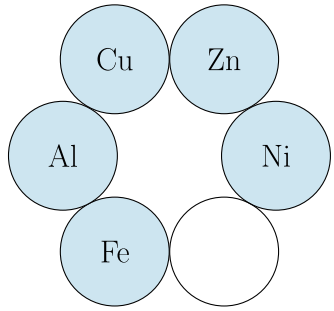

Figure 4. Sketch of the object for the experimental neutron data [13]. The experimental set-up consists of six aluminum cylinders where five of the cylinders have been filled with aluminum (Al), iron (Fe), copper (Cu), nickel (Ni) and zinc (Zn) powders. The sixth cylinder is empty.

Table 1. Experimental set-up for the spectral neutron data used for the numerical experiments.

Pixels (n)	Energies (m)	Detectors (r)	Projections (p)	Spectral flat fields (s)
460^2	339	460	120	8

Figure 4 shows a sketch of the object of interest. The object consists of six aluminum cylinders whereof five are filled with high-purity metal powders, i.e. aluminum (Al), iron (Fe), copper (Cu), nickel (Ni) and zinc (Zn) powders.

3.1. Data acquisition

Ametova et al. [23] describe the data acquisition and pre-treatment of data in detail, i.e. beam instabilities correction, overlap correction and spectral averaging. We consider the reconstruction-ready measurements and confine ourselves to reporting only the essential details of the experimental set-up. The data set contains $m = 339$ spectral projections acquired at $p = 120$ equidistant angles distributed from 0° to 180° with 1.5° angular increments. Each spectral projection arises from binning all neutrons with energies in a narrow energy range around a central energy level, ranging from 1.0576 Å to 5.0321 Å, with an energy bin width of 0.0115 Å for all but a few bins, as explained in [23]. Eight spectral flat fields were acquired, four prior to the scan and four after, i.e. $s = 8$. Each projection consists of 460×460 pixels with a pixel size of 0.055 mm resulting in a view of approximately $25 \times 25\,\text{mm}^2$ and $r = 460$ detector elements. We use a square reconstruction grid with $n = 460^2$ pixels, and we use the 127th vertical detector row for our experiments, i.e. we consider a two-dimensional set-up. The experimental set-up is summarized in Table 1. The attenuation coefficient for the neutron experiment is denoted $\Sigma_{\text{tot}}(\lambda)$ and has unit cm^{-1}, see [23] for further details.

4. Numerical experiments

We consider two reconstruction models: FBP and a weighted least squares (WLS) reconstruction model combined with TV regularization. We compare our method to the conventional flat-field correction and the existing ring reduction techniques described in

Table 2. Abbreviations used for the reconstruction models and ring reduction techniques.

	Reconstruction Model	
Ring reduction technique	FBP	WLS with TV
Conventional	FBP	TV
Preprocessing Münch et al. [17]	WF-FBP	WF-TV
Preprocessing Vo et al. [18]	NLM-FBP	NLM-TV
Low-rank spectral flat fields	LR-FBP	LR-TV

Section 2.2. Table 2 provides an overview of the reconstruction models and ring reduction techniques.

We use AIR TOOLS II [28] to generate the parallel-beam geometry of the experimental set-up. For the FBP reconstructions, we use the `fbp` function from AIR TOOLS II with the `Hann` filter to reduce the noise in the computed FBP reconstructions. For the TV reconstructions, we used the implementation of WLS with TV from [10] in MATLAB. All TV reconstructions have a maximum number of iterations of 1000 and a regularization parameter of 0.005. The regularization parameter was found by visual inspection of reconstructions for varying values of the regularization parameter. Note that we pick the same regularization parameter for all energies which will result in some reconstructions being a bit over-regularized whereas other reconstructions might be slightly under-regularized. The reason is that the SNR changes significantly as a function of energy and thus there is not a single regularization parameter that fits all energy channels. However, the purpose of the methods is ring reduction and thus we limit the experiments to considering the same regularization parameter for all reconstructions.

4.1. Error measures

The quality of the computed reconstructions is assessed by visual inspection combined with contrast-to-noise ratio (CNR). The CNR metric is used for evaluating the image contrast and noise properties for a selected region of interest (ROI) [3]. We use the method proposed by Bian et al. [29] where an ROI with a low-contrast structure is compared to a background ROI while taking the standard deviations of both the signal and background ROIs into account. The ROIs are clearly marked on the figures when applicable.

4.2. Experiment: different energies

There are too many energies to visualize all reconstructions, and the SNR varies significantly between the energies [23]. Thus, we select three energies based on the relative difference (RD) between the computed FBP and LR-FBP solution to ensure a representative visualization of energy channels. We define the RD as

$$\mathrm{RD}(k) = \frac{\|X_k^{\mathrm{FBP}} - X_k^{\mathrm{LR\text{-}FBP}}\|_2}{\|X_k^{\mathrm{LR\text{-}FBP}}\|_2}, \tag{7}$$

where X_k^{FBP} and $X_k^{\mathrm{LR\text{-}FBP}}$ denote the FBP and LR-FBP reconstructions for energy k, respectively. We select three energies corresponding to the minimum, median and maximum RD (3.1, 4.2 and 2.9 Å, respectively).

The FBP reconstructions in Figure 5 (first column) are all affected by ring artifacts. However, the severity of the ring artifacts increases from top to bottom, i.e. when the RD increases. The WF-FBP reconstructions (second column) seem to mitigate the ring artifacts at first glance. However, when carefully comparing the WF-FBP with NLM-FBP and LR-FBP, wave-like artifacts can be seen in the WF-FBP reconstructions. Thus, the ring artifacts have been reduced but not eliminated. No visible ring artifacts remain in the NLM-FBP and LR-FBP reconstructions (third and fourth columns). Considering the CNR listed in Table 3, we see that the CNR decreases from top to bottom for all four methods. Thus, when the SNR is low, our method differs most from the FBP, which can be explained by the fact that the ring artifacts are more dominating when the SNR is low. The CNR for all four methods is quite close in Figure 5. A possible explanation is that the dominating noise contribution comes from the measurement noise and not the ring artifacts since all FBP reconstructions suffer from a very low SNR. Note that the high noise level may conceal remaining ring artifacts.

Inspecting the TV reconstructions in Figure 6, we see that TV regularization reduces the noise level significantly, and consequently, the ring artifacts appear more severe for the TV reconstructions. In addition, we also note that the vague wave structures in the WF-FBP reconstructions are even clearer in the WF-TV reconstructions compared to the WF-FBP reconstructions. By carefully inspecting the NLM-TV reconstruction for 2.9 Å, one can see that the preprocessing has introduced a dark spot with negative values in the center. The other NLM-TV reconstructions show no introduced artifacts and closely resemble

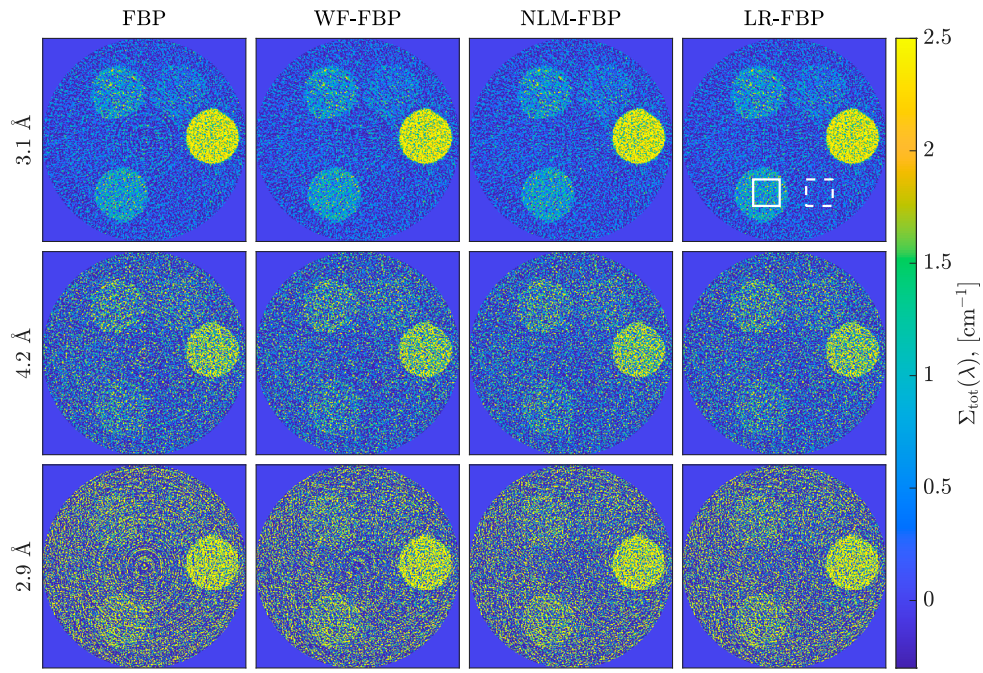

Figure 5. FBP (first column), WF-FBP (second column), NLM-FBP (third column) and LR-FBP (fourth column) reconstructions for three energies chosen by the RD measure. The white squares with full and dashed lines mark the structure and background ROIs for CNR, respectively.

Table 3. CNR for experiment depicted in Figures 5 and 6.

	FBP				TV			
Energy	Conv.	WF	NLM	LR	Conv.	WF	NLM	LR
3.1 Å	2.15	2.31	2.17	2.31	$1.02 \cdot 10^3$	$1.03 \cdot 10^3$	$1.01 \cdot 10^3$	$1.06 \cdot 10^3$
4.2 Å	0.28	0.30	0.29	0.30	$0.16 \cdot 10^2$	$0.16 \cdot 10^3$	$0.21 \cdot 10^3$	$0.28 \cdot 10^3$
2.9 Å	0.20	0.24	0.24	0.25	$0.03 \cdot 10^1$	$0.04 \cdot 10^3$	$0.07 \cdot 10^3$	$0.07 \cdot 10^3$

Note: The ROIs are marked in Figures 5 and 6.

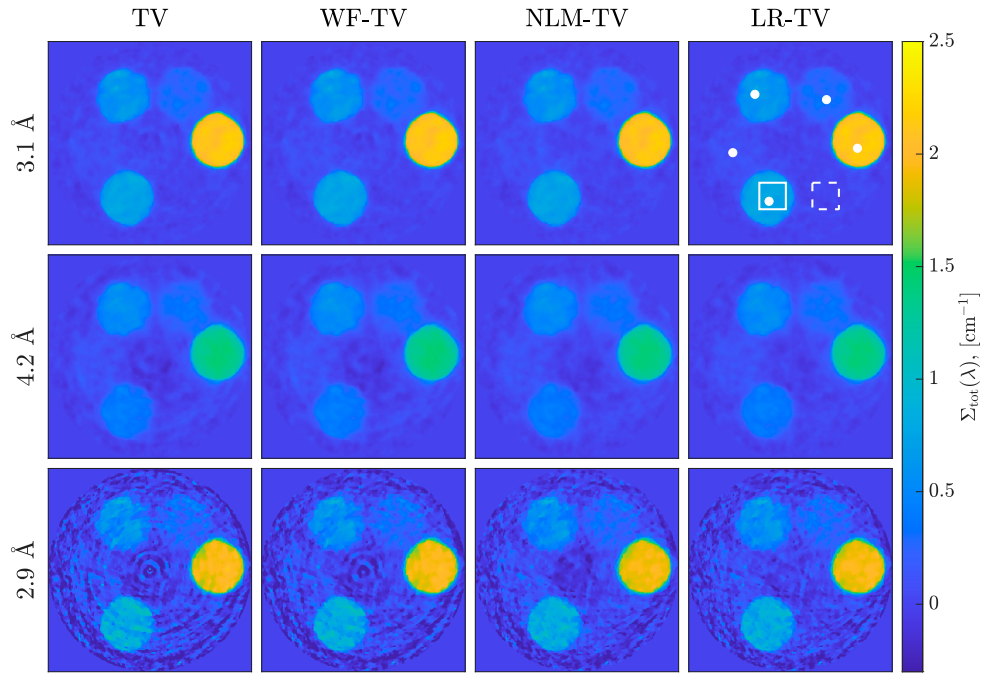

Figure 6. TV (first column), WF-TV (second column), NLM-TV (third column) and LR-TV (fourth column) reconstructions for three energies chosen by the RD measure. The white squares with full and dashed lines mark the structure and background ROIs for CNR, respectively. The white circles represent the pixels chosen for the spectral plot in Figure 7.

the LR-TV reconstructions (fourth column). The LR-TV reconstructions reveal no ring structure even though the noise level is reduced. The TV reconstructions yield significant improvements in the CNR compared to the FBP reconstructions, which can be seen in Table 3. For the TV experiment in Figure 6, we also see an increase in CNR when applying a ring reduction method, especially for NLM-TV and LR-TV. The LR-TV reconstruction achieves the highest CNR for two out of the three energies depicted in Figure 6.

Figure 7 illustrates the spectral dimension of the reconstructions by plotting spectral profiles for each material. The pixels chosen for the spectral profiles are marked by full circles on the upper right part in Figure 6. We generated spectral plots for both the FBP and TV experiments, but we only include the spectral plot for the TV reconstruction since the spectral profiles obtained with FBP were very noisy. Figure 7 shows minor improvement in the spectral domain by using one of the preprocessing methods. This minor improvement is quantified by computing the mean square error (MSE) between the estimated and

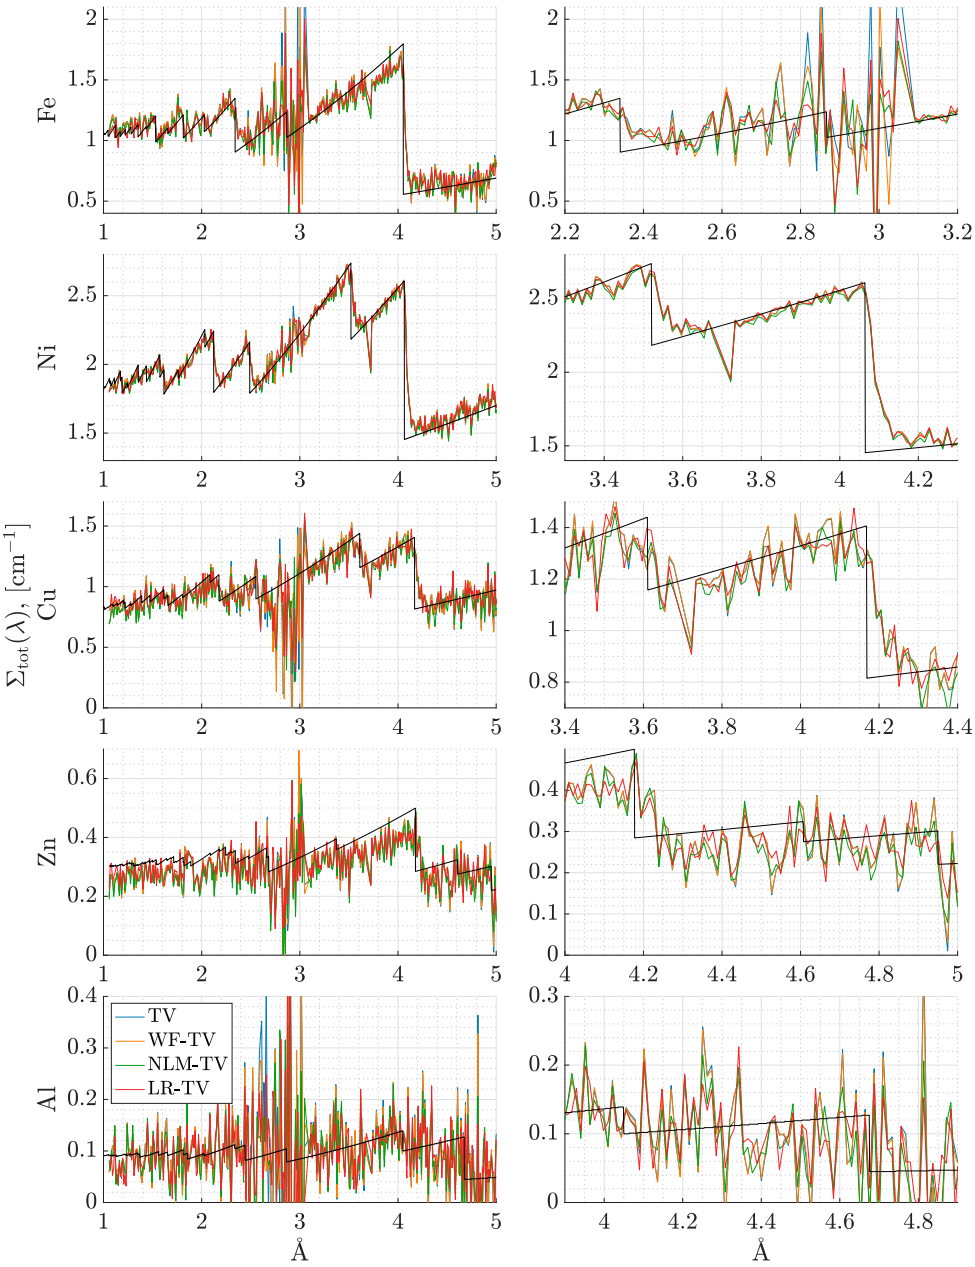

Figure 7. Spectral profiles for the TV reconstructions across all energies. The pixels chosen for the spectral profiles are marked with circles on the upper right subfigure in Figure 6.

theoretical spectral profile. The MSE measures are listed in Table 4 and show a decrease in MSE for the NLM and LR methods for all materials except nickel (Ni). It is notable that the significant improvement in the spatial domain has only a small effect in the spectral domain. Thus, for improved material decomposition, we need reconstruction methods which include spectral regularization.

Table 4. Mean square error (MSE) computed for each of the spectral plots in Figure 7.

	MSE			
Material	TV	WF-TV	NLM-TV	LR-TV
Fe	5.68	5.48	5.24	5.33
Ni	3.04	3.04	3.00	3.04
Cu	4.50	4.38	3.95	3.67
Zn	5.00	5.01	5.07	4.32
Al	22.6	18.2	15.3	16.5

4.3. Experiment: effect of number of spectral flat fields

We now perform an experiment with a varying number of spectral flat fields to examine the robustness of the proposed method. We consider the energy 2.1 Å for the experiment and use eight, four, two and one flat-field measurements for the flat-field correction, respectively. The FBP, WF-FBP, NLM-FBP and LR-FBP reconstructions are shown in Figure 8, and the TV reconstructions for the same experiment are shown in Figure 9. The CNR for the reconstructions are listed in Table 5.

In the first column of Figure 8, we see that there are visible ring artifacts in all FBP reconstructions, and the severity of the ring artifacts increases when the number of flat fields decreases (i.e. from top to bottom). The WF-FBP reconstructions are shown in the second column of the figure. Vague ring artifacts can be seen in the WF-FBP reconstruction using all eight flat fields. Like in the FBP reconstructions, the severity of the ring artifacts increases significantly when the number of flat fields decreases. The NLM-FBP reconstructions in the third column show less wave-like artifacts than the WF-FBP reconstructions. However, with four or fewer flat fields, ring artifacts start to become visible in the NLM-FBP reconstructions. The LR reconstructions in the fourth column show no sign of ring artifacts, not even in the case where a single flat field is used. This indicates that the low-rank approximation is quite robust, which is also supported by the CNR results reported in Table 5. Indeed, the CNR is nearly constant for the four reconstructions based on the proposed LR flat field, but it increases with the number of flat fields for both FBP, WF-FBP and NLM-FBP. The experiment was repeated using TV regularization in the reconstruction model, and the findings are consistent with those for FBP. Note that the use of TV regularization yields a significant increase in the SNR, which is easily seen by comparing the reconstructions in Figure 8 with those in Figure 9. Moreover, the TV-regularized reconstructions expose some remaining ring artifacts that were not visible in the FBP reconstructions due to noise.

