



TNF

Technisch-Naturwissenschaftliche
Fakultät

Problem-adapted Regularization for Inverse Problems in the Deterministic and Stochastic Setting

DISSERTATION

zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften

im Doktoratsstudium der

Naturwissenschaften

Eingereicht von:

Dipl.-Math. techn. Daniel Gerth

Angefertigt am:

Doktoratskolleg Computational Mathematics

Beurteilung:

O.Univ.-Prof. Dipl.-Math. Dr. Ronny Ramlau (Betreuung)

Prof. Dr. Samuli Siltanen

Linz, September, 2015

Acknowledgments

First and foremost I would like to thank my adviser Ronny Ramlau for the supervision during my time in Linz. His support, ideas and suggestions helped me greatly all the way until eventually writing this thesis. I am very thankful for the freedom he allowed and trustfulness I felt. The lively and friendly atmosphere he and all the members of his institute created made me feel comfortable very quickly. He takes the main share in making me feel part of the Inverse Problems community and seeking to continue working in the academic environment.

I am very grateful for being able to attend conferences, go abroad and work with other professors. Each of the stays not only broadened my scientific background, but also was a great time to discover new places and people. I would like to thank Prof. Samuli Siltanen from University of Helsinki, Prof. Lothar Reichel from Kent State University, Prof. Alfred K. Louis from Saarland University and Prof. Chein-Shan Liu from National Taiwan University for their hospitality and support during my research stays.

The many friendly people I met in Linz soon made me feel at home in the new town. Out of the many that deserve my gratitude, I can only name a few here. I would like to thank Gábor Kiss for his friendship and the countless games and teas we had together; Kirk Soodhalter for his exquisite taste in bad movies and good food and answering my questions about the English language; Daniela Saxenhuber for patiently explaining Adaptive Optics and, together with Roland Wagner and Markus Ableidinger, assisting me in handling the students of my lectures.

Without the support of my family and friends I would not be where I am today. The continuous and unconditional support from my parents Uwe and Andrea allowed me to pursue this road. The lasting friendship of Gregor and Michael Habeck is very much appreciated. Last, but certainly not least, I want to thank Meng-Chun Lee for her love and and truly being my “second half”.

This thesis was created during my employment at the Doctoral Program “Computational Mathematics” (W1214) at the Johannes Kepler University in Linz. The financial support by the Austrian Science Fund (FWF) under grant W1214-N15, project DK08, is gratefully acknowledged.

KURZFASSUNG

Die Modellierung des unvermeidlichen Messfehlers ist von größter Bedeutung für die theoretische und praktische Behandlung Inverser und schlecht-gestellter Probleme. Man betrachtet entweder einen deterministischen, größtmöglichen Fehler oder eine stochastische Formulierung. Beide Varianten haben Vor- und Nachteile. Einige frühere Veröffentlichungen haben Verbindungen zwischen dem Modellen angedeutet wenn in den Annahmen der deterministischen Theorie ein stochastisches Fehlermodell angenommen wird. Ein Hauptanliegen dieser Arbeit ist es, diese Verbindung weiter darzulegen. Als zweites Hauptthema dient eine spezielle Anwendung. Es werden zwei Regularisierungsmethoden für das Problem der atmosphärischen Tomografie für Adaptive Optik untersucht. Nach der Klärung des Konzepts Inverser und Schlecht-gestellter Probleme werden grundlegende Regularisierungsmethoden im deterministischen und stochastischen Modell diskutiert und ein Überblick über die vorhandene Literatur zur Verbindung der Modelle gegeben. Vor der Präsentation eigener Ergebnisse, werden einige stochastische Konzepte eingeführt, insbesondere die Ky-Fan-Metrik als Haupttechnik der späteren Analyse. Die ersten eigenen Resultate erfolgen für Filter-basierte Regularisierungsmethoden. Bevor Konvergenz und Konvergenzraten dieser Methoden unter dem stochastischen Modell gezeigt werden, werden zwei spezielle Filtermethoden im deterministischen Modell diskutiert. Diese, beide genannt fractional Tikhonov regularization, wurden eingeführt um das Überglätten der traditionellen Tikhonov-Regularisierung zu vermeiden. Es wird gezeigt, dass beide nur in Spezialfällen eine Verbesserung bringen obwohl beide sowohl mit einer a-priori Parameterwahl als auch mit dem Diskrepanzprinzip von optimaler Konvergenzordnung sind. Ausgewählte numerische Ergebnisse untermauern die theoretischen Erkenntnisse. Zur Vorbereitung späterer Kapitel wird dann das Problem der atmosphärischen Tomografie für Adaptive Optik eingeführt. Dies ist ein System zur Verbesserung der Bildqualität von auf der Erde stationierter Teleskope. Darauf folgend wird zu einer wavelet-sparsity basierten Regularisierungsmethode in Besov-Räumen übergeleitet. Das deterministische Tikhonov-Funktional wird aus einer rein stochastischen Formulierung motiviert, nämlich dem Bayes'schen Ansatz. Konvergenz und Konvergenzraten unter dem stochastischen Modell werden gezeigt. Durch die Bayes'sche Formulierung werden alle Größen als Zufallsvariablen modelliert. Dies erlaubt die Konstruktion einer neuartigen a-priori Parameterwahlregel. Nach einem akademischen numerischen Beispiel wenden wir die Theorie auf die Atmosphärische Tomografie an, da darin die Unbekannte als Zufallsvariable in einem Besov-Raum interpretiert werden kann. Danach werden die bisherigen Ergebnisse zur Verbindung zwischen deterministischem und stochastischen Fehlermodell zusammengefasst und verallgemeinert. Im numerischen Beispiel nutzen wir das Diskrepanzprinzip zur Bestimmung des Regularisierungsparameters für ein nichtlineares Faltungsproblem mit sparsity-Strafterm. Im letzten Teil der Arbeit stellen wir einen auf der Methode der Approximativen Inversen beruhenden Algorithmus für das atmosphärische Tomografieproblem vor. In dieser Regularisierungsstrategie wird der Großteil des Rechenaufwands zur Berechnung sogenannter Rekonstruktionskerne benutzt. Zur Laufzeit müssen nur innere Produkte zwischen Daten und Rekonstruktionskernen ausgewertet werden. Numerische Fallbeispiele zeigen die Konkurrenzfähigkeit dieser Methode.

ABSTRACT

In theory and application of Inverse and Ill-posed Problems, the modeling of the inevitable measurement noise is of utmost importance. One either works with a worst case deterministic error model or a stochastic one. Both approaches have their respective advantages and disadvantages. Several works have hinted at connections between both theories when the deterministic assumptions hold except for a stochastic noise model. A main concern of this thesis is to further elaborate this connection. A particular application is the second main thread of the work. Namely, we investigate two regularization approaches for the problem of Atmospheric Tomography in Adaptive Optics. After clarification of the concept of Inverse and Ill-posed Problems we discuss basic regularization strategies and results in both the deterministic and stochastic setting. We then proceed to review the existing literature in regard to connections between the two settings. Before presenting our own contributions, we introduce some stochastic concepts, in particular the Ky Fan metric which is a main ingredient for our stochastic analysis. We also discuss its connection to convergence in expectation. Our first results are on convergence properties of filter-based regularization methods. Before we show convergence and convergence rates of this type of methods under the stochastic noise assumption, we discuss two particular filter based methods, both called fractional Tikhonov regularization, in the deterministic setting. Being designed to reduce the oversmoothing of traditional Tikhonov regularization, we show that only in special cases they are superior to the standard form although both are of optimal convergence order with an a priori parameter choice as well as with the discrepancy principle. Selected numerical examples are presented to experimentally verify the theoretical findings. In preparation of later parts of the thesis we then introduce the topic of Atmospheric Tomography for Adaptive Optics, a hardware system used to correct for image perturbations in large earth-bound telescopes caused by atmospheric turbulence. Following this we move to the analysis of a wavelet-sparsity promoting regularization method in Besov spaces. The deterministic Tikhonov-type functional is motivated from a purely stochastic, namely Bayesian, point of view and convergence and convergence rates in the stochastic setting are proven. In particular, in this approach the unknown is modeled as a random variable. We use this construction to introduce a novel a priori parameter choice rule. After an academic numerical example, we apply to the results to Atmospheric Tomography as the unknown in this problem can be regarded as a Besov-space valued random variable which fits precisely in the theoretical setting. We proceed by summarizing and generalizing our results on the lifting of deterministic convergence results into the stochastic setting for general regularization methods. In our numerical example we use the discrepancy principle to find the regularization parameter for a nonlinear problem with sparsity constraints. In the last part of the thesis we present an algorithm for Atmospheric Tomography based on the method of the Approximate Inverse. In this regularization strategy, most of the computational effort is moved to the computation of so called reconstruction kernels. A fast algorithm is expected as the run-time procedure consists solely of evaluating inner products between the data and the precomputed reconstruction kernels. In our numerical case studies this method leads to results of the same quality as state-of-the-art techniques.

Contents

1	Regularization of Inverse and Ill-posed problems	1
1.1	Inverse and Ill-posed Problems	2
1.2	Deterministic noise model	4
1.2.1	The generalized inverse	4
1.2.2	Regularization theory	5
1.3	Stochastic regularization theory	8
1.3.1	On the noise model	8
1.3.2	The Bayesian approach	9
1.3.3	Other stochastic approaches	13
1.4	The Ky Fan metric	14
2	Introduction to Adaptive Optics	19
2.1	Introduction to Adaptive Optics	20
2.2	Optical imaging through an ideal telescope	21
2.3	Atmospheric turbulence processes	24
2.3.1	Exposure times	26
2.4	Principle of Adaptive Optics	27
2.5	Components in Adaptive Optics	28
2.5.1	Wave front sensors and wavefront reconstruction	28
2.5.2	Deformable Mirrors and Temporal Control	29
2.5.3	Guide Stars	31
2.6	Multi Conjugate Adaptive Optics	33
2.6.1	Introduction	33
2.6.2	Mathematical model of Atmospheric Tomography for MCAO	35
2.7	Remarks on Numerical Simulation	37
3	Filter-based reconstruction methods	39
3.1	Regularization methods and filter factors	40
3.2	Fractional Tikhonov regularization	43
3.2.1	Fractional Tikhonov regularization in the literature	43
3.2.2	Order optimality of fractional Tikhonov methods	45
3.2.3	Numerical examples	51
3.3	Filter-based regularization methods in the stochastic setting	59
3.3.1	Convergence results	60

3.3.2	Numerical Results	67
4	A sparsity promoting regularization method in the stochastic setting	71
4.1	Sparsity promoting regularization methods	72
4.1.1	Introduction and short survey	72
4.2	From Bayes to Tikhonov	73
4.2.1	Wavelets	75
4.2.2	Besov Spaces and random variables therein	77
4.3	Convergence of maximum a posteriori solutions with Besov priors	83
4.3.1	Review of a deterministic result as basis for the lifting approach	83
4.3.2	Convergence in the stochastic setting	85
4.4	Numerical examples for $p = 1$	91
4.5	Application to Atmospheric Tomography	94
5	Summary and generalization of the lifting strategy	101
5.1	Theoretical results	102
5.2	An example in theory and numerical results	105
6	The method of the Approximate Inverse for Adaptive Optics	109
6.1	Method of the Approximate Inverse	110
6.1.1	Reconstruction of scalar functions	110
6.1.2	The approximate inverse and filter-based methods	112
6.1.3	Reconstruction of vector-valued functions	113
6.2	Application to Atmospheric Tomography	115
6.2.1	Continuation of the operator A	116
6.2.2	Calculation of the reconstruction kernel	121
6.2.3	The extension error	123
6.3	Numerical Results	125
6.3.1	Reconstruction Method	125
6.3.2	Simulation results	126
6.3.3	Computational complexity	129
7	Conclusions	133
	Bibliography	135

Chapter 1

Regularization of Inverse and Ill-posed problems

1.1 Inverse and Ill-posed Problems

The study of Inverse and Ill-posed Problems is a branch of mathematics which emerged after the publications of A. Tikhonov [1, 2] in the middle of the 20th century. Inverse Problems are strongly connected to applications in science and technology. One is often interested in a certain quantity, the *cause*, that is not directly accessible, but of which indirect measurements, the *effect*, are available. For example, in computerized tomography (see, e.g., [3]) one images cross sections of a human body to help with the diagnosis of many illnesses. In order to not having to open the patient's body to directly observe the cause, one sends X-ray beams through the patient and measures the intensity of the beams after exiting the body. The task then is to recover the interior of the body from the measured data. The relation or *model* that connects cause and effect is assumed to be known. Throughout most parts of this work the cause will be denoted by x , the effect by y and the model by an operator A , such that the three quantities are connected by an equation

$$Ax = y. \tag{1.1}$$

Here, x and y are assumed to be elements of suitable spaces \mathcal{X} and \mathcal{Y} , respectively, which will be defined more precisely in the description of the individual sections. In most cases, we will consider Hilbert spaces and a bounded linear operator A mapping from \mathcal{X} into \mathcal{Y} , i.e., $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ fulfills

$$\begin{aligned} A(x_1 + x_2) &= Ax_1 + Ax_2 \quad \forall x_1, x_2 \in \mathcal{X}, \\ A(\lambda x) &= \lambda Ax \quad \forall x \in \mathcal{X}, \lambda \in \mathbb{R}, \\ \|Ax\|_{\mathcal{Y}} &\leq C\|x\|_{\mathcal{X}} \quad \forall x \in \mathcal{X} \text{ with } C > 0. \end{aligned}$$

An Inverse Problem, in its literal meaning, is to find the cause given the effect, in contrast to the direct or forward problem which is to determine the effect given the cause. Often, the Inverse Problem is much more difficult to solve than the forward problem due to the inversion of the model A . This issue becomes even more severe when the problem is *ill-posed*. The following definition is due to Hadamard [4].

Definition 1.1.1. An equation $Ax = y$ is called *well-posed*, if the following three conditions are satisfied:

- (Existence) for every $y \in \mathcal{Y}$ there exists an $x \in \mathcal{X}$ with $Ax = y$
- (Uniqueness) for every $y \in \mathcal{Y}$ there is exactly one $x \in \mathcal{X}$ with $Ax = y$
- (Stability) the solution $x \in \mathcal{X}$ depends continuously on the data $y \in \mathcal{Y}$

If any of the above is violated, the problem is called *ill-posed*.

Throughout this work we will drop the explicit reference to ill-posed problems and assume that, whenever we speak of an Inverse Problem, it is ill-posed. The difficulty of inversion of Inverse Problems is in particular due to the third condition of Definition 1.1.1. If the solution does not depend continuously on the data, small perturbations of the measured data may lead to arbitrarily large deviations in the solution. This is problematic as measured data naturally contains some kind of noise since any measurement device can only operate up to a certain precision. In some cases, the model of the forward problem A might only be known approximately, which can be interpreted as a kind of noise. Additionally, the measurements typically do not represent the effect y itself which often is an infinite dimensional quantity, but a discretized version. One may directly model the effect of discretization as additional noise [5, 6]. Returning to the example of Computerized Tomography, this corresponds to the fact that there is only a finite amount of sensors for the intensity of the X-ray beams, setting a limit to the amount of data one may acquire. The convention throughout this work is to assume an additive error to the true data, i.e.,

$$y^{\text{measured}} = y + \epsilon, \quad (1.2)$$

where y from (1.1) is the true data to the unknown object of interest x and ϵ represents noise in the measurements. Throughout the whole work we assume that $y \in \mathcal{R}(A)$ where $\mathcal{R}(A)$ denotes the range of the operator A such that for unperturbed data (1.1) is solvable. Although the actual value of ϵ is unknown, typically some information is available to estimate the magnitude of the error. In Inverse Problems, two approaches are common. On one hand, there is the *deterministic* approach where a worst case error bound on the magnitude of the error is assumed, i.e., there is a real number $\delta > 0$ such that $\|y^\delta - y\| \leq \delta$ where y^δ is the measured, noisy data and $\|\cdot\|$ a suitable norm. Due to this assumption, it is often possible to find strict upper bounds on the quality of the solutions. Analogously to y^δ , all error-dependent quantities will be denoted with an index δ whenever we are operating in the deterministic setting. In opposition to the worst-case scenario, in the *stochastic* setting noise models are used on the base of, as the name suggests, stochastic information. Typically one has an idea about the stochastic distribution of the noise and can approximately quantify properties like expectation and variance. The stochastic model allows arbitrarily large error magnitudes, but with low probability. As a result, also estimates about the quality of the solutions have to be given in a probabilistic formulation since it may be that the noise in the data is so large that it completely overshadows the true data y . In this work, the stochastic aspect is assumed to be purely in the noisy data in contrast to the genuine stochastic setting which also allows for randomness in other quantities such as the solution x , data y or the operator A . Stochastic algorithms often include repeated sampling of random variables which is computationally expensive. On the other hand, the error parameters in the stochastic setting are typically easier to obtain than a sharp worst-error bound for the deterministic theory. A main concern of this work is to use the stochastic information of

the noise with reconstruction methods from the deterministic setting. In order to clearly distinguish the different settings, all quantities involving stochastic noise models will be denoted by a superscript η , e.g., the noisy data is denoted by y^η .

In the following sections, we will address the two approaches separately.

1.2 Deterministic noise model

This section reviews definitions and properties of deterministic regularization methods; see, e.g., [7, 8, 9, 10] for further details. For simplicity, we will stick to the scenario that \mathcal{X} and \mathcal{Y} are Hilbert spaces.

1.2.1 The generalized inverse

Due to the ill-posedness of the Inverse Problem, a straight forward inversion of (1.1) is not feasible, in particular when the right hand side is contaminated with noise. In order to include elements $y \in \mathcal{Y}$ that are not in the range of the operator, it is natural to find an approximate solution to (1.1) by minimizing the residual functional

$$J(x) := \|Ax - y\|_{\mathcal{Y}}^2. \quad (1.3)$$

Denoting the orthogonal projection of \mathcal{Y} onto $\overline{\mathcal{R}(A)}$ by $P_{\overline{\mathcal{R}(A)}}$ one can split the residual into two parts,

$$J(x) = \|Ax - P_{\overline{\mathcal{R}(A)}}y\|_{\mathcal{Y}}^2 + \|y - P_{\overline{\mathcal{R}(A)}}y\|_{\mathcal{Y}}^2.$$

If $y \in \mathcal{R}(A) \oplus \mathcal{R}(A)^\perp$, the first part is solvable exactly while the second part is independent of x . Hence, the minimizer of $J(x)$ is given by the solution of

$$Ax = P_{\overline{\mathcal{R}(A)}}y.$$

If this equation has more than one solution, we select among all of them the one with minimal norm and denote it by x^\dagger , i.e.,

$$\|x^\dagger\| < \|\bar{x}\| \quad \text{for all } \bar{x} \neq x^\dagger : J(x^\dagger) = J(\bar{x}) = \min_x J(x).$$

The element x^\dagger is called *minimum norm solution*. It is unique and exists whenever $y \in \mathcal{R}(A) \oplus \mathcal{R}(A)^\perp$ [10]. Note that one may also define the minimum norm solution with respect to a fixed element $x^* \in \mathcal{X}$, for example when a reasonable estimate for the true solution is available; or an (x^*-) minimum norm solutions with respect to appropriate functionals of x other than the norm in \mathcal{X} .

Definition 1.2.1. The operator

$$A^\dagger : \mathcal{D}(A^\dagger) := \mathcal{R}(A) \oplus \mathcal{R}(A)^\perp \subset \mathcal{Y} \rightarrow \mathcal{X}$$

that maps y to x^\dagger is called the *generalized inverse* or *Moore-Penrose inverse* of A .

It can be shown that $x^\dagger = A^\dagger y$ is the unique solution of the normal equation

$$A^*Ax = A^*y \quad (1.4)$$

in $\mathcal{N}(A)^\perp$ where A^* is the Hilbert-space adjoint, i.e.,

$$\langle Ax, y \rangle = \langle x, A^*y \rangle \quad \forall x \in \mathcal{X}, y \in \mathcal{Y}.$$

Here $\mathcal{N}(A)^\perp$ denotes the orthogonal complement of the null space of A . We collect important properties of the generalized inverse in a theorem.

Theorem 1.2.1. *Let $A \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$. Then*

- $\mathcal{N}(A^\dagger) = \mathcal{R}(A)^\perp$,
- $\mathcal{R}(A^\dagger) = \mathcal{N}(A)^\perp = \overline{\mathcal{R}(A^*)}$,
- A^\dagger is linear,
- A^\dagger is continuous if and only if $\mathcal{R}(A)$ is closed.

The generalized inverse, by construction, deals with the first two conditions in Definition 1.1.1. Due to the last point of the previous theorem, however, the problem of stability remains whenever the range of A is not closed. We will assume that this is the case throughout the whole work. A particular class of operators for which this is true are compact operators. An operator A is called *compact* when it maps bounded sets in \mathcal{X} onto relatively compact sets in \mathcal{Y} . Integral operators are often compact between L_2 -spaces or the space of continuous functions. In the Chapters 3 and 4 we use compact integral operators in our numerical examples. A non-compact integral operator is the autoconvolution operator from Chapter 5. In order to deal with the issue of instability, the generalized inverse has to be approximated by another operator that allows continuous dependence of the solutions on the data. Such so called *regularizations* will be the topic of the next section.

1.2.2 Regularization theory

Definition 1.2.2. A *regularization method* for A^\dagger is a family of operators

$$\{R_\alpha\}_{\alpha>0}, \quad R_\alpha : \mathcal{Y} \rightarrow \mathcal{X}$$

with the following properties: There is a mapping $\alpha : \mathbb{R}_+ \times \mathcal{Y} \rightarrow \mathbb{R}_+$ such that for all $y \in \mathcal{D}(A^\dagger)$ and all $y^\delta \in \mathcal{Y}$ with $\|y - y^\delta\|_{\mathcal{Y}} \leq \delta$, it holds

$$\lim_{\delta \rightarrow 0} R_{\alpha(\delta, y^\delta)} y^\delta = A^\dagger y.$$

Here α is called the *regularization parameter* which has to fulfill

$$\lim_{\delta \rightarrow 0} \alpha(\delta, y^\delta) = 0.$$

If α does not depend on y^δ it is called an *a-priori* parameter choice, otherwise an *a-posteriori* parameter choice.

In the deterministic theory, the quality of a regularization method is determined by the asymptotics of $\|x^\dagger - R_\alpha y^\delta\|_{\mathcal{X}}$ as $\delta \rightarrow 0$. Convergence rates can only be achieved under additional assumptions on the solution as in general the convergence can be arbitrarily slow. For our analysis, we mostly assume a Hölder-type smoothness assumption, i.e., that the minimal norm solution x^\dagger of the error-free problem (1.1) satisfies a smoothness condition of the form

$$x^\dagger \in \mathcal{R}((A^*A)^{\nu/2}) \quad \text{with} \quad \|x^\dagger\|_\nu := \{\|z\|_{\mathcal{X}} : x^\dagger = (A^*A)^{\nu/2}z, z \in \mathcal{N}(A)^\perp\} \leq \varrho \quad (1.5)$$

for some constant $\varrho > 0$. An alternative representation of the ν -norm for a certain class of operators is given in (3.3). One can show that the worst case error

$$E_\nu(\delta, \varrho, R_\alpha) := \sup\{\|R_\alpha y^\delta - A^\dagger y\|_{\mathcal{Y}} : \|y - y^\delta\|_{\mathcal{Y}} \leq \delta, \|A^\dagger y\|_\nu \leq \varrho\}$$

is bounded by

$$E_\nu(\delta, \varrho, R_\alpha) \leq \delta^{\frac{\nu}{\nu+1}} \varrho^{\frac{1}{\nu+1}} \quad (1.6)$$

where there exists a sequence $\{\delta_k\}_{k \in \mathbb{N}}$, $\delta_k \rightarrow 0$ as $k \rightarrow \infty$ such that equality holds. In other words, (1.6) is the best possible guaranteed rate of decay of the worst case error. Finding a regularization method fulfilling (1.6) is difficult. In practice, one is usually content when the optimal rate is achieved up to some constant. A regularization method is said to be *order optimal* if there is a constant c independent of δ and ϱ such that

$$\|R_\alpha y^\delta - R_\alpha y^\delta\|_{\mathcal{X}} \leq c \delta^{\frac{\nu}{\nu+1}} \varrho^{\frac{1}{\nu+1}}. \quad (1.7)$$

The main task in Inverse Problems is to design regularization methods in the sense of Definition 1.2.2. One is always interested in guaranteeing a decay of the regularization error $\|R_\alpha y^\delta - R_\alpha y^\delta\|$ as fast as possible, i.e., to hit the convergence rate in (1.7) in the optimal case. We will now review some standard regularization methods in preparation of the following chapters.

Since for Inverse Problems the minimization of the residual (1.3) need not lead to meaningful solutions, a common remedy is to alter the problem and approximate (1.3) by a “close” problem

$$x_\alpha^\delta = \min_{x \in \mathcal{D}(A)} \|Ax - y^\delta\|_{\mathcal{Y}}^2 + \alpha \Theta(x) \quad (1.8)$$

where $\Theta(x)$ is a proper *penalty functional*, for example a norm in a function space (which not necessarily needs to be the one in \mathcal{X}). The residual term $\|Ax - y^\delta\|_{\mathcal{Y}}^2$ may be replaced

to fit certain noise types. For example in [11, 12] arbitrary norm powers $\|Ax - y^\delta\|_Y^q$, $q > 0$, are considered, in [13] the Kullback-Leibler distance is used. We will stick to the squared residual for the largest parts of this work. The regularization parameter α has to balance between the residual and the penalty functional: if α is too small, the penalty has little effect and the problem becomes unstable again; if α is too large, the penalty dominates and the residual becomes irrelevant. One can also interpret α as a Lagrange multiplier [7]. Tikhonov [1, 2] introduced the method (1.8) with $\Theta(x) = \|x\|_2^2$. Therefore regularization methods of type (1.8) are called *Tikhonov-type regularizations*. We will refer to Tikhonov's classical formulation

$$x_\alpha^\delta = \min_{x \in \mathcal{D}(A)} \|Ax - y^\delta\|^2 + \alpha \|x\|^2 \quad (1.9)$$

as *Tikhonov regularization in standard form*. The minimizer is given by the solution of a stabilized normal equation

$$(A^*A + \alpha I)x_\alpha^\delta = A^*y^\delta. \quad (1.10)$$

It is well known that Tikhonov regularization in standard form is an order optimal method, see, e.g., [8]. We will discuss this approach and some convergence properties in more detail in Chapter 3.

For some Tikhonov-type functionals with a penalty term other than the standard norm $\|\cdot\|_2$ a normal equation similar to (1.10) can be found, see the example of generalized Tikhonov regularization in Section 3.2. For other penalty terms, iterative methods are used to obtain a minimizer. An example can be found in Chapter 4 where we will use a Besov-space penalty term. In order to minimize the respective Tikhonov functional, we use an iterative soft-shrinking algorithm (Section 4.4). This algorithm is connected to the Landweber method which we shall briefly discuss now.

The Landweber method is an iterative algorithm which can be motivated from several points of view. One may, for example rewrite the normal equation (1.4) into a fixed point equation or construct a gradient method for the minimization of the residual (1.3) with the observation that $\nabla_x(\|Ax - y\|^2) = 2(A^*Ax - A^*y)$ [7, 10]. In any way, starting from an initial guess x_0 , one obtains the iterative procedure

$$x_{k+1} = x_k + \gamma A^*(y - Ax_k), \quad k = 0, 1, 2, \dots \quad (1.11)$$

where $0 < \gamma < 2/\|A\|^2$ is a stepsize parameter. For an initial value $x_0 \in \mathcal{X}$ the method converges to $A^\dagger y + P_{\mathcal{N}(A)}x_0$ for $y \in \mathcal{D}(A^\dagger)$. However, in the form given by (1.11) it is not a regularization method. When only noisy data is available, the algorithm at first converges to a meaningful approximation of the true solution but at some point starts to diverge. Stopping the iteration early enough, on the other hand, furnishes a regularization method that is even of optimal order as the following theorem from [7] shows. In this case, the stopping index takes the role of the regularization parameter.

Theorem 1.2.2. [7, Theorem 4.3.3.] *The Landweber method (1.11) with $0 < \gamma < 2/\|A\|^2$ is a linear regularization method. It is of optimal order (1.7) for all $\nu > 0$ if the iteration is stopped after k_* steps where*

$$k_* = \mathbf{floor} \left(\gamma \left(2 \frac{\gamma}{\nu} e \right)^{\frac{\nu}{\nu+1}} \left(\frac{\rho}{\delta} \right)^{\frac{2}{\nu+1}} \right).$$

Here, $\mathbf{floor}(\cdot)$ maps a real number to the next smallest integer.

In practice, an a-priori parameter choice is often impracticable and a-posteriori rules are used. A common strategy is the so called *discrepancy principle* which can also be used to determine a regularization parameter in non-iterative methods such as Tikhonov regularization in standard form. The idea is to find a regularization parameter such that the residual obtained with the regularized solution is of the same order as the magnitude of the noise. Namely, for iterative methods one chooses k_* such that

$$\|Ax_{k_*}^\delta - y^\delta\|_{\mathcal{Y}} \leq \tau\delta < \|Ax_k^\delta - y^\delta\|_{\mathcal{Y}}, \quad k = 0, 1, \dots, k_* - 1 \quad (1.12)$$

where $\tau > 1$ is fixed. This definition is common; a slightly different formulation is used in Chapter 3. One can show, c.f. [10], that the Landweber method with the discrepancy principle (1.12) is an order optimal regularization method for all $\nu > 0$.