5. Discussion

Our method relies on the assumption that the true flat field can be expressed as (or approximated by) a separable function. The robustness of the method is related to the fact that the detector elements provide redundant information about the spectrum. The data set used in our experiments includes eight flat fields with 339 energy channels, and our method

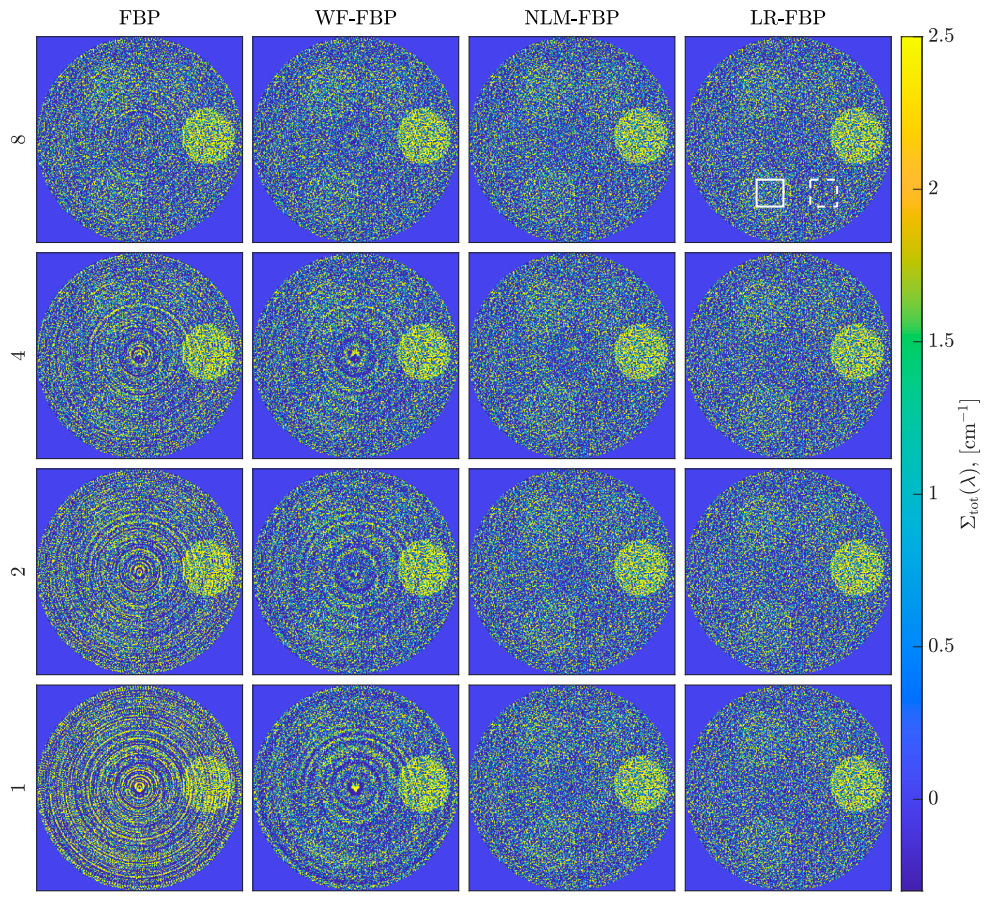

Figure 8. FBP, WF-FBP, NLM-FBP and LR-FBP reconstructions of energy 2.1 Å with eight, four, two and one flat field, respectively. The white squares with full and dashed lines mark the structure and background ROIs, respectively.

utilizes all of these channels to compute a rank-1 approximation. In contrast, the conventional method computes the sample mean, which ignores the spectral dimension. In the monochromatic case, our method and the conventional method are the same.

We emphasize that our proposed method is applied as a preprocessing operation prior to reconstruction and as such can be combined with any reconstruction method. We chose here to demonstrate it in combination with standard FBP, which does not assume a particular object composition, and TV regularization which is expected to perform well for an object with piecewise constant attenuation as considered here. Our results demonstrate clear improvements both visually and quantitatively with both reconstruction methods and we would expect comparable results with other reconstruction methods and for other types of objects.

As a simple extension of our method, the spectral flat-field measurements can be treated as a three-dimensional tensor of dimension $s \times r \times m$. Tensor decomposition methods such as the Tucker decomposition and parallel factors decomposition (PARAFAC) [30, 31] can then be used to compute a low-rank approximation. However, some preliminary

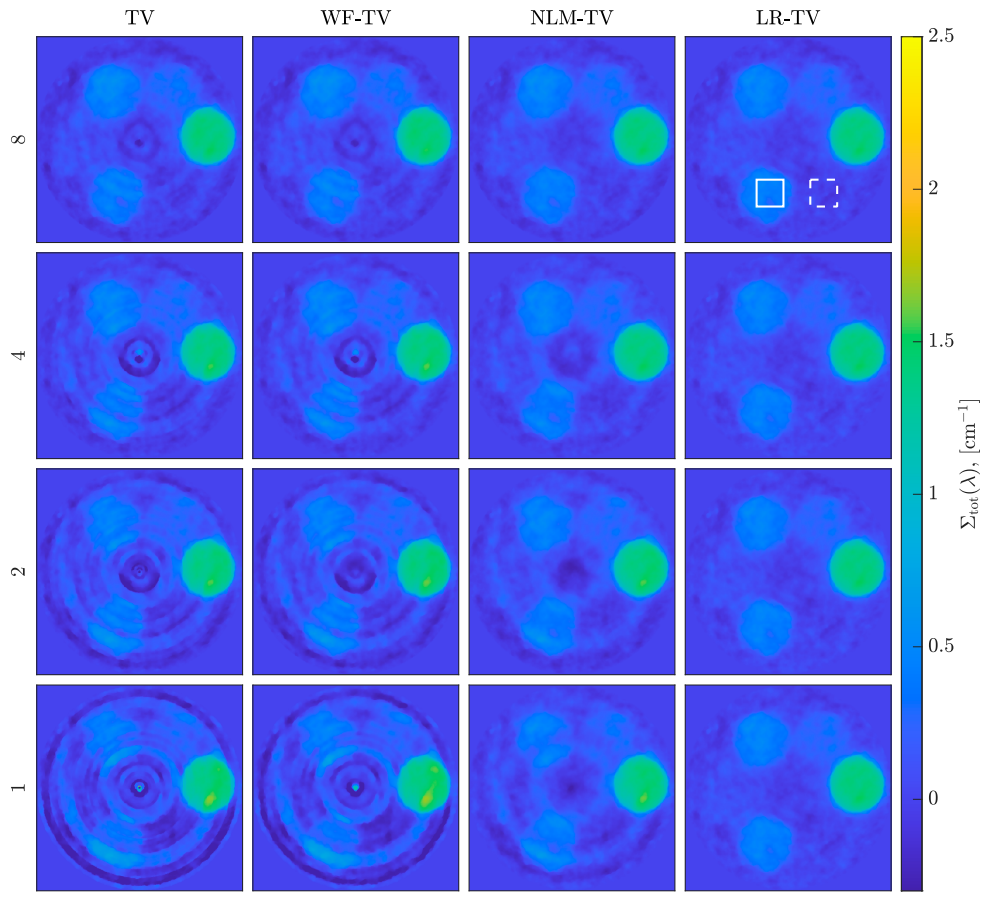

Figure 9. TV, WF-TV, NLM-TV and LR-TV reconstructions of energy 2.1 Å with eight, four, two and one flat field, respectively. The white squares with full and dashed lines mark the structure and background ROIs, respectively.

Table 5. CNR for experiment depicted in Figures 8 and 9 for energy 2.1 Å.

	FBP				TV			
	Conv.	WF	NLM	LR	Conv.	WF	NLM	LR
8	0.15	0.17	0.16	0.16	0.95×10^2	0.95×10^2	1.62×10^2	1.70×10^2
4	0.13	0.15	0.14	0.16	0.55×10^2	0.56×10^2	1.39×10^2	1.79×10^2
2	0.11	0.14	0.13	0.16	0.36×10^2	0.39×10^2	0.87×10^2	1.78×10^2
1	0.06	0.12	0.14	0.16	0.14×10^2	0.15×10^2	0.52×10^2	1.80×10^2

experiments with the tensor approach showed no significant difference between a low-rank tensor approximation and the low-rank approximation proposed in this paper.

Another extension of this work would be to jointly estimate a reconstruction and the spectral flat field, e.g. using an extension of the model proposed in [25]. The estimate obtained by the LR method would then be a natural initial guess for the spectral flat-field matrix.

6. Conclusion

We have proposed an extended flat-field model for spectral CT that exploits high correlation in the spectral flat fields. Our approach is to use a low-rank approximation of the spectral flat-field measurements to obtain a less noisy spectral flat-field estimate, thereby mitigating ring artifacts in the subsequent reconstruction. The proposed methodology can be combined with existing reconstruction methods, and it only depends on a single parameter (the approximation rank), which is easy to choose by inspection of the singular values of the spectral flat-field measurements. We have demonstrated the usefulness of the method based on a neutron CT data set and comparisons with conventional flat-field correction and two existing preprocessing methods for ring reduction. Our method successfully mitigated ring artifacts in all experiments, whereas the other methods struggled to suppress ring artifacts, especially in more challenging cases with severe ring artifacts.

Acknowledgments

We gratefully acknowledge beamtime RB1820541 (DOI: 10.5286/ISIS.E.100529645) at the IMAT Beamline of the ISIS Neutron and Muon Source, Harwell, UK.

Data availability

The code and data to compute the results and produce the figures can be obtained from [32] and [13], respectively.

Disclosure statement

The authors of this manuscript have no conflicts of interest to disclose.

Funding

This work was supported by The Villum Foundation (grant no. 25893). EA was partially funded by the EPSRC grant EP/V007742/1 'Rich Nonlinear Tomography for Advanced Materials' and partially by the Federal Ministry of Education and Research (BMBF) and the Baden-Württemberg Ministry of Science as part of the Excellence Strategy of the German Federal and State Governments.

ORCID

Katrine Ottesen Bangsgaard http://orcid.org/0000-0001-7491-0064
Genoveva Burca http://orcid.org/0000-0001-6867-9628
Evelina Ametova http://orcid.org/0000-0002-8867-3001
Martin Skovgaard Andersen http://orcid.org/0000-0002-4654-3946
Jakob Sauer Jørgensen http://orcid.org/0000-0001-9114-754X

References

[1] Heismann BJ, Schmidt BT, Flohr TG. Spectral computed tomography. Bellingham (WA): SPIE; 2012.
[2] McCollough CH. Methods for spectral CT imaging. Cham: Springer; 2020. p. 223–242. ISBN 978-3-030-26957-9.
[3] Warr R, Ametova E, Cernik RJ, et al. Enhanced hyperspectral tomography for bioimaging by spatiospectral reconstruction. Sci Rep. 2021;11:20818.

[4] Kazantsev D, Jørgensen JS, Andersen MS, et al. Joint image reconstruction method with correlative multi-channel prior for x-ray spectral computed tomography. Inverse Probl. 2018;34:20818.
[5] Schirra CO, Roessl E, Koehler T, et al. Statistical reconstruction of material decomposed data in spectral CT. IEEE Trans Med Imaging. 2013;32:1249–1257.
[6] Barber RF, Sidky EY, Schmidt TG, et al. An algorithm for constrained one-step inversion of spectral CT data. Phys Med Biol. 2016;61:3784–3818
[7] Whiting BR, Massoumzadeh P, Earl OA, et al. Properties of preprocessed sinogram data in x-ray computed tomography. Med Phys. 2006;33:3290–3303.
[8] Sijbers J, Postnov A. Reduction of ring artefacts in high resolution micro-CT reconstructions. Phys Med Biol. 2004;49:N247–N253.
[9] Goldman LW. Principles of CT and CT technology. J Nucl Med Technol. 2007;35:115–128.
[10] Aggrawal HO, Andersen MS, Rose SD, et al. A convex reconstruction model for x-ray tomographic imaging with uncertain flat-fields. IEEE Trans Comput Imaging. 2018;4:17–31.
[11] Di Trapani V, Brun F. Pre- and post-reconstruction digital image processing solutions for computed tomography with spectral photon counting detectors. Nucl Instrum Methods Phys Res Sec A Accel Spectrom Detect Assoc Equip. 2021;1010:165510.
[12] Zeng D, Huang J, Zhang H, et al. Spectral CT image restoration via an average image-induced nonlocal means filter. IEEE Trans Biomed Eng. 2016;63:1044–1057.
[13] Jorgensen J, Ametova E, Pasca E, et al. Neutron TOF imaging phantom data to quantify hyperspectral reconstruction algorithms. 2019. DOI:10.5286/ISIS.E.RB1820541
[14] Wu D, Zhang L, Xu X, et al. Spectral CT reconstruction with weighted non-local total-variation minimization. In: 2015 IEEE Nuclear Science Symposium and Medical Imaging Conference. NSS/MIC; 2015. p. 7582050.
[15] Lv X, Ren X, He P, et al. Image denoising and ring artifacts removal for spectral CT via deep neural network. IEEE Access. 2020;8:225594–225601.
[16] Fang W, Li L, Chen Z. Removing ring artefacts for photon-counting detectors using neural networks in different domains. IEEE Access. 2020;8:42447–42457.
[17] Münch B, Trtik P, Marone F, et al. Stripe and ring artifact removal with combined wavelet—Fourier filtering. Opt Express. 2009;17:8567.
[18] Vo NT, Atwood RC, Drakopoulos M. Superior techniques for eliminating ring artifacts in x-ray micro-tomography. Opt Express. 2018;26:28396.
[19] Hagemann J, Vassholz M, Hoeppe H, et al. Single-pulse phase-contrast imaging at free-electron lasers in the hard x-ray regime. J Synchrotron Radiat. 2021;28:52–63. DOI:10.1107/s160057752001557x
[20] Nieuwenhove VV, Beenhouwer JD, Carlo FD, et al. Dynamic intensity normalization using eigen flat fields in x-ray imaging. Opt Express. 2015;23:27975.
[21] Zhang Y, Xi Y, Yang Q, et al. Spectral CT reconstruction with image sparsity and spectral mean. IEEE Trans Comput Imaging. 2016;2:510–523.
[22] Gao H, Yu H, Osher S, et al. Multi-energy CT based on a prior rank, intensity and sparsity model (prism). Inverse Probl. 2011;27:115012.
[23] Ametova E, Burca G, Chilingaryan S, et al. Crystalline phase discriminating neutron tomography using advanced reconstruction methods. J Phys D Appl Phys. 2021;54:325502.
[24] Herman GT. Fundamentals of computerized tomography: image reconstruction from projections. London: Springer; 2009.
[25] Bangsgaard KO, Andersen MS. A statistical reconstruction model for absorption CT with source uncertainty. Inverse Probl. 2021;37:085009.
[26] Loan CF. The ubiquitous kronecker product. J Comput Appl Math. 2000;123:85–100.
[27] Halko N, Martinsson PG, Tropp JA. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. SIAM Rev. 2011;53:217–288. DOI:10.1137/090771806
[28] Hansen PC, Jørgensen JS. AIR tools II: algebraic iterative reconstruction methods, improved implementation. Numer Algorithms. 2017;79:107–137.

[29] Bian J, Siewerdsen JH, Han X, et al. Evaluation of sparse-view reconstruction from flat-panel-detector cone-beam CT. Phys Med Biol. 2010;55:6575–6599.
[30] Filipović M, Jukić A. Tucker factorization with missing data with application to low-n-rank tensor completion. Multidimens Syst Signal Process. 2015;26:677–692.
[31] Harshman RA, Lundy ME. PARAFAC: parallel factor analysis. Comput Stat Data Anal. 1994;18:39–72.
[32] Bangsgaard KO, Burca G, Ametova E, et al. Code for 'Low-rank flat-field correction for artifact reduction in spectral computed tomography'; 2023. DOI:10.5281/zenodo.7569942

APPENDIX C

Bayesian parameter estimation and uncertainty quantification for hemodialysis

Published in Mathematical Biosciences and Engineering, Volume 20, Issue 3: 4455-4492, DOI: 10.3934/mbe.2023207

Katrine O. Bangsgaard, Morten Andersen, James G. Heaf and Johnny T. Ottesen.

Mathematical Biosciences and Engineering

http://www.aimspress.com/journal/mbe

MBE, 20(3): 4455–4492.
DOI: 10.3934/mbe.2023207
Received: 08 September 2022
Revised: 13 December 2022
Accepted: 19 December 2022
Published: 26 December 2022

Research article

Bayesian parameter estimation for phosphate dynamics during hemodialysis

Katrine O. Bangsgaard [1], **Morten Andersen** [2], **James G. Heaf** [3] **and Johnny T. Ottesen** [2,*]

[1] Department of Applied Mathematics and Computer Science, Technical University of Denmark, Richard Petersens Plads, Building 324, 2800 Kongens Lyngby, Denmark

[2] Centre for Mathematical Modeling – Human Health and Disease, Department of Science and Mathematics, Roskilde University, 4000 Roskilde, Denmark

[3] Department of Medicine, Zealand University Hospital, 4000 Roskilde, Denmark

* **Correspondence:** Email: johnny@ruc.dk.

Abstract: Hyperphosphatemia in patients with renal failure is associated with increased vascular calcification and mortality. Hemodialysis is a conventional treatment for patients with hyperphosphatemia. Phosphate kinetics during hemodialysis may be described by a diffusion process and modeled by ordinary differential equations. We propose a Bayesian model approach for estimating patient-specific parameters for phosphate kinetics during hemodialysis. The Bayesian approach allows us to both analyze the full parameter space using uncertainty quantification and to compare two types of hemodialysis treatments, the conventional single-pass and the novel multiple-pass treatment. We validate and test our models on synthetic and real data. The results show limited identifiability of the model parameters when only single-pass data are available, and that the Bayesian model greatly reduces the relative standard deviation compared to existing estimates. Moreover, the analysis of the Bayesian models reveal improved estimates with reduced uncertainty when considering consecutive sessions and multiple-pass treatment compared to single-pass treatment.

Keywords: hemodialysis; phosphate kinetics; mathematical modeling; parameter estimation; uncertainty quantification

1. Introduction

Phosphate enables the body to perform vital processes such as construction of nucleic acids, energy transport and bone tissue formation [1]. The level of phosphate is tightly controlled, and excess phosphate is excreted by the kidneys [2]. However, for patients with renal failure, the control of phosphate homeostasis is impaired. An abnormal level of phosphate is associated with increased vascular calcification and mortality [3, 4].

About half of all dialysis patients suffer from hyperphosphataemia, and strategies to control phosphate levels include phosphate binders, low-phosphate diet and removal of phosphate by hemodialysis [5]. Hemodialysis (HD) is a conventional treatment for renal failure where a patient is coupled to a dialysis machine for four to eight hours. The blood plasma and dialysate fluid are passed through a filter that causes a diffusion process that removes toxic substances, e.g., phosphate, from the blood to the dialysate. The phosphate kinetics in HD is of particular interest because it differs the other removed toxins, e.g., urea, by the fact that hypophoshataemia is fatal for the patient [6]. Thus, the phosphate concentration should not be exhausted, but kept within the critical values.

1.1. Previous studies

The control of the phosphate concentration is a considerable clinical problem and has been studied extensively; The conventional hemodialysis treatment is the single-pass (SP) treatment. Agar et al. [7] and Debowska et al. [3] both study the SP treatment by considering a simple two-compartment ordinary differential equation (ODE) model for phosphate removal during HD. They present their results as an average of the measured patients to obtain confidence intervals for their parameters, however, these are not patient specific. Poleszczuk et al. [2] extend the model proposed by Debowska et al. [3] to include a time delay. The time delay is introduced to improve the fit at the later stage of the HD where a minor rebound is observed in some clinical experiments. Andersen et al. [8] analyze the same model analytically and estimate parameters using an optimization-driven approach. Here the parameters are estimated for each patient, but the uncertainty of the parameter estimates is not addressed. Laursen et al. [9, 10] propose a two- and three-compartment model for phosphate clearance during SP and find that the three-compartment model produces the most satisfying fit but does not address the uncertainty associated with the parameter estimates. Spalding et al. [11] propose a complicated four-compartment model where the fourth pool is a control pool for avoiding dangerously low phosphate concentrations. They argue that a simple two-compartment model cannot fit the relapse phase sufficiently. The relapse phase refers to the period after ended treatment where the phosphate concentration starts to increase. However, both Andersen et al. [8] and Debowska et al. [3] demonstrate that the simple two-compartment model can produce adequate fits for the relapse phase as well. A novel HD treatment called multiple pass (MP) [12–14] provides an alternative to the conventional SP. This novel treatment reduces the amount of dialysis fluid needed for a single session of HD. Andersen et al. [8] and Heaf et al. [14] analyse and compare the MP treatment and SP treatment.

However, none of the above-listed models address patient specific uncertainties associated with the parameter estimates. Moreover, the reported uncertainty of the parameter estimates for the average of the measured patients is very large, e.g., Debowska et al. [3] report a phosphate clearance with a relative standard deviation of 79% and Agar et al. [7] report a relative standard deviation of 47%, indicating that parameters of the two-compartment model are poorly identified. Common for all models is that they assume that the phosphate concentration in the inner-source compartment is known exactly through measurements at time zero. However, measurements are noisy and can potentially bias the results.

The Bayesian approach for parameter estimation for ODE modeling has gained attention in later years [15, 16] since it provides an elegant way of addressing the uncertainty associated with the estimated parameters and includes clinical knowledge. The Bayesian approach gives a complete image of the parameter estimation in terms of uncertainty quantification, i.e., posterior mean, credibility intervals and correlations. A Bayesian approach for patient-specific parameters for hemodialysis has been

proposed by Bianchi et al. [17] but does not consider the phosphate kinetics.

1.2. Contribution

We propose a Bayesian approach for estimating patient-specific parameters for phosphate dynamics during hemodialysis. Moreover, we include the phosphate concentration in the inner compartment as a parameter of the model. We use uncertainty quantification to assess the reliability of our parameter estimates and explore the full parameter space. We address the identifiability of the parameters for the SP, MP and the combination of the two, denoted combined-pass (CP). While CP has limited direct clinical impact it improves the parameter estimation slightly, since both SP and MP data are considered simultaneously. In addition, we also investigate how the parameter estimation can be improved by including relapse measurements and / or measure consecutive sessions.

1.3. Outline

Section 2 describes the phosphate kinetics during hemodialysis and introduces the single- and multiple-pass treatments. Section 3 introduces the Bayesian model and describes implementation and sampling diagnostics. In Section 4, we test and validate SP, MP and CP models on data sets and discuss findings from synthetic data which are found in the supplementary materials. Lastly, we conclude the paper in Section 5.

2. Hemodialysis modeling

About 85% of the total phosphate in the human body is stored in the bones [18]. We assume that we have an inexhaustible source (bone) that excretes phosphate to the blood, including extracellular fluid. The phosphate transport from source to blood is driven by diffusion. The diffusion process is governed by the diffusion coefficient (permeability) and concentration gradient. The blood compartment is coupled to the dialysate compartment through a semipermeable membrane which generates a flow of phosphate to the dialysate fluid. The flow of phosphate from blood to dialysate is mainly governed by diffusion and to an insignificant degree by a convection process. [9] However, comprehensive investigations have shown that the convective flow has a negligible effect on the model and parameters during the normal range of dialysis treatment, i.e., up to eight hours [8]. Thus, we exclude the convection term from the models. In this paper, we consider three types of models for HD for phosphate clearance in dialysis patients, the conventional SP, MP and the combination CP.

The value of this analysis for clinicians is twofold. Firstly, accurate modeling permits the prediction of phosphate removal during different forms of dialysis, e.g., short and long dialysis or use of filters with standard or high phosphate clearances. Secondly, it is possible to get insight into the underlying physiological causes of phosphate dynamics.

The model parameters are individually calibrated, so the parameter values are patient specific signatures. Hence, reliable estimates for these parameters are clinically important in order to assess the phosphate kinetics.

2.1. Single-pass dialysis

For the SP treatment, the dialysate is constantly replenished by fresh dialysate such that the phosphate concentration in the outflowing dialysate remains low. Data shows this phosphate concentration to be approximately constant throughout the treatment. SP requires excessive amounts of dialysate for each session. A conceptual diagram of the SP treatment is depicted in Figure 1 that illustrates the removal of phosphate by diffusion.

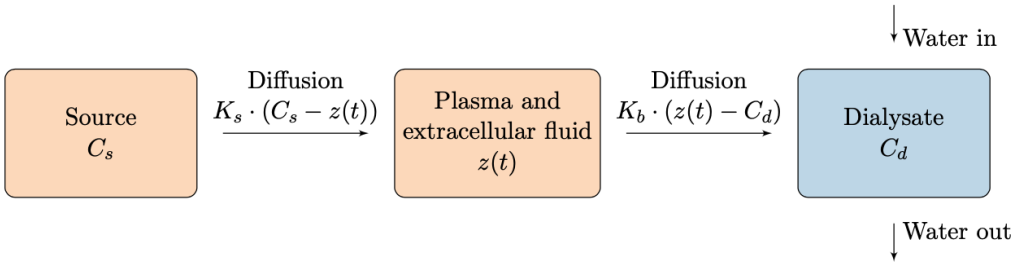

Figure 1. Conceptual diagram for single-pass (SP). In SP, blood and dialysate are passed through a filter which initiates a diffusion process that removes toxic substances from the blood (plasma and extracellular fluid). The outflowing dialysate is constantly replenished by fresh dialysate, and the concentration of phosphate in the dialysate is assumed constant.

Agar et al. [7] proposed a simple compartment model for SP consisting of a single linear autonomous ODE,

$$V_b \frac{dz(t)}{dt} = C_s K_s - (K_s + K_b) z(t) + K_b C_d, \tag{2.1}$$

where $z(t)$ is the concentration of phosphate in the blood compartment at time t, C_s is the constant concentration in the source compartment and C_d is the phosphate concentration in the dialysate assumed to be constant and measurable. K_s and K_b are diffusion rates from source to blood and from blood to dialysate, respectively. Lastly, V_b denotes the blood volume taken as the blood plasma and extracellular volume. For the system to have a unique solution, we equip the ODE with the initial condition $z(0) = z_0$. Notice that the system is not identifiable since V_b can be integrated in the remaining parameters and thus we assume that V_b is known through measurements for SP.

The assumption of a constant C_d is not crucial. If we allow the phosphate concentration to be a variable with initial value 0, then we can extend the model by an extra differential equation. This extension results in a fast transient in C_d toward the steady state value given by data shown in Table 1 with at doubling time of approximately 10-15 minutes (see supplementary, Figure E.14). Moreover, such extension does not affect the parameter estimates achieved. Hence we confine ourselves to consider C_d as a constant.

2.2. Multiple-pass dialysis

Contrary to SP where dialysate is constantly replenished, the dialysate for MP is recirculated, and consequently, the removed substances accumulate in the dialysate fluid over time. A conceptual diagram of the MP treatment is depicted in Figure 2.

MP is less effective than SP due to the accumulation of substances in dialysate. However, MP greatly reduces the amount of dialysate fluid needed for HD treatment, which makes a smaller clinical setting possible. Furthermore, it may ease HD treatment at home and treatment during travels, which can possibly greatly improve the quality of life for renal failure patients. [12–14]

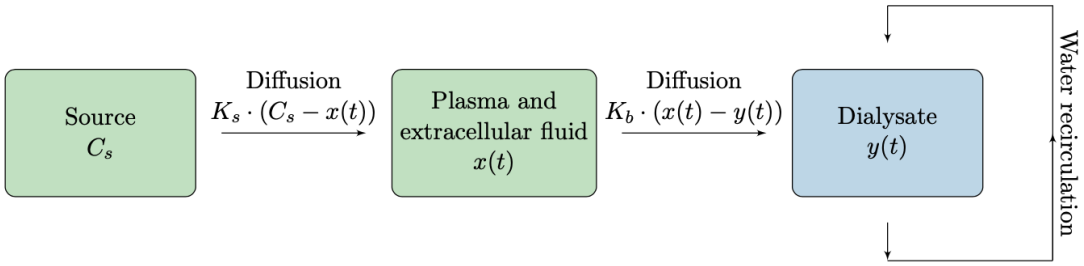

Figure 2. Conceptual diagram for multiple pass (MP) treatment. Like in conventional SP, blood and dialysate is passed through a filter that causes a diffusion process that removes toxic substances from the blood (plasma and extracellular fluid). The dialysate is recirculated and consequently, the removed substances accumulate in the dialysate, i.e., $y(t)$ changes as a function of time.

The MP model can be described by the following system of linear autonomous ODEs,

$$V_b \frac{dx(t)}{dt} = C_s K_s - (K_s + K_b)x(t) + K_b y(t), \tag{2.2a}$$

$$V_d \frac{dy(t)}{dt} = K_b(x(t) - y(t)), \tag{2.2b}$$

where $x(t)$ and $y(t)$ are the time-varying phosphate concentrations for the blood compartment and in the dialysate at time t, respectively, and V_d is the volume of the dialysate. The remaining parameters, i.e., V_b, C_s, K_s and K_b, have the same interpretation as for the SP model in (2.1). The initial conditions are $x(0) = x_0$ and $y(0) = y_0$ corresponding to the phosphate concentration in blood and dialysate at time $t = 0$, respectively. The phosphate concentration in the dialysate at time $t = 0$ is zero, i.e., we assume $y_0 = 0$ henceforth.

The MP model carries additional information compared to the SP model since the only new parameter, the dialysate volume V_d is accurately known from the dialysis equipment. Hence, we have an additional equation in the model but the same number of unknown parameters compared to SP. Thus, given sufficient data, the MP model allows for structural identifiability of the parameters due to the addition of (2.2b) since we cannot simply integrate V_b in the remaining parameters. However, since bioimpedance measurements of V_b are available, we will consider V_b to be known a priori, since it is not practically identifiable.

2.3. Combined-pass dialysis

The parameters for a single patient are shared for the two treatments. Thus, if a patient completes both SP and MP, we can utilize all available information by considering the CP model,

$$V_b \frac{dz(t)}{dt} = C_s K_s - (K_s + K_b)z(t) + K_b C_d, \tag{2.3a}$$

$$V_b \frac{dx(t)}{dt} = C_s K_s - (K_s + K_b)x(t) + K_b y(t), \tag{2.3b}$$

$$V_d \frac{dy(t)}{dt} = K_b(x(t) - y(t)), \tag{2.3c}$$

with $z(0) = z_0$, $x(0) = x_0$ and $y(0) = 0$, and the parameters as described for SP. The CP model, just as the MP model, allows for structural identifiability, and potentially even more precise estimation compared to the MP model due to the addition of the SP model.

2.4. Clinical data

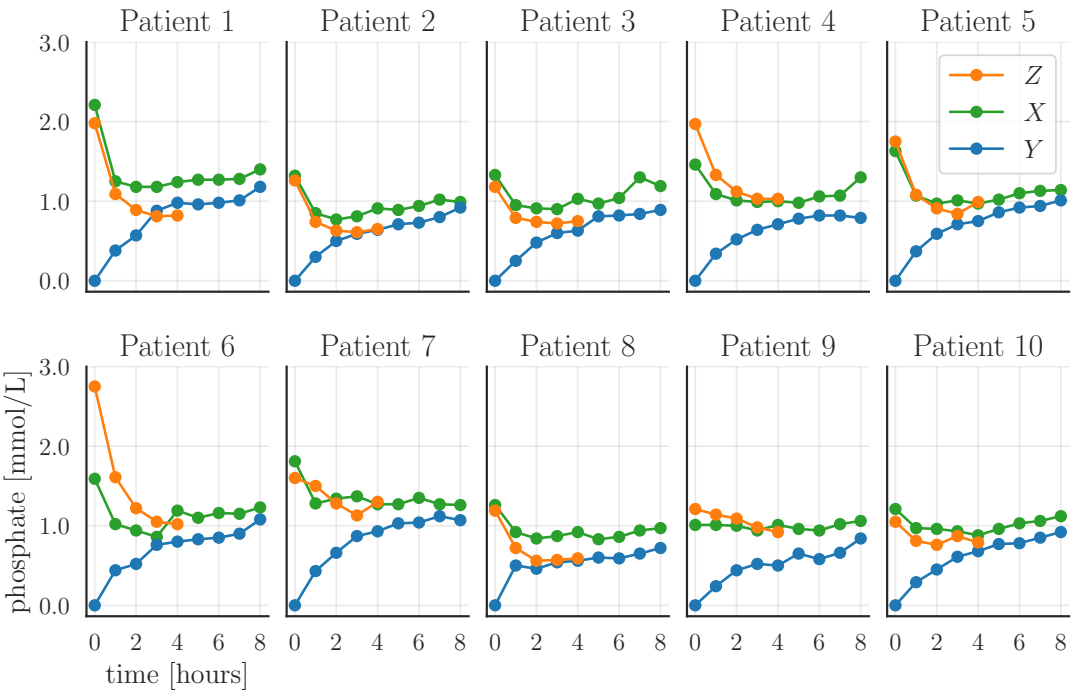

Figure 3. Visualization of the measured phosphate concentrations for SP and MP. The dots represent the measurements, and the full line is the linear interpolation of the measurements. The concentration of phosphate in dialysate in MP is denoted Y and the phosphate concentration in the blood is denoted Z and X for SP and MP, respectively.

We consider longitudinal data sets from 10 patients with renal failure that were measured during an SP session and an MP session. The measured phosphate concentrations for SP and MP (Z, X and Y)

are depicted in Figure 3. Measurements were once every hour for a total of four and eight hours for SP and MP, respectively. No measurements were taken in the relapse phase, i.e., after ended treatment.

Considering the SP measurements (orange dotted line) in Figure 3, we see an exponential-like decay in the measured phosphate concentration after two hours as predicted by (2.1). Thereafter, phosphate concentration seems to stabilize around a reduced concentration level. For the MP measurements (green and blue dotted lines), we see similar exponential-like decay for the phosphate concentration in agreement with the bi-exponential solution to (2.2). However, this drop in phosphate concentration is a bit slower for some patients and after two hours it starts to slowly increase due to the accumulation of phosphate in the dialysate. The phosphate concentration in the dialysate increases rapidly in the beginning of the treatment but slows down and approaches an equilibrium with the phosphate concentration in the blood. This behavior is expected according to the model in (2.2) since the concentration gradient vanishes.

3. Bayesian inference

We solve the parameter estimation problem using a Bayesian approach, where we consider the parameters, measurement noise and initial conditions as random variables. In Bayesian inference, we are interested in the posterior probability of the parameters. The posterior probability consists of two components: a prior probability reflecting our knowledge or beliefs about likely parameter values, and a likelihood function that expresses how likely it is to observe the data for a set of parameters. Thus, the posterior allows us to formally include clinical prior knowledge in the model. Moreover, the inclusion of the prior may have a regularizing effect on the parameter estimation problem in the sense that the parameter estimates become less sensitive to measurement noise.

We use uncertainty quantification to assess the reliability of the parameter estimates and the concentrations in terms of posterior statistics, i.e., mean, correlation and 95% credibility intervals (CI). A strength of the uncertainty quantification is that the solution is based on all probable outcomes instead of being solely based on a point estimate [19]. Uncertainty quantification can also be used for model analysis and improvement, e.g., revealing strong correlation or identifying potential measurements that could improve identifiability of the model [20]. Hence, uncertainty quantification is a flexible method to assess how certain we are of the parameter values and parameter-dependent solutions

3.1. Likelihood and prior modeling

We describe the Bayesian model for the SP and MP treatments and presume data for the relevant state variables (phosphate concentrations) are measured. Notice that the Bayesian formulation is trivially extended to CP by combining the SP and MP models.

3.1.1. Single-pass formulation

First, we consider the Bayesian formulation for SP. Let $\theta = [C_s, K_s, K_b]$ denote the vector of unknown parameters and z_{IC} denote the initial condition for SP. We assume that V_b and C_d are known to a sufficient degree a priori and do not estimate them based on the model. A justification of this assumption is given in Section 4.

The state variable $z(t, \theta, z_{\mathrm{IC}})$ is the solution to (2.1) and we wish to infer the model parameters θ and initial condition z_{IC} defining the state variable. Henceforth, we shorten notation such that $z(t) \equiv$

$z(t, \theta, z_{\text{IC}})$.

We assume that the measurement noise is normally distributed such that the state variable, $z(t)$ is inferred through the Gaussian likelihood function,

$$Z_i \sim \mathcal{N}\left(z(t_i), \sigma_{\text{d}}^2\right), \text{ for } i = 1, 2, ..., m, \tag{3.1}$$

where $Z \in \mathbb{R}^m$ is a vector with the measurement of the phosphate concentration in the blood at time $t = t_i$. The parameter $\sigma_{\text{d}}^2 \in \mathbb{R}_+$ is a hyperparameter describing the variance of the measurement noise. The hyperparameter σ_{d}^2 is not known a priori. Thus, we infer σ_{d}^2 as a parameter of the model and assign an inverse gamma prior [19]. We enforce non-negativity on the likelihood function by truncating it at 0 since the phosphate concentrations are non-negative.

We consider the initial condition z_{IC} to have mean equal to the phosphate concentration at time $t = 0$ and variance σ_{d}^2 equal to the measurement error, i.e.,

$$z_{\text{IC}} \sim \mathcal{N}\left(Z_0, \sigma_{\text{d}}^2\right). \tag{3.2}$$

This choice of prior for the initial condition can be interpreted as the initial measurement following the same measurement model as the measurements for time $t > 0$, i.e., we do not assume that the first measurement is more accurately measured than the subsequent ones.

We model the prior of the unknown parameters θ by the Gaussian distribution,

$$\theta \sim \mathcal{N}\left(\begin{bmatrix} \mu_{C_s} \\ \mu_{K_s} \\ \mu_{K_b} \end{bmatrix}, \begin{bmatrix} \sigma_{C_s}^2 & 0 & 0 \\ 0 & \sigma_{K_s}^2 & 0 \\ 0 & 0 & \sigma_{K_b}^2 \end{bmatrix}\right), \tag{3.3}$$

where μ_{C_s}, μ_{K_s} and μ_{K_b} represent the prior clinical knowledge, i.e., our prior belief about most likely parameter values and $\sigma_{C_s}^2$, $\sigma_{K_s}^2$ and $\sigma_{K_b}^2$ are the variances for C_s, K_s and K_b, respectively. As with the likelihood function, we impose constraints such that we only consider the parameters in a physiologically meaningful range.

As commonly done, we assume that at the start of the dialysis, i.e., $t = 0$, the patient's phosphate concentration is approximately in a steady state, i.e., we assume that Z_0 is close to C_s and we choose $\mu_{C_s} = Z_0$. The steady state assumption follows from (2.1) where $K_b = 0$ when the patient is not receiving dialysis treatment.

In previous publications [3, 7, 8] C_s is fixed to the value of the initial phosphate measurement. However, the data from Agar et al. [7] show large uncertainty for the first measurement point. Our choice of prior allows C_s to deviate from the initial measurement of the phosphate concentration and thereby our model is not oblivious to measurement errors for the initial phosphate measurement.

We base our values for μ_{K_s} and μ_{K_b} on literature and we choose $\mu_{K_s} = 8.06$ L/hour and $\mu_{K_b} = 7.56$ L/hour [3].

We initially considered $\sigma_{C_s}^2$, $\sigma_{K_s}^2$ and $\sigma_{K_b}^2$ to be parameters of the model. However, preliminary results showed that it greatly decreased the stability of the results. Thus, we choose $\sigma_{C_s}^2 = 0.2$, $\sigma_{K_s}^2 = 2.0$ and $\sigma_{K_b}^2 = 2.0$ based on visual inspection of the prior to incorporate adequate uncertainty about the prior mean. Modest increase of the prior variances did not lead to change in conclusions, see subsection E.2 and Table E.4. Thus, the chosen prior variances are robust with respect to the results and not sensitive to the prior assumptions.