1.3 Stochastic regularization theory

1.3.1 On the noise model

Before addressing regularization methods, we would like to discuss stochastic noise modeling and an intrinsic conflict with the deterministic model. Here, and throughout the rest of the work, assume

$$(\Omega, \mathcal{F}, \mathbb{P}) \quad (1.13)$$

to be a complete probability space with a set Ω of outcomes of the stochastic event, \mathcal{F} the corresponding σ -algebra and \mathbb{P} a probability measure, $\mathbb{P} : (\Omega, \mathcal{F}) \rightarrow [0, 1]$. In the Hilbert-space setting, the noise is typically modeled as follows, see for example [7, 14, 15]. Let $\xi : \Omega \rightarrow \mathcal{Y}$ be a stochastic process. Then for $y \in \mathcal{Y}$

$$\langle y, \xi \rangle \quad (1.14)$$

defines a real-valued random variable. Assuming that

$$\mathbb{E}(\langle \tilde{y}, \xi \rangle^2) < \infty \quad (1.15)$$

for all $\tilde{y} \in \mathcal{Y}$ and that this expectation is continuous in \tilde{y} ,

$$\mathbb{E}(\langle \tilde{y}, \xi \rangle \langle y, \xi \rangle)$$

defines a continuous, symmetric nonlinear bilinearform. In particular, there exists the *covariance operator*

$$\mathcal{C} : \mathcal{Y} \rightarrow \mathcal{Y}$$

with

$$\langle \mathcal{C}\tilde{y}, y \rangle = \mathbb{E}(\langle \tilde{y}, \xi \rangle \langle y, \xi \rangle).$$

For deterministic infinite dimensional problems, (1.14) is problematic. Namely, if $\{u_n\}_{n \in \mathbb{N}}$ is an orthonormal basis in \mathcal{Y} , the set $\{\langle u_n, \xi \rangle\}_{n \in \mathbb{N}}$ consists of infinitely many identically distributed random variables with $0 < \mathbb{E}|\langle u_n, \xi \rangle|^2 = \text{const} < \infty$ [7]. Thus

$$\mathbb{E} \left(\sum_{n=1}^{\infty} |\langle u_n, \xi \rangle|^2 \right) \quad (1.16)$$

is almost surely infinite. Therefore, a realization of the noise is an element of the Hilbert space \mathcal{Y} with probability zero. The problem is similar to the definition of the Besov-space valued random variable we present in Section 4.2.2. Here, we briefly discuss another example. Gaussian white noise can be described with the above construction. With $\mathbb{E}(\langle y, \xi \rangle) = 0 \forall y \in \mathcal{Y}$ and the covariance operator $\mathcal{C} = \eta^2 I$ where I is the identity and η the variance parameter, the Gaussian white noise is described [7, 15]. As explained for example in [15], a realization of such a Gaussian random variable is an element of an infinite dimensional L_2 -space with probability zero. It is therefore inappropriate to use an L_2 -norm for the residual. Since a realization of Gaussian white noise only lies (almost surely) in any Sobolev space H^s with $s < -d/2$ where d is the dimension of the domain, one should use such a norm instead. However, except for the paper [15] this issue seems not to have been addressed in the literature. For the practical solution of the Inverse Problem this is not a large issue since the measurements are finite dimensional and the unknown requires a finite dimensional representation to be handled by a computer. In this case, where the sum in (1.16) is finite, the noise lies within the finite dimensional space almost surely. It is problematic, however, if one seeks to investigate convergence of the discretized problem to its underlying infinite dimensional problem. Throughout this work we will assume that $\mathbb{E}||\epsilon|| < \infty$ or use the slightly weaker bound on the Ky-Fan metric (see Section 1.3). If we have a finite dimensional problem then this formulation is clear. For infinite dimensional problems, however, we have to assume that the noise is smooth enough to be in the respective spaces almost surely. Due to this, we will not discuss convergence of finite dimensional problems to their infinite dimensional representation.

1.3.2 The Bayesian approach

We divide the stochastic regularization theory in two subgroups: the Bayesian approach and other methods. We will start our review with a short explanation of the Bayesian

framework. For more detailed information, we refer to [6, 16, 17, 18, 19, 20]. In the Bayesian setting, the solution of the Inverse Problem is given as a distribution of the random variable of interest, the *posterior distribution* π_{post} , determined by Bayes formula

$$\pi_{post}(x|y^\eta) = \frac{\pi_{pr}(x)\pi_\epsilon(y^\eta|x)}{\pi_{y^\eta}(y^\eta)}. \quad (1.17)$$

That is, roughly spoken, all values x are assigned a probability of being a solution to (1.1) given the noisy data y^η . In (1.17), the *likelihood function* $\pi_\epsilon(y^\eta|x)$ represents the model for the measurement noise whereas the *prior distribution* π_{pr} represents a-priori information about the unknown. The data distribution $\pi_{y^\eta}(y^\eta)$ as well as the normalization constants are usually neglected since they only influence the normalization of the posterior distribution. In practice, however, one is often more interested in finding a single representation as solution instead of the distribution itself. Popular point estimates are the *conditional expectation* (conditional mean, CM)

$$\mathbb{E}(\pi_{post}(x|y^\eta)) = \int x\pi_{post}(x|y^\eta)dx \quad (1.18)$$

and the *maximum a-posteriori* (MAP) solution

$$x^{\text{MAP}} = \underset{x}{\operatorname{argmax}} \pi_{post}(x|y^\eta), \quad (1.19)$$

i.e., the most likely value for x . Both point estimators are widely used although they cannot capture the whole picture of the distribution of the posterior distribution. While the CM-estimate returns the value in the “middle” of the probability density, the MAP gives the value with maximal probability. If both likelihood and prior distribution are Gaussian, CM and MAP coincide. Since the computation of (1.18) often involves repeated evaluation of high dimensional integrals, we only consider the maximum a-posteriori approach because it essentially leads to a Tikhonov-type minimization problem which can typically be solved much faster. Namely, assuming $\pi_\epsilon(y^\eta|x) \propto \exp(-\|Ax - y^\eta\|_{\mathcal{C}_\eta}^2)$ where $\|\cdot\|_{\mathcal{C}}^2 = \langle \mathcal{C}^{-1}\cdot, \cdot \rangle$ with \mathcal{C}_η being the covariance of the noise and $\pi_{pr}(x) \propto \exp(-\alpha\Theta(x))$ for some proper functional $\Theta : \mathcal{X} \rightarrow \mathbb{R}$, then

$$\pi_{post}(x|y^\eta) \propto \exp(-\|Ax - y^\eta\|_{\mathcal{C}_\eta}^2) \exp(-\alpha\Theta(x)) = \exp(-\|Ax - y^\eta\|_{\mathcal{C}_\eta}^2 - \alpha\Theta(x)).$$

Thus

$$x^{\text{MAP}} = \min_x \|Ax - y^\eta\|_{\mathcal{C}_\eta}^2 + \alpha\Theta(x),$$

which is a Tikhonov-type functional. If, in particular, both likelihood and prior function are Gaussian, and (1.1) is a finite dimensional problem, then x^{MAP} has an explicit representation. Finite dimensional Gaussian random variables are defined as follows.

Definition 1.3.1. Let $x_0 \in \mathbb{R}^n$ and $\Gamma \in \mathbb{R}^{n \times n}$ be a symmetric positive definite matrix. A Gaussian n -variate random variable X with mean x_0 and covariance Γ is a random variable with the probability density

$$\pi(x) = \left(\frac{1}{2\pi \det(\Gamma)} \right)^{n/2} \exp \left(-\frac{1}{2} (x - x_0)^T \Gamma^{-1} (x - x_0) \right).$$

We use the notation $X \sim \mathcal{N}(x_0, \Gamma)$.

The following Theorem from [6] fully characterizes the solution to the Inverse Problem (1.1) in a finite dimensional Gaussian setting.

Theorem 1.3.1. *Assume that $X : \Omega \rightarrow \mathbb{R}^n$ and $\mathcal{E} : \Omega \rightarrow \mathbb{R}^m$ are mutually independent Gaussian random variables, $X \sim \mathcal{N}(x_0, \Gamma_{pr})$ and $\mathcal{E} \sim \mathcal{N}(e_0, \Gamma_\eta)$ where $\Gamma_{pr} \in \mathbb{R}^{n \times n}$ and $\Gamma_\eta \in \mathbb{R}^{m \times m}$ are positive definite. Assume further that we have a linear model for a noisy measurement Y (1.1),(1.2), where $A \in \mathbb{R}^{m \times n}$ is a known matrix. Then the posterior probability density of X given the measurement $Y = y$ is*

$$\pi(x|y) \propto \exp \left(-\frac{1}{2} (x - \bar{x})^T \Gamma_{post}^{-1} (x - \bar{x}) \right),$$

where

$$\bar{x} = x_0 + \Gamma_{pr} A^T (A \Gamma_{pr} A^T + \Gamma_\eta)^{-1} (y - A x_0 - e_0)$$

and

$$\Gamma_{post} = \Gamma_{pr} - \Gamma_{pr} A^T (A \Gamma_{pr} A^T + \Gamma_\eta)^{-1} A \Gamma_{pr}.$$

One easily sees that \bar{x} is nothing else than the MAP-solution. Assuming further that $X \sim \mathcal{N}(0, I)$ and $\mathcal{E} \sim \mathcal{N}(0, \eta^2 I)$, one obtains

$$x^{\text{MAP}} = (A^T A + \eta^2 I)^{-1} A^T y$$

which coincides with the solution of the regularized normal equation (1.10). Convergence of the MAP-estimator to the true solution in the setting of Theorem 1.3.1 when $\eta \rightarrow 0$ has been investigated in [21, 22, 23]. Recently, analysis of the MAP estimator in the infinite dimensional setting has emerged, see [24, 25]. There is plenty of literature on Inverse Problems in the Bayesian setting. We will give a more detailed review on sparsity-promoting regularization in Chapter 4. For a general overview we again refer to [6, 16] and the paper of Lasanen [26]. The latter reference gives a condensed review of the development of the Bayesian approach in Inverse Problems also in its historical development.

Since we already introduced the finite dimensional Gaussian random variable, we present the following theorem where we collect some properties which will be used later on.

Lemma 1.3.2. Let $e = (\epsilon_1, \epsilon_2, \dots, \epsilon_m)^T \in \mathbb{R}^m$ where ϵ_i , $i = 1, \dots, m$, are independent identically distributed Gaussian random variables with zero mean and variance η^2 . Then it holds that

i) for any $c > 0$

$$\mathbb{P}(\|e\| > c) = \frac{\Gamma(\frac{m}{2}, \frac{c^2}{2\eta^2})}{\Gamma(\frac{m}{2})}. \quad (1.20)$$

ii) for any $s \in \mathbb{N}$ it is

$$\mathbb{E}(\|e\|^s) = 2^{s/2} \eta^s \frac{\Gamma(\frac{s+m}{2})}{\Gamma(\frac{m}{2})}$$

iii) for any $\tau > 1$

$$\mathbb{P}(\|e\| \geq \tau \mathbb{E}(\|e\|)) = \frac{\Gamma(\frac{m}{2}, (\tau \Gamma(\frac{m+1}{2}) / \Gamma(\frac{m}{2}))^2)}{\Gamma(\frac{m}{2})} \quad (1.21)$$

with the gamma functions $\Gamma(\cdot)$ and $\Gamma(\cdot, \cdot)$,

$$\Gamma(a) = \int_0^\infty t^{a-1} e^{-t} dt, \quad \Gamma(a, z) = \int_z^\infty t^{a-1} e^{-t} dt.$$

Proof. Let $\epsilon_i \sim \mathcal{N}(0, \eta^2)$, $i = 1 \dots m$. Consider the real valued, non-negative random variable $X = \|e\|^2 = \sum_{i=1}^m \epsilon_i^2$. Then X is χ^2 -distributed and obeys the probability density function

$$f_X(\zeta) = \frac{1}{2^{\frac{m}{2}} \eta^m \Gamma(\frac{m}{2})} \zeta^{\frac{m}{2}-1} e^{-\frac{\zeta}{2\eta^2}},$$

see, e.g., [27]. Set $Y := \sqrt{X} = \|e\|$, $\xi^2 = \zeta$. Then the probability density is given by $f_Y(\xi) = f_X(\xi^2) \cdot \frac{d\xi^2}{d\xi}$ and hence

$$f_{\|e\|}(\xi) = \frac{1}{2^{\frac{m}{2}-1} \eta^m \Gamma(\frac{m}{2})} \xi^{m-1} \exp\left(-\frac{\xi^2}{2\eta^2}\right) \quad \xi \geq 0.$$

Now i) follows from evaluating

$$\mathbb{P}(\|e\| > c) = \int_c^\infty f_{\|e\|}(\xi) d\xi,$$

ii) is the result of

$$\mathbb{E}(\|e\|^s) = \int_0^\infty \xi^s f_{\|e\|}(\xi) d\xi,$$

and iii) holds due to i) and ii) with $s = 1$. □

1.3.3 Other stochastic approaches

We quickly mention some non-Bayesian approaches to stochastic Inverse Problems to show similarities between stochastic and deterministic approaches. Considering x, y, ϵ as realizations of the random variables X, Y, \mathcal{E} with

$$X : \Omega \rightarrow X, \quad Y, \mathcal{E} : \Omega \rightarrow Y$$

and

$$AX = Y + \mathcal{E}$$

as well as assuming $\mathbb{E}(X) = \mathbb{E}(Y) = 0$, that X and \mathcal{E} are uncorrelated and that $\mathcal{C}_{\mathcal{E}}$ is invertible, the best linear estimator, i.e., the linear operator L minimizing

$$\mathbb{E}(|\langle \xi, X - LY \rangle|^2) \quad \forall \xi \tag{1.22}$$

is given by

$$L_{min} = \mathcal{C}_X A^* (\mathcal{A} \mathcal{C}_X \mathcal{A} + \mathcal{C}_{\mathcal{E}})^{-1}$$

where \mathcal{C}_X and $\mathcal{C}_{\mathcal{E}}$ are the covariance operators of X and \mathcal{E} , respectively [7]. If $\mathcal{C}_X = I$ and $\mathcal{C}_{\mathcal{E}} = \eta^2 I$

$$L_{min} = (A^* A + \eta^2 I)^{-1} A^*,$$

i.e., Tikhonov regularization with regularization parameter $\alpha = \eta^2$. Thus, the best linear estimator is a particular Tikhonov method [7].

In the paper [14], the authors consider convergence of solutions to the linear inverse problems of type (1.1) in the *mean integrated square error (MISE)*,

$$\mathbb{E}(\|x_{\alpha}^{\eta} - x^{\dagger}\|^2). \tag{1.23}$$

Using conditions on their regularization method similar to the ones in Theorem 3.1.2 and Theorem 3.1.3, they derive convergence rates for the MISE assuming a combination of both stochastic and deterministic noise. Under similar conditions, in [28] the discrepancy principle was investigated with particular focus on the Conjugate Gradient Method, convergence in the MISE was shown. In contrast to the previous previous methods, where Gaussian noise was used, in [29], a Tikhonov-type regularization for Poisson noise in the deterministic and stochastic setting was discussed.

The Dissertation of Hofinger [30] was the first and, to the authors knowledge only, attempt of a rigorous analysis of Inverse Problems in a stochastic setting by lifting results from the deterministic theory. The author considers a purely stochastic problem

$$y^{\eta}(\omega) = A(\omega)x(\omega) + \varepsilon(\omega),$$

i.e., also the operator is allowed to be random. One of the main ingredients of his analysis is the Ky Fan metric, which we will also employ and introduce in the next section.

1.4 The Ky Fan metric

The Ky Fan metric (cf. [31]) will be the main tool for our stochastic convergence analysis. It is defined as follows.

Definition 1.4.1. Let X_1 and X_2 be random variables in a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with values in a metric space $(\mathcal{X}, d_{\mathcal{X}})$. The distance between X_1 and X_2 in the *Ky Fan metric* is defined as

$$\rho_K(X_1, X_2) := \inf_{\varepsilon > 0} \{\mathbb{P}(\{\omega \in \Omega : d_{\mathcal{X}}(X_1(\omega), X_2(\omega)) > \varepsilon\}) < \varepsilon\}. \quad (1.24)$$

This metric essentially allows to lift results from a metric space to the space of random variables as the connection to the deterministic setting is inherent via the metric $d_{\mathcal{X}}$ used in its definition. Throughout this work the deterministic metric will in most cases be induced by a norm $\|\cdot\|$. We will also implicitly assume that equation (1.1) is scaled appropriately since $\rho_K(X_1, X_2) \leq 1 \forall X_1, X_2$ by definition. An immediate consequence of (1.24) is that $\rho_K(X_1, X_2) = 0$ if and only if $X_1 = X_2$ almost surely. Convergence in the Ky Fan metric is equivalent to convergence in probability, i.e., for a sequence $\{X_k\}_{k \in \mathbb{N}} \in \mathcal{X}$ and $X \in \mathcal{X}$ one has

$$\rho_K(X_k, X) \xrightarrow{k \rightarrow \infty} 0 \quad \Leftrightarrow \quad \forall \varepsilon > 0 : \quad \mathbb{P}(\|X_k - X\|_{\mathcal{X}} > \varepsilon) \xrightarrow{k \rightarrow \infty} 0.$$

Hence convergence in the Ky Fan metric also leads to pointwise (almost sure) convergence of certain subsequences in the metric $d_{\mathcal{X}}$ [32].

A somewhat more intuitive and more frequently used metric is the expectation, or more general, a (stochastic) L_p metric. For random variables X_1 and X_2 with values in a metric space $(\mathcal{X}, d_{\mathcal{X}})$ and $\mathbb{E}(\|X_1\|^p), \mathbb{E}(\|X_2\|^p) < \infty$,

$$\mathbb{E}(\|X_1 - X_2\|_{\mathcal{X}}^p) = \int_{\Omega} \|X_1(\omega) - X_2(\omega)\|_{\mathcal{X}}^p d\mathbb{P}(\omega)$$

defines the p -th moment of $\|X_1 - X_2\|$ for $p \geq 1$. We will mostly use $p = 1$ and refer to it as *convergence in expectation*.

It is well-known that convergence in expectation implies convergence in probability, see for example [32]. Hence, convergence in the Ky Fan metric is implied by convergence in expectation (and also by convergence of higher moments). Namely, with *Markov's inequality* (see, e.g., [33]) one has, for an arbitrary nonnegative random variable X with $E(X) \leq \infty$

$$\mathbb{P}(X > C) \leq \frac{\mathbb{E}(X)}{C}. \quad (1.25)$$

Under an additional assumption, one can show that convergence in probability implies convergence in expectation. We have the following definition.

Definition 1.4.2 ([34], Definition A.3.1.). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space. A family $\mathcal{G} \subset L_1(\mathbb{P})$ is called *uniformly integrable* if

$$\lim_{C \rightarrow \infty} \sup_{x \in \mathcal{G}} \int_{|x| > C} |x(t)| \mathbb{P}(dt) = 0$$

Theorem 1.4.1 ([34], Theorem A.3.2.). Let $\{x_k\}_{k \in \mathbb{N}} \subset L_1(\mathbb{P})$ be a sequence convergent almost everywhere (or in probability) to a function x . If the sequence $\{x_k\}_{k \in \mathbb{N}}$ is uniformly integrable, then it converges to x in the norm of $L_1(\mathbb{P})$

We will make use of this result several times. From a practical point of view, uniform integrability of a sequence of regularized solutions to an Inverse Problem is a rather natural condition. Since Inverse Problems typically arise from some real-world application, it is to be expected that the true solution is bounded. For example, in Computer Tomography, the density of the tissue inside the body cannot be arbitrarily large. Although for an Inverse Problem with a stochastic noise model, boundedness of the regularized solutions can not be guaranteed due to the possibly huge measurement error, one can enforce the condition from a priori knowledge of the solution.

Assumption 1.4.2. Assume that the true solution x^\dagger fulfills $\|x^\dagger\| \leq \varrho$ and $|x^\dagger| \leq C$ globally for some fixed $\varrho, C > 0$.

Under this assumption, let $\{x_k^{\eta(k)}\}_{k \in \mathbb{N}}$ be a sequence of regularized solution with variance $\eta(k) \xrightarrow{k \rightarrow \infty} 0$. Let $C_1, C_2 > 1$ and define

$$\tilde{x}_k^\eta := \begin{cases} x_k^\eta, & \|x_k^\eta\| \leq C_1 \varrho, |x_k^\eta| \leq C_1 C \\ 0, & \text{otherwise} \end{cases}. \quad (1.26)$$

Then the sequence $\{\tilde{x}_k^\eta\}_{k \in \mathbb{N}}$ is uniformly integrable. In other words, by discarding solutions that must be far away from the true solution in regard of a priori knowledge, convergence in the Ky Fan metric implies convergence in expectation.

Due to the special structure of the Ky Fan metric (1.24) estimates of the Ky Fan distance between to random variables X_1, X_2 can easily be obtained from probability estimates in two ways. If one has a parameter dependent probability estimate of the form

$$\mathbb{P}(\|X_1 - X_2\| \geq C_1(\alpha)) \leq C_2(\alpha), \quad (1.27)$$

with, for simplicity, $\alpha > 0$, then, one may choose α^* such that $C_1(\alpha^*) = C_2(\alpha^*)$ (provided this problem is solvable) and it follows immediately from the definition of the Ky Fan metric that

$$\rho_K(X_1, X_2) \leq C_1(\alpha^*) = C_2(\alpha^*).$$

We will use this balancing principle in Theorem 4.3.8. If on the other hand

$$\mathbb{P}(\|X_1 - X_2\| \geq C_1) \leq C_2, \quad (1.28)$$

without any tuning parameter, then again (1.24) immediately implies

$$\rho_K(X_1, X_2) \leq \max\{C_1, C_2\}.$$

A main application of this property is the following Theorem by Hofinger and Pikkarainen which proves that pointwise estimates are essentially preserved in the Ky Fan metric.

Theorem 1.4.3. [21], Theorem 6 Let X_1, X_2 and Y_1, Y_2 be random variables on metric spaces $(\mathcal{X}, d_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}})$, respectively. Let

$$d_{\mathcal{X}}(X_1(\omega), X_2(\omega)) \leq \Phi(d_{\mathcal{Y}}(Y_1(\omega), Y_2(\omega)))$$

for almost all $\omega \in \Omega$, where Φ is a monotonically increasing right-continuous function. Then

$$\rho_K(X_1, X_2) \leq \max\{\rho_K(Y_1, Y_2), \Phi(\rho_K(Y_1, Y_2))\}.$$

The theorem raises the question how to calculate $\rho_K(Y_1, Y_2)$ for given random variables Y_1, Y_2 . To simplify notation let the distance $d_{\mathcal{Y}}$ be given by a norm $\|Y_1 - Y_2\|$. As pointed out by Hofinger [30], an upper bound for $\rho_K(Y_1, Y_2)$ can be derived from moments of $\|Y_1 - Y_2\|$.

Theorem 1.4.4. Let Y_1, Y_2 be random variables in a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and $\mathbb{E}(\|Y_1 - Y_2\|^s) < \infty$ for some $s \in \mathbb{N}$. Then

$$\rho_K(Y_1, Y_2) \leq \sqrt[s+1]{\mathbb{E}(\|Y_1 - Y_2\|^s)} \quad (1.29)$$

Proof. One has due to Markov's inequality (1.25) and the monotonicity of the mapping $z \mapsto z^s$ for $z \geq 0$ it is

$$\mathbb{P}(\|Y_1 - Y_2\| > C) = \mathbb{P}(\|Y_1 - Y_2\|^s > C^s) \leq \frac{\mathbb{E}(\|Y_1 - Y_2\|^s)}{C^s}$$

for $C \geq 0$. Solving $C = \frac{\mathbb{E}(\|Y_1 - Y_2\|^s)}{C^s}$ for C yields the claim. \square

From the above statement one can already guess that the Ky Fan distance between two random variables is larger than their expectation. This holds true even if the moments exist for all $s \in \mathbb{N}$. Let us investigate a prominent special case. In practice one often assumes a normal distribution of the measurement noise. Let therefore ϵ be a random vector in \mathbb{R}^m , $m \in \mathbb{N}$, $\epsilon \sim \mathcal{N}(0, \eta^2 I_m)$ (see Definition 1.3.1) such that each component is normally distributed with zero mean and variance η^2 . In this case, a direct estimate for the Ky Fan distance was derived by Hofinger and Pikkarainen [21] and refined by Neubauer and Pikkarainen [23].

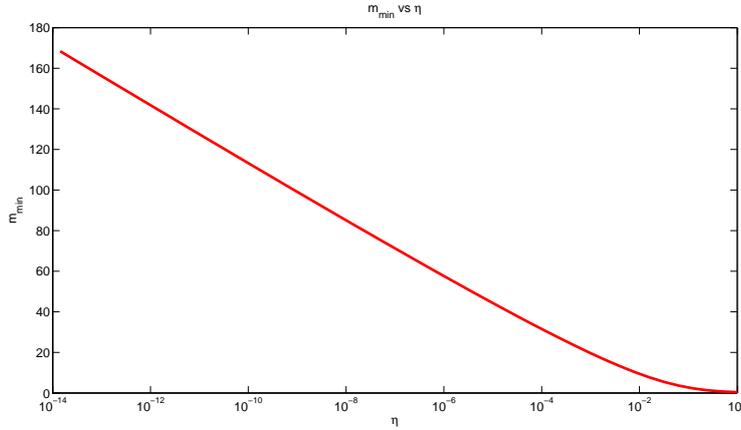


Figure 1.1: m_{\min} from (1.31) vs η . Even for $\eta = 10^{-10}$, only 114 measurement points are necessary to keep the logarithm in (1.30) inactive, i.e., $\ln^-(\cdot) = 0$.

Proposition 1.4.5. *Let ϵ be a random variable with values in \mathbb{R}^m . Assume that the distribution of ϵ is $\mathcal{N}(0, \eta^2 I)$ with $\eta > 0$. Then it holds in $(\mathbb{R}^m, \|\cdot\|_2)$ that*

$$\rho_K(\epsilon, 0) \leq \min \left\{ 1, \sqrt{2}\eta \sqrt{m - \ln^- \left(\eta^2 2\pi m^2 \left(\frac{e}{2} \right)^m \right)} \right\}, \quad (1.30)$$

where $f^-(h) := \min\{0, f(h)\}$.

The smallest m for which the \ln^- -term vanishes is at the zero of $\ln^- \left(\eta^2 2\pi m^2 \left(\frac{e}{2} \right)^m \right)$. It is given by

$$m_{\min} = \mathbf{ceil} \left(\frac{2}{1 - \ln 2} W \left(\frac{1 - \ln 2}{2\sqrt{2\pi}} \frac{1}{\eta} \right) \right), \quad (1.31)$$

where W is the Lambert W-function defined by $W(z)e^{W(z)} = z$ (cf. [35]), and $\mathbf{ceil}(\cdot)$ the function which maps a real number to the smallest following integer. In practice, when a real-world problem is solved with fixed m and η , $m > m_{\min}$ typically is fulfilled, see Figure 1.1. Then $\rho_K(\xi, 0) \leq \min \{1, \sqrt{2}\eta\sqrt{m}\}$ and the Ky-Fan distance is of the same order as the expectation of the error.

Chapter 2

Introduction to Adaptive Optics

Although the theoretical results presented in this work can be used for a huge variety of applications in Inverse Problems, a particular one will be presented in more detail. Namely, we will discuss the problem of Atmospheric Tomography in the context of Adaptive Optics, a strategy needed to improve the image quality of large earth-bound telescopes. This chapter is meant to introduce the reader to this topic in a condensed way, aimed at explaining basic concepts and quantities without going too deep into the matter. For more information, several works are available. We refer to, e.g., [36, 37, 38] and in particular [39].