3.1.2. Multiple-pass formulation

The main difference between the MP formulation and the SP formulation is the inclusion of an additional state variable through equation (2.2b). Hence, the likelihood function for MP is

$$\begin{bmatrix} X_i \\ Y_i \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} x(t_i) \\ y(t_i) \end{bmatrix}, \sigma_{\mathrm{d}}^2 I \right), \text{ for } i = 1, 2, ..., n, \tag{3.4}$$

where $X \in \mathbb{R}^n$ and $Y \in \mathbb{R}^n$ are vectors with the measurements of the phosphate concentration in the blood and dialysate at time $t = t_i$, respectively, and I is the 2×2 identity matrix. The initial condition for the phosphate concentration in the dialysate is set to zero, i.e., $y_{IC} = 0$, and the initial condition for the phosphate concentration in the blood is assigned a prior with mean X_0 and variance equal to the measurement variance, i.e.,

$$x_{\mathrm{IC}} \sim \mathcal{N}\left(X_0, \sigma_{\mathrm{d}}^2\right). \tag{3.5}$$

Lastly, we choose the prior for the parameters θ to be (3.3) with the exception that $\mu_{C_s} = X_0$.

3.2. Implementation and diagnostics

We use sampling-based techniques to approximate the posterior [20]. Markov Chain Monte Carlo (MCMC) is a sampling technique that generates a Markov chain of samples that converges to the posterior distribution of the parameters [21]. Hence, we can compute posterior statistics, i.e., mean, 95% CI and correlation from the Markov chain.

The simple MCMC techniques such as random walk Metropolis Hastings and the Gibbs sampler are plagued by inefficient exploration of the parameter space via random walks and are highly sensitive to correlated parameters. Hamiltonian Monte Carlo (HMC) is an MCMC method that avoids random walk behavior by taking a series of first-order gradient informed steps in the simulation and explores the parameter space well even in the case of correlated parameters. The performance of the HMC sampler is highly sensitive to the choice of user-specified parameters. However, the No-U-Turn Sampler (NUTS) is an HMC method where the user-specified parameters are automatically estimated. [22] We use Runge-Kutta 45 (RK45) to solve the ODE system [23] and the `PySTAN` implementation of NUTS [24] with default choice for all associated parameters to compute the samples that approximate the posterior distribution.

For each simulation, we generate four sample chains from random initializations, and we consider the potential scale reduction statistic, the so-called $\hat{R}$ value for sampling diagnostics [25]. The $\hat{R}$ value measures the ratio of the average variance of samples within each chain to the variance of the pooled samples across chains, and if all chains are at equilibrium, then the $\hat{R}$ value will be one.

4. Results

In this section, we consider two data sets for dialysis patients during hemodialysis. For each patient, we generate 4000 samples and visualize the results in terms of posterior mean and 95% CIs for the estimated parameters and phosphate concentrations during and after hemodialysis. In addition, we also visualize the pairwise correlation for the parameters by scatter plots of the samples and compute the relative standard deviation. All presented results returned an $\hat{R}$ value of one, indicating convergence

of the sample chains. In addition, we visually inspected the sample chains, which appeared well mixed. Tables with estimated posterior means, 95% CIs and relative standard deviations are found in Appendix B, and RMSE is listed in Table 2. The assumptions of normal distributed measurement errors are investigated by the empirical error distributions and QQ-plots which are shown in Appendix D, Figure D.1 and Figure D.2. We may not reject the hypothesis about normal distributed measurement errors based on the investigations.

We have also investigated the models using synthetic data to confirm the findings of the results with real data. These synthetic experiments can be found in the supplementary. Here we present the results obtained by the Bayesian model described in Section 3 for the data depicted in Figure 3.

4.1. Single-pass and multiple-pass

First, we consider the hemodialysis data for the ten patients shown in Figure 3. Beside phosphate concentrations in the blood and dialysate depicted, we have hourly measurements of the phosphate concentration in the dialysate (C_d) for SP, the volume of the blood compartment (V_b) for both SP and MP, and the dialysate volume (V_d) for MP. C_d was measured when exiting the dialysate compartment after initializing the dialysis process. We assume that the concentration of phosphate in the dialysis for SP is constant as suggested by data, and for each patient, we compute C_d as the spatial average of the concentration of phosphate from inlet to outlet of the dialysis machine. Table 1 lists C_d, V_d and V_b estimated directly from available data and Figure A.1 and Figure A.2 in Appendix A provide exploratory statistics of the corresponding data.

Table 1. The mean concentration of phosphate in the dialysate for SP, C_d, the mean dialysate volume for MP, V_d and the mean extracellular volume, V_b for both SP and MP.

	Estimate	Patient									
		1	2	3	4	5	6	7	8	9	10
SP	C_d [mmol/L]	0.16	0.12	0.11	0.16	0.18	0.21	0.14	0.09	0.14	0.09
	V_b [L]	16.88	17.74	16.92	21.20	18.20	14.74	15.32	13.04	20.20	18.00
MP	V_b [L]	16.99	17.80	17.43	21.20	18.39	15.15	15.48	13.76	20.21	18.25
	V_d [L]	22.61	23.00	26.57	31.93	28.51	20.15	23.42	14.24	28.59	23.25

4.1.1. Estimation

The estimated phosphate concentrations obtained for SP are depicted in Figure 4 along with the predicted relapse. The solid line represents the posterior mean, the full circles are data points and the transparent region indicates the 95% CI i.e., the region that contains 95% of the samples. Considering the estimated phosphate concentrations for SP, we see that the sampler has computed a decent fit in terms RMSE in Table 2 and posterior mean with a narrow 95% CI for the treatment phase. However, there is a large 95% CI for the relapse phase.

The corresponding parameter estimates with 95% CI for SP are visualized in Figure 5 and listed in

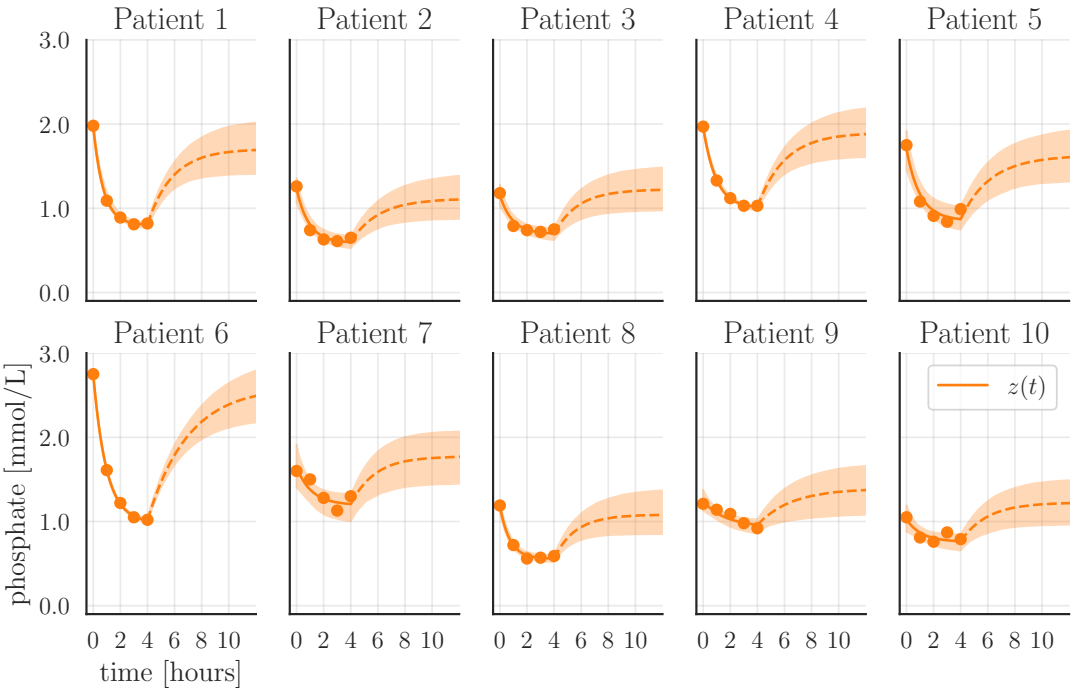

Figure 4. Estimated treatment and relapse for SP. The full lines are the posterior mean of the samples whereas the transparent regions represent the 95% CI. The full circles are measurements, and the dashed line is the posterior mean of the estimated relapse. For RMSE, see Table 2.

Table B.1 where the average relative standard deviation is 10.3%, 18.4% and 18.6% for C_s, K_s and K_b, respectively. The full posterior density for the parameters for patient 2 is shown in Figure 6. We have chosen to only include a correlation plot for patient 2 in this section since it shows the general trend of the estimated parameters. The correlation plots for the remaining patients are found in Figure C.1-C.9 in Appendix C.

Figure 7 shows the estimated phosphate concentrations and predicted relapse phase for MP. Figure 7 and Table 2 show that the parameter estimation has found a satisfying fit both visually and in terms of RMSE for MP as for SP. However, the width of the 95% CIs is smaller for the relapse phase. The reduced uncertainty in the relapse can be explained by the reduced 95% CI for C_s in MP compared to SP which is shown in Figure 5 and Figure 6 and quantified by the decreased average standard deviation of 7.3% in Table B.2, i.e., a reduction of 3%.

Moreover, Figure 5 shows a great reduction in the uncertainty about K_b as expected from the addition of equation (2.2b) with a relative standard deviation of 9.6%, i.e., a reduction of 9% compared to SP. However, the uncertainty about K_s remains largely unaffected by the additional knowledge utilized by the MP model and the uncertainty actually increases on average with an average relative standard deviation of 24.8%.

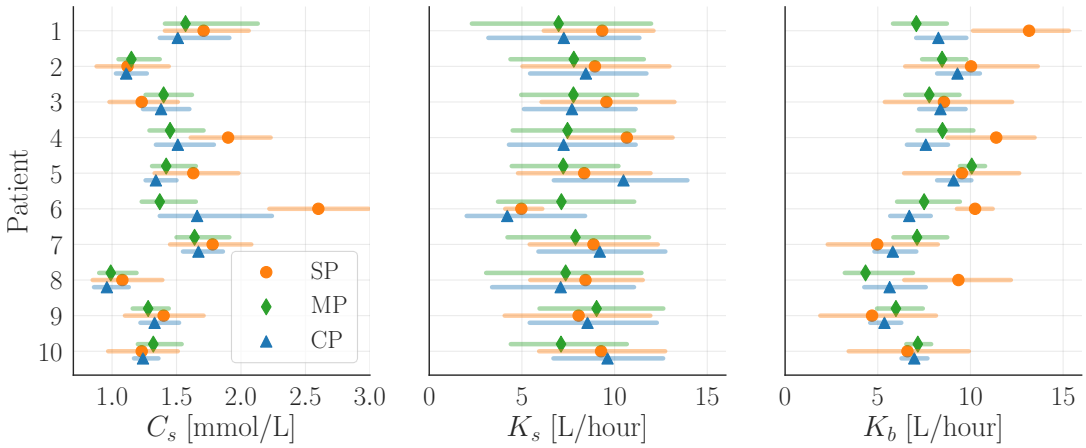

Figure 5. Visualization of the parameter estimates. The dots, diamonds and triangles represent the posterior mean for SP, MP and CP, respectively. The transparent region is the 95% CI. The full posterior of the parameters for patient 2 is shown in 6 and for the remaining patients in Figure C.1-C.9 in Appendix C.

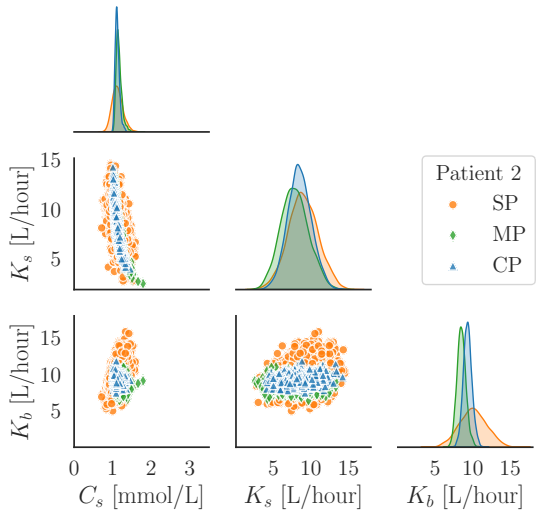

Figure 6. Plot of the posterior density and correlation of the parameters estimated for patient 2 for SP, MP and CP. The density plots show the posterior density functions, and the scatter plots show the posterior samples.

Considering the CP results in Figure 8, we see that CP finds a unified set of parameters that describe the SP and MP sessions for each patient. Moreover, the CP estimates a satisfying fit both visually and in terms of RMSE in Table 2. The parameter estimates are very similar to the ones obtained by MP as seen in Figure 5 except for patient 6 and with only a slight reduction compared to MP in average relative

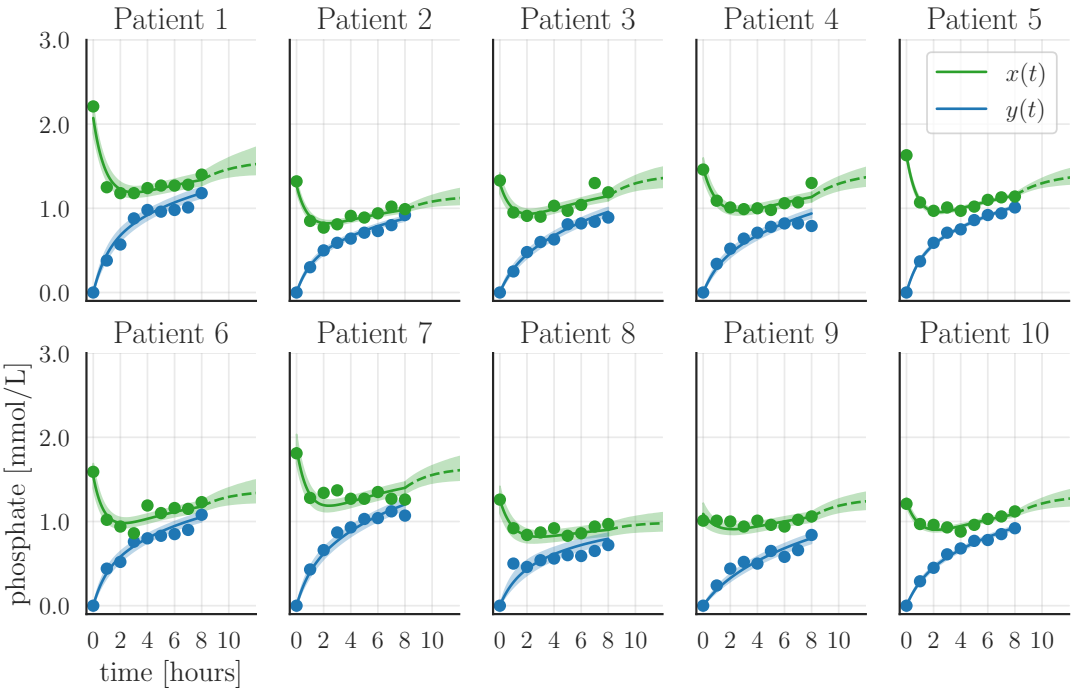

Figure 7. Estimated treatment and relapse for MP. The full lines are the posterior mean of the samples whereas the transparent regions represent the 95% CI. The full circles are measurements, and the dashed line is the posterior mean of the estimated relapse. For RMSE, see Table 2.

standard deviation, 6.9%, 22.9% and 8.2% for C_s, K_s and K_b, respectively. A possible explanation for the difference in K_s for patient 6 is the large difference in initial measured phosphate concentration, indicating that steady state had not been reached before treatment onset.

The synthetic results in Figure E.2 and E.3 in the supplementary materials show that with fixed C_s and an uniform prior on K_s and K_b (mimicking the parameter estimation in [3, 7, 8]), we have a very limited identifiability of K_s and K_b for SP, whereas MP and CP recover values very close to the true parameters with significantly lower uncertainty. In addition, the parameter estimates for K_s and K_b were highly correlated and this correlation was significantly reduced by MP and CP. We also considered the full Bayesian model with priors on the synthetic data and the results are depicted in Figure E.5 and E.6 in the supplementary materials. The results showed that MP and CP in general came closer to the true parameters with smaller 95% CI and showed similar results in terms of relative standard deviation.

In summary, the uncertainty associated with the SP results is reduced significantly by using the Bayesian model with priors compared to the standard parameter estimation without the clinical knowledge incorporated. For the Bayesian models, we see that MP and CP are superior to SP in estimating patient-specific parameters C_s and K_b, but that the gain of considering CP compared to MP is limited.

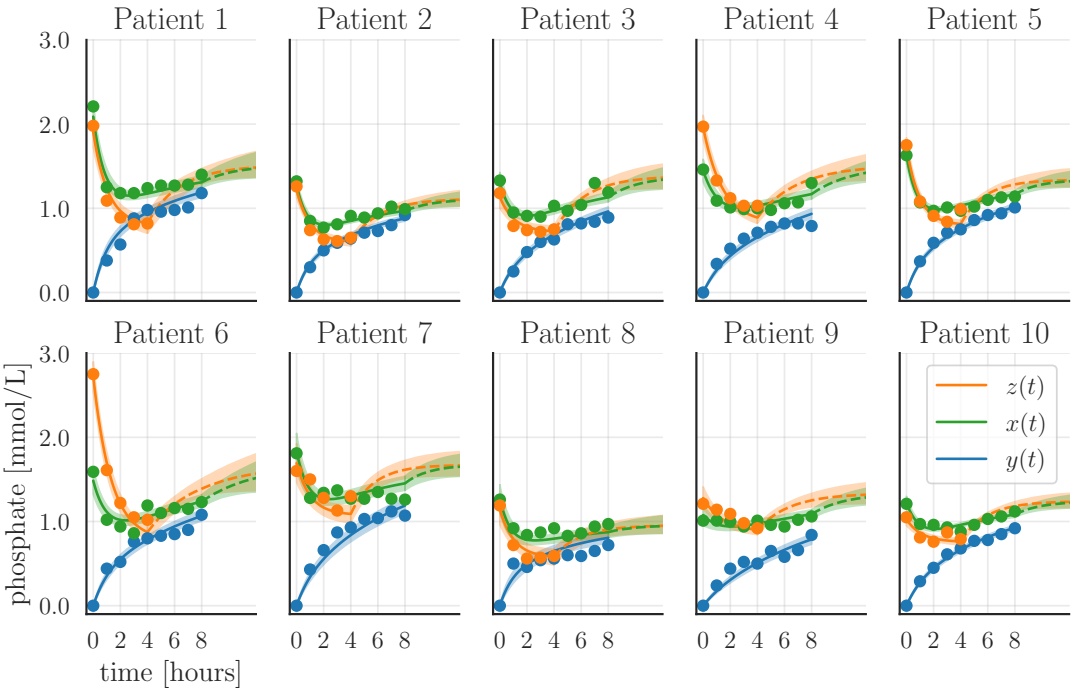

Figure 8. Estimated treatment and relapse for CP. The full lines are the posterior mean of the samples whereas the transparent regions represent the 95% CI. The full circles are measurements, and the dashed line is the posterior mean of the estimated relapse. For RMSE, see Table 2.

However, we see that the uncertainty about K_s is large even when using all available data with the CP model. These findings are further supported by the synthetic results in the supplementary materials, where the estimates obtained by MP and CP are closer to the true parameter value and with less uncertainty. Thus, based on the estimation results, it seems that the SP data without relapse data or consecutive sessions are not sufficient for estimating the parameters reliably.

4.2. Consecutive SP sessions

Debowska et al. [3] present a data set consisting of 25 patients that were examined during three consecutive SP sessions of a one-week dialysis treatment cycle. They present the data as the average of the measurement for the 25 patients and we have read off the data from the figures. Measurements were obtained hourly for a total duration of four hours with the addition of a measurement 45 minutes after ended treatment, i.e., we have five SP measurements and a relapse measurement for each of the three consecutive SP sessions. We choose $V_b = 20$ and $C_d = 0$.

Table 2. Computed RMSE for Figure 4, 7 and 8. We compute RMSE by the formula RMSE = $\sqrt{\frac{1}{n}\sum_{i=1}^{n}(\hat{X}_i - X_i)^2}$ where $\hat{X}$ and X are the estimated and measured phosphate concentrations, respectively.