2.1 Introduction to Adaptive Optics

Mankind has always been fascinated with the observation of the stars. With the invention of telescopes in the late 17th century people were able to observe more and more details of the night sky, leading to better and better understanding of astronomical phenomena. Since often a new discovery leads to further questions, more and more detailed observations of the sky are necessary. As will be explained in the next section, the resolution of a telescope is, even in the ideal case, limited by its size. Hence, the size of telescopes has increased up to a mirror diameter of currently around 10m. The next generation of telescopes will feature mirrors of up to 40m in diameter. After a long planning and design phase, the construction of the European Extremely Large Telescope (E-ELT), see Figure 2.1, was authorized by the European Southern Observatory (ESO) in 2012. With a primary mirror diameter of (to this point planned) 39.3m it will be the worlds largest telescope. First light is expected in 2024. However, as with any earth-bound telescope, the image quality is negatively affected by the turbulence of the air in the atmosphere. Space based telescopes such as the Hubble Space Telescope, with a main mirror of 2.4m in diameter, are free from this effect. However they still suffer from the image degradation from diffraction which can only be improved by increasing the telescope size. The costs for space telescopes exceed those of earth-bound by far. While the build cost of the E-ELT is currently estimated at 1.1 billion euro, the construction of the Hubble telescope was about 2.5 billion dollars. Additionally, service of the telescope in space is very expensive. The overall costs of the Hubble mission until 2010 are estimated around 10 billion dollars [40]. Hence it is desirable to improve the telescopes on earth which requires to be able to remove the influence of the atmosphere. This can, up to a certain extend, be realized with *Adaptive Optics* (AO). Outside the atmosphere of the earth, the light from a distant object outside the solar system can be assumed to be planar, i.e., a lens focuses the light in a single point. However, as the light travels through the atmosphere of the earth, turbulences of the air cause small local changes in the refractive index of the air, leading to the light rays traveling at different speed. At the telescope, the lens no longer focuses the rays in a single point and hence the observed image is blurred and lacks details. The principle of Adaptive Optics is to “undo” this



Figure 2.1: Render of the planned E-ELT in comparison with the VLT (Very Large Telescope, largest mirror diameter 8.2m) and St. Stephen's Cathedral in Vienna (Source: ESO)

process using deformable mirrors such that the incoming light can be focused properly again, see Figure 2.2. With this approach, significantly more details of the astronomical objects are visible, see again Figure 2.2 for two examples.

2.2 Optical imaging through an ideal telescope

The light observed by a telescope is modeled as a complex field

$$u(\mathbf{x}, t) = A(\mathbf{x}, t)e^{i\varphi(\mathbf{x}, t)},$$

where $A(\mathbf{x}, t)$ is the *amplitude* and $\varphi(\mathbf{x}, t)$ the *phase* of the signal. Both depend on the spatial coordinate $\mathbf{x} = (x_1, x_2) \in \mathbb{R}^2$ and, in general, the time t . We will usually drop the explicit dependence on the time as we consider a sequence of observations, each frozen in time, rather than a time-continuous process. The level sets of the phase, i.e., the sets $\{\mathbf{x} \in \mathbb{R}^2 : \varphi(\mathbf{x}) = C\}$ for constants $C \in \mathbb{R}$ are called *wavefronts*. Light propagates along rays normal to these sets. Although optical devices are only sensitive to the intensity $I = |A(\mathbf{x})|^2$, a non-planar phase also influences the intensity: if the phase is not constant, the light waves interfere with each other and, according to the superposition principle, add up or cancel out. However, even without the effects of the turbulent atmosphere, the quality of an observed image is determined and limited essentially by the diameter of the telescope aperture. The image of a point light source observed with a telescope is shaped by *diffraction*. Mathematically, this can be described with a *point spread function* (also called *Airy function*, see Figure 2.3), modeling the light intensity in the focal plane as a

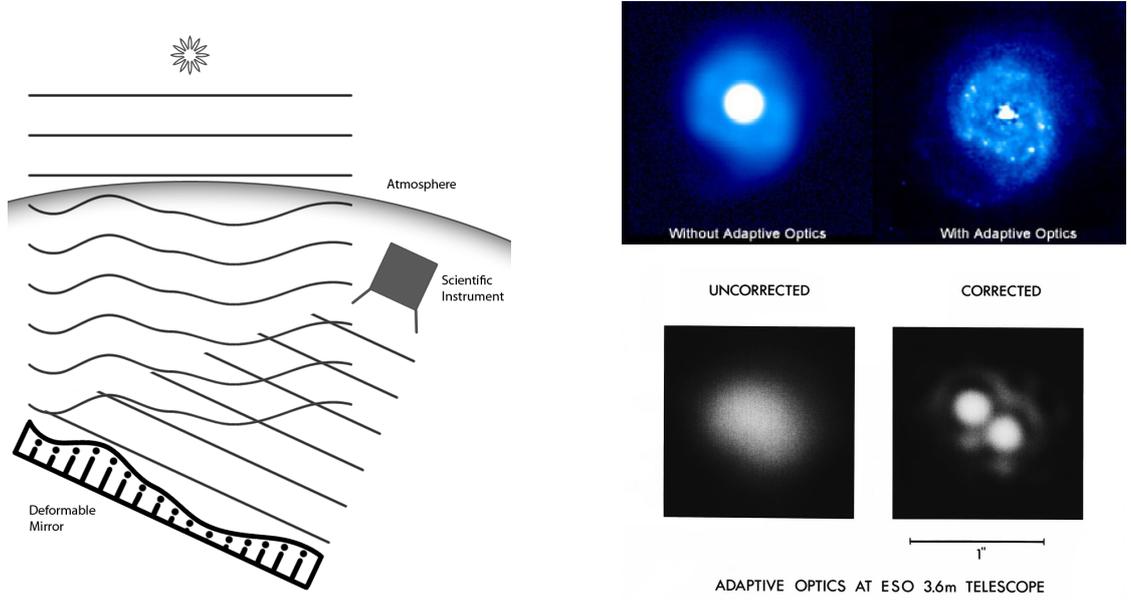


Figure 2.2: Left: The former planar light is distorted when travelling through the atmosphere. A deformable mirror can correct this. Right: Two examples of images taken from earth bound telescopes, each without and with Adaptive Optics. The corrected images are much more detailed. (Sources: [41], ESO)

function of the angular coordinate \mathbf{x} . With the wavelength of light λ and the diameter of the (throughout this work always circular) telescope aperture D , the point spread function PSF_0 for diffraction limited imaging is given by

$$PSF_0 = \frac{\pi D^2}{4\lambda^2} \left[\frac{2J_1(\pi D|\mathbf{x}|)}{\pi D|\mathbf{x}|/\lambda} \right]^2, \quad (2.1)$$

where $J_1(\cdot)$ is the Bessel function of the first kind. The effective focal spot size is $1.22\lambda L/D$, where L is the focal distance. It is often taken as a measure of resolution in an ideal telescope. For any fixed wavelength it depends solely on the geometry of the telescope. Hence, this quantity, corresponding to the first dark ring in Figure 2.3, is the reason for building larger and larger telescopes.

The Airy function describes the imaging process of a telescope. Let $I(\mathbf{x})$ be the image of an astronomical object $O(\mathbf{x})$. Then each point of the object is spread into an airy function, i.e., the observed image is blurred and its resolution is degraded. Mathematically, this corresponds to the convolution of the object with the point spread function,

$$I(\mathbf{x}) = \int O(\mathbf{y})PSF_0(\mathbf{x} - \mathbf{y}) d\mathbf{y}. \quad (2.2)$$

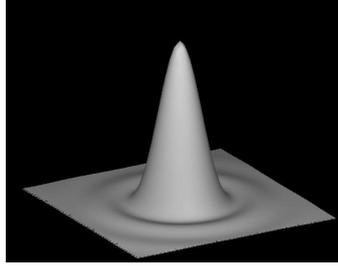


Figure 2.3: Example of an Airy function, taken from [39]

This is called the *imaging equation*. PSF_0 corresponds to the least possible degradation of an image for a given telescope with diameter D , not including additional distortions, for example from atmospheric processes. Such an ideal image is called *diffraction-limited*. Equation (2.2) also holds for real astronomical images, although with a different point spread function PSF . Due to PSF_0 being the best obtainable point spread function, such a general PSF must be “worse” than PSF_0 , i.e., it may be broader, with smaller peak value and also it may not be smooth anymore. Hence, the observed image is degraded further. Implicitly, we assumed so far the the point spread functions are such that

$$\int_{\mathbb{R}^2} PSF(\mathbf{x}) d\mathbf{x} = 1. \quad (2.3)$$

They can be interpreted as image of a star of unit intensity. Therefore, the imaging equation preserves the total flux of light coming from an astronomical object, it only distributes it differently between, in praxis, the pixels of a sensor. The width of a point spread function is a measure of the resolution of the telescope, i.e., the minimum distance between distinguishable objects in an image. Usually, real point spread functions are irregular, and other numeric measures are taken to compare the resolution of different telescopes. In this work we use the *Strehl ratio* S which is the quotient of the central intensities of the telescopes PSF and the ideal, diffraction-limited PSF_0 , i.e., $S = PSF(0)/PSF_0(0)$. Since diffraction-limited imaging is best possible, one always has $0 \leq S \leq 1$. The Strehl value can be calculated also if no ground-truth solution is available, i.e., for real data. We refer to, e.g., [42] on further information on the definition of the Strehl ratio. Other measures are the *Full Width at Half Maximum (FWHM)* and the *Encircled energy*, see e.g. [39].

Since the imaging equation is a convolution, it can easily be expressed with the Fourier transform. Denote with $\mathcal{F}(A)$ the Fourier transform of a bivariate function A ,

$$\mathcal{F}(A)(\mathbf{f}) = \frac{1}{2\pi} \int_{\mathbb{R}^2} A(\mathbf{x}) \exp(-i\langle \mathbf{x}, \mathbf{f} \rangle) d\mathbf{x},$$

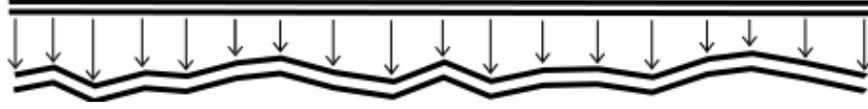


Figure 2.4: Due to differences in the air refractive index, the light rays propagate with different speed. A once planar wavefront is therefore deformed.

where $\langle \cdot, \cdot \rangle$ is the \mathbb{R}^2 euclidean scalar product. Then, due to the Fourier convolution theorem, (2.2) is equivalent to

$$\mathcal{F}(I)(\mathbf{f}) = \mathcal{F}(O)(\mathbf{f}) \cdot \mathcal{F}(PSF)(\mathbf{f}). \quad (2.4)$$

The function $\mathcal{F}(PSF)(\mathbf{f})$ is called *Optical Transfer Function* (OTF). It describes the change of the modulus and phase of the Fourier Transform of the object O in the imaging process. For any optical system, it holds $|\mathcal{F}(PSF)(\mathbf{f})| = 0$ for $|\mathbf{f}| > \omega_c$ where ω_c is the *cut-off frequency*, given by $\omega_c = D/\lambda$. Any information at frequencies larger than ω_c is lost in the image, i.e., the smaller an astronomical object is, the larger a telescope is needed in order to observe it, again motivating the need for very large telescopes.

2.3 Atmospheric turbulence processes

The main cause of wavefront distortions are small changes in the refractive index of the air, i.e., changes in the propagation properties of light through that medium. With the speed of light in vacuum c_0 and the phase velocity c_φ of light in an optical medium, the refractive index n is given by

$$n = \frac{c_0}{c_\varphi}. \quad (2.5)$$

The refractive index of air is influenced, for example, by temperature, humidity and pressure. Small local changes in these properties are caused by atmospheric turbulences. The fluctuations in the refractive index $n = n(\mathbf{x})$ cause local differences in the light propagation speed (reformulating (2.5), one has $c_\varphi(\mathbf{x}) = \frac{c_0}{n(\mathbf{x})}$). Due to the inhomogeneities, individual light rays travel slightly faster or slower than the neighbouring rays, causing the once planar wavefronts to be distorted (imagine a formerly flat sheet of paper being crumpled), see Figure 2.4. The path of a ray through the medium (air) is called *optical path* $l = \int n(z) dz$, thus the differences in travel distance appear as local differences in the optical path $\Delta l(\mathbf{x})$. The dependence of $l(\mathbf{x})$ on the wavelength is small and hence in practice neglected [36]. The phase of the wave $\varphi(\mathbf{x}) = \frac{2\pi}{\lambda} \Delta l(\mathbf{x})$, however, strongly depends the wavelength.

The turbulences in the atmosphere can be modelled as a random process. In order to correct the distortions of the wavefronts caused by the turbulences, it is necessary to statistically characterize the turbulence. We shall do so in the following. For a more

detailed discussion, see [37, 36]. In the description of the atmosphere, a random process u is modelled by the *structure function*

$$D_u(\Delta \mathbf{x}) = \mathbb{E}([u(\mathbf{x} + \Delta \mathbf{x}) - u(\mathbf{x})]^2). \quad (2.6)$$

The structure function is the expected (average) difference between two values of a random process. The covariance

$$C_u(\Delta \mathbf{x}) = \mathbb{E}(u(\mathbf{x})u(\mathbf{x} + \Delta \mathbf{x}))$$

and structure functions are connected via

$$D_u(\Delta \mathbf{x}) = 2[C_u(0) - C_u(\Delta \mathbf{x})] \quad (2.7)$$

as long as $|\mathbf{x}|$ is finite [38]. Another fundamental quantity is the *power spectrum*, given by the Fourier transform of the covariance

$$\phi_u := \mathcal{F}(C_u).$$

It describes how the square of a stationary quantity (i.e., the statistical properties are independent of shifts in space as well as time) is distributed in frequency.

The fundamental model for the atmospheric turbulence is the *Kolmogorov model* due to Kolmogorov [43], see also, e.g., [44, 38]. In this model, the statistical differences between the air refractive index (and thus phase) fluctuations depends, for any given two points in space, only on the distance between these points. The structure function of the refractive index is given by

$$D_n(\Delta \mathbf{x}) = C_n^2(h)|\Delta \mathbf{x}|^{\frac{2}{3}}, \quad (2.8)$$

where C_n^2 is the *refractive index structure constant*. Roughly speaking, $C_n^2(h)$ describes the strength of the turbulence at height h above ground. In practice it is obtained experimentally with weather balloons [37]. Speaking of the *turbulence profile* often refers to C_n^2 in dependence of the altitude. We will again refer to the C_n^2 profile in Section 2.6. The power spectrum of the refractive index in the Kolmogorov model is given by

$$\phi_n(\mathbf{f}) = 0.033(2\pi)^{-\frac{2}{3}}C_n^2|\mathbf{f}|^{-\frac{11}{3}}.$$

It is important to note that the Kolmogorov model only holds in a certain frequency interval called the *inertial range* $1/L_0 \leq |\mathbf{f}| \leq 1/\ell_0$, where ℓ_0 and L_0 describe the length of the smallest and largest *eddies* (i.e., small pockets of air with homogeneous refractive index) in the turbulence, respectively. From the respective quantities of the refractive index, one can calculate the structure function of the phase and obtains

$$D_\varphi(\Delta \mathbf{x}) = 6.88 \left(\frac{|\Delta \mathbf{x}|}{r_0} \right)^{\frac{5}{3}} \quad (2.9)$$

where r_0 is the *Fried parameter* which will be introduced soon. The corresponding power spectrum reads

$$\phi_\varphi(\mathbf{f}) = 0.023r_0^{-\frac{5}{3}}|\mathbf{f}|^{-\frac{11}{3}}. \quad (2.10)$$

The van Karman model (see, e.g., [37]) generalizes this for arbitrary outer scales. Namely, for all $|\mathbf{f}| \geq 0$ one has

$$\phi_n(\mathbf{f}) = 0.033(2\pi)^{-\frac{2}{3}}C_n^2 \left(|\mathbf{f}|^2 + \frac{1}{\tilde{L}_0} \right)^{-\frac{11}{6}} \exp\left(-\frac{|\mathbf{f}|^2}{|\mathbf{f}_m|^2}\right), \quad (2.11)$$

where $|\mathbf{f}_m| = \frac{5.92}{\ell_0}$ and \tilde{L}_0 stands for the outer scale of the spatial coherence of the turbulence.

The Fried parameter r_0 [45], which occurred in (2.9) and (2.10), can be seen to characterize the size of the so called *coherence spots*. In coherent spots, parts of the wavefronts remain locally coherent despite the global turbulence. Typical values for r_0 are between 5cm and 20cm for optical wave lengths and 0.5m–1m in the infrared range.

2.3.1 Exposure times

According to the *frozen turbulence model*, inhomogeneities in the air refractive index are fixed. The turbulence originates from air flow movements (wind, convection) which move these inhomogeneities through the atmosphere. In particular, the speed with which a light beam is crossed corresponds directly to the air flow velocity. Hence, a turbulence is modeled as fixed layers of inhomogeneities (*phase screens*) which are driven by the wind. The temporal behavior not only depends on the wind velocity, but also on the spatial behavior of the phase distortions. It can be assumed, however, that during very short time intervals the turbulences are stationary. This is described by the *atmospheric time constant* τ_0 , by

$$\tau_0 = 0.31 \frac{r_0}{\bar{v}}, \quad (2.12)$$

where r_0 is the Fried parameter and \bar{v} the averaged wind velocity over a given altitude. Typical values for \bar{v} are around $20 \frac{m}{s}$, leading to values of τ_0 from few to tens of milliseconds. This means that during a time interval of length τ_0 the atmosphere can be assumed to be in a fixed state. Hence, τ_0 corresponds to the time in which an Adaptive Optics system has to operate in order to constantly correct for the turbulent atmosphere, leading to required operation frequencies up to several tens or hundreds of Hz. Images of astronomical objects are called *short exposure images* (SE) when the exposure time was at most τ_0 . If the exposure time is (significantly) longer, several images of the same objects are taken with varying atmospheric distortions. Hence, the individual aberrations are averaged, and a *long exposure* (LE) point spread function is obtained. Returning to the notion of optical transfer functions (see (2.4)), the LE PSF can be described by

$$\mathcal{F}(PSF_{LE})(\mathbf{f}) = \mathcal{F}(PSF_0)(\mathbf{f}) \cdot \mathcal{F}(PSF_a)(\mathbf{f}), \quad (2.13)$$

where PSF_0 is the OTF of the telescope (2.1) and $\mathcal{F}(PSF_a)(\mathbf{f})$ the *Atmospheric Optical Transfer Function* (Atmospheric OTF). The atmospheric OTF is related to the phase structure function (2.9) by

$$\mathcal{F}(PSF_a)(\mathbf{f}) = \exp\left(-\frac{1}{2}D_\varphi(\lambda\mathbf{f})\right).$$

Inserting the structure function of the Kolmogorov model, one obtains

$$\mathcal{F}(PSF_a)(\mathbf{f}) = \exp\left(-3.44\left(\frac{\lambda|\mathbf{f}|}{r_0}\right)^{\frac{5}{3}}\right).$$

The LE PSF is given by the inverse Fourier transform. The Strehl number of the atmospheric PSF is exactly the same as in an ideal telescope of diameter r_0 . Thus, for a large telescope with diameter $D \gg r_0$, the Strehl ratio of the atmospheric PSF equals $S = \left(\frac{r_0}{D}\right)^2$. Assume $r_0 = 10\text{cm}$, $D = 40\text{m}$, the Strehl ratio is $S = 0.0025$, i.e., far off the best possible value 1. For large telescopes with good optical quality the resolution is almost entirely defined by the atmosphere. Approximating (2.13) by $\mathcal{F}(PSF_{LE})(\mathbf{f}) \approx \mathcal{F}(PSF_a)(\mathbf{f})$ yields the imaging equation for long exposure imaging

$$\mathcal{F}(I)(\mathbf{f}) = \mathcal{F}(O)(\mathbf{f}) \cdot \mathcal{F}(PSF_a)(\mathbf{f}). \quad (2.14)$$

Another interpretation of the aim of Adaptive Optics is to manipulate the atmospheric PSF such that it is as close as possible to the diffraction limited PSF (2.1).

Since the structure function of the turbulence is statistically the same everywhere, the atmospheric PSF is independent of the viewing direction. However, at a fixed moment of time the phase aberrations depend on the viewing direction (they are *anisoplanatic*). This is a crucial effect for Adaptive Optics, as it sets a limit to the distance between guide star and the object of interest, since then a phase correction in direction of the guide star does not correct the phase in direction of the object sufficiently. The characteristic angular difference is called *isoplanatic angle* θ_0 or *isoplanatic patch*, defined by

$$\theta_0 = 0.31\frac{r_0}{\bar{h}}, \quad (2.15)$$

where \bar{h} is some characteristic average turbulence altitude with typical values $\bar{h} \approx 5\text{km}$.

2.4 Principle of Adaptive Optics

The principle of an Adaptive Optics system is shown in Figure 2.5, for the simplest case of one Deformable Mirror (see Section 2.5.2) and one Wavefront Sensor (cf. Section 2.5.1). Before entering the atmosphere, the light from the astronomical object is planar. Due to changes in the refractive index of the air, the phase is then perturbed. The

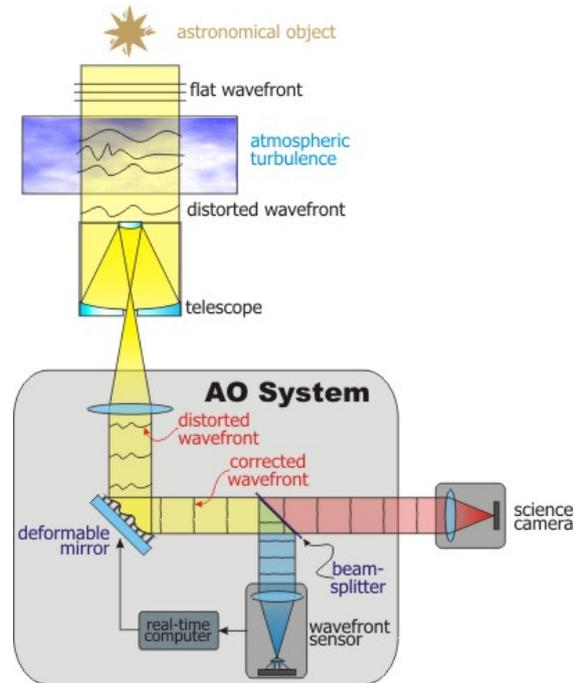


Figure 2.5: Principle of Adaptive Optics for one wavefront sensor and one deformable mirror.

telescope collects the perturbed light (that is, the sum of all individual perturbations in the atmosphere), which is then reflected to the science camera by a deformable mirror. A wavefront sensor measures the incoming wavefront, and a real time computer controls the shape of the mirror such that after reflection at the mirror the light is planar again. Thus, the observed image is sharpened.

2.5 Components in Adaptive Optics

2.5.1 Wave front sensors and wavefront reconstruction

Throughout this work, we will assume that we are provided with measurements in form of incoming wavefronts. The incoming wavefronts cannot be measured directly. Reconstructing them from the actual measurements is a non-trivial problem itself, as it is in fact an Inverse Problem [36]. Although stable when considered in suitable subspaces of L^2 [46], several additional restrictions on the wavefront sensors and reconstruction algorithms in particular for large telescopes complicate the problem.

Although there exist several different types of wavefront sensors, we will assume that our telescope uses *Shack-Hartman* wavefront sensors. There, an array of small identical

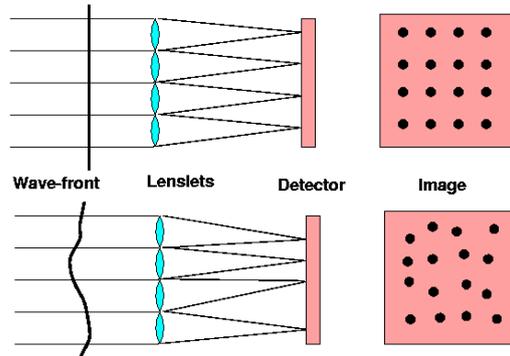


Figure 2.6: Scheme of the Shack-Hartman WFS. Image taken from [39].

lenses (*lenslets*) is placed in the aperture plane of the telescope. Each lens focuses light in a small part of the aperture, called *subaperture*. In each subaperture a CCD sensor, located in the focal plane of the lenslet array, is used to read the image. In order to keep the read-out noise minimal, often 4 CCD-pixels are used for each subaperture in a 2×2 array. Geometries with more CCD pixels are also successfully used. The idea behind the Shack-Hartmann WFS is as follows, compare Figure 2.6. If the incident wavefront is planar, the images formed by the lenslet array create a regular grid in the focal plane. If, however, the wavefront is distorted, the grid becomes non-uniform; in each subaperture the image is deflected from its original position in the $x - y$ plane. The size of the deflection is measured from the CCD data separately in x and y direction. From this, the wavefront can be reconstructed. Many reconstruction methods exist for the Shack-Hartmann WFS. In this work, we employ the so called *Cumulative Reconstructor with Domain Decomposition* (CuReD) [46, 47, 48]. The idea behind CuRe is to abuse the fact that the Shack-Hartmann operator approximates the gradient of the wavefronts in a discrete setting and inverting it by integration of the measurements. CuRe is fast, parallelizable, pipeline-able and provides results that can compete with other methods. Adding a domain decomposition step stabilizes the results for large telescope diameters [48].

2.5.2 Deformable Mirrors and Temporal Control

A Deformable Mirror (DM) is a device with a, as the name suggests, surface that is able to change its shape up to a certain extent. If a planar wave front is reflected by a flat mirror, the reflected phase is again planar. If the mirror surface is not flat, however, the phase of the light is distorted according to the mirror distortions. The idea behind Adaptive Optics is to revert this process: If the wavefronts of the incoming light are known, the surface of the mirror is adjusted such that the reflected light is planar

again. Since the wavefronts change over time, so must the mirror change its shape. Segmented DMs are quite expensive and gaps between the segments may further distort the wavefronts. This problem does not occur for continuous DMs. Continuous face-sheet DM's, see Figure 2.7, therefore have a significantly lower fitting error. The figure of the DM is shaped by the movement actuators of typically piezoelectric type. Other types of Continuous Deformable Mirrors are Bimorph Deformable Mirrors, Micro-Opto-Electro-Mechanical-Systems or Voice-Coil Deformable Mirrors, see, e.g., [36, 39].

For most types of DMs the phase correction φ_{DM} can be obtained by the action of a linear operator H on the actuator commands a ,

$$\varphi_{DM} = Ha. \quad (2.16)$$

Since the incoming phase φ_{tur} is (approximately) known from the WFS data, a standard approach [49] is to find the actuator commands a^* by minimizing a least squares functional

$$a^* = \min_a \|Ha - \varphi_{tur}\|^2. \quad (2.17)$$

In the numerical simulations we use a *pseudo open-loop control* (POLC) (see below). POLC has been successfully used in Adaptive optics, see, e.g., [50]. By adding the current figure of the DM $\varphi_{DM}(\mathbf{x}, t)$ to the measured residual phase $\varphi_{res}(\mathbf{x}, t)$, *open loop* data is created artificially, i.e., we can assume that the incoming phase $\varphi_{turb}(\mathbf{x}, t)$ is measured. In our simplified model, time delays are neglected. Since we mainly consider Multi-Conjugate Adaptive Optics (see Section 2.6), we assume that the mirror fitting step (2.17) is given by $H = -I$ where I is the identity operator. Thus, the minimizer of (2.17) is simply given by $-\varphi_{tur}$.

As explained in previous sections, new wavefront measurements are produced in short time intervals and the shape of the DMs has to be adjusted accordingly. Each individual update will be referred to as a *time step*. In order to do this in a stable way, a control algorithm is needed. Adaptive Optics systems often work in closed loop: instead of measuring the incoming wavefront, the WFSs measure the difference between incoming phase and the current figure of the DM. Several individual components are involved in the process of controlling the DM shape, see Figure 2.7. The main components are the wavefront sensors (WFS), a sampler (analog-to-digital converter, ADC), real time computer (RTC) doing the computations, digital-to-analog-converter (DAC), a power amplifier (PA) and the Deformable Mirrors themselves. Mathematically, the temporal behavior of all these components can be described by *transfer functions*. We do not include this lengthy discussion here, as for the numerical simulations presented in later chapters we use a simplified control algorithm that ignores most of the effects occurring in practice. For more information on the transfer function we refer the interested reader to [36].

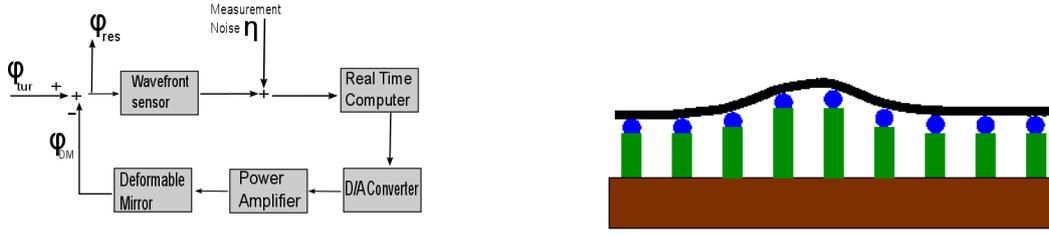


Figure 2.7: Left: Block diagram of the AO control loop. Right: Continuous Deformable Mirror. Image taken from [39].