Patient	SP	MP	CP
1	0.29	0.08	0.1
2	0.17	0.04	0.05
3	0.11	0.06	0.08
4	0.32	0.06	0.09
5	0.26	0.02	0.04
6	0.59	0.07	0.1
7	0.27	0.08	0.12
8	0.16	0.08	0.11
9	0.19	0.06	0.08
10	0.1	0.03	0.04

4.2.1. Simulations and estimates

The aim of this subsection is to investigate the improvement of information obtained by including relapse measurement and / or consecutive sessions in the SP model. We investigate the four following scenarios, Scenario 1 (S1) where we consider the first SP treatment only, Scenario 2 (S2) with the first SP treatment with the addition of a measured relapse point, Scenario 3 (S3) where we consider the first SP treatment with relapse point and the first measured data point of the second SP, and Scenario 4 (S4) where we include the data from all three SP consecutive sessions.

The results for the four scenarios are depicted in Figure 9. The measurements included in each parameter estimation are marked with colored dots, whereas the measurements not included in the model estimation are marked with black open circles. The posterior statistics for the parameters are shown in Figure 10 and listed in Table B.4. Correlation of the parameters is shown in Figure 11.

Figure 9a shows estimation without relapse measurement for a single SP session, the phosphate concentration has a quite large 95% CI and undershoots the relapse. If we consider the uncertainty in the correlation plot for the parameters in Figure 11 and Figure 10, we see a large 95% CI for the parameter estimates and relative standard deviation in Table B.4 which is similar to the uncertainty associated with the estimate for the SP estimation in Section 4.1.

A model estimation including the measured relapse 45 minutes after ended treatment is depicted in Figure 9b. The 95% CIs for the phosphate concentration is slightly reduced, but the 95% CIs for the parameters have barely changed as seen in Figure 10, Figure 11 and Table B.4. Hence, including a measurement after 45 minutes relapse has limited effect on the uncertainty of the parameter estimates. This can also be seen by considering the correlation plot in Figure 11, where the width of the distribution is only slightly changed. It is noteworthy that the addition of the relapse point has such limited effect on the estimation. However, this limited effect is due to the very rapid dynamics in the initial relapse phase. The initial relapse is not very sensitive to small changes, whereas a relapse point measured later e.g., after two hours, will have a larger effect on the estimation process due to the slower change in the concentration.

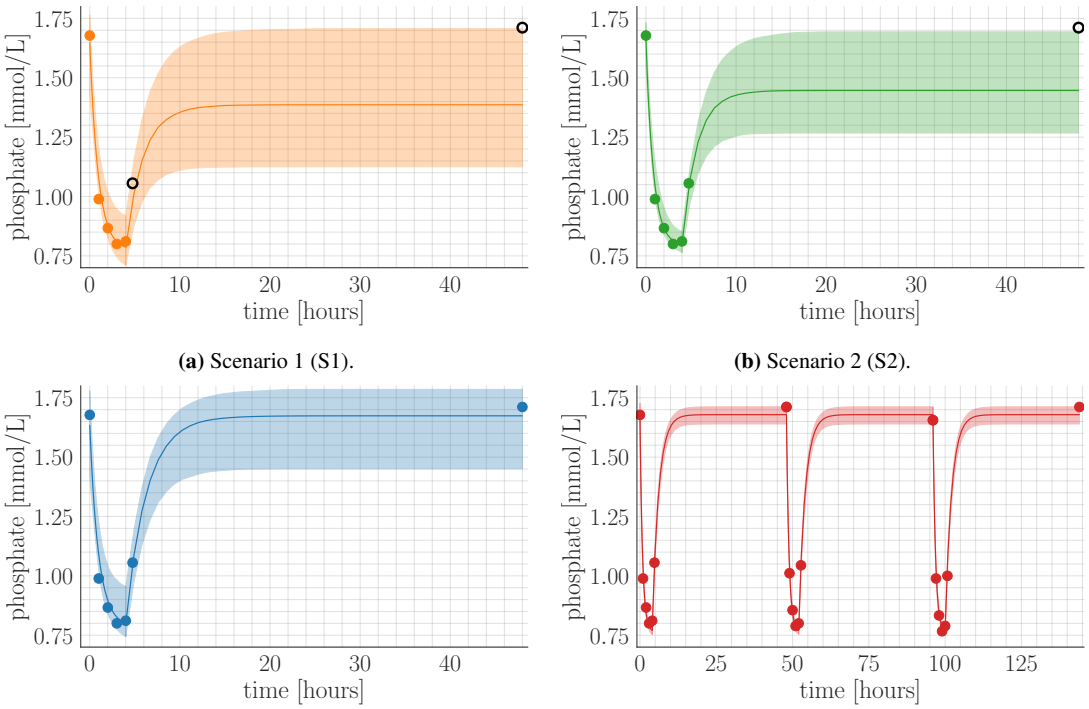

(a) Scenario 1 (S1).

(b) Scenario 2 (S2).

(c) Scenario 3 (S3).

(d) Scenario 4 (S4).

Figure 9. Four scenarios for the relapse data. (a) first SP treatment with no relapse data (S1), (b) first SP treatment with a single relapse data point after 4.75 hours (S2), (c) first SP treatment with relapse data after 4.75 and 48 hours (S3). Lastly (d) shows the fit when including all three consecutive SP treatment with relapse data (S4). The measurements are shown with colored circles. The open black circles in (a) and (b) indicate that the measurements are not used for estimation. RMSE is S1 =0.05, S2=0.03, S3=0.06 and S4=0.03, respectively.

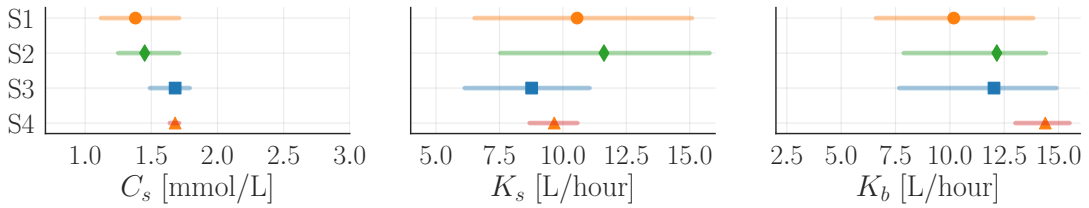

Figure 10. Posterior mean and 95% CI for the parameter estimates for the four scenarios, S1, S2, S3, and S4 depicted in Figure 9. The figure shows that the uncertainty about the parameter estimates decreases as the number of measurements increases.

Considering the full relapse in Figure 9c, we see the effect of having a relapse measurement several hours after ended treatment. The estimated steady state for the phosphate concentration has an in-

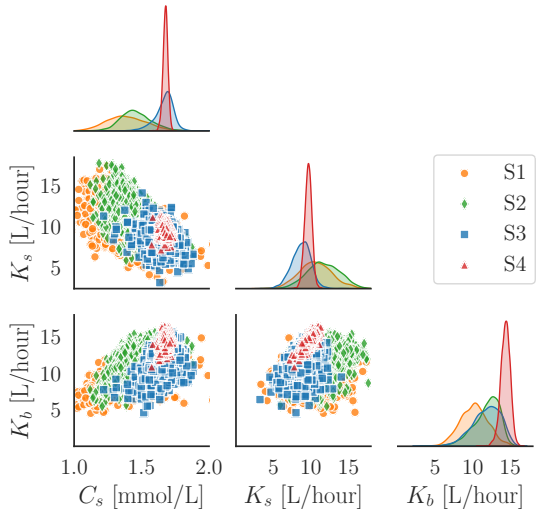

Figure 11. Plot of the posterior density and correlation for the parameter estimates for the four scenarios, S1, S2, S3, and S4. The density plots show the posterior density functions, and the scatter plots show the posterior samples. The uncertainty associated with the posterior mean of the parameters decreases as more information is included in terms of relapse measurement and /or consecutive sessions.

creased posterior mean and narrower 95% CI compared to Figure 9a and 9b. This increase is explained by the increase for C_s which can be seen in Figure 10 and Figure 11. There is also a slight narrowing of the 95% CI for K_s whereas the effect on K_b is limited as the relative standard deviation actually increases from 16% to 18% compared to the partial relapse. Hence, including relapse measurements has limited effect on the identifiability of K_b, but reduces the uncertainty associated with the estimates for C_s and K_s. This observation is expected based on the model (2.1), since we have $K_b = 0$ in the relapse phase.

Lastly, if we have three consecutive SP treatments for the same patient, we can reduce the uncertainty even further, as shown in Figure 9d. The three consecutive SP treatments carry significant information since the repetition makes the estimates less sensitive to fluctuations in the data, which can also be seen in Figure 10 and Figure 11. Considering the relative standard deviation for K_s in Table B.4, we find that it decreases from 20% to 5 % by considering the consecutive sessions compared to a single session. However, even in the case of a single session, our Bayesian approach has significantly smaller relative standard deviation compared to the estimates found by Debowska et al. [3] and Agar et al. [7], who report a relative standard deviation of 79% and 47%, respectively. Even for K_b, we see a significant narrowing of the 95% CI. Thus, measuring consecutive sessions greatly increases the identifiability of all three model parameters as the relative standard deviation decreases significantly for all three parameter estimates, as seen in Table B.4.

We also investigated the effect of including relapse measurements for the synthetic data for SP, MP and CP. The results including relapse measurements are shown in Figure E.8-E.11 and results for two consecutive sessions are shown in Figure E.12 and E.13. Here we found that the consecutive sessions

were more effective than relapse measurements to reduce the uncertainty of the parameters which aligns with the findings in Figure 9. In general for the synthetic data, we found that MP compared to SP had less uncertainty and came closer to the true parameter values.

5. Conclusion

Phosphate clearance with hemodialysis is crucial for patients with renal failure since abnormal levels of phosphate are associated with increased vascular calcification and mortality. We propose a Bayesian approach to parameter estimation for patients undergoing hemodialysis treatments (SP, MP and CP). The Bayesian approach allows us to formally include clinical knowledge in the model and to use uncertainty quantification to assess how reliably we can estimate the three model parameters: phosphate concentration in the bones, phosphate clearance from bone to blood and from blood to dialysate.

We validated and tested our Bayesian model on two data sets for patients with renal failure. The results showed that the uncertainty for the parameter estimates is greatly reduced by considering MP and CP compared to SP while CP is not significantly better than MP. However, for the parameter governing the diffusion rate between bone phosphate and blood, the uncertainty remained unchanged. We also investigated the impact of including relapse data and consecutive treatments. The results showed that including an early relapse measurement (after 45 minutes) had little effect on the estimation process if not combined with a measurement in the later relapse phase. The relapse measurements taken more than 45 minutes after ended treatment had significant impact on the reliability of the model parameters. Moreover, the results showed that we can reduce the relative standard deviation for the phosphate clearance from blood to bone from 20% to 5% by including consecutive sessions in the estimation process compared to estimation based on a single session.

Numerical results on synthetic data confirmed the findings obtained from the real data, and showed that the parameters were poorly identified for SP if no prior information was included. The uncertainty of the estimates greatly decreased when using the Bayesian model incorporating clinical knowledge, and the MP model generally was closer to the true parameter values of the model. Compared to existing parameter estimates of the phosphate clearance from bone to blood, our Bayesian model can estimate a parameter associated with significantly lower uncertainty for both SP and MP. As the consecutive SP sessions may also be used to reduce the uncertainty of the estimated parameters significantly and such action comes without any costs, it seems very straight forward to implement clinically. Hence, pooling data, e.g., from three consecutive SP sessions to estimate patient specific parameters, is recommendable.

Acknowledgments

This work was supported by The Villum Foundation (grant no. 25893).

Conflict of interest

The authors declare there is no conflict of interest.

References

1. C. Basile, D. Schneditz, Haemodialysis adequacy monitoring for phosphate: An old problem with new solutions?, *Nephrol. Dial. Transpl.*, **30** (2014), 9–11. https://doi.org/10.1093/ndt/gfu351
2. J. Poleszczuk, M. Debowska, A. Wojcik-Zaluska, A. Ksiazek, W. Zaluska, Phosphate kinetics in hemodialysis: Application of delayed pseudo one-compartment model, *Blood Purificat.*, **42** (2016), 177–185. https://doi.org/10.1159/000445934
3. M. Debowska, J. Poleszczuk, A. Wojcik-Zaluska, A. Ksiazek, W. Zaluska, Phosphate kinetics during weekly cycle of hemodialysis sessions: Application of mathematical modeling, *Artif. Organs*, **39** (2015), 1005–1014. https://doi.org/10.1111/aor.12489
4. B. Kestenbaum, J. N. Sampson, K. D. Rudser, D. J. Patterson, S. L. Seliger, B. Young, et al., Serum phosphate levels and mortality risk among people with chronic kidney disease, *J. Am. Soc. Nephrol.*, **16** (2004), 520–528. https://doi.org/10.1681/asn.2004070602
5. J. P. Gutzwiller, Estimating phosphate removal in haemodialysis: An additional tool to quantify dialysis dose, *Nephrol. Dial. Transpl.*, **17** (2002), 1037–1044. https://doi.org/10.1093/ndt/17.6.1037
6. R. Subramanian, R. Khardori, Severe hypophosphatemia: Pathophysiologic implications, clinical presentations and treatment, *Medicine*, **79** (2000), 1–8. https://doi.org/10.1097/00005792-200001000-00001
7. B. U. Agar, A. Akonur, A. K. Cheung, J. K. Leypoldt, A simple method to estimate phosphorus mobilization in hemodialysis using only predialytic and postdialytic blood samples, *Hemodial. Int.*, **15** (2011), S9–S14. https://doi.org/10.1111/j.1542-4758.2011.00596.x
8. M. Andersen, K. O. Bangsgaard, J. T. Ottesen, Analytical solution of phosphate kinetics during and after single pass or multi pass dialysis, *submitted for publication*.
9. S. H. Laursen, A. Buus, M. H. Jensen, P. Vestergaard, O. K. Hejlesen, Distribution volume assessment compartment modelling: Theoretic phosphate kinetics in steady state hemodialys patients, *Int. J. Artif. Organs*, **38** (2015), 580–587. https://doi.org/10.5301/ijao.5000449
10. S. H. Laursen, P. Vestergaard, O. K. Hejlesen, Phosphate kinetic models in hemodialysis: A systematic review, *Am. J. Kidney Diseases*, **71** (2018), 75–90. https://doi.org/10.1053/j.ajkd.2017.07.016
11. E. M. Spalding, P. W. Chamney, K. Farrington, Phosphate kinetics during hemodialysis: Evidence for biphasic regulation, *Kidney Int.*, **61** (2002), 655–667. https://doi.org/10.1046/j.1523-1755.2002.00146.x
12. S. Eloot, R. Vanholder, W. Van Biesen, Less water for haemodialysis: Is multiple pass the future pace to go?, *Nephrol. Dial. Transpl.*, **27** (2012), 3975–3978. https://doi.org/10.1093/ndt/gfs435
13. S. Eloot, W. V. Biesen, M. Axelsen, G. Glorieux, R. S. Pedersen, J. G. Heaf, Protein-bound solute removal during extended multipass versus standard hemodialysis, *BMC Nephrol.*, **16** (2015). https://doi.org/10.1186/s12882-015-0056-y
14. J. G. Heaf, M. Axelsen, R. S. Pedersen, Multipass haemodialysis: A novel dialysis modality, *Nephrol. Dial. Transpl.*, **28** (2012), 1255–1264. https://doi.org/10.1093/ndt/gfs484

15. H. Huang, A. Handel, X. Song, A Bayesian approach to estimate parameters of ordinary differential equation, *Comput. Stat.*, **35** (2020), 1481–1499. https://doi.org/10.1007/s00180-020-00962-8

16. J. Russell-Buckland, C. P. Barnes, I. Tachtsidis, A Bayesian framework for the analysis of systems biology models of the brain, *PLOS Comput. Biol.*, **15** (2019), e1006631. https://doi.org/10.1371/journal.pcbi.1006631

17. C. Bianchi, E. Lanzarone, G. Casagrande, M. L. Costantino, A Bayesian approach for the identification of patient-specific parameters in a dialysis kinetic model, *Stat. Methods Med. Res.*, **28** (2018), 2069–2095. https://doi.org/10.1177/0962280217745572

18. O. S. Iheagwara, T. S. Ing, C. M. Kjellstrand, S. Q. Lew, Phosphorus, phosphorous, and phosphate, *Hemod. Int.*, (2012). https://doi.org/10.1111/hdi.12010

19. R. C. Smith, *Uncertainty quantification: Theory, implementation and applications*, Society for Industrial and Applied Mathematics, USA, 2014.

20. J. Vanlier, C. A. Tiemann, P. A. J. Hilbers, N. A. W. van Riel, Parameter uncertainty in biochemical models described by ordinary differential equations, *Math. Biosci.*, **246** (2013), 305–314. https://doi.org/10.1016/j.mbs.2013.03.006

21. C. P. Robert, G. Casella, *Monte Carlo Statistical Methods*, Springer, New York, 2004.

22. M. D. Hoffman, A. Gelman, The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo, *J. Mach. Learn. Res.*, **15** (2014), 1593–1623.

23. R. J. LeVeque, *Finite difference methods for ordinary and partial differential equations: Steady-state and time-dependent problems*, Society for Industrial and Applied Mathematics, USA, 2007.

24. A. Gelman, D. Lee, J. Guo, Stan: A probabilistic programming language for Bayesian inference and optimization, *J. Educ. Behav. Stat.*, **40** (2015), 530–543. https://doi.org/10.3102/1076998615606113

25. A. Gelman, D. B. Rubin, Inference from iterative simulation using multiple sequences, *Stat. Sci.*, **7** (1992). https://doi.org/10.1214/ss/1177011136

26. T. W. Anderson, D. A. Darling, Asymptotic theory of certain "Goodness of Fit" criteria based on stochastic processes, *Ann. Math. Stat.*, **23** (1952), 193–212. https://doi.org/10.1214/aoms/1177729437

27. R. B. D'Augostino, M. A. Stephens, Goodness-of-Fit Techniques, Marcel Dekker, *Statistics: textbooks and monographs*, **68** (1986).

Appendix

A. Summary statistics

B. Tables

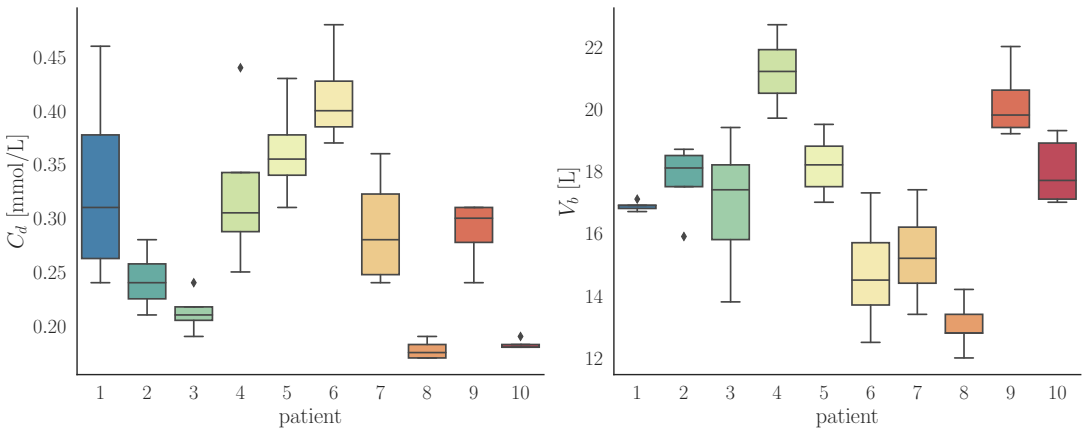

Figure A.1. Boxplots of the measured parameters C_d and V_b for the SP sessions in Figure 3.