2.5.3 Guide Stars

In order to measure the incoming wavefronts, proper sources of light crossing the atmosphere are needed. Optimally, a bright star is available on the sky close to the observation area. Since the light of such a star travels an enormous distance in space, the wavefronts can be assumed to be planar prior to entering the atmosphere of the earth. Unfortunately, the coverage of the sky with suitable, so called natural guide stars (NGS), is rather low. A NGS has to be within the isoplanatic patch (see Section 2.3.1) of the area of observation, i.e., the difference in the turbulence profile in direction between NGS and astronomical object has to be small. This guarantees that a correction for the turbulence in direction of the NGS also sufficiently corrects for the turbulences in direction of the object. Additionally, the probability of finding a suitable guide star close to an astronomical object (called *sky coverage*) is related to the brightness (stellar magnitude) of the star. Namely, the photon noise error in the wavefront measurement is inversely proportional to the photon flux. The brighter the star, the smaller the photon noise error. At optical wavelengths, the sky coverage with NGS is only a few percent of the sky. In order to increase the sky coverage, artificial guide stars are used. These are created by shooting a laser beam to the sky (cf. Figure 2.8), hence being called laser guide stars (LGS). A part of the light is backscattered to the telescope and can be used to measure the wavefronts. There exist two different sources of backscattering, hence also two implementations of LGS. We will only use Sodium-LGS here. At about 90 km above ground, there is a layer of sodium with a thickness of about 10 km. The sodium atoms are excited by the laser beam and the returning light can be used to measure the wavefront. For a more in-depth introduction see, e.g. [36, 51]. With laser guide stars, the sky coverage is increased. However, additional physical effects have to be taken into account. Throughout this work, LGS are modelled as a point light source at altitude 90 km. Since the telescope itself is circular, the propagation of the light from this point to

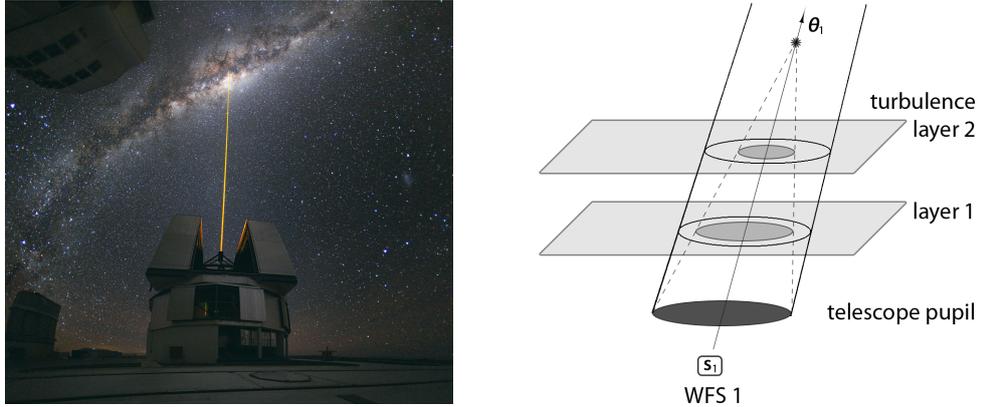


Figure 2.8: Left: Laser Guide Star for the very Large Telescope (VLT). Right: The cone effect for laser guide stars in comparison with cylindrical light propagation for natural guide stars.

the telescope pupil has the form of a cone rather than a cylinder, as would be the case for natural guide stars; see Figure 2.8 for a comparison. As result of this so called *cone effect* only, LSG AO is deteriorated form NGS AO by three distinct effects: Firstly, turbulences above the height $H = 90000m$ of the LGS can not be seen by the telescope. Second, turbulences outside the cone are not visible for the telescope (although some still would be visible for NGS); and third, the wavefronts are scaled differently: the diameter of a slice of the cone at height $h \leq H$ is reduced by $\mu_l = 1 - h/H$ compared to the telescope diameter. The cone effect states a serious limitation of LGS AO in particular for large telescopes and at short wavelength. Several ideas have been developed to cope with the cone effect, textcolorredsee REF for an overview. The preferred option is to use multiple laser guide stars and reconstruct an approximation to the 3-dimensional turbulence profile. This will be the task in *atmospheric tomography*, see Section 2.6. Two additional effects of LGS further decrease the expected image quality. In this work we will not consider them as part of our model. *Tip-tilt indetermination* [52, 41] makes it impossible to compensate the global tilt of the light path as the tilts upward and downward compensate completely as the light travel through the same atmosphere twice. Among various approaches to deal with the incorrect tip/tilt is the approach to measuring the tip/tilt from a NGS, see the note at the end of section 6.2.1. *Spot elongation* [53, 54, 55] takes into account that in reality the light focused by the lenslets of a Shack-Hartman WFS is shaped like an ellipsis rather that a circle.

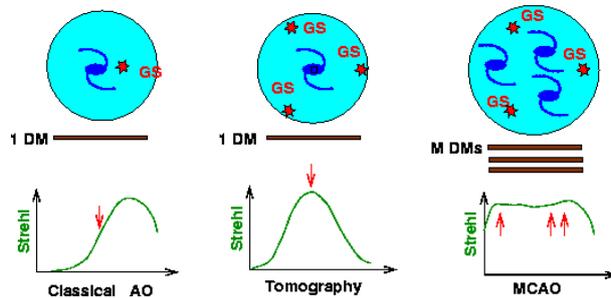


Figure 2.9: Strehl in a cross section of the field of view for SCAO, LTAO and MCAO. Image taken from [39]

2.6 Multi Conjugate Adaptive Optics

2.6.1 Introduction

There are several different AO strategies. In the most simple one, one deformable mirror is used to correct for the atmosphere in the direction of one single guide star. This method, called *Single Conjugate Adaptive Optics (SCAO)* suffers from the need of a (preferably) natural guide star in close proximity to the object of interest. Due to the cone effect and other issues, laser guide stars are not a satisfying solution to this problem. One idea to remedy this is to use several (laser) guide stars around the object, and to create a 3-dimensional turbulence profile from which the optimal shape of the deformable mirror can be calculated. This allows to achieve a peak in Strehl ratio exactly at the point of interest, see Figure 2.9. This approach is called *Laser Tomography Adaptive Optics (LTAO)*. We will discuss the term *tomography* in more detail shortly. Since only one DM is used, however, a good compensation of the wavefront perturbations can only be achieved in a certain direction. In *Multi Conjugate Adaptive Optics (MCAO)* systems, several deformable mirrors and several guide stars are used in order to achieve high imaging quality over a large field of view, see Figure 2.9. This allows in particular to observe several astronomical objects simultaneously or the observation of very large single objects, respectively. MCAO uses several laser guide stars as the main source of wavefront measurements. Additionally, 3 natural guide stars are needed to obtain wavefront information that can not be recovered with LGS [36]. The combined wavefront measurements allow to reconstruct the turbulence profile of the atmosphere in the tomography step mentioned previously for LTAO, which we will now present in more detail.

Recall the X-ray tomography example from Section 1.1: the measurement of the intensity of several X-rays which have passed through a body in different angles/positions allows to reconstruct the inside of the body (that is, one reconstructs location and

“density” of the material inside the body). Similarly, as the wavefront measurements are taken from several directions, one can reconstruct the density inside the medium, i.e., the air refractive index in the atmosphere. We later compare the mathematical models in (2.20) and (2.21). Atmospheric Tomography for telescopes is a variant of limited angle tomography, which is known to be a severely ill-posed problem, see, e.g., [56, 57]. However, a common assumption in atmospheric tomography is that turbulences in the atmosphere are restricted to certain heights. As introduced in Section 2.3, the turbulence strength is measured in the C_n^2 -profile. In practice there are often only a few altitudes where most of the turbulence is concentrated. This leads to the layered model for atmospheric tomography, which is standard in Adaptive Optics [37, 41, 50, 58]: Instead of one 3-dimensional object, the atmosphere is modeled by several 2-dimensional *layers* which are fixed at a certain altitude. The C_n^2 profile then describes the “thickness” of the layer, or, in other words, the percentage of overall atmospheric turbulence concentrated in a certain layer.

As mentioned above, MCAO uses several deformable mirrors for the correction. The mirrors are conjugated to different heights (that is, they are positioned in the optical path of the telescopes to achieve the same effect as if they were positioned at the conjugated altitude in the atmosphere) and are deformed based on the reconstructed turbulence profile, see Figure 2.10. In the simplest case an MCAO system has as many DMs as layers. In this case, a single DM is used for correcting the influence of a single layer. In reality, however, one usually uses a model with more layers than deformable mirrors (e.g., the MCAO system at the E-ELT will use 3 DMs). Therefore, the shape of the mirrors has to be determined by an additional optimization routine, cf. [59, 60], such that the deformable mirrors correct the turbulence profiles of the atmosphere as good as possible.

Algorithms for the computation of an MCAO system can be separated into two groups: Those tackling the whole problem at once, considering the operator that maps the wavefront sensor data to the commands that drive the deformable mirror, and those that split it into sub problems. The MCAO problem can be written by a matrix-vector equation $Ax = b$ of high dimensionality. Thus the inversion process involves either the inversion of large matrices [62, 63] or the solution of large linear systems, which can be achieved, e.g., by CG methods [58, 50, 64, 65] or preconditioned CG methods [66, 67, 68, 69, 70]. Also, Fourier transform based reconstruction methods have been proposed [71, 72, 57, 73]. Iterative methods using a specific representation of the covariance matrix of the statistics of the turbulence have been investigated to further speed up the computations, e.g., the Fractal Iterative Method (*FrIM*) [74] and a wavelet based method (*FEWHA*), [75, 76].

Splitting the problem into substeps allows to exploit individual properties of each operation, thus possibly speeding up the whole process. For MCAO, three such substeps have to be performed: First, the incoming wavefronts have to be reconstructed from the WFS data. Then, the atmospheric tomography problem has to be solved. In the last step, the optimal configuration of the deformable mirrors has to be calculated. Step

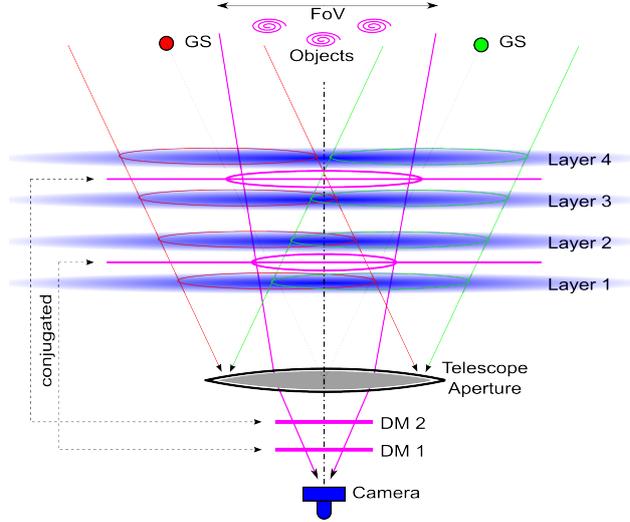


Figure 2.10: Mirrors in an MCAO system, courtesy of [61]

one and three have been discussed in principle in Sections 2.5.1 and 2.5.2, respectively. From the so called *three step methods* we would like to mention in particular a Kaczmarz based method [60, 52, 77, 41].

In this work, we will mainly be concerned with the tomography step, see next section for a mathematical description. We assume that the wavefronts have already been reconstructed. The sought-after atmosphere consists of as many layers as available deformable mirrors. The temporal control is realized with POLC, hence the shape of each DM is set as the negative turbulence profile of its corresponding atmospheric layer as explained in Section 2.5.2.

2.6.2 Mathematical model of Atmospheric Tomography for MCAO

As mentioned previously, the atmosphere itself is modelled as a vector of two dimensional functions rather than one three dimensional object. Each of these functions is called a layer of the atmosphere. All layers together represent an approximation of the turbulence of the atmosphere. Hence, for MCAO, we aim at the reconstruction of a finite number L of turbulent layers $\Phi^{(l)}$, located at known heights h_l , $l = 1, \dots, L$, each corresponding to a deformable mirror conjugated to height h_l . Available data are the incoming wavefronts φ_{β_g} , $g = 1, \dots, G$ which have been reconstructed from the measurements of the wavefronts sensors in an intermediate step. We identify the guide stars by their direction from the projection of the guide star onto the plane in which the telescope pupil is situated to the center of the telescope pupil, given by a vector $\beta_g \in \mathbb{R}^2$. Let $\bar{\beta}_g = (\bar{\beta}_{g,1}, \bar{\beta}_{g,2}, h)^T \in$

\mathbb{R}^3 be the three dimensional vector from the telescope pupil to the guide star. Then $\beta_g \in \mathbb{R}^2$ corresponds to the first two components of the vector $\bar{\beta}_g$ at altitude $h = 1$, i.e., $\beta_g = (\beta_{g,1}, \beta_{g,2})^T$ where $\bar{\beta}_g = (\beta_{g,1}, \beta_{g,2}, 1)^T$. Due to the cone effect (see Figure 2.8), the circular part of the atmosphere visible from the telescope in a certain direction shrinks to a point with increasing height. Therefore we define the shrinkage parameters $\mu_{l,\beta_g} := \frac{H_g - h_l}{H_g}$ with values between zero and one, where H_g denotes the height of the LGS point source (typically 90km). In this notation, we allow a different scaling not only for each layer l , but also in each direction β_g . The case of a natural guide star is covered by setting $\mu_{l,\beta_g} := 1$.

At altitude h_l , we will only be able to reconstruct what is “seen” by the sensors, i.e. the layer $\Phi^{(l)}$ will only be reconstructed within the area

$$\Omega_l = \bigcup_{g=1}^G \Omega_D^{\mu_{l,\beta_g}}(h_l \beta_g),$$

where

$$\Omega_D^{\mu_{l,\beta_g}}(h_l \beta_g) := \{\rho \in \mathbb{R}^2 : \mu_{l,\beta_g}^{-1}(\rho - h_l \beta_g) \in \Omega_D\}, \quad (2.18)$$

i.e., the shifted and scaled projections of the telescope pupil onto the layers in each direction. In this notation, Ω_D represents the telescope pupil with radius D , i.e.,

$$\Omega_D = \{r \in \mathbb{R}^2 : \|r\| \leq D\}.$$

We consider $\Phi^{(l)} \in L_2(\Omega_l)$ and collect all the layers in a vector

$$\Phi := (\Phi^{(1)}, \dots, \Phi^{(L)})^T \in \bigotimes_{l=1}^L L_2(\Omega_l).$$

The symbol $\bigotimes_{l=1}^L$ denotes the tensor product of the respective spaces. On this tensor product space, an inner product is defined via

$$\langle \Phi, \Psi \rangle_{\bigotimes_{l=1}^L L_2(\Omega_l)} := \sum_{l=1}^L \frac{1}{c_l} \langle \Phi^{(l)}, \Psi^{(l)} \rangle_{L_2(\Omega_l)}, \quad (2.19)$$

where c_l is the (known) relative strength of a layer in the atmosphere described in the C_n^2 profile. The wavefronts, when traveling through the atmosphere, are perturbed by the turbulences in the air. Each layer of turbulence the wavefront passes through adds to the previous perturbations. Additionally, for each guide star direction β_g only the part $\Omega_D^{\mu_{l,\beta_g}}(h_l \beta_g)$ of the layer is visible, c.f. Figure 2.8 and (2.18). Therefore, assuming geometric propagation of the light, the incoming wavefronts are given as sums of parts

of the layers [59],

$$A : \bigotimes_{l=1}^L L_2(\Omega_l) \longrightarrow (L_2(\Omega_D))^G,$$

$$A\varphi(r) := \left(\sum_{l=1}^L \Phi^{(l)}(\mu_{l,\beta_g} r + h_l \beta_g) \right)_{g=1,\dots,G}, \quad r \in \Omega_D. \quad (2.20)$$

We can interpret this as a Radon transform operator over a medium that is discretized in a special way. The Radon transform, see, e.g., [3], of a function f of n variables is given by

$$\mathbf{R}f(\theta, s) = \int_{\theta^\perp} f(s\theta + y) dy \quad (2.21)$$

i.e., the integral of f over the hyperplane perpendicular to $\theta \in S^{n-1}$ with signed distance $s \in \mathbb{R}$ from the origin. In (2.20), f corresponds to the complete three dimensional turbulence profile of the atmosphere. Instead of integrals along a certain direction θ , the atmospheric tomography considers sums over the parts of the layers visible from the telescope in the direction of each guide star. A detailed discussion justifying the term “atmospheric tomography” from a mathematical point of view can be found in [57, 49]. In sum, the goal is the reconstruction of Φ from the measured data $\varphi = (\varphi_{\beta_1}, \dots, \varphi_{\beta_G})^T$ as solution of the equation

$$A\Phi = \varphi.$$

2.7 Remarks on Numerical Simulation

Whenever we conduct numerical experiments, the computations are performed with the simulation tool MOST developed by the Austrian Adaptive Optics Team [78]. Since our study is aimed specifically at the E-ELT telescope, we consider a telescope with 42m mirror diameter, as originally proposed for the E-ELT. The guide stars are positioned in a circle of radius 3.75 arcmin for laser guidestars and 5 arcmin for natural guide stars, respectively. For each guide star, a Shack-Hartmann wavefront sensor with 84×84 subapertures is used. The wavefronts φ_{β_g} , $g = 1, \dots, G$ are, if not stated otherwise, reconstructed with CuReD [47, 48] and used as input for our methods. In case only laser guide stars are considered, the data is tip/tilt removed. Otherwise, the (full) natural guide stars are used for the tip/tilt estimation. The simulated field of view is 10 arcmin. The error on the measurements of the wavefront sensor is assumed to be low due to a high number of photons per subaperture. The atmosphere is simulated with 9 layers, moving at a speed of roughly 15 m/s with a Fried parameter $r_0 = 0.129\text{m}$ at 500nm in K-band. The isoplanatic angle is $\theta_0 = 2''$, the outer scale for the van Karmann turbulence model is $\tilde{L}_0 = 20\text{m}$. This is a standard model employed by ESO. For the

clarification of the parameters we refer to Sections 2.3. In MCAO, three deformable mirrors are simulated, conjugated to heights of 0m, 4000m and 12700m, corresponding to the MAORY configuration [79]. Hence, we seek to reconstruct an artificial atmosphere consisting of 3 layers at heights $h_1 = 0\text{m}$, $h_2 = 4000\text{m}$ and $h_3 = 12700\text{m}$ and obtain the shrinkage parameters $\mu_1 = 1$, $\mu_2 = 0.956$ and $\mu_3 = 0.859$ for laser guide stars. The relative strength of the layers are $c_1 = 0.6$, $c_2 = 0.2$ and $c_3 = 0.2$, respectively.

Chapter 3

Filter-based reconstruction methods

Many linear regularization methods can be described via the filtering of the singular value expansion which will be the topic of this chapter. After reviewing some well-known deterministic convergence results we discuss two particular methods both introduced as *fractional Tikhonov regularization*. We prove convergence of both variants and discuss their applicability in comparison to standard Tikhonov regularization both theoretically and in numerical examples. This part is largely taken from the paper [80]. In the last section of the chapter, we derive convergence results for general filter methods with the stochastic noise model.

3.1 Regularization methods and filter factors

We start in the deterministic setting of Section 1.2 and consider a linear operator equation $Ax = y$ (1.1) between Hilbert spaces \mathcal{X} and \mathcal{Y} where instead of the correct data we are only provided with noisy data y^δ satisfying $\|y - y^\delta\| \leq \delta$ for $\delta > 0$. Assume that A is compact. Then the *singular system* of A can be used to construct regularization methods. It is derived as follows, see for example [7, 8]. With A^* being the adjoint of A , A^*A is a self-adjoint linear operator and admits a singular value decomposition

$$(A^*A)x = \sum_{n \in \mathbb{N}} \sigma_n^2 \langle x, v_n \rangle v_n \quad (3.1)$$

where σ_n^2 and v_n , $n \in \mathbb{N}$, are the eigenvalues and eigenvectors of A^*A satisfying

$$(A^*A)v_n = \sigma_n^2 v_n.$$

Let the eigenvalues be arranged in decreasing order and set $u_n := \sigma_n^{-1} A v_n$. Then the triple $\{\sigma_n, u_n, v_n\}$ is called the *singular system* of A , the σ_n are called *singular values*. The following proposition collects some important properties of the singular system. The proofs can be found, e.g., in [7, 8].

Proposition 3.1.1. *Let $\{\sigma_n, u_n, v_n\}$ be the singular system of the compact operator $A \in \mathcal{L}(X, Y)$. Then it holds*

- $\{v_n\}$, $n \in \mathbb{N}$, is a complete orthonormal system for $\overline{\mathcal{R}(A^*)} = \mathcal{N}(A)^\perp$
- $\{u_n\}$, $n \in \mathbb{N}$, is a complete orthonormal system for $\overline{\mathcal{R}(A)} = \mathcal{N}(A^*)^\perp$
- for $x \in \mathcal{D}(A)$ it is

$$Ax = \sum_n \sigma_n \langle x, v_n \rangle u_n$$

- for $y \in \mathcal{D}(A^\dagger)$ it is

$$A^\dagger y = \sum_{\sigma_n > 0} \sigma_n^{-1} \langle y, u_n \rangle v_n \quad (3.2)$$

The singular vector calculus also allows another characterization of the norm $\|\cdot\|_\nu$ from (1.5). Namely, it is

$$\|x\|_\nu^2 = \sum_n \sigma_n^{-2\nu} |\langle x, v_n \rangle|^2. \quad (3.3)$$

For the solution of the Inverse Problem (1.1) in particular the last property of Proposition 3.1.1 is of interest. Given the singular system, the construction of the generalized inverse of A is straight forward. Since the singular values go to zero for compact operators, their inverse grows unbounded. If noisy data $y^\delta = y + \epsilon$ has to be inserted into (3.2), one obtains

$$A^\dagger y^\delta = \sum_{\sigma_n > 0} \sigma_n^{-1} \langle y, u_n \rangle v_n + \sum_{\sigma_n > 0} \sigma_n^{-1} \langle \epsilon, u_n \rangle v_n,$$

i.e., the noise is amplified and the sum does not converge unless the noise decays fast enough (which is most often not the case in practice). The idea of filter methods is to control the noise amplification by introducing a real valued *filter function* $F_\alpha(\sigma)$ as additional factor in (3.2). Here, we excluded the possibility of defining the filter functions also in dependence of the data y^δ since then we no longer obtain linear regularization methods. Typically, the $F_\alpha(\sigma)$ depend on the regularization parameter α which has to be chosen appropriately. Thus,

$$R_\alpha y^\delta := \sum_{\sigma_n > 0} F_\alpha(\sigma_n) \sigma_n^{-1} \langle y^\delta, u_n \rangle v_n \quad (3.4)$$

furnishes an approximation to the generalized inverse (3.2). Under appropriate conditions on the filter function, $R_\alpha : \mathcal{Y} \rightarrow \mathcal{X}$ from (3.4) defines a regularization method in the sense of Definition 1.2.2.

Theorem 3.1.2. [7, Theorem 3.3.3.] *Let the filter F_α fulfill*

$$\begin{aligned} \sup_n |F_\alpha(\sigma_n) \sigma_n^{-1}| &= c(\mu) < \infty, \\ \lim_{\mu \rightarrow 0} F_\alpha &= 1 \quad \text{pointwise in } \sigma_n, \\ |F_\alpha(\sigma_n)| &\leq c \quad \text{for all } \alpha \text{ and } \sigma_n. \end{aligned}$$

Then the operators R_α from (3.4) are regularizations of A^\dagger , i.e., it holds for $y \in \mathcal{D}(A^\dagger)$

$$\lim_{\alpha \rightarrow 0} \|A^\dagger y - R_\alpha y\| = 0.$$

The first of the three conditions ensures that the filter does not allow uncontrolled noise amplification, the second and third one ensure convergence of R_α to A^\dagger as $\alpha \rightarrow 0$. Stricter requirements on the filter functions F_α allow to deduce order optimality (see Section 1.2.2) of the induced regularization method.

Theorem 3.1.3. [7, Theorem 3.4.3] Let $y \in \mathcal{R}(A)$ and $\|y - y^\delta\|_Y \leq \delta$. Assume that, in addition to the conditions of Theorem 3.1.2, it holds $\|x^\dagger\|_\nu \leq \varrho$ and for $0 \leq \nu \leq \nu^*$,

$$\sup_{0 < \sigma \leq \sigma_1} \sigma^{-1} |F_\alpha(\sigma)| \leq c\mu^{-\beta} \quad (3.5)$$

$$\sup_{0 < \sigma \leq \sigma_1} |1 - F_\alpha(\sigma)| \sigma^{\nu^*} \leq c_{\nu^*} \mu^{\beta\nu^*}, \quad (3.6)$$

where $\beta > 0$ and c, c_{ν^*} are constants independent of δ . Then with the a priori parameter choice

$$\mu = C \left(\frac{\delta}{\varrho} \right)^{1/\beta(\nu+1)}, \quad C > 0 \text{ fixed}, \quad (3.7)$$

the method induced by the filter F_α is order optimal for all $0 \leq \nu \leq \nu^*$.

Proof. We shall include the proof from [7, Theorem 3.4.3] for later reference. It is

$$\|A^\dagger y - R_\alpha y^\delta\| \leq \|(A^\dagger - R_\alpha)y\| + \|R_\alpha(y - y^\delta)\|. \quad (3.8)$$

For the noise term we obtain with (3.5)

$$\begin{aligned} \|R_\alpha(y - y^\delta)\| &= \left(\sum_{\sigma_n > 0} (F_\alpha(\sigma_n) \sigma_n^{-1})^2 |\langle y - y^\delta, u_n \rangle|^2 \right)^{1/2} \\ &\leq \sup_{\sigma_n > 0} |F_\alpha(\sigma_n) \sigma_n^{-1}| \delta \leq c\alpha^{-\beta} \delta. \end{aligned} \quad (3.9)$$

The filter error can be estimated via

$$\begin{aligned} \|(A^\dagger - R_\alpha)y\| &= \left(\sum_{\sigma_n > 0} (1 - F_\alpha(\sigma_n))^2 \sigma_n^{-2} |\langle y, u_n \rangle|^2 \right)^{1/2} \\ &= \left(\sum_{\sigma_n > 0} (1 - F_\alpha(\sigma_n))^2 |\langle x, v_n \rangle|^2 \right)^{1/2} \\ &= \left(\sum_{\sigma_n > 0} (F_\alpha(\sigma_n) - 1)^2 \sigma_n^{2\nu} \sigma_n^{-2\nu} |\langle x, v_n \rangle|^2 \right)^{1/2} \\ &\leq \sup_{\sigma_n > 0} |(F_\alpha(\sigma_n) - 1) \sigma_n^\nu| \|x\|_\nu \leq c_\nu \alpha^{\beta\nu} \varrho. \end{aligned} \quad (3.10)$$

using (3.6) and $\|x^\dagger\|_\nu \leq \varrho$ from (3.3). Inserting the parameter choice (3.7) into (3.10) and (3.9) in combination with (3.8) yields the order optimality since then

$$\begin{aligned} \|A^\dagger y - R_\alpha y^\delta\| &\leq c \left(C \left(\frac{\delta}{\varrho} \right)^{1/\beta(\nu+1)} \right)^{-\beta} \delta + c_\nu \left(C \left(\frac{\delta}{\varrho} \right)^{1/\beta(\nu+1)} \right)^{\beta\nu} \varrho \\ &\leq \tilde{C} \delta^{\frac{\nu}{\nu+1}} \varrho^{\frac{1}{\nu+1}} \end{aligned}$$

with a constant $\tilde{C} > 0$ independent of δ and ρ . \square

To close this section, we shortly mention three widely used filter-based regularization methods. The first one, called *truncated singular value decomposition*, follows directly from the observation that the smaller the singular values become, the more the noise is amplified. This effect is counteracted by dropping all elements in the sum (3.4) whenever the corresponding singular value is below a threshold $\alpha > 0$. The parameter α corresponds to the regularization parameter and has to be chosen appropriately in order to find an appropriate balance between the noise amplification and the cutting off of important parts of the solution. The filter function is given by

$$F_{\alpha}^{trunc}(\sigma) = \begin{cases} \sigma & \sigma \geq \alpha \\ 0 & \sigma < \alpha \end{cases}. \quad (3.11)$$

The other two filter-based methods presented here have been mentioned Chapter 1 already. Namely, Tikhonov regularization in standard form (1.9) can be written in the form (3.4) with

$$F_{\alpha}^{tikh}(\sigma) = \frac{\sigma^2}{\sigma^2 + \alpha}, \quad (3.12)$$

[7], and the Landweber method corresponds to a filtering with the function

$$F_k^{LW}(\sigma) = 1 - (1 - \gamma\sigma^2)^k \quad (3.13)$$

where k is the iteration number, corresponding to the regularization parameter $\alpha \sim \frac{1}{k}$ [81] and γ the stepsize parameter in 1.11.

3.2 Fractional Tikhonov regularization

3.2.1 Fractional Tikhonov regularization in the literature

It is well known that Tikhonov regularization in standard form (1.9) typically determines a regularized solution x_{α}^{δ} that is too smooth, i.e., many details of the desired solution x^{\dagger} are not represented by x_{α}^{δ} . This shortcoming led Klann and Ramlau [81] to introduce the fractional Tikhonov regularization method. Subsequently another approach, also referred to as fractional Tikhonov regularization, was investigated by Hochstenbach and Reichel [82]. The latter approach fits both into the framework of generalized Tikhonov regularization introduced by Louis [7, Chapter 4] and into the framework presented in [83]. A more detailed comparison of the three approaches is presented in Section 3.2.2. Application of the fractional approach of [82] to Lavrentiev regularization is discussed in [84].

Let us quickly recall that Tikhonov regularization in standard form can be expressed via the normal equation

$$(A^*A + \alpha I)x_{\alpha}^{\delta} = A^*y^{\delta}, \quad (3.14)$$

or equivalently by the optimization problem

$$x_\alpha^\delta = \min_{x \in \mathcal{D}(A)} \|Ax - y^\delta\|_{\mathcal{Y}}^2 + \alpha \|x\|_{\mathcal{X}}^2. \quad (3.15)$$

The method in [82] can be derived by replacing the Y -norm in the fidelity term in (3.15) by a weighted seminorm

$$\|y\|_W := \|W^{1/2}y\|_{\mathcal{Y}}$$

with

$$W = (AA^*)^{(\mu-1)/2} \quad (3.16)$$

for some parameter $0 \leq \mu \leq 1$, where W is defined with the aid of the Moore–Penrose pseudoinverse of AA^* when $\mu < 1$. We obtain the minimization problem

$$\min_{x \in X} \|Ax - y^\delta\|_W^2 + \alpha \|x\|_{\mathcal{X}}^2 \quad (3.17)$$

and denote its solution by \tilde{x}_α^δ . It can be computed by solving the associated normal equations

$$((A^*A)^{(\mu+1)/2} + \alpha I)x = (A^*A)^{(\mu-1)/2}A^*y^\delta. \quad (3.18)$$

Oversmoothing in Tikhonov regularization in standard form (which corresponds to $\mu = 1$) is caused by the fact that y^δ is multiplied by A^* . Letting $0 < \mu < 1$ reduces oversmoothing. We will show that, by the theory presented in [7, 83], choosing $-1 < \mu \leq 1$ is possible.

Klann and Ramlau [81] propose another approach to reduce oversmoothing. They advocate that an approximation of x^\dagger be computed by solving

$$(A^*A + \alpha I)^\mu x = (A^*A)^{\mu-1}A^*y^\delta \quad (3.19)$$

for some $0 < \mu \leq 1$, where $(A^*A)^{\mu-1}$ is defined with the Moore–Penrose pseudoinverse when $\mu < 1$. This leads to an interpolation between standard Tikhonov regularization and the generalized inverse. We denote the solution by \hat{x}_α^δ . Also this method simplifies to Tikhonov regularization in standard form when $\mu = 1$. An associated minimization problem is not known.

The normal equations (3.18) and (3.19) already resemble the one of Tikhonov regularization in standard form (3.14). The resemblance increases by considering the associated filter functions. Namely, Tikhonov regularization in standard form can be characterized by the filter function (3.12). Thus, the minimizer of (1.9),(3.15) can be computed as

$$x_\alpha^\delta = \sum_{\sigma_n > 0} \frac{\sigma_n}{\sigma_n^2 + \alpha} \langle y^\delta, u_n \rangle v_n. \quad (3.20)$$

as well as by (3.14). The filter function associated with the fractional Tikhonov regularization method (3.17) is given by

$$\tilde{F}_{\alpha,\mu}(\sigma) = \frac{\sigma^{\mu+1}}{\sigma^{\mu+1} + \alpha} \quad (3.21)$$

and gives the associated approximation

$$\tilde{x}_{\alpha}^{\delta} = \sum_{\sigma_n > 0} \frac{\sigma_n^{\mu}}{\sigma_n^{\mu+1} + \alpha} \langle y^{\delta}, u_n \rangle v_n \quad (3.22)$$

of x^{\dagger} . This expression is provided in [82] with slightly different notation. The fractional Tikhonov method (3.19) can be written in terms of a filter function in a similar fashion. We have

$$\hat{F}_{\alpha,\mu}(\sigma) := (F_{\alpha}^{\text{Tikh}}(\sigma))^{\mu} = \left(\frac{\sigma^2}{\sigma^2 + \alpha} \right)^{\mu}. \quad (3.23)$$

The corresponding approximation of x^{\dagger} is given by

$$\hat{x}_{\alpha}^{\delta} = \sum_{\sigma_n > 0} \left(\frac{\sigma_n^{2-1/\mu}}{\sigma_n^2 + \alpha} \right)^{\mu} \langle y^{\delta}, u_n \rangle v_n. \quad (3.24)$$

Graphs of the filter functions for various values of the free parameters can be found in [81] and [82], respectively. From now on we refer to the method (3.17) of [82] by (3.22), the method (3.19) of [81] by (3.24) and Tikhonov regularization in standard form by (3.20).

3.2.2 Order optimality of fractional Tikhonov methods

As mentioned previously, the method (3.22), introduced here as described in [82], can be treated as a special case of the methods proposed by Louis [7] and Mathé–Tautenhahn [83]. We will show in this section that the results of [7] and [83] coincide for the special case of fractional Tikhonov regularization (3.22). It is worth mentioning that neither in [7] nor [83] the applicability of the methods was discussed; no numerical experiments were shown. One of the major concerns of this work is to show that fractional Tikhonov methods are only in certain situations preferable to Tikhonov regularization in standard form.

Generalized Tikhonov regularization as in [7] is obtained by replacing the penalty term in (3.15) by $\|Bx\|_{\mathcal{X}}^2$, where $B : \mathcal{N}(A)^{\perp} \rightarrow \mathcal{X}$ is an operator whose domain $\mathcal{D}(B)$ is dense in $\mathcal{N}(A)^{\perp}$ and $(B^*B)^{-1} : \mathcal{N}(A)^{\perp} \rightarrow \mathcal{X}$ is continuous. The associated functional is

$$J_{\mu,B}(x) := \|Ax - y^{\delta}\|_{\mathcal{Y}}^2 + \alpha \|Bx\|_{\mathcal{X}}^2. \quad (3.25)$$

Certain conditions on the operator B allow for results on optimality and order optimality of generalized Tikhonov regularization; see [7]. Louis considers a special case of generalized Tikhonov regularization (3.25) and obtains with $B^*B = (A^*A)^{-\zeta}$ the normal equations

$$(A^*A + \alpha(A^*A)^{-\zeta})x_\alpha^\delta = A^*y^\delta, \quad (3.26)$$

where $(A^*A)^{-\zeta}$ is replaced by the Moore–Penrose pseudoinverse if A is not invertible. He [7, Satz 4.2.3] establishes that this method is order optimal for $\zeta \geq -1/2$ under an a priori parameter choice rule if x^\dagger fulfills $\|x^\dagger\|_\nu \leq \varrho$ with $\nu \leq 2\zeta + 2$. Multiplying (3.26) with $(A^*A)^\zeta$ and setting $\zeta = \frac{\nu-1}{2}$ we obtain (3.18). Hence, with this correlation between the parameters, we can immediately conclude that the fractional Tikhonov method (3.22) is order optimal for $\mu \geq 0$ and $\nu < \mu + 1$. Results from Louis's book [7] make it possible to extend the range of μ to $-1 < \mu \leq 1$ using Theorem 3.1.3.

Lemma 3.2.1. *The regularizing filter $\tilde{F}_{\alpha,\mu}(\sigma)$ from (3.21) with $-1 < \mu \leq 1$ fulfills (3.5),(3.6) with $\beta = \frac{1}{\mu+1}$ and $\nu^* = \mu + 1$.*

Proof. The filter $\tilde{F}_{\alpha,\mu}(\sigma)$ is continuous on $(0, \infty)$. The regularizing properties of the Filter $F_{\alpha,\mu}$ are easily verified. One sees that

$$\begin{aligned} \lim_{\sigma \rightarrow 0} \sigma^{-1} |\tilde{F}_{\alpha,\mu}(\sigma)| &= \lim_{\sigma \rightarrow \infty} \sigma^{-1} |\tilde{F}_{\alpha,\mu}(\sigma)| = 0, \\ \lim_{\sigma \rightarrow 0} |1 - F_{\alpha,\mu}(\sigma)| \sigma^\nu &= 0, \\ \lim_{\sigma \rightarrow \infty} |1 - F_{\alpha,\mu}(\sigma)| \sigma^\nu &= \begin{cases} \infty & \nu > \mu + 1, \\ 1 & \nu = \mu + 1, \\ 0 & \nu < \mu + 1. \end{cases} \end{aligned}$$

Hence, as long as $\nu < \mu + 1$, the suprema in (3.5),(3.6) are attained as local maxima, which can be derived by simple calculus. One obtains

$$\sup_{0 < \sigma \leq \sigma_1} \sigma^{-1} |F_{\alpha,\mu}(\sigma)| \leq c\alpha^{-\frac{1}{\mu+1}} \quad \text{and} \quad \sup_{0 < \sigma \leq \sigma_1} |1 - F_{\alpha,\mu}(\sigma)| \sigma^{\nu^*} \leq c_{\nu^*} \alpha^{\frac{1}{\mu+1}\nu}.$$

□

Consequently, by Theorem 3.1.3 the fractional method (3.22) is order optimal for $-1 < \mu \leq 1$ and $0 \leq \nu \leq \mu + 1$.

Now let us investigate the connection of the fractional Tikhonov regularization (3.22) to [83]. There, without mentioning the previous work by Louis [7], the authors considered the solution of the equation

$$((A^*A)^{s+1} + \mu I)x_\mu^\delta = (A^*A)^s A^*y^\delta, \quad (3.27)$$

for $s > -1$, assuming that

$$\|(A^*A)^{q/2}(y - y^\delta)\| \leq \delta \quad \text{and} \quad x^\dagger \in \{x \in X : (A^*A)^{p/2}v, \|v\| \leq E\}. \quad (3.28)$$

Setting $s = \frac{\mu-1}{2}$, $q = 0$, $p = \nu$ and $E = \varrho$, we arrive exactly at the fractional method (3.22) in the setting introduced in Section 3.2. The following result is shown in [83].

Theorem 3.2.2. *Under the assumptions (3.28) and with $\alpha = \left(\frac{\delta}{E}\right)^{(2s+2)/(p+q+1)}$, the method (3.27) is of optimal order provided that $s \geq \max(\frac{p-2}{2}, \frac{q-1}{2})$.*

Translating this theorem into our notation of the fractional Tikhonov method (3.22), we obtain order optimality provided that $\mu \geq \max(\nu - 1, 0)$, i.e., $\nu \leq \mu + 1$ in case $\nu > 1$. We summarize the results in the following theorem.

Theorem 3.2.3. *Let $A : \mathcal{X} \rightarrow \mathcal{Y}$ be a linear compact operator between Hilbert spaces \mathcal{X} and \mathcal{Y} . Let $x^\dagger := A^\dagger y$ satisfy $\|x^\dagger\|_\nu \leq \varrho$. Then for all $-1 < \mu \leq 1$ and $0 \leq \nu \leq \mu + 1$ the fractional Tikhonov method (3.22) is of optimal order under the a-priori parameter choice rule*

$$\alpha = C \left(\frac{\delta}{\varrho} \right)^{(\mu+1)/(\nu+1)}.$$

While the method (3.22) is order optimal for all $-1 < \mu \leq 1$ and appropriate ν , this is not the case for the fractional Tikhonov method (3.24). We have the following result.

Proposition 3.2.4. *[81, Proposition 3.2] Let $A : \mathcal{X} \rightarrow \mathcal{Y}$ be a compact operator with singular system $\{\sigma_n, u_n, v_n\}_{n \geq 0}$, and let $x^\dagger := A^\dagger y$ satisfy $\|x^\dagger\|_\nu \leq \varrho$ for some constant ϱ and the ν -norm defined by (3.3). Then for $\alpha \in (1/2, 1]$, the fractional Tikhonov method defined (3.24) is order optimal with the parameter choice rule*

$$\alpha = C \left(\frac{\delta}{\varrho} \right)^{1/2(\nu+1)}$$

for all $0 < \nu < 2$. Here C is a positive constant independent of δ and ϱ .

Klann and Ramlau [81, Theorem 4.4] show that after appropriate presmoothing of the error-contaminated data y^δ , fractional powers $0 < \mu \leq 1/2$ together with a suitable choice of the regularization parameter α yield quasi-optimal convergence rates.

The above approaches to determine α generally are not very useful for the solution of specific problems. When an accurate estimate of the norm of the error in the data $\|y^\delta - y\|_{\mathcal{Y}}$ is known, the *discrepancy principle*, discussed, e.g., in [8, 85], can be applied to determine a suitable value of α . We already briefly introduced the concept in Section 1.2.2. Here we will make use of the filter properties and, instead of (1.12) choose the regularization parameter $\alpha = \alpha(\delta, y^\delta)$ such that

$$\|Ax_\alpha^\delta - y^\delta\|_{\mathcal{Y}} = \tau\delta, \quad (3.29)$$

where $\tau > 1$ is again a user-supplied constant independent of δ . This is a nonlinear equation for α . Its solution can be calculated by finding the positive zero of

$$G_\mu(\alpha) := \sum_{n \in \mathbb{N}} (1 - F_{\alpha, \mu}(\sigma)) \langle y^\delta, u_n \rangle^2 - (\tau\delta)^2, \quad (3.30)$$

for example with Newton's method; see, e.g., [82] for further details. This reference discusses linear discrete ill-posed problems, but the results carry over to the setting of the present work.

Convergence of regularized approximate solutions determined by filtered regularization methods using the discrepancy principle has been analyzed in the works of Louis [7] and Mathé and Tautenhahn [83]. As for the a-priori parameter choice, the order optimality of the fractional Tikhonov method (3.22) is derived immediately from both [7] and [83] by identifying corresponding parameters as done in the first part of this section. Thus without going into detail again, the following theorem follows independently from [7, Theorem 3.5.2] and [83, Theorem 2.4].

Theorem 3.2.5. *Let $A : \mathcal{X} \rightarrow \mathcal{Y}$ be a linear compact operator between Hilbert spaces \mathcal{X} and \mathcal{Y} . Let $x^\dagger := A^\dagger y$ satisfy $\|x^\dagger\|_\nu \leq \varrho$. Then for all exponents $\mu > 0$ and $0 < \nu \leq \mu$, the fractional Tikhonov method (3.22) is order optimal with the regularization parameter α determined by the discrepancy principle (3.29).*

Remark. It might appear appealing to substitute the standard norm in (3.29) by the weighted norms from (3.17) (or (3.28) for $-1 < q < 0$, respectively). Then with again $W = (A^*A)^{(\mu-1)/2}$,

$$\|A\tilde{x}_\alpha^\delta - y^\delta\|_W^2 = \sum_{\sigma_n > 0} (1 - \tilde{F}_{\alpha, \mu}(\sigma_n))^2 \sigma_n^{\mu-1} \langle y^\delta, u_n \rangle^2. \quad (3.31)$$

However, since $\lim_{\sigma_n \rightarrow 0} \tilde{F}_{\alpha, \mu}(\sigma_n) = 0$, the sum typically will converge very slowly since for large n the inner products $\langle y^\delta, u_n \rangle$ generally are dominated by the error in y^δ . In a discrete setting using finite-precision arithmetic, the quantities $\langle y^\delta, u_n \rangle$ also are contaminated by propagated round-off errors introduced during the computations. The residual (3.31) therefore will be very large due to error amplification, and the equation

$$\|A\tilde{x}_\alpha^\delta - y^\delta\|_W = \tau\delta$$

is not guaranteed to have a solution. Hence, the weighted residual norm is in general not useful for determining the regularization parameter in actual computations.

Order optimality for the fractional Tikhonov method (3.24) can be shown with the help of the results in [7].

Theorem 3.2.6. *Let $A : \mathcal{X} \rightarrow \mathcal{Y}$ be a linear compact operator between Hilbert spaces X and Y . Let $x^\dagger := A^\dagger y$ satisfy $\|x^\dagger\|_\nu \leq \varrho$. Then for all exponents $\mu \in (1/2, 1]$ and $0 < \nu \leq 1$, the fractional Tikhonov method of (3.24) is order optimal with the regularization parameter α given by the discrepancy principle (3.29).*

Proof. Klann and Ramlau showed in [81] that the filter function (3.23) fulfills (3.5), (3.6) with $\beta = \frac{1}{2}$ and $\nu^* = 2$ for $\mu > \frac{1}{2}$. Hence, by [7, Theorem 3.5.2], the claim follows. \square

Approximations of x^\dagger determined by fractional Tikhonov regularization typically are closer to x^\dagger in the X -norm than approximations obtained with Tikhonov regularization in standard form; see [82] for computed examples. However, a smaller error does not always correspond to a more pleasing approximation of x^\dagger , because the fractional Tikhonov approximation may be more oscillatory than the approximation determined by Tikhonov regularization in standard form. We would like to elucidate in which situations fractional methods yield more pleasing approximations. The following lemma is helpful. A similar result has been shown in [82].

Lemma 3.2.7. *Let the mappings $\alpha \mapsto F_{\alpha, \mu}(\sigma)$ and $\mu \mapsto F_{\alpha, \mu}(\sigma)$ be continuous and monotonically decreasing for $\alpha > 0$ and μ in an interval $\underline{\mu} < \mu < \bar{\mu}$. Let $\alpha = \alpha(\mu)$ be determined by the discrepancy principle (3.29). Then $\frac{d\alpha(\mu)}{d\mu} < 0$.*

Proof. We can write (3.30) in the form $G(\mu, \alpha(\mu)) = 0$. Since G is differentiable, we have

$$\frac{dG}{d\alpha} = \sum_{\sigma_n > 0} 2(1 - F_{\alpha, \mu}(\sigma_n)) \cdot (-1) \cdot \frac{dF_{\alpha, \mu}}{d\alpha} \cdot \langle y^\delta, u_n \rangle^2 > 0,$$

because $1 - F_{\alpha, \mu}(\sigma) > 0$ $\frac{dF_{\alpha, \mu}}{d\alpha} = < 0$. Analogously, one finds that $\frac{dG}{d\mu} > 0$. Hence, by the implicit function theorem,

$$\frac{d\alpha}{d\mu} = - \left(\frac{dG}{d\alpha} \right)^{-1} \frac{dG}{d\mu} < 0.$$

\square

An immediate consequence of the above lemma is that decreasing μ results in an increase of the regularization parameter α . It is therefore inappropriate to compare fractional methods with the standard Tikhonov filter using the same regularization parameter.

We are now in position to take a closer look at the computed approximations. Again we will make use of the explicit representation of the solution in terms of the singular system of A . Let as in (1.2) $\epsilon = y^\delta - y$. Since

$$\sigma_n \langle x^\dagger, v_n \rangle = \langle x^\dagger, A^* u_n \rangle = \langle y, u_n \rangle,$$

cf. [8], and

$$\langle y^\delta, u_n \rangle = \langle y, u_n \rangle + \langle \epsilon, u_n \rangle,$$

the approximation error $e(\delta, \alpha, \mu) := x^\dagger - R_\alpha y^\delta$ is given by

$$e(\delta, \alpha, \mu) = \sum_{\sigma_n > 0} (1 - F_{\alpha, \mu}(\sigma_n)) \langle x^\dagger, v_n \rangle v_n + \sum_{\sigma_n > 0} F_{\alpha, \mu}(\sigma_n) \frac{1}{\sigma_n} \langle -\epsilon, u_n \rangle u_n. \quad (3.32)$$

Let ϵ be fixed. The quality of the computed solution then is determined by the positive coefficients $1 - F_{\alpha, \mu}(\sigma)$ and $F_{\alpha, \mu}(\sigma)$. One immediately sees that the filter $F_{\alpha, \mu}(\sigma)$ has to achieve two contradicting properties: $F_{\alpha, \mu}(\sigma)$ should be close to one to give a small deviation of the reconstruction from x^\dagger , and $F_{\alpha, \mu}(\sigma)$ should be close to zero in order to effectively reducing propagation of the error ϵ into the computed approximation.

It is not obvious from (3.32) in which situations letting $\mu < 1$ improves the quality of the computed approximation of x^\dagger . We can shed some light on this by studying the derivative $\frac{d}{d\mu} F_{\alpha, \mu}(\sigma)$. We first consider the filter function (3.21). Since α depends on μ (the function (3.30) used to calculate α depends on μ via the filter functions $F_{\alpha, \mu}(\sigma)$), we get

$$\frac{d}{d\mu} \tilde{F}_{\alpha, \mu}(\sigma) = -\frac{d}{d\mu} \left(1 - \tilde{F}_{\alpha, \mu}(\sigma) \right) = h(\sigma, \mu, \alpha(\mu)) \left(\ln \sigma - \frac{\alpha'(\mu)}{\alpha(\mu)} \right), \quad (3.33)$$

where $h(\sigma, \mu, \alpha(\mu))$ is a positive function. The sign of the derivative is determined by the factor $\ln \sigma - \frac{\alpha'(\mu)}{\alpha(\mu)}$. When μ and the error norm δ are fixed, so is α , and the sign only depends on σ . By Lemma 3.2.7, $\alpha'(\mu) < 0$. Therefore, the derivative (3.33) changes sign at some $0 < \tilde{\sigma}_0 < 1$. Only for n with $\sigma_n < \tilde{\sigma}_0$, the coefficient of $\langle x^\dagger, v_n \rangle$ in (3.32) will be reduced by decreasing μ , since then $\frac{d}{d\mu} (1 - \tilde{F}_{\alpha, \mu}(\sigma_n)) > 0$. Hence, the coefficient of $\langle \epsilon, u_n \rangle$ increases. The opposite holds true for the coefficients of the terms associated with the propagated error. Whereas for large σ_n the propagated error is damped, it is amplified for all $\sigma_n < \tilde{\sigma}_0$.

The result for the fractional filter (3.23) is analogous. Similarly to (3.33), one has

$$\begin{aligned} \frac{d}{d\mu} \hat{F}_{\alpha, \mu}(\sigma) &= -\frac{d}{d\mu} \left(1 - \hat{F}_{\alpha, \mu}(\sigma) \right) \\ &= \hat{h}(\sigma, \mu, \alpha(\mu)) \left(-\ln \left(\frac{\sigma^2 + \alpha(\mu)}{\sigma^2} \right) - \alpha \frac{\alpha'(\mu)}{\sigma^2 + \alpha(\mu)} \right) \end{aligned} \quad (3.34)$$

with $\hat{h}(\sigma, \mu, \alpha(\mu)) > 0$. The logarithm is positive and $\alpha'(\mu) < 0$. Therefore, the sign of (3.34) changes at some $\sigma = \hat{\sigma}_0 > 0$. Hence, the above discussion also applies to this filter function. However, it is not clear whether the operator A has singular values that satisfy $\sigma_n > \hat{\sigma}_0$. If this is not the case, then decreasing μ will result in error amplification in all components of the computed approximate solution.

Although it is an open problem how to determine a value of μ that yields the best approximation of x^\dagger , we can identify two situations in which fractional Tikhonov methods outperform Tikhonov regularization in standard form:

- a) the problem is severely ill-posed, i.e., the singular values of A decrease rapidly to zero, and
- b) the error in y^δ is concentrated to low frequencies.

In case the problem is severely ill-posed, $\tilde{\sigma}_0$ and $\hat{\sigma}_0$ are likely to be large enough for the propagated error to be damped. A slight loss in accuracy of terms in (3.22) and (3.24) associated with large singular values is typically acceptable, since they are much larger than the error and therefore usually are recovered quite accurately. On the other hand, if there is only little error in the high frequency components in (3.22) and (3.24), the amplification of the error in b^δ is largely avoided, while the reconstruction is improved. In other cases, both fractional methods do not perform significantly better than Tikhonov regularization in standard form. The reason for this can again be found in the dependency of the filter factors $F_{\mu,\alpha}(\sigma)$ on the parameters α and μ . By decreasing α , the $F_{\mu,\alpha}(\sigma)$ increase. At the same time, decreasing α leads to increasing regularization parameter μ as shown in Lemma 3.2.7. From the definition of the filter factors (3.21) and (3.23), respectively, one sees that this leads to decreasing values of the filter factors. Hence, both effects cancel each other out to some extent. Although α is decreased below one, the filter factors corresponding to larger singular values stay almost constant. The following section provides some illustrative computed examples.

3.2.3 Numerical examples

Several computed examples that show the performance of the fractional Tikhonov regularization methods discussed in this paper are provided in [82, 81]. These examples demonstrate that it may be attractive to use fractional Tikhonov methods instead of standard Tikhonov regularization. In this paper, we present a few examples that show the relative performance of the fractional methods (3.22) and (3.24), and that illustrate the comments in the last paragraph of the previous section.

Our first example is a severely ill-posed Fredholm integral equation of the first kind given by

$$y_1(s) = [A_1x](s) = \int_0^1 \sqrt{s^2 + t^2}x(t)dt, \quad 0 \leq s \leq 1, \quad (3.35)$$

with error-free data $y_1(s) = \frac{1}{3}((1+s^2)^{3/2} - s^3)$ and solution $x_1^\dagger(t) = t$. This equation was first introduced by Fox and Goodwin, cf. [86]. Numerically, the singular values decrease exponentially until they stagnate around attainable computational precision. We used the discretization of (3.35) provided in Regularization Tools [87]. This gave a 1000×1000 matrix.

The second example is the mildly ill-posed Volterra integral equation of the first kind

$$y_2(s) = [A_2x](s) = \int_0^s x(t)dt, \quad 0 \leq s \leq 1, \quad (3.36)$$

with error-free data

$$y_2(s) = \begin{cases} -s & 0 \leq s \leq 0.5, \\ s-1 & 0.5 < s \leq 1, \end{cases}$$

and solution

$$x_2^\dagger(t) = \begin{cases} -1 & 0 \leq t \leq 0.5, \\ 1 & 0.5 < t \leq 1. \end{cases}$$

This example has been used in [81]. The coefficients $\langle x_1^\dagger, v_n \rangle$ decrease slowly to zero. The singular system $\{\sigma_n; u_n, v_n\}_{n \geq 1}$ of A_2 (without discretization) is given in [7]. The integration problem (3.36) has been discretized with the Nyström method based on the trapezoidal rule with 1000 equidistant nodes.

In all experiments shown, we equipped both the domain and range of the discretized operators with the Euclidean vector norm. We added noise with prescribed noise level $\theta := \frac{\|y-y^\delta\|}{\|y\|}$ to the error-free data y .

We compare approximate solutions obtained by the two fractional methods (3.22) (denoted by $\tilde{x}_{\alpha,\mu}^\delta$) and (3.24) (denoted by $\hat{x}_{\alpha,\mu}^\delta$) with the approximate solution determined by Tikhonov regularization in standard form (3.20) (denoted by \bar{x}_{α}^δ). The key ingredient for this is a comparable choice of the regularization parameter. For all approximate solutions, the regularization parameter is determined by the discrepancy principle (3.29). It turns out, that the relative performance of the fractional methods when compared to Tikhonov regularization in standard form varies significantly with the choice of the free parameter τ in (3.29). We therefore conduct the following experiment. For each $\tau \in \{1, 1.05, 1.1, 1.2, 1.3, 1.4, 1.5\}$, we compute the standard Tikhonov solution $\bar{x}_{\alpha(\tau)}^\delta$ and solutions for the fractional methods (3.22) and (3.24) for the fractional parameters μ from discrete sets $\mathcal{A}_1 := \{-1 + 0.05i, i = 0, 1, 2, \dots, 40\}$ and $\mathcal{A}_2 := \{0.05i, i = 0, 1, 2, \dots, 20\}$, respectively. Out of all these solutions, we select for each fractional method and each τ the best approximation. In other words, we choose $\mu^*(\tau)$ such that it gives the minimum deviation from the true solution x^\dagger over all μ -values considered,

$$\tilde{\mu}^* := \min_{\mu \in \mathcal{A}_1} \|\tilde{x}_{\alpha(\tau),\mu}^\delta - x^\dagger\| \quad \text{and} \quad \hat{\mu}^* := \min_{\mu \in \mathcal{A}_2} \|\hat{x}_{\alpha(\tau),\mu}^\delta - x^\dagger\|.$$

We compute for both fractional methods the relative deviation when compared to Tikhonov regularization in standard form,

$$\tilde{\text{re}}(\tau) = \frac{\|\tilde{x}_{\alpha(\tau),\tilde{\mu}^*}^\delta - x^\dagger\|}{\|\bar{x}_{\alpha(\tau)}^\delta - x^\dagger\|} \quad \text{and} \quad \hat{\text{re}}(\tau) = \frac{\|\hat{x}_{\alpha(\tau),\hat{\mu}^*}^\delta - x^\dagger\|}{\|\bar{x}_{\alpha(\tau)}^\delta - x^\dagger\|}. \quad (3.37)$$

The quantities $\widetilde{\text{re}}(\tau)$, $\widehat{\text{re}}(\tau)$, and the optimal fractional parameters are plotted as functions of τ in Figure 3.1 for the Fox–Goodwin problem (3.35). Only one realization of perturbed data is used for this experiment. Figure 3.1 is typical for many experiments. The fractional methods can be seen to yield reconstruction errors that are smaller than those obtained with Tikhonov regularization in standard form. For $\tau \geq 1.05$, the error obtained with the method (3.22) is about half the error obtained with standard Tikhonov regularization. Also direct inspection of the computed solutions, shown in Figure 3.6, strongly favors this fractional method. However, for $\tau = 1$, the difference in performance is almost negligible. In fact, in a direct comparison of the three approximate solutions computed with the fractional and standard Tikhonov methods, the difference between these solutions is barely visible. Moreover, for some noisy right-hand sides, with the same noise characteristics as in Figure 3.1 but different realizations of the noise, standard Tikhonov regularization with $\tau = 1$ gives the best reconstruction. The regularization parameters for the fractional methods and Tikhonov regularization in standard form are shown in Figure 3.2. For all $\tau \geq 1$, the regularization parameter for the latter method is much smaller than for the fractional methods.