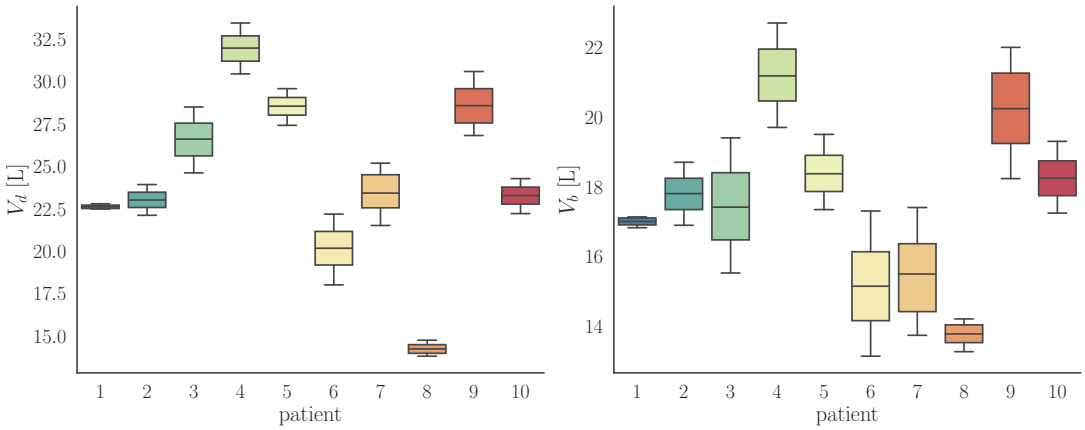

Figure A.2. Boxplots of the measured parameters V_b and V_d for the MP sessions in Figure 3.

Table B.1. Median (m), lower (l) and upper (u) 95% CI and relative standard deviation ($\frac{\text{std}}{\text{mean}}$) for the parameters for the SP estimation.

Patient	C_s				K_s				K_b			
	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$
1	1.41	1.71	2.06	9%	6.18	9.33	12.1	16%	10.14	13.16	15.31	10%
2	0.88	1.12	1.44	13%	5.03	8.94	12.96	22%	6.48	10.03	13.63	17%
3	0.98	1.23	1.51	11%	6.06	9.56	13.23	19%	5.38	8.57	12.25	21%
4	1.61	1.90	2.23	8%	7.50	10.66	13.13	13%	8.75	11.39	13.46	11%
5	1.33	1.63	1.98	10%	4.79	8.36	11.94	22%	6.42	9.54	12.63	17%
6	2.22	2.60	2.99	8%	4.10	4.97	6.10	10%	9.27	10.25	11.20	5%
7	1.45	1.78	2.08	9%	5.42	8.86	12.33	20%	2.30	4.97	8.24	30%
8	0.85	1.08	1.39	13%	5.45	8.43	11.52	18%	6.43	9.35	12.18	16%
9	1.10	1.4	1.71	11%	4.07	8.06	11.93	25%	1.90	4.69	8.14	34%
10	0.97	1.23	1.51	11%	5.93	9.27	12.73	19%	3.43	6.60	9.91	25%

Table B.2. Median (m), lower (l) and upper (u) 95% CI and relative standard deviation ($\frac{\text{std}}{\text{mean}}$) for the parameters for MP estimation.

Patient	C_s				K_s				K_b			
	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$
1	1.41	1.57	2.13	12%	2.30	6.98	11.97	38%	5.82	7.08	8.73	10%
2	1.05	1.15	1.37	7%	4.38	7.80	11.58	24%	7.39	8.47	9.78	7%
3	1.26	1.40	1.62	7%	4.96	7.78	11.22	21%	6.48	7.78	9.41	10%
4	1.29	1.45	1.71	7%	4.50	7.46	11.06	23%	7.14	8.49	10.16	9%
5	1.31	1.42	1.65	6%	4.45	7.23	10.21	22%	9.43	10.07	10.79	3%
6	1.23	1.37	1.65	8%	3.71	7.13	11.06	26%	6.04	7.50	9.43	11%
7	1.50	1.64	1.91	6%	4.22	7.89	11.86	24%	5.84	7.12	8.75	11%
8	0.90	0.99	1.19	7%	3.07	7.36	11.45	28%	3.22	4.34	6.9	20%
9	1.16	1.28	1.44	6%	5.95	9.03	12.63	19%	4.96	5.98	7.43	10%
10	1.2	1.32	1.54	7%	4.40	7.11	10.66	23%	6.54	7.16	7.88	5%

Table B.3. Median (m), lower (l) and upper (u) 95% CI and relative standard deviation ($\frac{\text{std}}{\text{mean}}$) for the parameters for CP estimation.

Patient	C_s				K_s				K_b			
	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$
1	1.37	1.51	1.91	8%	3.19	7.26	11.34	28%	7.07	8.27	9.78	8%
2	1.03	1.11	1.27	5%	5.44	8.45	11.72	19%	8.2	9.30	10.51	6%
3	1.24	1.38	1.60	7%	5.11	7.70	11.14	19%	7.23	8.38	9.74	8%
4	1.34	1.51	1.79	8%	4.30	7.25	11.14	24%	6.57	7.59	8.79	7%
5	1.26	1.34	1.50	5%	6.72	10.48	13.93	18%	8.20	9.09	10.05	5%
6	1.37	1.66	2.24	14%	2.02	4.21	8.41	39%	5.67	6.70	7.86	8%
7	1.55	1.67	1.86	5%	5.89	9.20	12.74	19%	4.82	5.81	7.08	10%
8	0.86	0.96	1.13	7%	3.39	7.09	11.05	27%	4.27	5.64	7.59	16%
9	1.22	1.33	1.52	6%	5.42	8.54	12.28	20%	4.59	5.36	6.27	9%
10	1.17	1.24	1.36	4%	6.70	9.62	12.62	16%	6.28	6.97	7.68	5%

Table B.4. Median (m), lower (l) and upper (u) 95% CI and relative standard deviation ($\frac{\text{std}}{\text{mean}}$) for the relapse data.

	C_s				K_s				K_b			
	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$
NR	1.11	1.37	1.74	11%	6.44	10.51	14.84	20%	6.48	10.04	13.53	18%
PR	1.25	1.45	1.71	8%	7.57	11.70	15.61	18%	8.18	12.17	14.42	14%
FR	1.49	1.68	1.79	4%	5.97	8.75	11.21	15%	7.77	11.94	14.88	16%
FTR	1.64	1.68	1.71	1%	8.64	9.66	10.57	5%	13.01	14.4	15.51	4%

C. Correlation plots for SP and MP

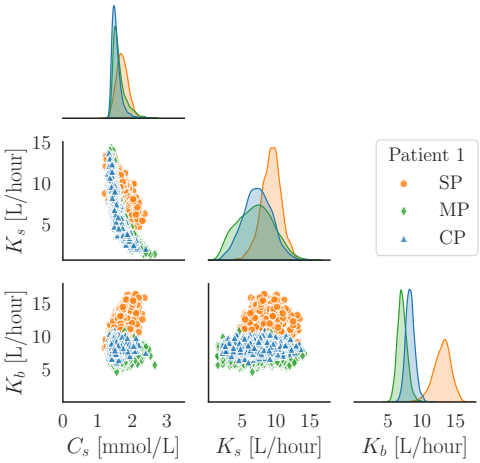

Figure C.1. Correlation and posterior density for the patient 1.

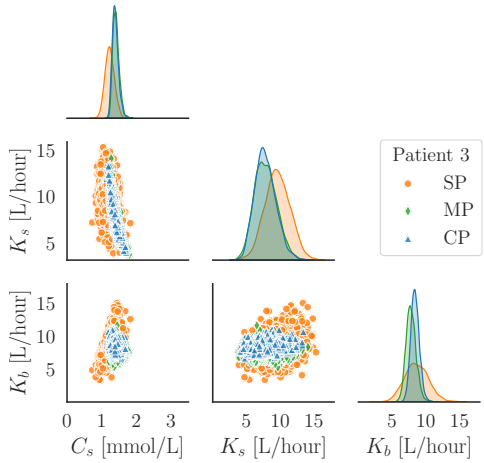

Figure C.2. Correlation and posterior density for the patient 3.

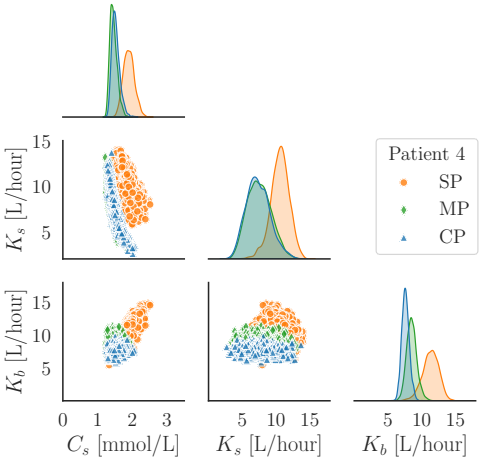

Figure C.3. Correlation and posterior density for the patient 4.

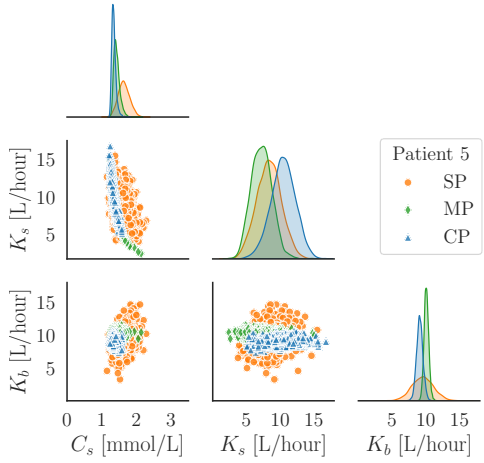

Figure C.4. Correlation and posterior density for the patient 5.

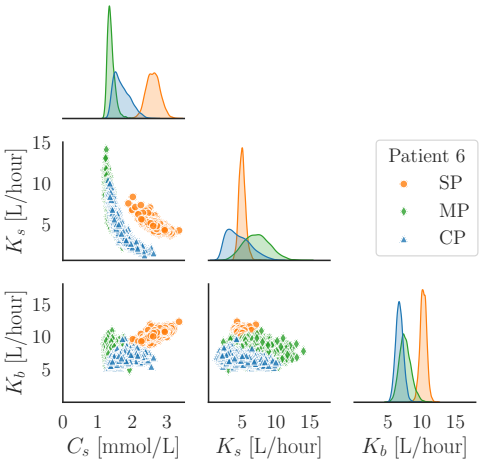

Figure C.5. Correlation and posterior density for the patient 6.

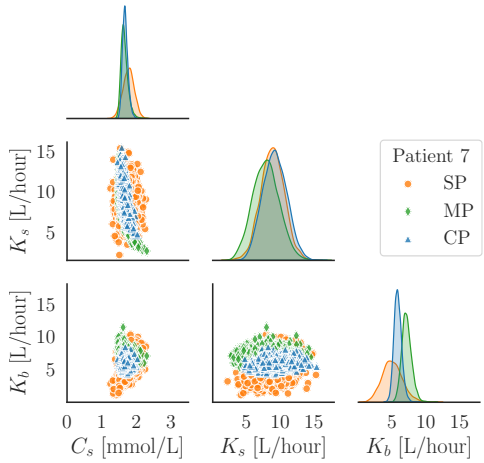

Figure C.6. Correlation and posterior density for the patient 7.

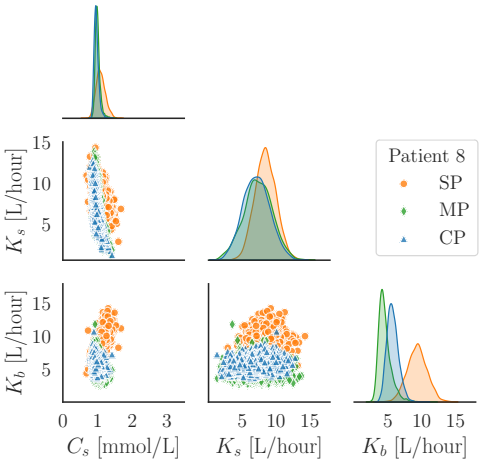

Figure C.7. Correlation and posterior density for the patient 8.

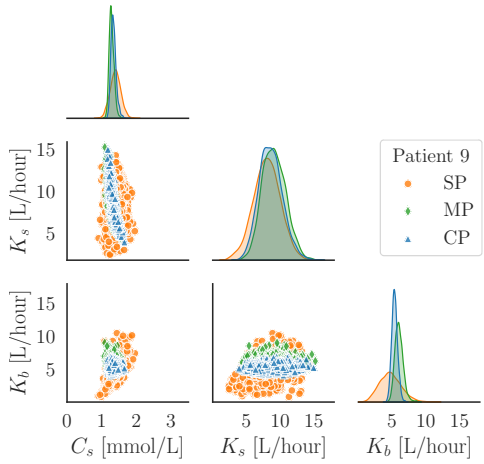

Figure C.8. Correlation and posterior density for the patient 9.

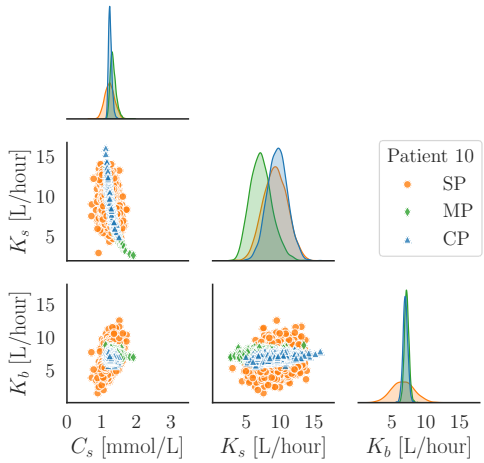

Figure C.9. Correlation and posterior density for the patient 10.

D. Validation of model assumption

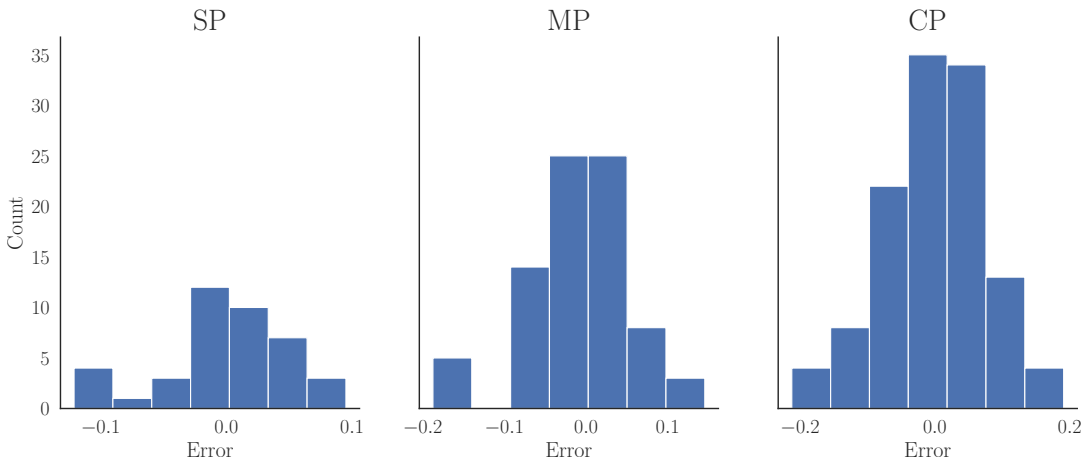

Figure D.1. Histogram of the computed errors for the experiment in Figures 4, 7 and 8.

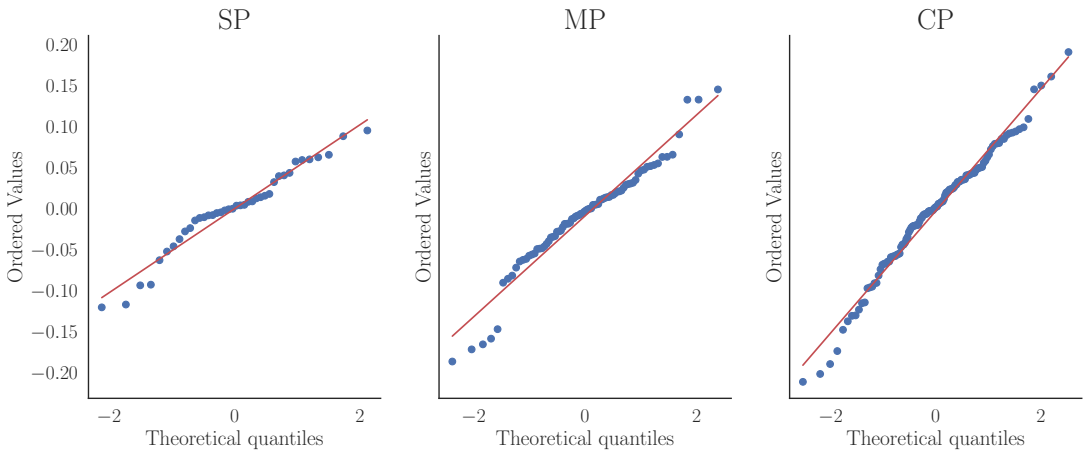

Figure D.2. QQ plot of the computed errors for the experiment in Figures 4, 7 and 8.

Due the thin tails in the QQ-plots we supplement the visual inspections in Figure D.1 and Figure D.2 with an additional analytic test, the Anderson-Darling test [26]. We determine the p-values for the Anderson-Darling statistic [27]. The Anderson-Darling test gives p-values of p=0.08, p=0.04, and p=0.07 for SP, MP, and CP, respectively. With a significant level of $\alpha = 0.05$, we see that all p-values are close to α but those for SP and CP are just above while that for MP is just below it. Thus, the conclusions are very sensitive to the choice of significant level, and one should be careful to make definitive conclusions until more data are collected. Altogether, we do not reject the hypothesis that the residuals are normal distributed but acknowledge that more data is needed to reach a robust conclusion.

Supplementary

E. Synthetic data experiments

We generate synthetic data for a patient with renal failure to investigate the identifiability of the hemodialysis model parameters. We base the true parameters on the estimates obtained for patient 2 and simulate both SP and MP treatments with relapse. We add Gaussian noise with mean equal to the "true" trajectory of the state variables and with variance $\sigma^2 = 0.05$. Figure E.1 shows simulated data for a renal failure patient during hemodialysis with added Gaussian noise. The true parameters for the simulation are listed in Table E.1. Note that we have used the true values for the parameters estimated from data, i.e., V_b, V_d and C_d and thus the only uncertainty introduced arise from the Gaussian measurement noise.

We will use the simulated data to investigate the following four cases:

- Reduced Bayesian model with fixed C_s and uniform priors
- Full Bayesian model
- Full Bayesian model with relapse measurements
- Full Bayesian model with consecutive sessions

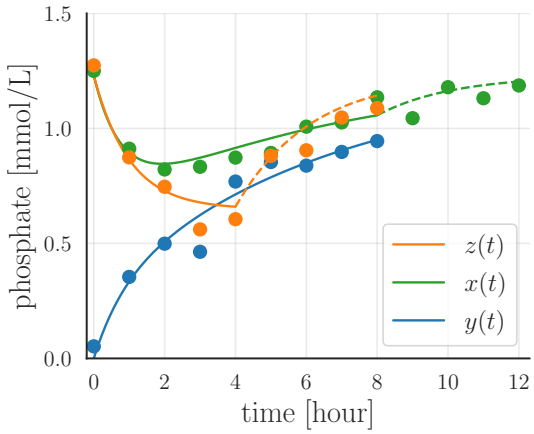

Figure E.1. Synthetic data with Gaussian noise for a dialysis patient during SP and MP treatments. The full lines are the true phosphate concentrations during hemodialysis and the dashed lines are the true relapse phase. The dots represent the hourly measurements with Gaussian noise.

Table E.1. The true model parameters chosen for the simulation experiment.

C_s	K_s	K_b	V_b	V_d	C_d	Z_0	X_0
1.23	8.39	9.23	17.77	23.0	0.12	1.23	1.23

E.1. Reduced Bayesian model with fixed C_s and uniform priors

We fix the parameter C_s and reduce the priors to uniform priors. The uniform priors correspond to simply having bounds on the parameters but with equal probability for all outcomes in the specified interval. Hence this reduced Bayesian model mimics the simple model suggested by Debowska et al. [3] and Agar et al. [7]. The results for the reduced Bayesian model are shown in Figure E.2, the parameter estimates are visualized in Figure E.3 and listed in Table E.2. Correlation plot is shown in Figure E.4.

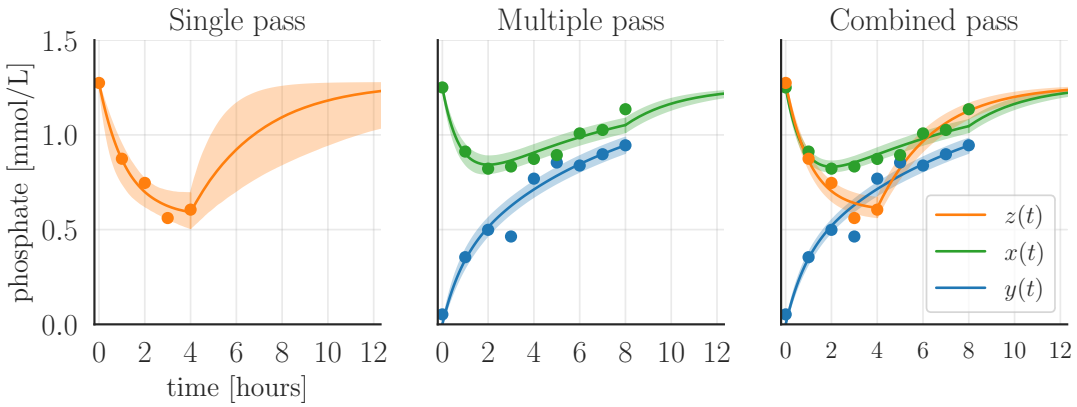

Figure E.2. Results for the simulated experiment for SP, MP and CP. The dots represent the measurements, the solid line is the estimated mean and the transparent region shows the 95% CI. RMSE is SP= 0.04, MP= 0.05 and CP= 0.11.

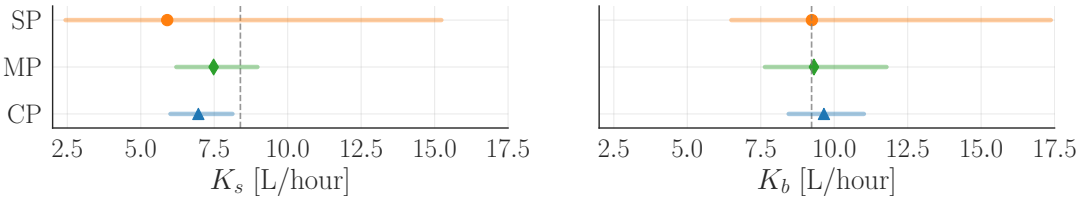

Figure E.3. Visualization of the parameter estimates and 95% for the estimated phosphate concentrations in Figure E.2. The vertical gray dotted line represents the true parameter value.

Remarkably, even though we have ignored the uncertainty in C_s, the SP model is very uncertain about K_s and K_b, as seen in Figure E.3 and in Table E.2. Moreover, we see a pronounced correlation between K_s and K_b in Figure E.4. Hence, the parameters of the SP model are poorly identified when no prior knowledge is included in the model. The parameters for MP and CP are, on the other hand, more certain about their estimates and closer to the true parameter values with lower relative standard deviation. However, this certainty comes at a cost in terms of validity of the available knowledge of C_s since the estimates for K_s and K_b may be biased if this estimate is not sufficiently close to the true parameter.

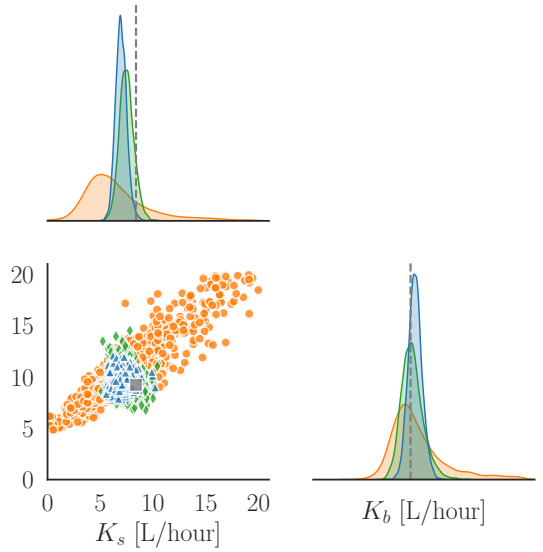

Figure E.4. Correlation plot for simulated experiment in Figure E.2.

Table E.2. Median (m), lower (l) and upper (u) 95% CI and relative standard deviation ($\frac{\text{std}}{\text{mean}}$) for Figure E.2.

HD	K_s				K_b			
	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$
SP	1.83	5.93	15.62	50%	6.22	9.28	17.8	28%
MP	6.22	7.48	9.09	10%	7.62	9.27	11.82	11%
CP	6.01	6.97	8.06	8%	8.44	9.64	11.07	7%

E.2. Full Bayesian model

We consider C_s to be a parameter of the model such that we investigate the same estimation problem as presented in Section 4 for the SP and MP data. The results are shown in Figure E.5, the parameter estimates and corresponding 95% CI are visualized in Figure E.6 and listed in Table E.3. Correlation plot shown in Figure E.7.

Figure E.5 shows the same behavior as when using real data in Section 4. All three models find a satisfying fit to the data and all three models estimate posterior means of the parameters that are close to the true parameter values. Figure E.6 and E.4 show that the uncertainty in the parameter space is much larger for SP than for MP and CP for the parameters C_s and K_b, whereas the uncertainty is only slightly reduced for MP and CP compared to SP for the estimate of K_s. Hence, the identifiability of C_s and K_b is greatly increased by considering MP measurements since the mean is closer to the true value and the uncertainty associated with the estimate is greatly reduced compared to SP.

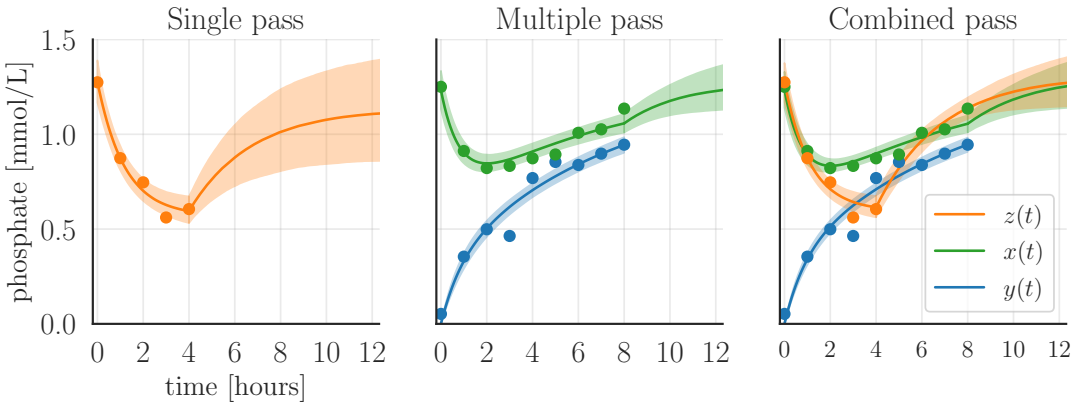

Figure E.5. Results for the simulated experiment for SP, MP and CP. The dots represent the measurements, the solid line is the estimated mean and the transparent region shows the 95% CI. RMSE is SP= 0.04, MP= 0.05 and CP= 0.1.

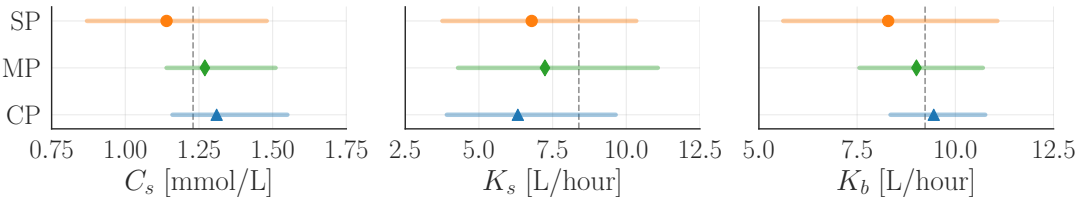

Figure E.6. Visualization of the parameter estimates and 95% CI for the estimated phosphate concentrations in Figure E.5. The vertical gray dotted line represents the true parameter value.

Table E.3. Median (m), lower (l) and upper (u) 95% CI and relative standard deviation ($\frac{\text{std}}{\text{mean}}$) for Figure E.5.

HD	C_s				K_s				K_b			
	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$
SP	0.86	1.13	1.48	14%	3.94	6.85	10.52	24%	5.76	8.32	11.19	17%
MP	1.14	1.27	1.5	7%	4.39	7.23	10.94	23%	7.51	8.99	10.78	9%
CP	1.16	1.31	1.55	8%	3.97	6.36	9.7	22%	8.31	9.47	10.79	7%

To ensure that the priors for K_s and K_b are not too strong, we redid the investigation with an 10% increase of the variance, i.e., $\sigma_{K_s} = 2.2$ and $\sigma_{K_b} = 2.2$. The results are shown in Table E.4 and show no significant change in the results.

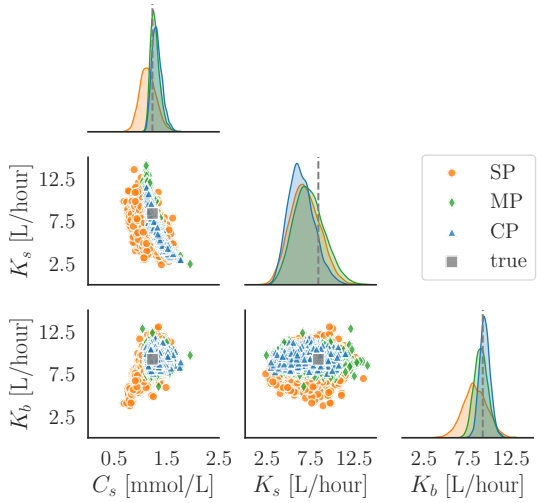

Figure E.7. Correlation plot for simulated experiment in Figure E.5.

Table E.4. Median (m), lower (l) and upper (u) 95% CI and relative standard deviation ($\frac{\text{std}}{\text{mean}}$) for the experiment in Figure E.5 with 10% increased variance of the priors for K_s and K_b, i.e., $\sigma_{K_s} = 2.2$ and $\sigma_{K_b} = 2.2$.

HD	C_s				K_s				K_b			
	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$
SP	0.86	1.15	1.5	14 %	3.49	6.63	10.75	28%	5.61	8.37	11.4	18%
MP	1.14	1.27	1.51	7 %	4.24	7.21	11.09	24 %	7.54	9.05	10.9	9 %
CP	1.16	1.31	1.56	8 %	3.89	6.21	9.59	23 %	8.32	9.47	10.75	6 %

E.3. Full Bayesian model with relapse

In the third simulated experiment, we investigate the effect of including relapse measurements 1, 2 and 4 hours after ended treatment for both SP and MP. The estimated phosphate concentrations are depicted in Figure E.8 and Figure E.9 for SP and MP, respectively. The corresponding parameter estimates are visualized in Figure E.11 and listed in Table E.5.

Figure E.11 shows that the uncertainty about C_s is reduced for all methods by including relapse measurements, especially by including a relapse measurement after 4 hours.

The uncertainty for K_s is only slightly reduced for SP and MP by including relapse measurements after 1,2 and 4 hours. This observation is quantified in Table E.5 where the relative standard deviation is reduced by 5% and 6% by including relapse up to four hours for SP and MP, respectively. Considering the estimate for K_s for CP we obtain an reduction in relative standard deviation of 10% by including the relapse measurements.

For K_b, we see a limited effect of including relapse measurements for MP and CP, as expected,

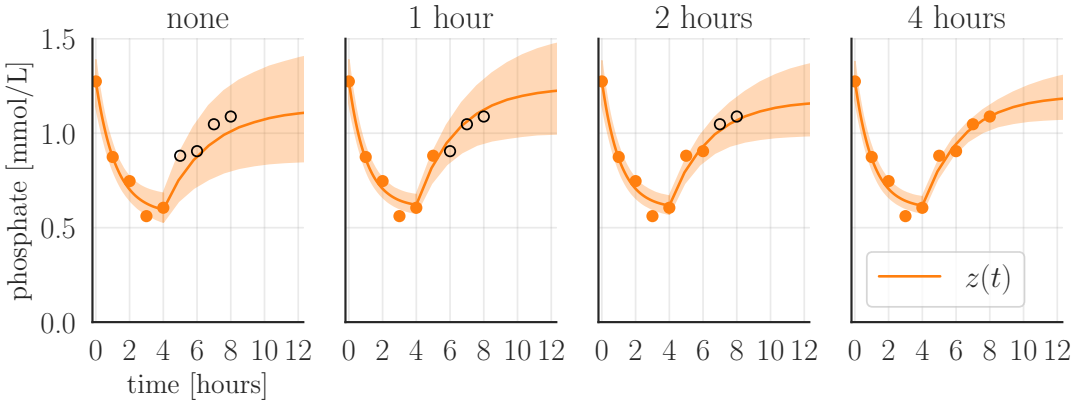

Figure E.8. Results for the simulated experiment for SP with included relapse measurements after 1, 2 and 4 hours, respectively. The measurements used for estimation are marked by full dots and the measurements left out are marked by black open circles. RMSE is None = 0.04, 1 hour = 0.04, 2 hours = 0.04 and 4 hours = 0.04.

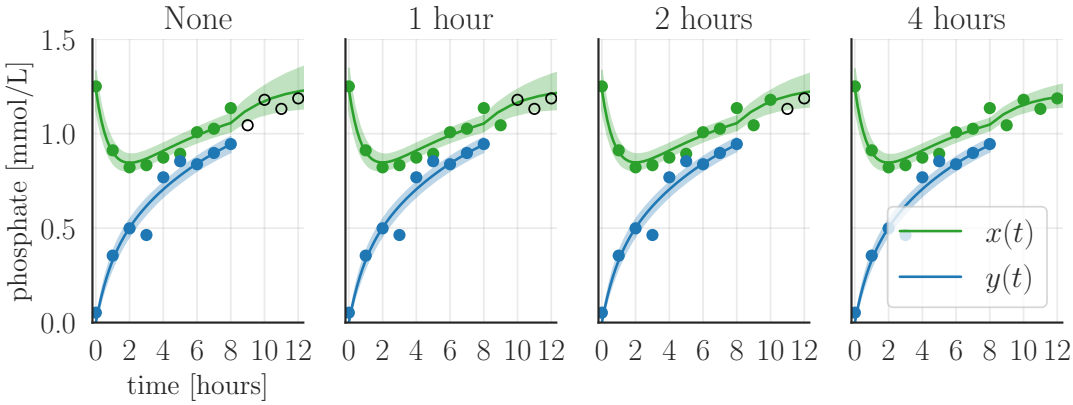

Figure E.9. Results for the simulated experiment for MP with included relapse measurements after 1, 2 and 4 hours, respectively. The measurements used for estimation are marked by full dots and the measurements left out are marked by black open circles. RMSE is None = 0.05, 1 hour = 0.05, 2 hours = 0.05 and 4 hours = 0.05.

since K_b is not active in the equations during relapse since $y(t)$ and C_d are zero. However, we do see a reduction in uncertainty for K_b for SP. These findings align very well with the data experiment including relapse presented in Section 4. Overall, it seems that SP has the greatest reduction in uncertainty by including relapse measurements, which might be explained by SP having the most uncertain starting point.

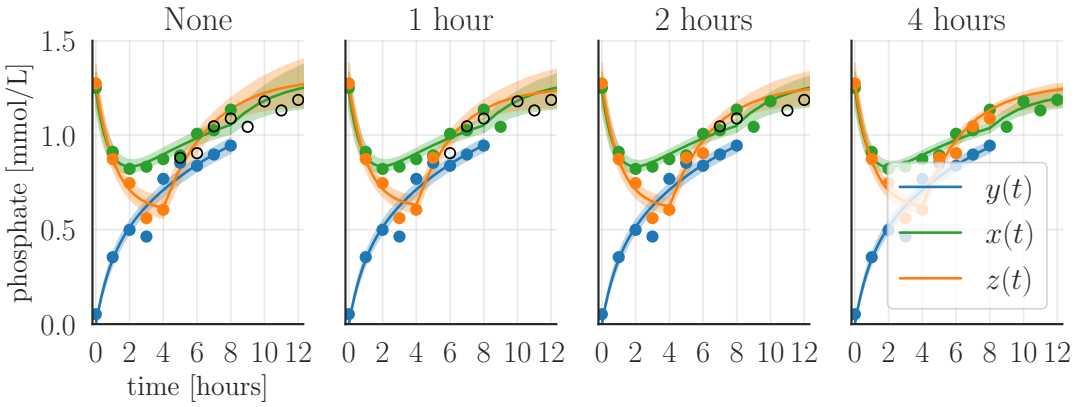

Figure E.10. Results for the simulated experiment for CP with included relapse measurements after 1, 2 and 4 hours, respectively. The measurements used for estimation are marked by full dots and the measurements left out are marked by black open circles. RMSE is None = 0.1, 1 hour = 0.1, 2 hours = 0.1 and 4 hours = 0.1.

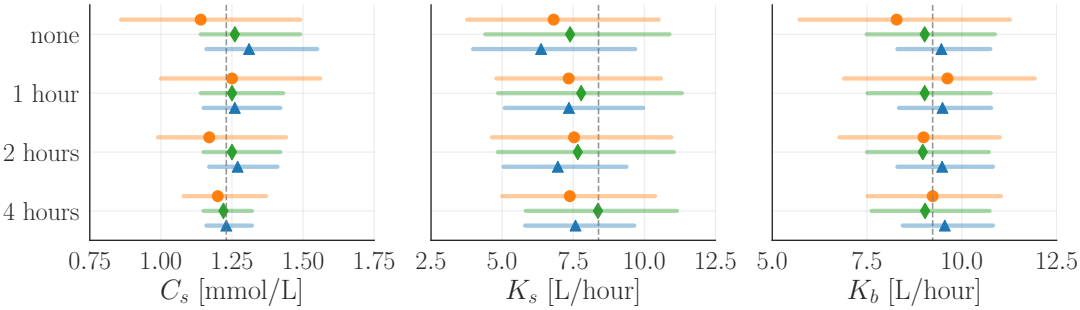

Figure E.11. Visualization of the parameter estimates and 95% CI for the estimated phosphate concentrations with relapse measurements in Figure E.8, E.9 and E.10. The vertical gray dotted lines represent the true parameter values.

E.4. Full Bayesian model with consecutive sessions

Lastly, we investigate the effect of measuring a patient for two consecutive sessions. We simulated that the patient is measured during two sessions with 18 hours between start of the first and second session. The phosphate concentration is in a steady state after 18 hours and thus, it does not affect the results if the time between the starts is larger, e.g., 48 hours. The results are shown in Figure E.12, Figure E.13 and Table E.6.

Figure E.12 shows that the uncertainty for the relapse phase, especially for SP, is greatly reduced compared to the results in Figure E.5. Similar for the parameter estimates, we see a great reduction in uncertainty for C_s and K_b for SP by comparing Figure E.6 and E.13 whereas there is only a slight reduction in the uncertainty for K_s. For MP and CP, we see a great reduction in uncertainty compared to SP for all three parameters, but that the mean estimate of K_s is slightly under estimated.

Table E.5. Median (m), lower (l) and upper (u) 95% CI and relative standard deviation ($\frac{\text{std}}{\text{mean}}$) for Figure E.8, E.9 and E.10.