We now carry out the same experiment as above with the integration problem (3.36). The result is shown in Figure 3.3. Observe that for this problem Tikhonov regularization in standard form with the regularization parameter obtained for $\tau = 1$ always gives the best solution. Although for larger τ -values the fractional methods give a slightly smaller error, one would nevertheless typically prefer the solution obtained with Tikhonov regularization in standard form since it is smoother than the fractional solutions. A typical case is shown in Figure 3.7. The plot of the regularization parameters is similar to the corresponding plot for the Fox–Goodwin problem; see Figure 3.2. We should mention that the relative performance of the methods in our comparison does not change significantly if a different true solution x^\dagger is chosen. The properties of the operator A are far more important than the choice of true solution for the performance of the methods.

Our experience with the Fox–Goodwin problem leads us to conclude that it is important to carefully choose the parameter τ in a comparison of the methods in order to avoid a bias towards fractional Tikhonov methods. To make this point even more evident, we repeat the above experiment, but now calculate the relative reconstruction errors

$$\widetilde{\text{re}}_1(\tau) = \frac{\left\| \widetilde{x}_{\alpha(\tau), \widetilde{\mu}^*}^\delta - x^\dagger \right\|}{\left\| \widetilde{x}_{\alpha(1)}^\delta - x^\dagger \right\|} \quad \text{and} \quad \widehat{\text{re}}_1(\tau) = \frac{\left\| \widehat{x}_{\alpha(\tau), \widehat{\mu}^*}^\delta - x^\dagger \right\|}{\left\| \widehat{x}_{\alpha(1)}^\delta - x^\dagger \right\|}, \quad (3.38)$$

with respect to the solution obtained with Tikhonov regularization in standard form for fixed $\tau = 1$; see Figure 3.4. For the integration problem the relative reconstruction errors do not change much, but for the Fox–Goodwin problem they increase by a large factor. A comparison with Figure 3.1 illustrates that, although for each value of τ the fractional methods are better than Tikhonov regularization in standard form, this claim

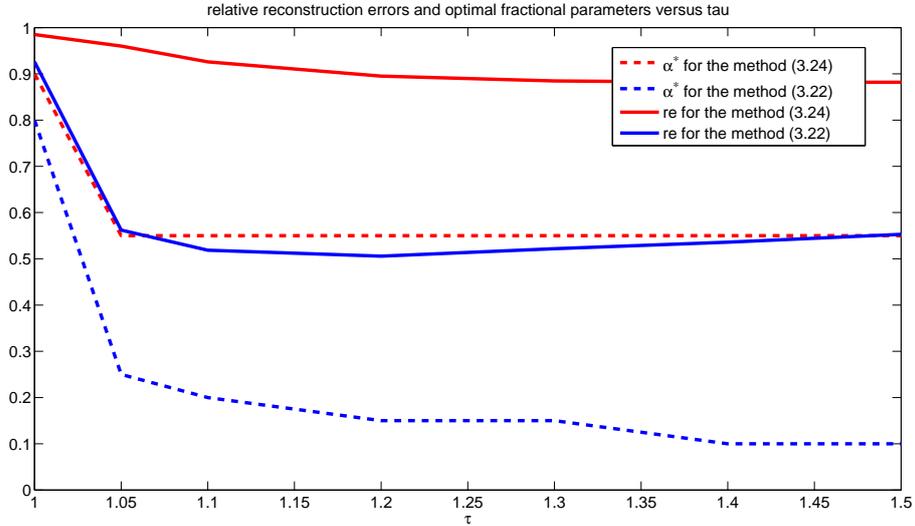


Figure 3.1: Relative errors from (3.37) and optimal fractional parameters μ^* as functions of τ for the Fox–Goodwin problem (3.35) with white noise of level $\theta = 0.05$. The larger τ is chosen, the more favorable in particular the fractional method (3.22) becomes. For $\tau = 1$, the difference is almost negligible.

does not hold anymore when $\tau = 1$ is used for Tikhonov regularization in standard form. Hence, it is important to choose τ carefully when comparing methods.

In the above examples the data was perturbed by white Gaussian noise. The observation of Section 3.2.2 lead us to repeat the experiments with low-frequency noise. An example of this kind of noise and a comparison with white noise is displayed in Figure 3.8. Figure 3.9 shows the fractional methods to give more accurate approximations of x^\dagger than Tikhonov regularization in standard form for low-frequency noise. The improved performance for low-frequency noise becomes even more evident if instead of a one-dimensional signal a two-dimensional image is take as the true solution. Figure 3.5 compares he solutions obtained with the methods (3.24), (3.22) and (3.20) for a chessboard-like image x^\dagger of size 75×75 pixels which has been blurred using a blurring matrix that is block Toeplitz with Toeplitz blocks with bandwidth 6 and approximates a Gaussian point-spread function with variance $\sigma = 0.7$. The regularization parameter was determined by the discrepancy principle (3.29) with $\tau = 1$. For this example, the method (3.24) produces the best approximation of the exact image. The reconstruction error is only 36% of the reconstruction error determined by Tikhonov regularization in standard form.

To further illustrate the different behaviors of the methods in our comparison in the

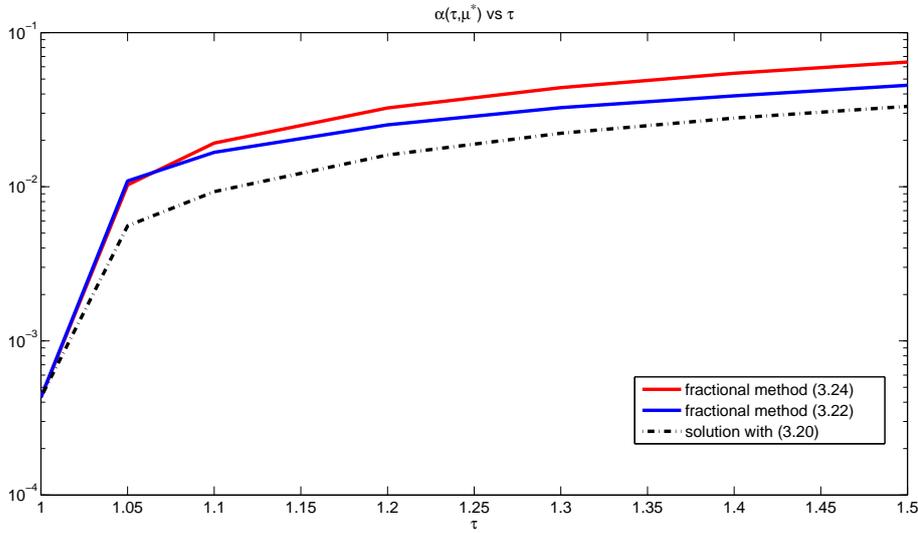


Figure 3.2: Regularization parameter for the solutions with the methods (3.22), (3.24) and (3.20) for the Fox–Goodwin problem (3.35) with white noise of level $\theta = 0.05$, obtained by the discrepancy principle (3.29). For $\tau = 1$ the regularization parameters almost coincide whereas for larger τ the regularization parameter of Tikhonov regularization in standard form is significantly smaller than for the fractional methods. The same quality of the curves is obtained for the integration problem (3.36).

settings introduced above, we include tables in which we give regularization parameters and approximation errors relative to those obtained with Tikhonov regularization in standard form (3.37). All errors are averages over 20 experiments with different error-realizations. Table 3.1 shows results for the Fox–Goodwin problem (3.35). In agreement with Figure 3.6, the fractional method (3.22) performs the best. For the problem (3.36) with Gaussian white noise, the error in the approximate solutions determined by the fractional methods is only slightly smaller than the error obtained with Tikhonov regularization in standard form, as shown in Table 3.2. However, using the same problem with low-frequency error instead of white Gaussian error, the fractional methods yield a much better approximations of x^\dagger than Tikhonov regularization in standard form; see Table 3.3.

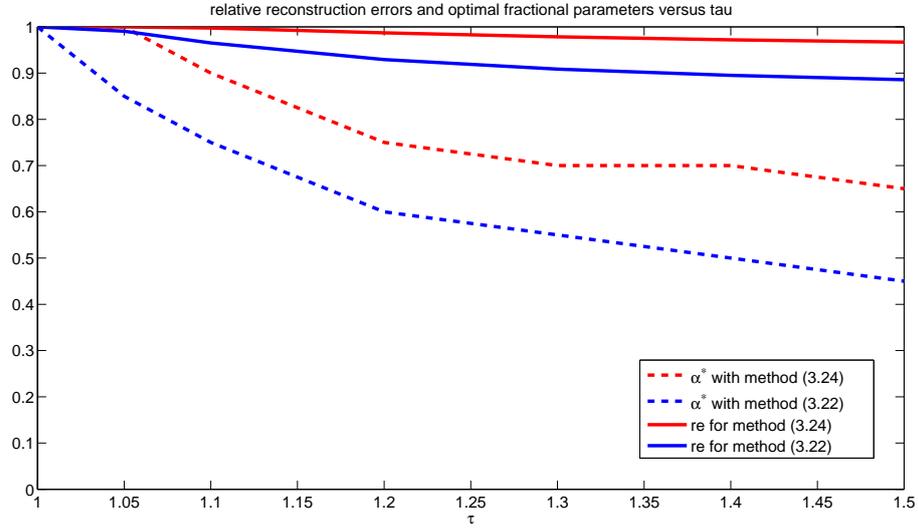


Figure 3.3: Relative errors from (3.37) and optimal fractional parameters μ^* as functions of τ for the integration problem (3.36) with white noise of level $\theta = 0.05$. Even for large values of τ , there is not a big difference between the fractional methods and Tikhonov regularization in standard form.

μ	0.05	0.1	0.3	0.5	0.6	0.7	0.9	1
$\tilde{\alpha}$	6.1e-3	5.9e-3	5.2e-3	4.3e-3	3.8e-3	3.3e-3	2.4e-3	2.0e-3
$\hat{\alpha}$	1.1e-1	3.9e-2	8.2e-3	4.4e-3	3.5e-3	2.9e-3	2.2e-3	2.0e-3
\tilde{r}_e	10.3	5.2	1.5	0.60	0.61	0.72	0.92	1
\hat{r}_e	2.4e15	4.7e13	1e7	4.2	0.93	0.91	0.97	1

Table 3.1: Regularization parameter and relative reconstruction error for both fractional filters, tilde standing for (3.22), hat for (3.24); and the Fox–Goodwin problem (3.35). In both cases μ grows monotonically with decreasing μ . The reconstruction errors (3.37) are shown in the two bottom rows. For the method (3.22), there is a minimum clearly below one. Hence, the reconstructions are significantly improved. Since for $\mu < 0.5$ the filter (3.23) is not regularizing anymore, the reconstruction error explodes.

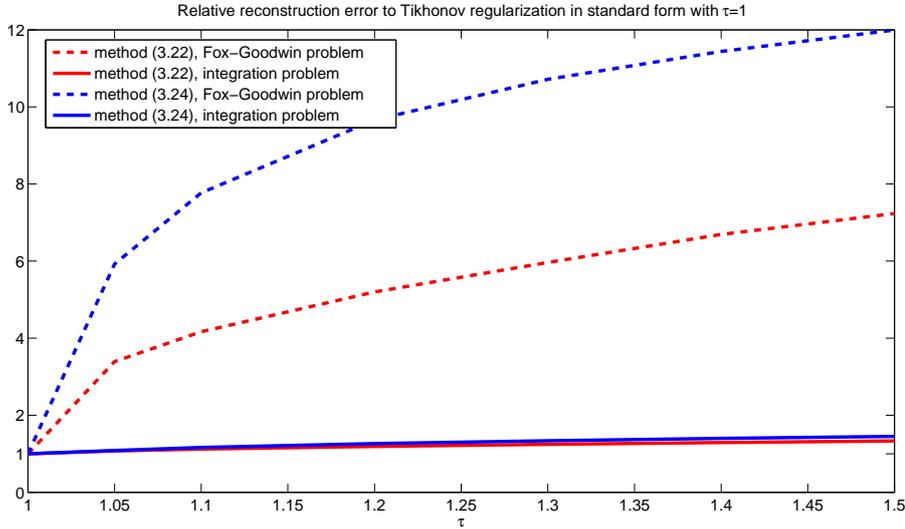


Figure 3.4: Comparison of relative reconstruction errors (3.38). For the integration problem (3.36), increasing τ does not change the error much. However, for the Fox–Goodwin problem (3.35), the difference becomes much more evident.

μ	0.05	0.1	0.3	0.5	0.6	0.7	0.9	1
$\tilde{\alpha}$	4.7e-3	4.3e-3	2.7e-3	1.6e-3	1.3e-3	1.0e-3	0.6e-3	0.4e-3
$\hat{\alpha}$	3.4e-2	1.1e-2	2.3e-3	1.1e-3	0.8e-3	0.7e-3	0.5e-3	0.4e-3
\tilde{r}_e	2.7	2.3	1.4	1.07	1.02	0.991	0.990	1
\hat{r}_e	21.8	12.5	2.3	1.13	1.03	1.007	0.998	1

Table 3.2: Regularization parameter and relative reconstruction error for both fractional filters and the integration problem (3.36). In both cases α grows monotonically with decreasing μ . The reconstruction errors (3.37) grow nearly monotonically, only for μ close to one it is slightly below one, i.e., the fractional methods give a slightly lower residual than Tikhonov regularization in standard form.

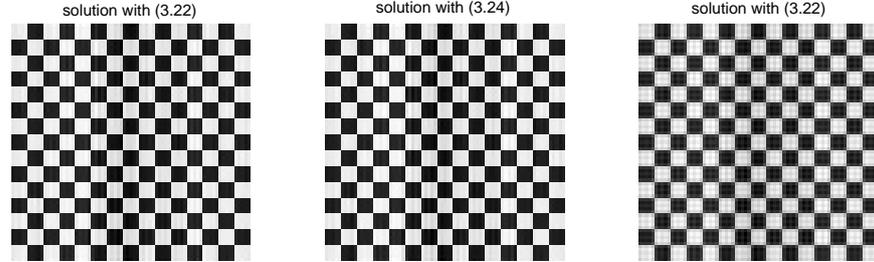


Figure 3.5: Comparison of fractional Tikhonov methods and Tikhonov regularization in standard form for a two dimensional deblurring problem of size 75×75 with low-frequency noise of noise level $\theta = 0.05$. The solutions obtained with the fractional methods provide much sharper edge recovery. The best fractional parameters were $\mu^* = 0.05$ for the method (3.24) and $\mu^* = -1$ for the method (3.22). The relative reconstruction error to Tikhonov solution in standard form (see (3.37)) are $\widehat{r}_e = 0.36$ and $\widetilde{r}_e = 0.48$, respectively.

μ	0.05	0.1	0.3	0.5	0.6	0.7	0.9	1
$\widetilde{\alpha}$	6.0e-3	5.5e-3	3.8e-3	2.6e-3	2.1e-3	1.7e-3	1.1e-3	0.9e-3
$\widehat{\alpha}$	3.9e-2	1.4e-2	3.5e-3	1.9e-3	1.5e-3	1.3e-3	1.0e-3	0.9e-3
\widetilde{r}_e	0.62	0.65	0.75	0.84	0.88	0.91	0.97	1
\widehat{r}_e	0.54	0.60	0.81	0.91	0.94	0.96	0.99	1

Table 3.3: Regularization parameter and relative reconstruction error for both filters and the integration problem (3.36) in presence of low frequency noise (cf. Figure 3.8). The reconstruction errors (3.37) are shown in the two bottom rows. Both fractional filters give a much better result than Tikhonov regularization in standard form.

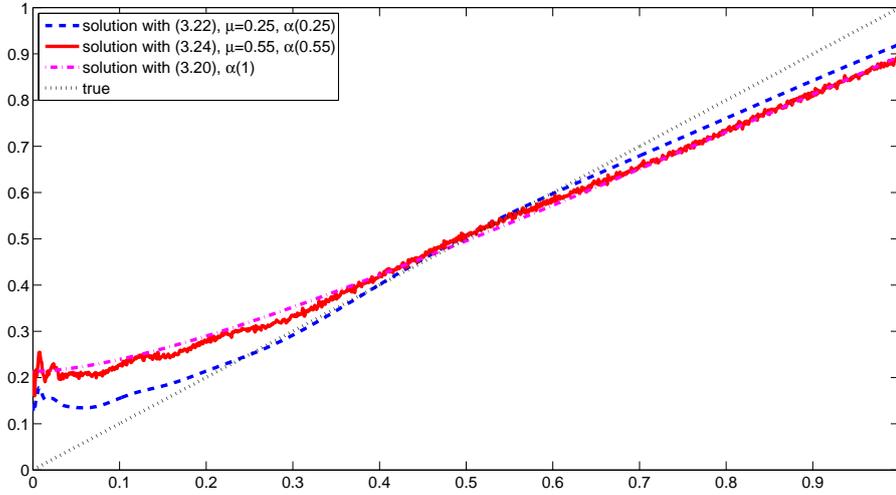


Figure 3.6: Comparison of solutions for the severely ill-posed Fox–Goodwin problem (3.35) with Gaussian white noise of level $\theta = 0.05$, α according to (3.29), $\tau = 1.1$. For the fractional methods the solutions with smallest reconstruction error are shown. The solution for the method (3.22) is plotted with the dashed line. For this type of problems it is to be preferred over the other two methods. Those are the method (3.24) (solid) and Tikhonov regularization in standard form (3.20) (dash-dotted).

3.3 Filter-based regularization methods in the stochastic setting

We now turn to the stochastic setting introduced in Chapter 1. That is, we adopt the deterministic setting from the previous sections, but substitute the worst case error bound $\|y - y^\delta\| \leq \delta$ with a stochastic assumption $\mathbb{E}(\|y - y^\eta\|) = \mathbb{E}(\|\epsilon\|) = f(\eta)$ where η is typically the variance of the error ϵ . The theory here is presented for filter-based regularization methods in general. Application of the results in particular to the fractional Tikhonov methods will be straight forward. We assume that the filters fulfill the conditions of Theorem 3.1.2 and thus furnish a regularization for A^\dagger . Note that this definition is independent of the error of the data. For the convergence analysis we use convergence in expectation as well as the Ky Fan metric as introduced in Section 1.4.

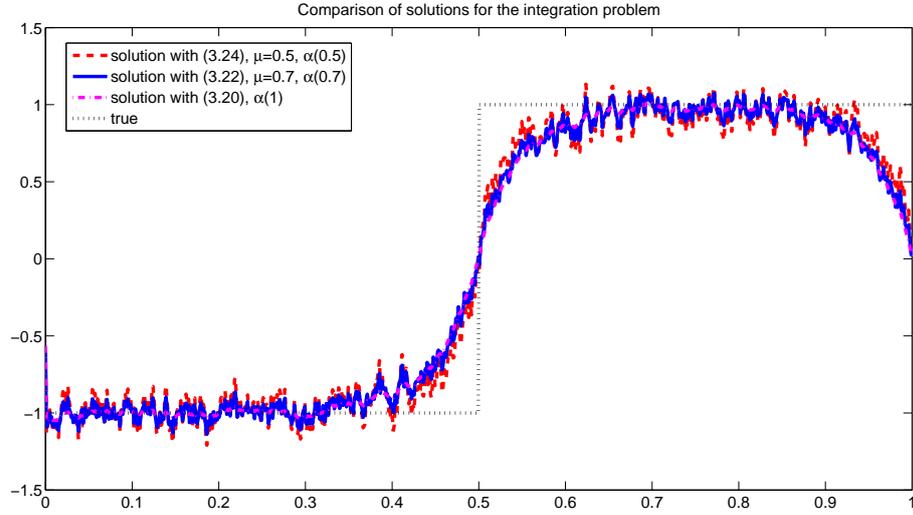


Figure 3.7: Comparison of solutions for the mildly ill-posed integration problem (3.36) with Gaussian white noise of level $\theta = 0.05$, α according to (3.29), $\tau = 1.1$. Upper plot: discontinuous solution, lower plot: smooth solution. In this case the fractional methods (3.22) and (3.24) do not perform better than Tikhonov regularization in standard form. On the contrary, the noise is amplified even more.

3.3.1 Convergence results

In case of an a priori parameter choice, we immediately obtain a result for the stochastic setting similar to the known deterministic theory by simply considering the expectation of the norms instead of the norms themselves. Theorem 3.1.3 translates directly into the new setting.

Theorem 3.3.1. *Let $y \in \mathcal{R}(A)$, $y^\eta = Ax + \epsilon$ and $\mathbb{E} \|\epsilon\| = f(\eta) < \infty$. For all $\sigma \in (0, \sigma_1]$, let $\alpha \mapsto |1 - F_{\alpha, \mu}|$ be continuous and monotonically increasing, and assume that the filter fulfils (3.5), (3.6). Let $\alpha = \alpha(\eta)$ be chosen as*

$$\alpha = C \left(\frac{\mathbb{E}(\|\epsilon\|)}{\varrho} \right)^{1/\mu(\nu+1)}, \quad C > 0 \quad \text{const.} \quad (3.39)$$

Then there is a constant d_ν independent of y , η , ϱ such that

$$\mathbb{E} \left(\left\| A^\dagger y - R_\alpha y^\eta \right\| \right) \leq d_\nu \mathbb{E}(\|\epsilon\|)^{\nu/(\nu+1)} \varrho^{1/(\nu+1)}. \quad (3.40)$$

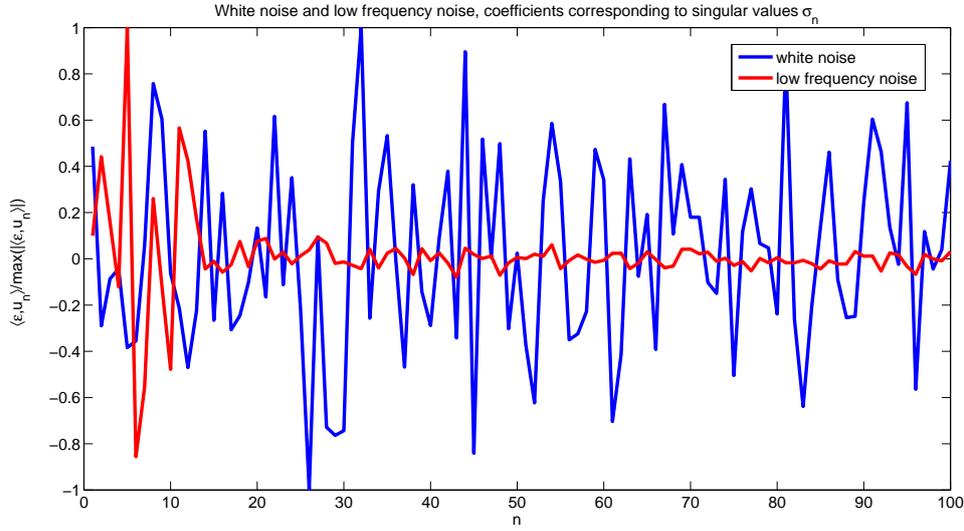


Figure 3.8: Comparison of typical random draws of white noise and low-frequency noise w.r.t. the singular values. For white noise, the coefficients are equally distributed over all singular values. The low-frequency noise decreases with growing n .

Proof. The proof follows the lines of the deterministic one, see, Theorem 3.1.3. Analogously to (3.8) it is, due to the linearity of the expectation,

$$\begin{aligned} \mathbb{E}(\|A^\dagger y - R_\alpha y^\eta\|) &\leq \mathbb{E}(\|(A^\dagger - R_\alpha)y\| + \|R_\alpha(y - y^\eta)\|) \\ &\leq \|(A^\dagger - R_\alpha)y\| + \mathbb{E}(\|R_\alpha(y - y^\eta)\|) \end{aligned} \quad (3.41)$$

since the first part of the sum is deterministic. Hence, for this term we obtain the same result as in (3.10). For the term featuring the stochastic noise, one has analogously to (3.9)

$$\begin{aligned} \mathbb{E}(\|R_\alpha(y - y^\eta)\|) &= \mathbb{E} \left(\left(\sum_{\sigma_n > 0} (F_\alpha(\sigma_n) \sigma_n^{-1})^2 |y^\eta - y, u_n|^2 \right)^{1/2} \right) \\ &\leq c \alpha^{-\mu/2} \mathbb{E} \|\epsilon\| \end{aligned} \quad (3.42)$$

due to Jensens inequality for concave function. A key ingredient of the proof is that due to the a priori choice the regularization parameter is not a stochastic quantity, but in view of the linearity of the expectation, merely a prescribed positive number. Choosing α according to (3.39) and inserting it into (3.41) with (3.10) and (3.42) again

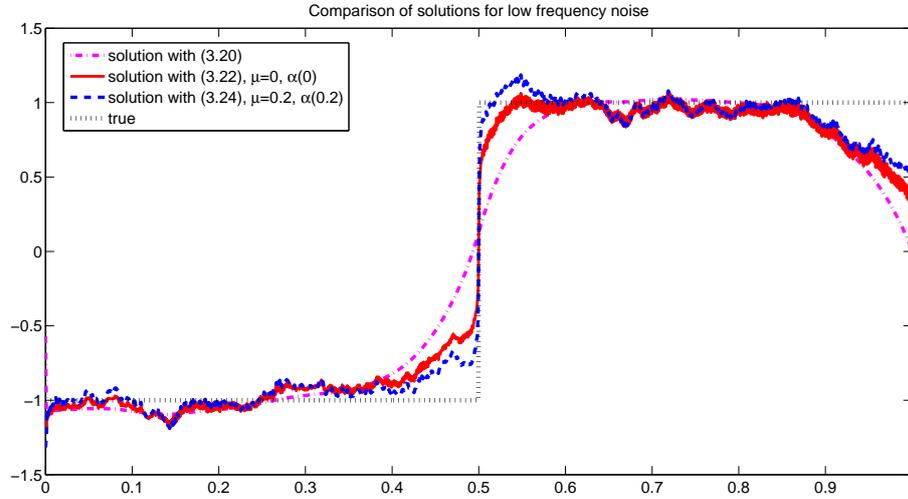


Figure 3.9: Comparison of solutions for the mildly ill-posed integration problem (3.36) with low-frequency noise (c.f. Figure 3.8) of level $\theta = 0.05$, α according to (3.29), $\tau = 1.1$. The solution of the fractional methods (3.22) (dashed) and (3.24) (solid) with appropriate μ approximate the discontinuity much better than the results of Tikhonov regularization in standard form (3.20) (dash-dotted).

results in the estimate

$$\mathbb{E}(\|A^\dagger y - R_\alpha y^\eta\|) \leq d_\nu \mathbb{E}(\|\epsilon\|)^{\nu/(\nu+1)} \varrho^{1/(\nu+1)},$$

where d_ν takes all the previous constants and does not depend on $\mathbb{E}(\|\epsilon\|)$, y and ϱ . \square

Although a-priori parameter choices of the above type may be more viable when the expectation of the noise is used rather than the worst case bound δ , the constant d_ν in (3.39) is still unknown. If by some method, for example through additional a priori knowledge or a posteriori methods a reasonable value for $d_n u$ is found, this explicit parameter choice is viable in practice, see Section 3.3.2 for numerical experiments. If $d_n u$ remains dubious, a-posteriori methods, in particular the discrepancy principle, may again be more relevant in application. In the deterministic setting, we defined the discrepancy principle in (3.29). The obvious translation of this definition to the stochastic setting is to simply replace δ with $\mathbb{E}(\|\epsilon\|)$. However, this leads to a fundamental change in the properties of α as ϵ is a stochastic quantity, i.e., $\epsilon = \epsilon(\eta, \omega)$ is one specific realization of a random variable $\mathcal{E}(\eta, \omega)$ in the complete probability space $(\Omega, \Sigma, \mathbb{P})$, see Chapter 1.3. Therefore, solving

$$\|Ax_\alpha^\eta - y^\eta\| = \tau E(\|\epsilon\|), \quad \tau > 1 \tag{3.43}$$

for α implicitly defines a function $g : (\Omega, \Sigma) \rightarrow (\Omega, \mathcal{B}(\mathbb{R}))$ such that $g(\epsilon(\omega)) = \alpha$, where $\mathcal{B}(\mathbb{R})$ denotes the Borel σ -algebra of \mathbb{R} . Hence, α is the result of a measurable function applied to a random variable and thus becomes a random variable itself. Consequently, the regularization operator R_α becomes a random operator.

Before addressing the convergence under such a random regularization parameter, we discuss a slightly different formulation of the discrepancy principle (3.29). Namely, if one finds α such that

$$\mathbb{E}(\|Ax_\alpha^\eta - y^\eta\|) = \tau E(\|\epsilon\|), \quad \tau > 1, \quad (3.44)$$

α does no longer depend on ω , i.e., α again becomes a deterministic quantity. The reason for this is that in this concept one alters the whole distribution of the residual $\|Ax_\alpha^\eta - y^\eta\|$ rather than working with realizations. As a result, convergence results can be derived from the deterministic results in a straight forward way again. We have the following theorem.

Theorem 3.3.2. *Let $y \in \mathcal{R}(A)$, $y^\eta = Ax + \epsilon$ and $\mathbb{E} \|\epsilon\| = f(\eta) < \infty$. For all $\sigma \in (0, \sigma_1]$, let $\alpha \mapsto |1 - F_{\alpha, \mu}|$ be continuous and monotonically increasing, and assume that the filter fulfills (3.5), (3.6). Let $\alpha = \alpha(\eta)$ be determined via (3.44). Then there is a constant d_ν independent of y , η , ϱ such that*

$$\mathbb{E} \left(\left\| A^\dagger y - R_\alpha y^\eta \right\| \right) \leq d_\nu \mathbb{E}(\|\epsilon\|)^{\nu/(\nu+1)} \varrho^{1/(\nu+1)}. \quad (3.45)$$

Proof. Since α is, by construction, a deterministic quantity, the remainder of the proof can be retraced from the deterministic one in [7, Theorem 3.5.2] or, in slightly different formulation of the filter approach, [8, Theorem 4.17]. Again, all inequalities hold in expectation due to its linearity and the construction of the regularization parameter. Since the proof is rather long and does not provide any additional insight, we omit it here. \square

Although equation (3.44) allows for simple theoretical treatment of the regularization method, it is not obvious how to solve the problem in practice as the formulation requires to adjust the distribution of $\|Ax_\alpha^\eta - y^\eta\|$. Since we do not know the true solution x or exact data y , this is not possible in practice. Instead, we suggest to approximate the expectation statistically. Assume that several measurements $y_i^\eta = y + \epsilon(\omega_i)$, $i = 1, \dots, N$, are available. Then we replace $\mathbb{E}(\|Ax_\alpha^\eta - y^\eta\|)$ by the mean of the available data, i.e.

$$\mathbb{E}(\|Ax_\alpha^\eta - y^\eta\|) \approx \frac{1}{N} \sum_{i=1}^N \|Ax_\alpha^\eta - y_i^\eta\|.$$

Since $\|Ax_\alpha^\eta - y_i^\eta\|$ is decreasing for $\alpha \rightarrow 0$ due to the properties of the filter functions, we start with a sufficiently large α_0 such that $\frac{1}{N} \sum_{i=1}^N \|Ax_{\alpha_0}^\eta - y_i^\eta\| > \tau \mathbb{E}(\|\epsilon\|)$. From

there, we decrease the regularization parameter slowly until we find a first α_* satisfying

$$\frac{1}{N} \sum_{i=1}^N \|Ax_{\alpha_*}^\eta - y_i^\eta\| < \tau \mathbb{E}(\|\epsilon\|). \quad (3.46)$$

The numerical experiments in Section 3.3.2 suggest that this is indeed a viable approach.

Let us now return to the case of finding α according to (3.43). Assume we have a fixed realization ϵ of the noise with $\|\epsilon\| < \infty$. We already discussed convergence properties in the deterministic setting. Whenever the norm of the given realization of the error is below the threshold, $\|\epsilon\| \leq \tau \mathbb{E}(\|\epsilon\|)$, we can write $\tau \mathbb{E}(\|\epsilon\|) = \tilde{\tau} \|\epsilon\|$ for some $\tilde{\tau} > 1$. Hence, (3.43) turns into (3.29) and according to deterministic theory we obtain

$$\|A^\dagger y - R_\alpha y^\eta\| \leq d_{\nu, \tilde{\tau}} \mathbb{E}(\|\epsilon\|)^{\nu/(\nu+1)} \varrho^{1/(\nu+1)}$$

where $d_{\nu, \tilde{\tau}} = d_\nu \tilde{\tau}^{\nu/(\nu+1)}$. If, however, $\|\epsilon\| > \tau \mathbb{E}(\|\epsilon\|)$, we can not deduce information about the reconstructed solution as the error may be arbitrarily large. Therefore,

$$\mathbb{P}\left(\|A^\dagger y - R_\alpha y^\eta\| \geq d_{\nu, \tau} \mathbb{E}(\|\epsilon\|)^{\nu/(\nu+1)} \varrho^{1/(\nu+1)}\right) \leq \mathbb{P}(\|\epsilon\| > \tau \mathbb{E}(\|\epsilon\|)). \quad (3.47)$$

Whenever the probability on the right goes to zero we have convergence of the regularized to the true solution in probability. The actual value depends on the distribution of the noise. However, in general it follows from Markov's inequality (1.25) that

$$\mathbb{P}(\|\epsilon\| \geq \tau \mathbb{E}(\|\epsilon\|)) \leq \frac{1}{\tau}$$

independent of the variance of the noise. Lemma 1.3.2 iii) shows that even after an exact calculation the probability $\mathbb{P}(\|\epsilon\| \geq \tau \mathbb{E}(\|\epsilon\|))$ might be independent of the variance. Since we require $\mathbb{P}(\|\epsilon\| \geq \tau \mathbb{E}(\|\epsilon\|))$ to converge to zero we therefore alter the formulation and allow τ to depend on the variance of the noise. If $\tau(\eta) \rightarrow \infty$ for $\eta \rightarrow 0$ slowly enough such that still $\tau(\eta) \mathbb{E}(\|\epsilon\|) \rightarrow 0$, then

$$\mathbb{P}(\|\epsilon\| \geq \tau(\eta) \mathbb{E}(\|\epsilon\|)) \leq \frac{1}{\tau(\eta)} \rightarrow 0 \quad (3.48)$$

without specifying the type of noise. The speed of convergence might be slow. However, when better estimates of the probability are available, the situation improves significantly. Consider the Gaussian distribution as in Lemma 1.3.2. The term $\Gamma(\frac{m}{2}, \frac{m\tau^2}{2})/\Gamma(\frac{m}{2})$ is, for not too small values of m and τ , close to zero and decays exponentially, see Figure 3.10. Therefore, already moderate values of τ lead to a small probability. This agrees with the intuition that τ should not be too close to 1 to compensate the uncertainty between expectation and actual value of the norm of the noise.

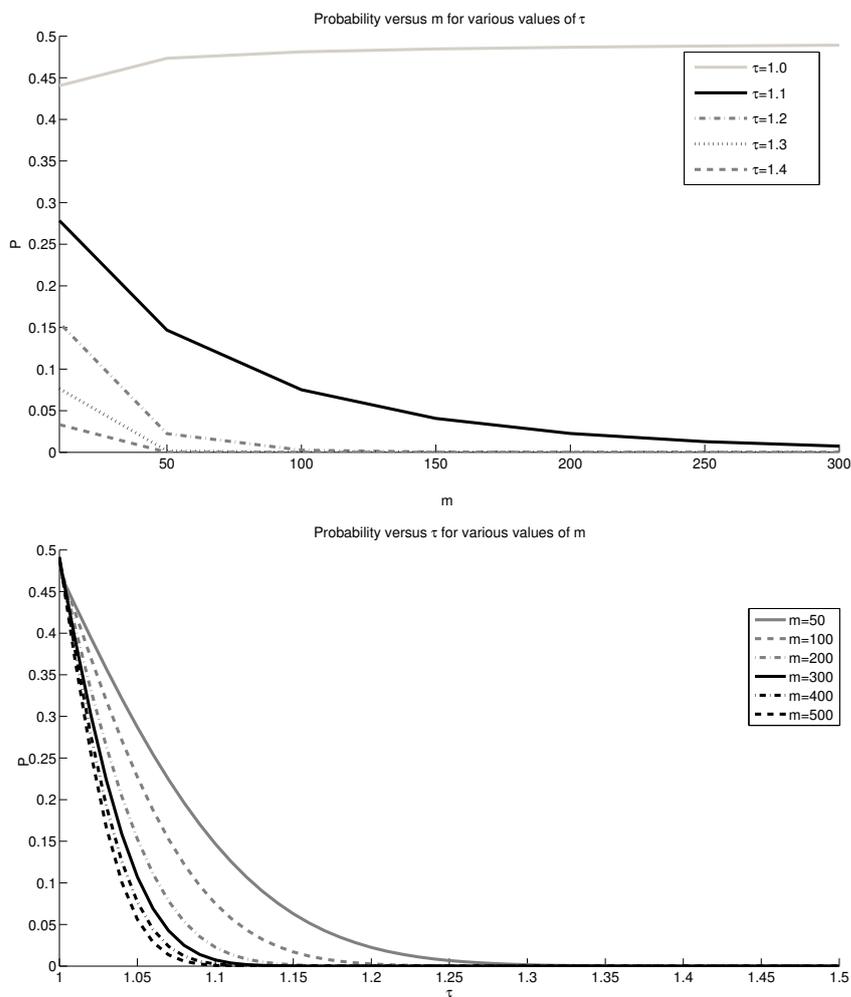


Figure 3.10: Plots of $\mathbb{P}(\|\epsilon\| \geq \tau \mathbb{E}(\|\epsilon\|))$ for i.i.d Gaussian noise $\epsilon \in \mathbb{R}^m$, $\epsilon_i \sim \mathcal{N}(0, \eta^2)$, $i = 1, \dots, m$. Top: \mathbb{P} versus m for various values of τ . Bottom: \mathbb{P} versus τ for various values of m .

Using this strategy, we implement the variance-dependent τ in the discrepancy principle to obtain the equation

$$\|Ax_\alpha^\eta - y^\eta\| = \tau(\eta)E(\|\epsilon\|) \quad (3.49)$$

where

$$\tau(\eta) \rightarrow \infty \quad \text{as } \eta \rightarrow 0, \quad \tau(\eta) > 1 \quad \forall 0 \leq \eta \leq \eta_0 \quad \text{and} \quad \lim_{\eta \rightarrow 0} \tau(\eta)E(\|\epsilon\|) = 0. \quad (3.50)$$

The probability estimate (3.47) then leads directly to a convergence Theorem with the maximum principle of Section 1.4.

Theorem 3.3.3. *Let $b \in \mathcal{R}(\cdot)A$, $y^\eta = Ax + \epsilon$ and $\mathbb{E}\|\epsilon\| = f(\eta) < \infty$. Let $\alpha = \alpha(\eta, y^\eta)$ be chosen via (3.49) with $\tau(\eta)$ fulfilling (3.50). Then*

$$\rho_K(A^\dagger y, R_\alpha y^\eta) \leq \max \left\{ d_\nu(\tau(\eta)\mathbb{E}(\|\epsilon\|))^{\nu/(\nu+1)} \varrho^{1/(\nu+1)}, \mathbb{P}(\|\epsilon\| > \tau(\eta)\mathbb{E}(\|\epsilon\|)), 1 \right\}. \quad (3.51)$$

In particular it holds that

$$\lim_{\eta \rightarrow 0} \rho_K(A^\dagger y, R_\alpha y^\eta) = 0.$$

If additionally Assumption 1.4.2 holds and one sets $x_\alpha^\eta := 0$ if any of the two conditions is not satisfied, one has

$$\lim_{\eta \rightarrow \infty} \mathbb{E}(\|A^\dagger y - R_\alpha y^\eta\|) = 0.$$

Proof. For any noise model with $\mathbb{E}(\|\epsilon\|) < \infty$ Markov's inequality gives

$$\mathbb{P}(\|\epsilon\| > \tau(\eta)\tau_0\mathbb{E}(\|\epsilon\|)) \leq \frac{\mathbb{E}(\|\epsilon\|)}{\tau(\eta)\tau_0\mathbb{E}(\|\epsilon\|)} = \frac{1}{\tau(\eta)\tau_0}. \quad (3.52)$$

Since $\tau(\eta) \rightarrow \infty$ for $\eta \rightarrow 0$ the probability goes to zero, i.e., $R_\mu b^\eta$ converges to $A^\dagger b$ in probability and hence in the Ky Fan metric. Enforcing boundedness of the solutions implies convergence in expectation, see Section 1.4 \square

Remark. If operator and noise allow a bias-variance decomposition

$$\mathbb{E}(\|x^\dagger - x_\alpha^\eta\|^2) = \|x^\dagger - x_\alpha\|^2 + \mathbb{E}(\|x_\alpha - x_\alpha^\eta\|^2),$$

see for example [28, 14], then all convergence theorems including convergence rates can be stated for stochastic L_2 convergence $\mathbb{E}(\|x^\dagger - x_\alpha^\eta\|^2)$

Lastly, we want to discuss the feasibility of the discrepancy principle in the stochastic setting. Again, the results follow from the deterministic theory. Let y^δ be a given

realization of noisy measured data. Then, for simplicity assuming the mapping $\alpha \mapsto |1 - F_\alpha(\sigma)|$ is continuous, the equation

$$\|AR_\alpha y^\delta - y^\delta\| = C, \quad C > 0,$$

is solvable for

$$\|(I - \mathcal{P}_{\overline{\mathcal{R}(A)}})y^\delta\| \leq C \leq \|y^\delta\|, \quad (3.53)$$

where $\mathcal{P}_{\overline{\mathcal{R}(A)}}$ is the projection of y^δ onto $\overline{\mathcal{R}(A)}$. Namely, for $\alpha \rightarrow \infty$ one has $\|AR_\alpha y^\delta - y^\delta\| \rightarrow \|y^\delta\|$ because $\lim_{\alpha \rightarrow \infty} F_\alpha(\sigma) = 0$ and $\sup_{\sigma, \alpha} |F_\alpha(\sigma)| \leq \infty$, c.f. Theorem 3.1.3 or [7]. In case $C > \|y^\delta\|$ one chooses $\alpha = \infty$. On the other hand, for $\alpha \rightarrow 0$ the filter functions go to one (see again Theorem 3.1.3), and hence the regularization operator R_α turns in the limit into the generalized inverse A^\dagger , compare (3.4) and (3.2). Since $AA^\dagger y^\delta = \mathcal{P}_{\overline{\mathcal{R}(A)}} y^\delta$, the lower bound in (3.53) follows. In the stochastic setting it may occur that a given realization of the noise is so large that the discrepancy principle based on that single realization fails. It is therefore advisable to choose τ larger than usual in the deterministic theory in order to decrease the probability of overly large noise, see Figure 3.10 for an example with Gaussian noise.

3.3.2 Numerical Results

In this section we will verify our results numerically. Our test problem is the inverse heat equation

$$y(t) = \int_0^t \frac{(s-t)^{-\frac{3}{2}}}{2\sqrt{\pi}} \exp\left(-\frac{1}{4(s-t)}\right) x(s) ds. \quad (3.54)$$

We used the discretization of (3.54) provided in Regularization Tools [87]. The true solutions x_ν^\dagger , depending on smoothness parameter ν , were created by multiplying a vector $h \in \mathbb{R}^n$, $h_i = 1$ for $i = 1, \dots, n$, with the operator $(A^*A)^{\nu/2}$. Thus $x_\nu^\dagger \in \mathcal{R}((A^*A)^{\nu/2})$ and $\varrho = \|h\|$.

We chose to show results for a discretization of the solution on $n = 200$ points, resulting in a 200×200 matrix A . The reason for not displaying too fine discretization levels is that in this case the difference between stochastic and deterministic curves becomes basically invisible, while at the same time the quality of the curves remains the same. In all experiments shown, we equipped both the domain and range of the discretized operators with the Euclidean vector norm and added Gaussian noise $\epsilon \sim \mathcal{N}(0, \eta^2 I_n)$ to the error-free data y .

In a first experiment we address the a priori parameter choice rule (3.39) for Tikhonov regularization in standard form where we set $C := 1$. We do not consider fractional Tikhonov regularization in order to avoid having another free parameter. Tikhonov regularization in standard form fulfills (3.5) and (3.6) with $\beta = \frac{1}{2}$, see Lemma 3.2.1 with $\mu = 1$.

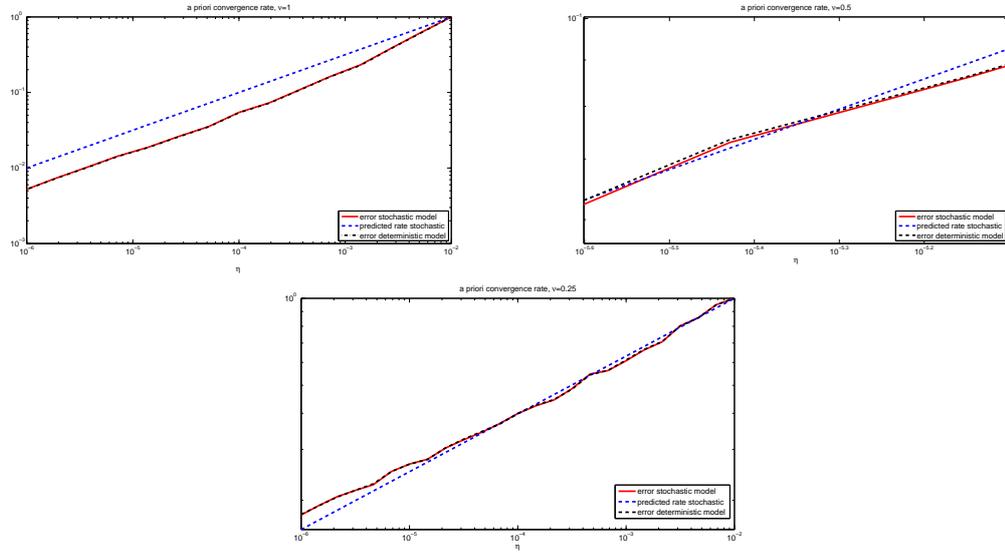


Figure 3.11: Log-log plots of a priori parameter choice with $\nu = 1$ (top left), $\nu = 0.5$ (top right) and $\nu = 0.25$ (bottom). Shown are the observed stochastic rate (red), the deterministic rate (black) and the predicted rate (blue). The convergence rates are normalized since the constant is unknown. Therefore, the slope of the curves is of primary interest.

In order to approximate the expectation of the reconstruction error, the reconstruction error has been averaged over 50 individual noise realizations for each variance η . We plot, for three values of ν , the predicted bound (3.40), averaged reconstruction error using the stochastic a-priori parameter choice rule and, for comparison, the reconstruction error under the well established deterministic a priori parameter choice rule. The values of the variance η are chosen from a log-uniform grid of 20 values between 0.000001 and 0.01, see Figure 3.11.

In the second experiment we seek to verify the convergence results for the discrepancy principle based on (3.44). As suggested, we approximate the expectation of the residual by (3.46). The expectation of the residual is approximated by averaging $N = 10$ residuals obtained from different realizations of the noise. Again we simply employ Tikhonov regularization in standard form and as before, plot for three values of ν the predicted bound (3.40), averaged reconstruction error using the stochastic parameter choice rule (3.44) and, for comparison, the reconstruction error under the well established deterministic discrepancy principle (3.29). The value of τ is identical for both. The values of the variance η are chosen from a log-uniform grid of 20 values between 0.000001 and 0.01, see Figure 3.12.

Finally the discrepancy principle (3.43) is addressed. We use Landweber's method in

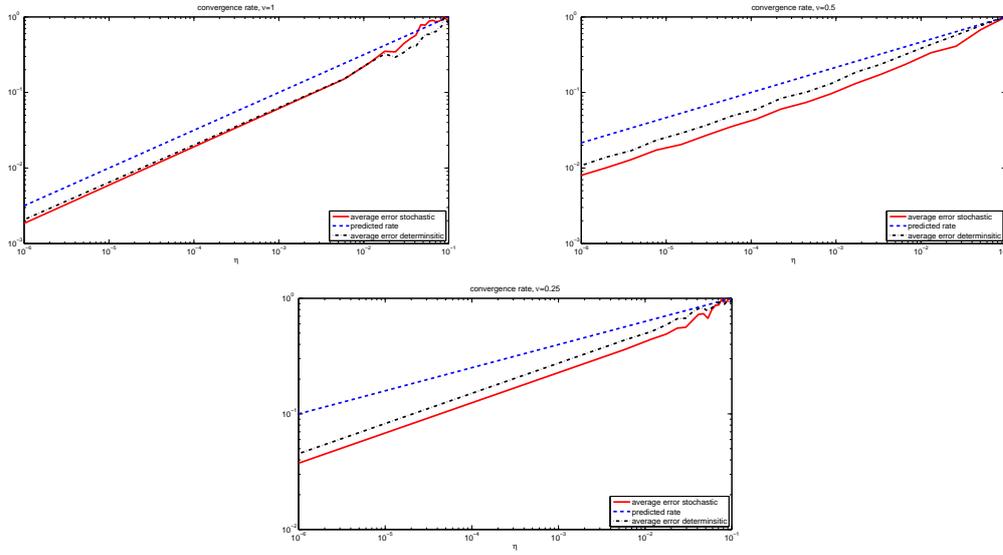


Figure 3.12: Log-log plots of a priori parameter choice with $\nu = 1$ (top left), $\nu = 0.5$ (top right) and $\nu = 0.25$ (bottom). Shown are the observed stochastic rate (red), the deterministic rate (black) and the predicted rate (blue). The convergence rates are normalized since the constant is unknown. Therefore, the slope of the curves is of primary interest.

order to give another example of a filter based regularization method. The Landweber method has been introduced in Section 1.2.2. The discrepancy principle is used to stop the iteration early enough. That is, we stop at the first solution for which

$$\|Ax_k^\eta - y^\eta\| \leq \tau(\eta)\mathbb{E}(\|\epsilon\|).$$

In the experiments we let

$$\tau(\eta) = 0.3|\log(\eta)|\eta^{-0.015}$$

For Gaussian noise it is $\tau(\eta)\mathbb{E}(\|\epsilon\|) \leq 0.3|\log(\eta)|\eta^{-0.015}\eta\sqrt{m} \rightarrow 0$ as $\eta \rightarrow 0$.

As before, we compare the predicted convergence rate with the observed rate in both stochastic and deterministic setting for three particular values of ν , see 3.13. The slope of the observed stochastic convergence rate in the log-log plots is slightly lower compared to the deterministic one. This coincides with the analytic result of Theorem 3.3.3 due to the parameter τ .

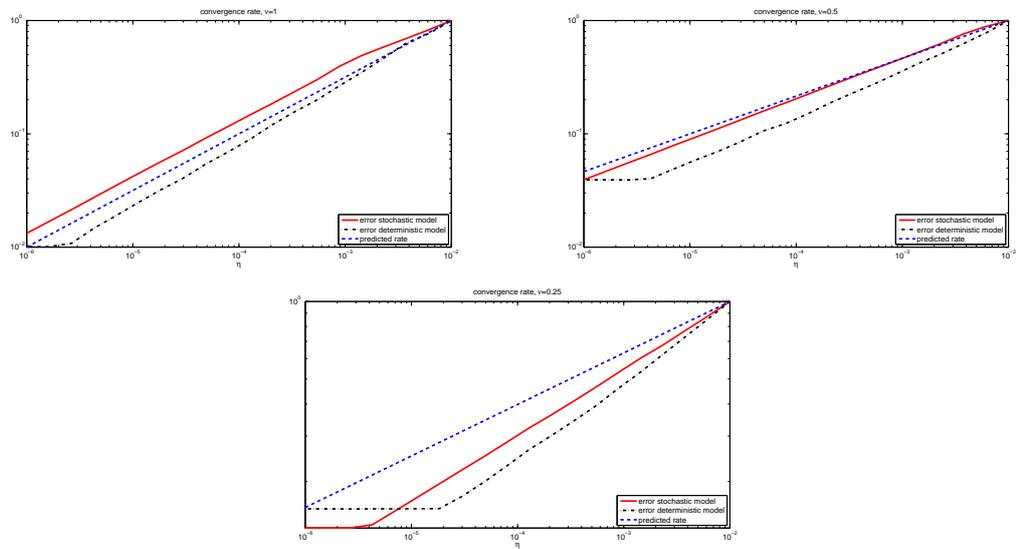


Figure 3.13: Log-log plots of a priori parameter choice with $\nu = 1$ (top left), $\nu = 0.5$ (top right) and $\nu = 0.25$ (bottom). Shown are the observed stochastic rate (red), the deterministic rate (black) and the predicted rate (blue). The convergence rates are normalized since the constant is unknown. Therefore, the slope of the curves is of primary interest.

Chapter 4

A sparsity promoting regularization method in the stochastic setting

In Chapter 3 we have shown how deterministic convergence results for a certain type of linear regularization methods for linear ill-posed problems can be lifted to fit stochastic error models. In this chapter, we stick with linear ill-posed problems, but address a particular nonlinear regularization method that we seek to apply to the problem of Atmospheric Tomography. We derive a Tikhonov-type functional from a pure Bayesian approach and show during the convergence analysis that the lifting arguments in particular allow the construction of a non-standard a-priori parameter choice rule. Large parts of this chapter are taken from the paper [88].

4.1 Sparsity promoting regularization methods

4.1.1 Introduction and short survey

Recently, sparsity promoting regularization methods have gained much attention in the Inverse Problems community. In this framework, it is assumed that the unknown x can be approximated well with only few coefficients of its expansion with respect to a preassigned dictionary Λ in \mathcal{X} . Let, for example, $\{\phi_j\}_{j \in \mathbb{N}}$, be an orthonormal basis (ONB) for the (infinite dimensional) Hilbert space \mathcal{X} . Then a function

$$\mathbf{x} = \sum_{j \in \mathbb{N}} \langle \mathbf{x}, \psi_j \rangle \psi_j \in \mathcal{X}$$

is sparse in that basis if only finitely many coefficients $\langle \mathbf{x}, \phi_j \rangle$ are nonzero. At this point, let us remark that in this chapter, infinite dimensional quantities will be denoted by bold-face letters. In the deterministic case, where usually Tikhonov-type functionals

$$\|\mathbf{A}\mathbf{x} - \mathbf{y}^\delta\|_{\mathcal{Y}}^2 + \hat{\alpha} \Theta_{\mathbf{w},p}(\mathbf{x}) \quad (4.1)$$

are minimized, it is known that penalties of the form

$$\Theta_{\mathbf{w},p}(\mathbf{x}) = \sum_{\lambda \in \Lambda} w_\lambda |\langle \mathbf{x}, \psi_\lambda \rangle|^p, \quad 1 \leq p < 2, \quad (4.2)$$

$\mathbf{w} = \{w_\lambda\}_{\lambda \in \Lambda}$ with $w_\lambda \geq c > 0 \forall \lambda \in \Lambda$, indeed lead to sparse reconstructions, i.e., the amount of nonzero coefficients $\langle \mathbf{x}, \psi_\lambda \rangle$, $\lambda \in \Lambda$, is small. Pushed by the seminal paper [89], where the authors studied the so-called iterative soft-shrinking algorithm to calculate the minimizer of a special Tikhonov-functional under the deterministic error assumption, sparsity has become a widely used regularization strategy both in the deterministic and the stochastic setting. We will base our theory on the paper [89], from where we also employ the numerical algorithm of *iterative soft shrinkage*, see Section 4.4. Although it is known to be rather slow, it serves our purpose. The results of [89] have been refined [90, 91, 92] and alternative minimization procedures for the solution of (4.1) have been introduced, based on for example, hard thresholding [93] projection

[94, 95], Newton-type methods [96] or domain decomposition [97]. The idea of sparsity has been generalized to nonlinear problems [98, 99, 100, 11], Banach spaces [101], or nonconvex penalties [102, 103, 104]. Sparsity promoting regularization may even be beneficial when the sparsity assumption fails [105, 106]. In [107] it has been shown how Tikhonov-Regularization in standard form can be used to find relevant coefficients for sparse regularization. Tomography is a particular application of sparsity that is often used in numerical examples. Other applications include, for example, inpainting [99, 108], inverse interface problems [109] or parameter identification, see [110] for a topical review. Sparse regularization also gained attention in the stochastic setting, c.f. [111, 112, 113]. A particular choice for $\Theta_{\mathbf{w},p}$ are Besov space norms which have already been used as sparsity constraints, see for example [89, 98, 114, 115] in the deterministic setting or [112, 113, 116, 117, 118, 119] in the stochastic setting. In Section 4.2.1 we introduce a wavelet-ONB which is later used to characterize Besov spaces in Section 4.2.2.