HD		C_s				K_s				K_b			
		l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$
SP	None	0.87	1.13	1.49	14%	3.82	6.81	10.36	24%	5.80	8.31	11.14	16%
	1 hour	1.00	1.25	1.56	9%	4.81	7.31	10.78	22%	6.87	9.63	12.04	11%
	2 hour	1.00	1.18	1.46	10%	4.55	7.46	10.96	22%	6.88	8.99	11.07	12%
	4 hour	1.08	1.21	1.39	6%	4.84	7.26	10.35	19%	7.55	9.20	11.01	9%
MP	None	1.14	1.27	1.50	7%	4.34	7.27	10.97	23%	7.54	9.02	10.76	9%
	1 hour	1.14	1.25	1.44	6%	4.68	7.72	11.33	21%	7.5	8.97	10.81	9%
	2 hour	1.15	1.26	1.44	6%	4.70	7.55	10.99	21%	7.57	8.98	10.75	9%
	4 hour	1.15	1.22	1.33	4%	5.72	8.34	11.24	17%	7.57	8.99	10.71	9%
CP	None	1.15	1.31	1.55	8%	3.92	6.26	9.54	23%	8.34	9.46	10.66	6%
	1 hour	1.15	1.26	1.41	5%	5.16	7.38	10.16	15%	8.28	9.46	10.73	7%
	2 hour	1.17	1.27	1.41	5%	5.04	6.96	9.39	16%	8.35	9.47	10.76	7%
	4 hour	1.16	1.23	1.32	3%	5.82	7.59	9.68	13%	8.42	9.56	10.83	6%

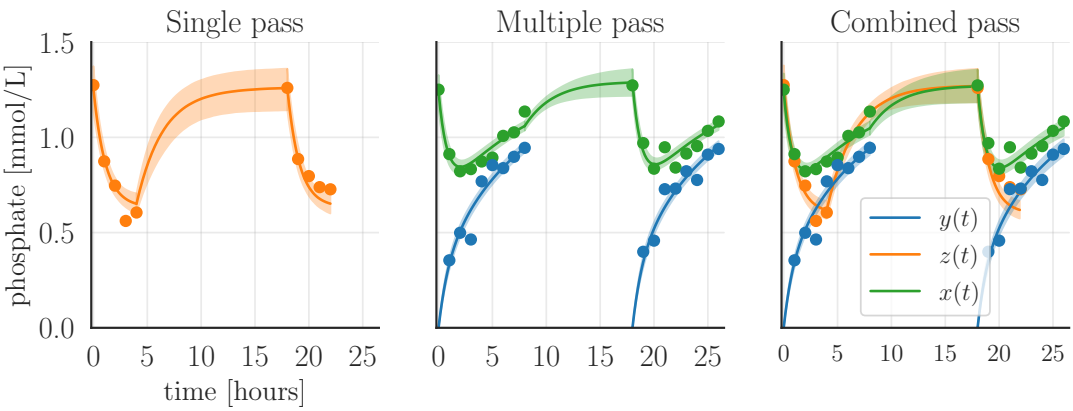

Figure E.12. Results for the simulated experiment for SP with included relapse measurements after 1, 2 and 4 hours, respectively. The measurements used for estimation are marked by full dots and the measurements left out are marked by black open circles. RMSE is SP= $8.30 \cdot 10^{-6}$, MP= $6.53 \cdot 10^{-6}$ and CP= $8.28 \cdot 10^{-6}$.

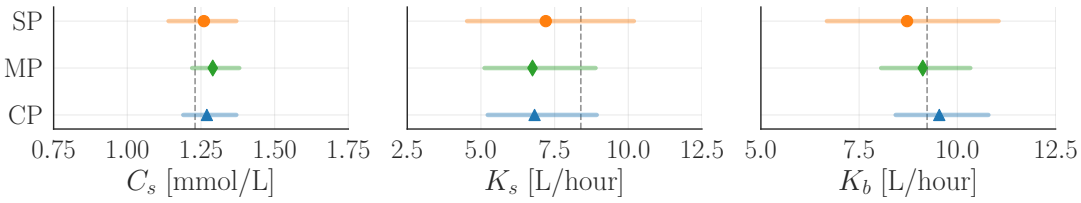

Figure E.13. Visualization of the parameter estimates and 95% CI for the estimated phosphate concentrations with consecutive sessions. The vertical gray dotted lines represent the true parameter values.

Table E.6. Median (m), lower (l) and upper (u) 95% CI and relative standard deviation ($\frac{\text{std}}{\text{mean}}$) for Figure E.12.

HD	C_s				K_s				K_b			
	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$	l	m	u	$\frac{\text{std}}{\text{mean}}$
SP	1.15	1.26	1.38	5%	4.58	7.1	10.15	20%	6.59	8.73	11.11	13%
MP	1.21	1.29	1.38	3%	5.14	6.77	9.18	15%	8.06	9.11	10.36	6%
CP	1.19	1.27	1.37	4%	5.26	6.83	8.79	13%	8.43	9.56	10.81	6%

E.5. Single-pass with non-constant phosphate concentration in the dialysate

We assume that the phosphate concentration in the dialysate for SP is constant. However, the assumption of a constant C_d is not crucial. Allowing it to be a variable with initial value 0 extends the model by an extra differential equation. The result of the extended SP model is shown in Figure E.14. This extension results in a fast transient in C_d toward the steady state value given by data with at doubling time of approximately 10-15 minutes.

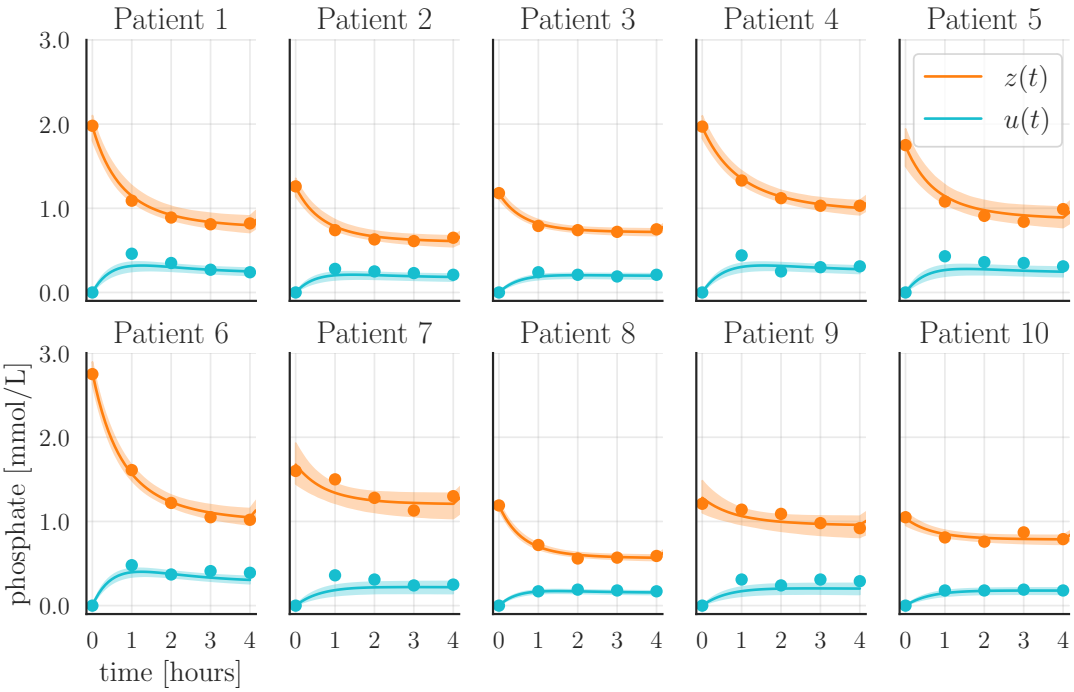

Figure E.14. Estimation of the phosphate concentration in blood ($z(t)$) for the extended SP model with non-constant phosphate concentration in the dialysate ($u(t)$). The estimation shows a fast transient in C_d toward the steady state.

Bibliography

[1] M. Andersen, K. O. Bangsgaard, J. G. Heaf, and J. T. Ottesen. "Analytical solution of phosphate kinetics for hemodialysis." In: *Journal of Mathematical Biology* 87.1 (June 2023). DOI: 10.1007/s00285-023-01942-4.

[2] Richard C. Aster, Brian Borchers, and Clifford H. Thurber. *Parameter Estimation and Inverse Problems*. Second Edition. Boston: Academic Press, 2013. ISBN: 978-0-12-385048-5. DOI: 10.1016/C2009-0-61134-X.

[3] Katrine O. Bangsgaard and Martin S. Andersen. "A statistical reconstruction model for absorption CT with source uncertainty." In: *Inverse Problems* 37.8 (July 2021), page 085009. DOI: 10.1088/1361-6420/ac11c7.

[4] Katrine O. Bangsgaard, Morten Andersen, James G. Heaf, and Johnny T. Ottesen. "Bayesian parameter estimation for phosphate dynamics during hemodialysis." In: *Mathematical Biosciences and Engineering* 20.3 (2022), pages 4455–4492. DOI: 10.3934/mbe.2023207.

[5] Katrine O. Bangsgaard, Morten Andersen, Vibe Skov, Lasse Kjær, Hans C. Hasselbalch, and Johnny T. Ottesen. "Dynamics of competing heterogeneous clones in blood cancers explains multiple observations - a mathematical modeling approach." In: *Mathematical Biosciences and Engineering* 17.6 (2020), pages 7645–7670. DOI: 10.3934/mbe.2020389.

[6] Katrine O. Bangsgaard, Genoveva Burca, Evelina Ametova, Martin Skovgaard Andersen, and Jakob Sauer Jørgensen. "Low-rank flat-field correction for artifact reduction in spectral computed tomography." In: *Applied Mathematics in Science and Engineering* 31.1 (March 2023). DOI: 10.1080/27690911.2023.2176000.

[7] Johnathan M. Bardsley. *Computational Uncertainty Quantification for Inverse Problems*. Philadelphia: Society for Industrial and Applied Mathematics, 2018. ISBN: 9781611975376. DOI: 10.1137/1.9781611975383.

[8] Johnathan M. Bardsley and Colin Fox. "An MCMC method for uncertainty quantification in nonnegativity constrained inverse problems." In: *Inverse Problems in Science and Engineering* 20.4 (2012), pages 477–498. DOI: 10.1080/17415977.2011.637208.

[9] Carlo Basile and Daniel Schneditz. "Haemodialysis adequacy monitoring for phosphate: an old problem with new solutions?" In: *Nephrology Dialysis Transplantation* 30.1 (November 2014), pages 9–11. DOI: 10.1093/ndt/gfu351.

[10] Amir Beck and Marc Teboulle. "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems." In: *SIAM Journal on Imaging Sciences* 2.1 (January 2009), pages 183–202. DOI: 10.1137/080716542.

[11] Michael Betancourt. "A Conceptual Introduction to Hamiltonian Monte Carlo." In: (2017). DOI: 10.48550/ARXIV.1701.02434.

[12] Stephen P. Boyd and Lieven Vandenberghe. *Convex optimization*. Padstow Cornwall: Cambridge University Press, 2004. ISBN: 978-0-521-83378-3.

[13] Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng, editors. *Handbook of Markov Chain Monte Carlo*. New York: Chapman and Hall/CRC, May 2011. DOI: 10.1201/b10905.

[14] Thorsten M. Buzug. *Computed Tomography*. Heidelberg: Springer Berlin, 2008. DOI: 10.1007/978-3-540-39408-2.

[15] Daniela Calvetti and Erkki Somersalo. *Introduction to Bayesian scientific computing : ten lectures on subjective computing*. New York: Springer, 2007. ISBN: 978-0-387-73393-7. DOI: 10.1007/978-0-387-73394-4.

[16] George Casella and Edward I. George. "Explaining the Gibbs Sampler." In: *The American Statistician* 46.3 (August 1992), pages 167–174. DOI: 10.1080/00031305.1992.10475878.

[17] William P. Dillon. "50th anniversary of computed tomography: past and future applications in clinical neuroscience." In: *Journal of Medical Imaging* 8.05 (October 2021). DOI: 10.1117/1.jmi.8.5.052112.

[18] Carl Eckart and Gale Young. "The approximation of one matrix by another of lower rank." In: *Psychometrika* 1.3 (September 1936), pages 211–218. DOI: 10.1007/bf02288367.

[19] Leah Edelstein-Keshet. *Mathematical Models in Biology*. Philadelphia: Society for Industrial and Applied Mathematics, 2005. ISBN: 978-0-898715-545. DOI: 10.1137/1.9780898719147.

[20] Sunny Eloot, Wim Van Biesen, Mette Axelsen, Griet Glorieux, Robert Smith Pedersen, and James Goya Heaf. "Protein-bound solute removal during extended multipass versus standard hemodialysis." In: *BMC Nephrology* 16.1 (April 2015). DOI: 10.1186/s12882-015-0056-y.

[21] Bjarne Kjær Ersbøll and Knut Conradsen. *Multivariate statistics: an introduction*. Lyngby: DTU Compute, 2016. ISBN: 9788764313154.

[22] Andrew Gelman and Donald B. Rubin. "Inference from Iterative Simulation Using Multiple Sequences." In: *Statistical Science* 7.4 (1992), pages 457–472.

[23] John Geweke. "Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments." In: (2009).

[24] Jean-Pierre Gutzwiller, Andreas R. Huber Daniel Schneditz, Christian Schindler, Felix Gutzwiller, and Carlos E. Zehnder. "Estimating phosphate removal in haemodialysis: an additional tool to quantify dialysis dose." In: *Nephrology Dialysis Transplantation* 17.6 (June 2002), pages 1037–1044. DOI: 10.1093/ndt/17.6.1037.

[25] Jacques Hadamard. *Lectures on Cauchy's Problem in Linear Partial Differential Equations*. New Haven: Yale University Press, 1923.

[26] Per Christian Hansen. *Discrete Inverse Problems*. Society for Industrial and Applied Mathematics, January 2010. ISBN: 978-0-898716-96-2. DOI: 10.1137/1.9780898718836.

[27] James Goya Heaf, Mette Axelsen, and Robert Smith Pedersen. "Multipass haemodialysis: a novel dialysis modality." In: *Nephrology Dialysis Transplantation* 28.5 (November 2012), pages 1255–1264. ISSN: 0931-0509. DOI: 10.1093/ndt/gfs484.

[28] Matthew D. Hoffman and Andrew Gelman. "The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo." In: *Journal of Machine Learning Research* 15 (2014), pages 1593–1623.

[29] Frank C. Hoppensteadt and Charles S. Peskin. *Modeling and Simulation in Medicine and the Life Sciences*. New York: Springer, 2002. ISBN: 978-0-387-95072-3. DOI: 10.1007/978-0-387-21571-6.

[30] O. Susan Iheagwara, Todd S. Ing, Carl M. Kjellstrand, and Susie Q. Lew. "Phosphorus, phosphorous, and phosphate." In: *Hemodialysis International* 17.4 (December 2012), pages 479–482. DOI: 10.1111/hdi.12010.

[31] Jakob Jorgensen, Evelina Ametova, Edoardo Pasca, Gemma Fardell, Alex Liptak, Daniil Kazantsev, William Lionheart, Martin Turner, Genoveva Burca, and Vaggelis Papoutsellis. *Neutron TOF imaging phantom data to quantify hyperspectral reconstruction algorithms*. 2019. DOI: 10.5286/ISIS.E.RB1820541.

[32] Anosh Joseph. "Hybrid (Hamiltonian) Monte Carlo." In: *Markov Chain Monte Carlo Methods in Quantum Field Theories: A Modern Primer*. Cham: Springer International Publishing, 2020, pages 59–69. ISBN: 978-3-030-46044-0. DOI: 10.1007/978-3-030-46044-0_7.

[33] Winfried Kockelmann, Triestino Minniti, Daniel Pooley, Genoveva Burca, Ranggi Ramadhan, Freddie Akeroyd, Gareth Howells, Chris Moreton-Smith, David Keymer, Joe Kelleher, Saurabh Kabra, Tung Lee, Ralf Ziesche, Anthony Reid, Giuseppe Vitucci, Giuseppe Gorini, Davide Micieli, Raffaele Agostino, Vincenzo Formoso, Francesco Aliotta, Rosa Ponterio, Sebastiano Trusso, Gabriele Salvato, Cirino Vasi, Francesco Grazzi, Kenichi Watanabe, Jason Lee, Anton Tremsin, Jason McPhate, Daniel Nixon, Nick Draper, William Halcrow, and Jim Nightingale. "Time-of-Flight Neutron Imaging on IMAT-@ISIS: A New User Facility for Materials Science." In: *Journal of Imaging* 4.3 (February 2018), page 47. DOI: 10.3390/jimaging4030047.

[34] Martin K. Kuhlmann. "Phosphate Elimination in Modalities of Hemodialysis and Peritoneal Dialysis." In: *Blood Purification* 29.2 (2010), pages 137–144. DOI: 10.1159/000245640.

[35] Randall J. LeVeque. *Finite Difference Methods for Ordinary and Partial Differential Equations*. Philadelphia: Society for Industrial and Applied Mathematics, 2007. DOI: 10.1137/1.9780898717839.

[36] Jennifer L. Mueller and Samuli Siltanen. *Linear and Nonlinear Inverse Problems with Practical Applications*. Philadelphia: Society for Industrial and Applied Mathematics, 2012. ISBN: 978-1-611972-33-7. DOI: 10.1137/1.9781611972344.

[37] Beat Münch, Pavel Trtik, Federica Marone, and Marco Stampanoni. "Stripe and ring artifact removal with combined wavelet—Fourier filtering." In: *Optics Express* 17.10 (2009), page 8567. DOI: 10.1364/OE.17.008567.

[38] Art B. Owen. *Monte Carlo theory, methods and examples*. 2013. URL: https://artowen.su.domains/mc/.

[39] Lawrence Perko. *Differential Equations and Dynamical Systems*. New York: Springer, 2001. ISBN: 978-0-387-95116-4. DOI: 10.1007/978-1-4613-0003-8.

[40] Jim Pitman. *Probability*. New York: Springer, 1993. ISBN: 0-387-97974-3.

[41] Johann Radon. "On the determination of functions from their integral values along certain manifolds." In: *IEEE Transactions on Medical Imaging* 5.4 (December 1986), pages 170–176. DOI: 10.1109/tmi.1986.4307775.

[42] Christian P. Robert and George Casella. *Monte Carlo Statistical Methods*. New York: Springer, 2004. ISBN: 978-1-4419-1939-7. DOI: 10.1007/978-1-4757-4145-2.

[43] Ronald D Ruth. "A canonical integration technique." In: *IEEE Trans. Nucl. Sci.* 30 (1983), pages 2669–2671. DOI: 10.1109/TNS.1983.4332919.

[44] Adrian F. M. Smith and Gareth O. Roberts. "Bayesian Computation Via the Gibbs Sampler and Related Markov Chain Monte Carlo Methods." In: *Journal of the Royal Statistical Society: Series B (Methodological)* 55.1 (September 1993), pages 3–23. DOI: 10.1111/j.2517-6161.1993.tb01466.x.

[45] Ralph C. Smith. *Uncertainty Quantification: Theory, Implementation, and Applications*. Philadelphia: Society for Industrial and Applied Mathematics, 2014. ISBN: 978-1-611973-21-1.

[46] Rajesh Subramanian and Romesh Khardori. "Severe Hypophosphatemia: Pathophysiologic Implications, Clinical Presentations, and Treatment." In: *Medicine* 79.1 (January 2000), pages 1–8. DOI: 10.1097/00005792-200001000-00001.

[47] Aixin Tan and James P. Hobert. "Block Gibbs Sampling for Bayesian Random Effects Models With Improper Priors: Convergence and Regeneration." In: *Journal of Computational and Graphical Statistics* 18.4 (2009), pages 861–878. DOI: 10.1198/jcgs.2009.08153.

[48] Lloyd N. Trefethen and David Bau. *Numerical Linear Algebra*. Philadelphia: Society for Industrial and Applied Mathematics, 1997.

[49] Nghia T. Vo, Robert C. Atwood, and Michael Drakopoulos. "Superior techniques for eliminating ring artifacts in X-ray micro-tomography." In: *Optics Express* 26.22 (2018), page 28396. DOI: 10.1364/OE.26.028396.

{2.7182818284}