4.2 From Bayes to Tikhonov

In contrary to the previous chapter we will start here directly in the stochastic setting. As before, we consider an equation

$$\mathbf{A}\mathbf{x} = \mathbf{y} \quad (4.3)$$

between Hilbert spaces \mathcal{X} and \mathcal{Y} . Now, however, we want to pay more attention to the fact the (typically) infinite dimensional problem has to be approximated by a finite dimensional one when doing calculations on a computer. Therefore, we first introduce a linear orthogonal projection operator $P_m : \mathcal{Y} \rightarrow \mathbb{R}^m$, modeling the mapping of the object \mathbf{y} on an m -dimensional vector y . The projection depends on the actual measurement device. One might for example think of measured function values or coefficients with respect to certain basis functions in \mathcal{Y} . As before $\epsilon \in \mathbb{R}^m$ denotes the typically unavoidable measurement noise. We thus have the *practical measurement model*

$$P_m \mathbf{y} = P_m \mathbf{A}\mathbf{x} + \epsilon.$$

Throughout this chapter we assume for simplicity that each component of the error is normally distributed with zero mean and variance η^2 , $\epsilon_i \sim \mathcal{N}(0, \eta^2)$ for $i = 1 \dots m$, $\eta > 0$. Let $\{\psi_\lambda : \lambda \in \Lambda\}$ be an orthonormal basis in \mathcal{X} , where Λ is an appropriate index set. In order to characterize the unknown \mathbf{x} by its coefficients with respect to $\{\psi_\lambda : \lambda \in \Lambda\}$, we introduce a second operator

$$T : \mathcal{X} \rightarrow \ell_2 \quad \text{via} \quad \mathbf{x} \mapsto \{\langle \mathbf{x}, \psi_\lambda \rangle\}_{\lambda \in \Lambda}. \quad (4.4)$$

T and its adjoint T^* ,

$$T^* : \ell_2 \rightarrow \mathcal{X} \quad \text{via} \quad \mathbf{g} \mapsto \sum_{\lambda \in \Lambda} \mathbf{g}_\lambda \psi_\lambda \quad (4.5)$$

which allow us to switch between function \mathbf{x} and the coefficients $\mathbf{x}_\lambda := \langle \mathbf{x}, \psi_\lambda \rangle$. Here $\langle \cdot, \cdot \rangle$ denotes the L_2 -inner product. Computations on a computer require a finite dimensional representation of \mathbf{x} . Therefore we restrict the index set Λ to a finite set Λ_n , where $n \in \mathbb{N}$ is the number of basis functions used for the discretization. The truncated projectors with respect to Λ_n are defined analogously to (4.4) and (4.5), respectively. We have

$$\begin{aligned} T_n : \mathcal{X} &\rightarrow \ell_2(\mathbb{R}^n), & \mathbf{x} &\mapsto \{\langle \mathbf{x}, \psi_\lambda \rangle\}_{\lambda \in \Lambda_n}, \\ T_n^* : \ell_2(\mathbb{R}^n) &\rightarrow \mathcal{X}, & \mathbf{g} &\mapsto \sum_{\lambda \in \Lambda_n} \mathbf{g}_\lambda \psi_\lambda. \end{aligned} \quad (4.6)$$

Our choice of $\{\psi_\lambda : \lambda \in \Lambda\}$ will be a wavelet basis as introduced in Section 4.2.1. We introduce the basis in more detail in Section 4.2.2. The operator T_n then corresponds to taking the first n coefficients of the wavelet expansion of \mathbf{x} and T_n^* corresponds to using a sequence of n real numbers as coefficients of the series expansion. Although our theory can be expanded to frames instead of a basis in \mathcal{X} , we restrict ourselves to the latter case for simplicity. Thus we arrive at the *computational model*

$$P_m \mathbf{y} = P_m \mathbf{A} T_n^* T_n \mathbf{x} + \epsilon. \quad (4.7)$$

To simplify the notation we denote $y := P_m \mathbf{y}$, $y^\eta := y + \epsilon$, $x := T_n \mathbf{x}$, $A := P_m \mathbf{A} T_n^*$ and obtain the linear model

$$y^\eta = Ax + \epsilon. \quad (4.8)$$

Since we use the Bayesian approach (see Section 1.3.2), the variables x , y^η and ϵ are realizations of the corresponding random variables in the equation

$$Y^\eta = AX + \mathcal{E}. \quad (4.9)$$

The solution of the Inverse Problem is given, as explained in Chapter 1, by Bayes formula

$$\pi_{post}(x|y^\eta) = \frac{\pi_{pr}(x)\pi_\epsilon(y^\eta|x)}{\pi_{y^\eta}(y^\eta)}. \quad (4.10)$$

In order to get a single solution rather than the whole distribution, we use the maximum a-posteriori (MAP) solution (1.19). Because of the normally distributed error we simply have

$$\pi_\epsilon \propto \exp\left(-\frac{1}{2\eta^2} \|Ax - y^\eta\|^2\right).$$

The *prior distribution* π_{pr} for the Besov space prior can formally be written as

$$\pi_{pr}(x) \propto \exp\left(-\frac{\alpha}{2} \|T_n^* x\|_{B_{p,p}^s(\mathbb{R}^d)}^p\right) \quad (4.11)$$

in the discretized setting. We will define it in more detail in Definition 4.2.2. Bayes' formula (4.10) thus yields

$$\pi_{post}(x|y^\eta) \propto \exp\left(-\frac{1}{2\eta^2}\|Ax - y^\eta\|^2\right) \exp\left(-\frac{\alpha}{2}\|T_n^*x\|_{B_{p,p}^s(\mathbb{R}^d)}^p\right).$$

The maximum a posteriori solution is thus given by

$$x^{\text{MAP}} = \operatorname{argmin}_{x \in \mathbb{R}^n} \frac{1}{2\eta^2}\|Ax - y^\eta\|^2 + \frac{\alpha}{2}\|T_n^*x\|_{B_{p,p}^s(\mathbb{R}^d)}^p.$$

Setting $\hat{\alpha} = \alpha\eta^2$ we arrive at

$$x_{\hat{\alpha}}^{\text{MAP}} = \operatorname{argmin}_{x \in \mathbb{R}^n} \|Ax - y^\eta\|^2 + \hat{\alpha}\|T_n^*x\|_{B_{p,p}^s(\mathbb{R}^d)}^p, \quad (4.12)$$

which is exactly a discretized version of the Tikhonov functional (4.1) known from the deterministic setting. Consequently, the same techniques can be used to calculate the minimizer. In order to avoid confusion with well-known results from deterministic theory, it is important to carefully distinguish between the parameter α originating from the prior distribution (4.11) and the actual regularization parameter $\hat{\alpha} = \alpha\eta^2$ in the Tikhonov functional (4.1), as they show different asymptotic behavior: in the following chapters we will see that $\hat{\alpha}$ goes to zero, whereas α has to grow to infinity.

4.2.1 Wavelets

In this work we focus on a particular type of basis for $L_2(\mathbb{R}^d)$, namely, the wavelet basis. The wavelet basis is constructed via a nested sequence of subspaces for $L_2(\mathbb{R})$ called *multiresolution analysis* (MRA), see for example [120]. In order to generate a MRA, a family of subspaces $\{V_j\}_{j \in \mathbb{Z}}$ of $L_2(\mathbb{R})$ has to fulfil the following conditions.

- (1) $\cdots \subset V_{-1} \subset V_0 \subset V_1 \subset \cdots$
- (2) $\bigcap_{j \in \mathbb{Z}} V_j = \{0\}$ and $\overline{\bigcup_{j \in \mathbb{Z}} V_j} = L_2(\mathbb{R})$
- (3) $x \in V_0 \Leftrightarrow x(\cdot - k) \in V_0 \forall k \in \mathbb{Z}$
- (4) $x \in V_j \Leftrightarrow x(2^{-j}\cdot) \in V_0$
- (5) There exists a function $\phi \in V_0$ such that $\{\phi(\cdot - k) : k \in \mathbb{Z}\}$ is an orthonormal basis of V_0

In particular,

$$\phi_{jk}(t) := 2^{j/2}\phi(2^j t - k)$$

is an orthonormal basis for V_j . The function ϕ is called the *scaling function*. Define W_j as the orthonormal complement of V_j in V_{j+1} ,

$$V_{j+1} = V_j \oplus W_j \quad \text{for all } j \in \mathbb{Z}. \quad (4.13)$$

The $\{W_j\}_{j \in \mathbb{Z}}$ inherit the scaling and translation properties (3) and (4) of the $\{V_j\}_{j \in \mathbb{Z}}$. Let $\psi \in W_0$, the *mother wavelet*, be a compactly supported function such that $\{\psi(\cdot - k) : k \in \mathbb{Z}\}$ constitutes an ONB in W_0 . Then the functions

$$\psi_{j,k} := 2^{j/2} \psi(2^j t - k), \quad k \in \mathbb{Z}$$

are an ONB for W_j . Since (4.13) holds for all $j \in \mathbb{Z}$, one can show that the W_j span $L_2(\mathbb{R})$, $\overline{\bigoplus_{j \in \mathbb{Z}} W_j} = L_2(\mathbb{R})$. Due to this and (1), one can represent $L_2(\mathbb{R}) = V_0 \oplus_{j \geq 0} W_j$. The construction of the functions ϕ and ψ is rather complex and will not be discussed here. In particular, smoothness properties play a big role in the design process. We shall just assume that both functions are given and that they are smooth enough, i.e., $\phi \in C^{\tilde{s}}(\mathbb{R})$ and $\psi \in C^{\tilde{s}}(\mathbb{R})$ for some $\tilde{s} > 0$. For more information see, for example, [120, 121, 122].

So far we introduced a MRA only for $L_2(\mathbb{R})$. Following Meyer [121] we expand this to a wavelet basis in $L_2(\mathbb{R}^d)$ by the tensor product method. Let

$$\phi_{j,k}^d(t) := 2^{dj/2} \phi(2^j t_1 - k_1) \dots \phi(2^j t_d - k_d)$$

with the same ϕ as previously. For each $j \in \mathbb{Z}$, the $\{\phi_{j,k}^d : k \in \mathbb{Z}^d\}$ span a subspace of $L_2(\mathbb{R}^d)$ which we denote by V_j^d . One can show that the set of these spaces fulfils the conditions (1)-(5) and hence defines a MRA for $L_2(\mathbb{R}^d)$. Let E denote the set of all $2^d - 1$ sequences $\nu = (\nu_1, \nu_2, \dots, \nu_d)$ with $\nu_j \in \{0, 1\}$ for all $j = 1, \dots, d$ and $\sum_j \nu_j > 0$. The spaces W_j^d defined by

$$V_{j+1}^d = V_j^d \oplus W_j^d, \quad j \in \mathbb{Z}$$

analogously to (4.13) are then spanned by the tensor-wavelets

$$\psi_{j,k}^\nu(t) := 2^{dj/2} \psi^{\nu_1}(2^j t_1 - k_1) \dots \psi^{\nu_d}(2^j t_d - k_d)$$

with the convention that $\psi^0 = \phi$, $\psi^1 = \psi$ and vectors $k = (k_1, k_2, \dots, k_d) \in \mathbb{Z}^d$. Since again $\overline{\bigoplus_{j \in \mathbb{Z}} W_j^d} = L_2(\mathbb{R}^d)$ the $\psi_{j,k}^\nu$ constitute an orthonormal basis for $L_2(\mathbb{R}^d)$. One also has

$$L_2(\mathbb{R}^d) = V_0^d \bigoplus_{j > 0} W_j^d.$$

Thus, defining an orthonormal basis of V_j^d via

$$\phi_{j,k}^d(t) := 2^{dj/2} \phi(2^j t_1 - k_1) \dots \phi(2^j t_d - k_d),$$

a function $\mathbf{x} \in L_2(\mathbb{R}^d)$ has the form

$$\mathbf{x} = \sum_{k \in \mathbb{Z}^d} \langle \mathbf{x}, \phi_{0,k}^d \rangle \phi_{0,k}^d + \sum_{j=0}^{\infty} \sum_{k \in \mathbb{Z}^d} \langle \mathbf{x}, \psi_{j,k}^\nu \rangle \psi_{j,k}^\nu. \quad (4.14)$$

To simplify the notation we follow [89] and denote the set of all functions $\phi_{0,k}^d$ and $\psi_{j,k}^\nu$, $j = 1, 2, \dots$, $k \in \mathbb{Z}$, $\nu \in E$, in (4.14) by Ψ_λ . However, in our framework we have to restrict the index set Λ such that on each scale we have only finitely many wavelets. This is guaranteed by the following assumptions on the wavelet expansion of \mathbf{x} :

- $\exists k_\phi^0 \in \mathbb{Z}^d, k_\phi \in \mathbb{Z} : \langle \mathbf{x}, \phi_{0,k}^d \rangle = 0$ for all $k \in \mathbb{Z}^d : \|k - k_\phi^0\|_{l_1} > k_\phi$. Define $\ell_\phi := \{\#k : \|k - k_\phi^0\|_{l_1} \leq k_\phi\}$.
- On each scale $j \geq 0$, $\exists k_\psi^0 \in \mathbb{Z}^n, k_\psi^j \in \mathbb{Z}$, $k_\psi^{j-} \leq k_\psi^{j+} : \langle \mathbf{x}, \psi_{j,k}^d \rangle = 0$ for all $k \in \mathbb{Z}^d : \|k - k_\psi^0\|_{l_1} > k_\psi^j$. Define $\ell_\psi^j := \{\#k : \|k - k_\psi^0\|_{l_1} \leq k_\psi^j\}$.
- There exists $\ell_\psi \in \mathbb{N}$ such that $\ell_\psi^j \leq 2^{jd} \ell_\psi$ for all $j \geq 0$.

These assumptions are for example satisfied for compactly supported functions or functions which are truly sparse, i.e., the number of nonzero inner products in (4.14) is finite. Thus (4.14) reads

$$\mathbf{x} = \sum_{k_\phi^- \leq k \leq k_\phi^+} \langle \mathbf{x}, \phi_{0,k}^d \rangle \phi_{0,k}^d + \sum_{j=0}^{\infty} \sum_{k_\psi^{j-} \leq k \leq k_\psi^{j+}} \langle \mathbf{x}, \psi_{j,k}^\nu \rangle \psi_{j,k}^\nu. \quad (4.15)$$

We denote the corresponding index set by Λ_f . Hence $\mathbf{x} = \sum_{\lambda \in \Lambda_f} \langle \mathbf{x}, \psi_\lambda \rangle \psi_\lambda$.

4.2.2 Besov Spaces and random variables therein

The $\Psi_\lambda = \{\psi_\lambda : \lambda \in \Lambda\}$ as introduced in the previous section are not only an orthonormal basis in $L_2(\mathbb{R}^d)$ but also a (Riesz) basis for other function spaces including Besov spaces. The Besov spaces $B_{p,q}^s(\mathbb{R}^d)$ are function spaces on \mathbb{R}^d consisting of, roughly spoken, functions which have s derivatives in $L_p(\mathbb{R}^d)$, where q provides some additional fine-tuning. They form Banach spaces whenever $p \geq 1$ and $q < \infty$. Besov spaces can be introduced in various ways. From function space theory, they can be motivated as interpolation spaces between the classical Sobolev spaces H^s [123]. Besov spaces are often used to measure smoothness of functions, which can be seen from the following common definition of Besov spaces, see [123, 124, 125].

Definition 4.2.1. Let $\Delta_h x(t) := x(t-h) - x(t)$ and define the *modulus of continuity* by

$$c_l(x, r)_p = \sup_{\|h\|_2 \leq r} \|\Delta_h^l x\|_{L_p}, \quad r > 0, l \in \mathbb{N}.$$

Let $m > s > 0$. A function x is an element of the Besov space $B_{p,q}^s(\mathbb{R}^d)$ if

$$|x|_{B_{p,q}^s(\mathbb{R}^d)} := \begin{cases} \int_0^\infty \left| \frac{c_m(x,r)_p}{r^s} \right|^q \frac{dr}{r} < \infty & 0 < q < \infty \\ \sup_{r \geq 0} \frac{c_m(x,r)_p}{r^s} & q = \infty \end{cases}. \quad (4.16)$$

Adding the norm $\|x\|_{L_p}$ to the seminorm (4.16) yields a norm in $B_{p,q}^s(\mathbb{R}^d)$ [123, 125]. In this work we restrict ourselves to the case $p = q < \infty$ for the Besov spaces. In this case the spaces $B_{p,p}^s$ coincide with the Sobolev spaces $W_p^s(\mathbb{R}^d)$ whenever s is not an integer [123]. For $p = 2$, $B_{2,2}^s$ even coincides with the classical the Sobolev spaces H^s . Throughout this work we will always assume that

$$\varsigma := s + d \left(\frac{1}{2} - \frac{1}{p} \right) \geq 0 \quad (4.17)$$

to ensure $B_p^s(\mathbb{R}^d)$ is a subset of $L_2(\mathbb{R}^d)$ [89]. With this convention, the wavelet basis, allows for a much less technical definition of the Besov spaces as in Definition 4.2.1. Namely, using the convention $|\lambda| = j$ to denote the scale of the wavelets, the norm

$$\|\mathbf{x}\|_{B_{p,p}^s(\mathbb{R}^d)} = \left(\sum_{\lambda \in \Lambda} 2^{\varsigma p |\lambda|} |\langle \mathbf{x}, \Psi_\lambda \rangle|^p \right)^{\frac{1}{p}} \quad (4.18)$$

is equivalent the traditional Besov space norm $\|x\|_{B_{p,p}^s} = (\|x\|_{L_p} + |x|_{B_{p,p}^s(\mathbb{R}^d)})^{1/p}$. In this formulation of the norm, the smoothness of the functions is controlled by the decay of the wavelet coefficients against the weight $2^{\varsigma p |\lambda|}$. A more detailed discussion of Besov spaces can be found, for example, in [120, 121, 123, 124, 125].

Since the convergence analysis of this chapter is performed in the Bayesian setting (see Section 1.3.2), all occurring quantities have to be understood as random variables. Defining Besov space random variables will directly lead to the prior distribution necessary in the Bayesian approach. With the wavelet basis Ψ_λ from the previous section, we have the following definition, adapted from [112].

Definition 4.2.2. Consider functions on \mathbb{R}^d , $d \in \mathbb{N}$. Let $1 \leq p < \infty$ and Λ_f from (4.15). Take $s \in \mathbb{R}$ such that $\varsigma := s + d(\frac{1}{2} - \frac{1}{p}) > 0$. Let $(X_\lambda^\alpha)_{\lambda \in \Lambda_f}$ be independent identically distributed real-valued random variables with probability density function

$$\pi_{X_\lambda^\alpha}(\tau) = c_p^\alpha \exp\left(-\frac{\alpha|\tau|^p}{2}\right), \quad \tau \in \mathbb{R}, \quad c_p^\alpha = \left(\frac{\alpha}{2}\right)^{\frac{1}{p}} \frac{p}{2\Gamma(\frac{1}{p})} \quad (4.19)$$

for all $\alpha > 0$. Let \mathbf{X} be the random function

$$\mathbf{X}(t) = \sum_{\lambda \in \Lambda_f} 2^{-\varsigma |\lambda|} X_\lambda^\alpha \psi_\lambda(t), \quad t \in \mathbb{R}^d.$$

Then we say \mathbf{X} is distributed according to a $B_{p,p}^s$ -prior with parameter α .

The following Lemma characterizes the random variables $|X_\lambda^\alpha|^p$ on which the stochastic properties of $\|\mathbf{X}\|$ essentially depend.

Lemma 4.2.1. *Let X_λ^α be defined as in Definition 4.2.2. Then the random variables $|X_\lambda^\alpha|^p$, $0 < p < \infty$, are distributed according to the probability density function*

$$\pi_{|X_\lambda^\alpha|^p}(\xi) = \left(\frac{\alpha}{2}\right)^{\frac{1}{p}} \frac{\xi^{\frac{1}{p}-1}}{\Gamma(\frac{1}{p})} \exp\left(-\frac{\alpha\xi}{2}\right), \quad \xi \geq 0 \quad (4.20)$$

and satisfy

$$\mathbb{E}(|X_\lambda^\alpha|^p) = \frac{2}{\alpha p}. \quad (4.21)$$

Proof. Let X_λ^α be defined as in (4.19). We are interested in the probability density of $Y := |X_\lambda^\alpha|^p$. Denote $\mathbb{F}_{X_\lambda^\alpha}(\tau)$ and $\mathbb{F}_Y(\xi)$ the cumulative distribution functions of X_λ^α and Y , respectively. Since $Y \geq 0$, also $\xi \geq 0$. For all $\xi > 0$,

$$\begin{aligned} \mathbb{F}_Y(\xi) &= \mathbb{P}(Y \leq \xi) = \mathbb{P}(|X_\lambda^\alpha|^p \leq \xi) = \mathbb{P}(-\sqrt[p]{\xi} \leq X_\lambda^\alpha \leq \sqrt[p]{\xi}) \\ &= \mathbb{F}_{X_\lambda^\alpha}(\sqrt[p]{\xi}) - \mathbb{F}_{X_\lambda^\alpha}(-\sqrt[p]{\xi}) \end{aligned}$$

and since $\sqrt[p]{\cdot}$ is continuously differentiable on $(0, \infty)$ for $0 < p < \infty$,

$$\begin{aligned} \pi_Y(\xi) &= \frac{d}{d\xi} \mathbb{F}_Y(\xi) = \frac{d}{d\xi} (\mathbb{F}_{X_\lambda^\alpha}(\sqrt[p]{\xi}) - \mathbb{F}_{X_\lambda^\alpha}(-\sqrt[p]{\xi})) \\ &= \pi_{X_\lambda^\alpha}(\sqrt[p]{\xi}) \cdot \left(\frac{1}{p} \xi^{\frac{1}{p}-1}\right) - \pi_{X_\lambda^\alpha}(-\sqrt[p]{\xi}) \cdot \left(-\frac{1}{p} \xi^{\frac{1}{p}-1}\right) \\ &= \left(\frac{\alpha}{2}\right)^{\frac{1}{p}} \frac{\xi^{\frac{1}{p}-1}}{\Gamma(\frac{1}{p})} \exp\left(-\frac{\alpha\xi}{2}\right), \quad \xi > 0. \end{aligned}$$

For $\xi = 0$, $\mathbb{F}_Y(\xi) = \mathbb{P}(Y \leq \xi) = \mathbb{P}(|X_\lambda^\alpha|^p \leq \xi) = 0$. Now (4.21) is given by

$$\mathbb{E}(|X_\lambda^\alpha|^p) = \int_0^\infty \xi \pi_Y(\xi) d\xi.$$

□

Since Λ_f contains infinitely many basis functions, a realization of such a Besov space random variable is an element of the space of definition with probability zero. To guarantee finiteness of the norm, the functions have to be defined in a Besov space which is smoother than the one where the realizations are measured. The following Lemma was adopted from [112, Lemma 2], but there the authors considered functions on a d -dimensional torus instead of \mathbb{R}^d , i.e., they used a different definition of the wavelet basis. Also, their proof does not include the parameter α .

Lemma 4.2.2. *Let \mathbf{X} be defined in $B_{p,p}^r(\mathbb{R}^d)$ as in Definition 4.2.2 for some $r > 0$ and $2 < \alpha < \infty$. Then the following three conditions are equivalent:*

- (i) $\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)} < \infty$ almost surely,
- (ii) $\mathbb{E} \exp\left(\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)}^p\right) < \infty$,
- (iii) $s < r - \frac{d}{p}$.

Proof. Let $(X_\lambda^\alpha)_{\lambda \in \Lambda_f}$ be as in Definition 4.2.2. First consider the expectation of $\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)}^p$. Because of (4.21) we have

$$\begin{aligned}
\mathbb{E}\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)}^p &= \mathbb{E} \sum_{\lambda \in \Lambda_f} 2^{(s+d(\frac{1}{2}-\frac{1}{p}))p|\lambda|} \left| 2^{-(r+d(\frac{1}{2}-\frac{1}{p}))|\lambda|} X_\lambda^\alpha \right|^p \\
&= \mathbb{E} \sum_{\lambda \in \Lambda_f} 2^{-(r-s)p|\lambda|} |X_\lambda^\alpha|^p = \sum_{\lambda \in \Lambda_f} 2^{-(r-s)p|\lambda|} \mathbb{E}|X_\lambda^\alpha|^p \\
&= \frac{2}{\alpha p} \sum_{\lambda \in \Lambda_f} 2^{-(r-s)p|\lambda|}. \tag{4.22}
\end{aligned}$$

Because of the construction of Λ_f there are ℓ_ϕ scaling functions and ℓ_ψ wavelets on the coarsest scale. Additionally, on scale $j > 0$ we have at most $2^{jd}\ell_\psi$ wavelets. Hence, the summation in (4.22), which is actually a double sum over all wavelets and scales, reduces to a simple sum and

$$\begin{aligned}
\mathbb{E}\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)}^p &= \frac{2}{\alpha p} \left(\ell_\phi + \sum_{j=0}^{\infty} 2^{-(r-s)pj} \cdot 2^{jd}\ell_\psi \right) \\
&= \frac{2}{\alpha p} \left(\ell_\phi + \ell_\psi \sum_{j=0}^{\infty} 2^{-j((r-s)p-d)} \right) \tag{4.23}
\end{aligned}$$

The sum converges if and only if $(r-s)p-d > 0$. Since finiteness of the expectation of a positive random variable implies almost sure finiteness of the random variable itself, $\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)}^p < \infty$ a.s. and also $\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)} < \infty$ a.s., hence (i) \Leftrightarrow (iii). Now we turn to condition (ii). It is

$$\begin{aligned}
\mathbb{E} \exp\left(\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)}^p\right) &= \mathbb{E} \exp\left(\sum_{\lambda \in \Lambda_f} 2^{-(r-s)p|\lambda|} |X_\lambda^\alpha|^p\right) \\
&= \prod_{\lambda \in \Lambda_f} \mathbb{E} \exp\left(2^{-(r-s)p|\lambda|} |X_\lambda^\alpha|^p\right) \\
&= \prod_{\lambda \in \Lambda_f} \left(1 - \frac{2^{-(r-s)p|\lambda|+1}}{\alpha}\right)^{-1/p} \\
&= \left(1 - \frac{2}{\alpha}\right)^{-\frac{\ell_\phi}{p}} \cdot \left(\prod_{j=0}^{\infty} \left(1 - \frac{2^{-(r-s)p|\lambda|+1}}{\alpha}\right)^{2^{jd}\ell_\psi}\right)^{-\frac{1}{p}} \tag{4.24}
\end{aligned}$$

where we used that the X_λ^α are independent and $\mathbb{E} \exp(c|X_\lambda^\alpha|^p) = (1 - \frac{2c}{\alpha})^{-1/p}$ if $\alpha > 2c$ (which is why we have to require $\alpha > 2$). Since $\prod_{l=0}^{\infty} (1 + a_l)$ converges if and only if $\sum_{l=0}^{\infty} \log(1 + a_l)$ converges we find that $\mathbb{E} \exp\left(\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)}^p\right) < \infty$ if

$$\sum_{j=0}^{\infty} 2^{jd} \ell_\psi \log\left(1 - \frac{2^{-(r-s)p|\lambda|+1}}{\alpha}\right) < \infty. \quad (4.25)$$

The root test yields

$$\lim_{j \rightarrow \infty} \left(2^{jd} \ell_\psi \log\left(1 - \frac{2^{-(r-s)p|\lambda|+1}}{\alpha}\right)\right)^{\frac{1}{j}} = 2^{-(r-s)p+d}. \quad (4.26)$$

Hence the sum and by that (4.24) converges if (iii) holds. Since (ii) obviously implies (i) the proof is complete. \square

The Lemma shows that, although we define the random variable in the Besov space $B_{p,p}^r(\mathbb{R}^d)$, its realizations will be elements of $B_{p,p}^r(\mathbb{R}^d)$ with probability zero. Instead we have to reduce the smoothness of the space in which the realizations are measured to $B_{p,p}^s(\mathbb{R}^d)$, $s < r - d/p$, so that the tail of the wavelet series is under control. If such a combination is used for spaces of definition and measurement of the random variables, finiteness of the norms in the latter space is ensured if condition (iii) is fulfilled. We will refer to this as the *infinite model* (MI). A second possibility is to consider a finite dimensional model: Let T_n, T_n^* be defined as in (4.6). Then for a function $\mathbf{x} \in L_2(\mathbb{R}^d)$ and arbitrary, but fixed $n \in \mathbb{N}$, $T_n^* T_n \mathbf{x} = \sum_{\lambda \in \Lambda_n} x_\lambda \psi_\lambda$ is an element of $B_{p,p}^s(\mathbb{R}^d)$ with probability one if the wavelet is smooth enough. This allows in particular to measure the realizations of random variables in the same norm as was used in the definition of the random function. This will be referred to as the *finite model* (MII). In order to derive convergence rates we need to calculate $\mathbb{P}(\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)} \geq \varrho)$ for given $\varrho > 0$. The following Corollary shows how this can be done for model (MI) using Lemma 4.2.2.

Corollary 4.2.3. *Consider model (MI). Let \mathbf{X} be defined in $B_{p,p}^r(\mathbb{R}^d)$ according to Definition 4.2.2 with $2 < \alpha < \infty$. Let $s < r - \frac{d}{p}$ and $\varrho > 0$. Then*

$$\mathbb{P}(\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)} > \varrho) \leq \frac{2}{\alpha p \varrho^p} \left(\ell_\phi + \ell_\psi \sum_{j=0}^{\infty} 2^{-j((r-s)p-d)} \right). \quad (4.27)$$

Proof. According to Markov's inequality, for any nonnegative random variable ξ with $\mathbb{E}\xi < \infty$, $\mathbb{P}(\xi > \varrho) \leq \frac{1}{\varrho} \mathbb{E}\xi$. Since the mapping $z \mapsto z^p$ is bijective for $z \geq 0$ and $1 \leq p \leq 2$, we have for given $\varrho > 0$

$$\mathbb{P}(\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)} > \varrho) = \mathbb{P}(\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)}^p > \varrho^p) \leq \frac{1}{\varrho^p} \mathbb{E} \|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)}^p.$$

The expectation of $\|\mathbf{X}\|_{B_{p,p}^s(\mathbb{R}^d)}^p$ is given by (4.23). \square

Using the finite model, we get the following result.

Lemma 4.2.4. *Consider model (MII). Let \mathbf{X} be defined as $B_{p,p}^s(\mathbb{R}^d)$ random function according to Definition 4.2.2, T_n as in (4.6) and take $\varrho > 0$. Denote $X_n := T_n^* T_n \mathbf{X}$. Then*

$$\mathbb{P}(\|X_n\|_{B_{p,p}^s(\mathbb{R}^d)} > \varrho) = \frac{\Gamma(\frac{n}{p}, \frac{\alpha \varrho^p}{2})}{\Gamma(\frac{n}{p})} \quad (4.28)$$

with the Gamma functions

$$\Gamma(a) = \int_0^\infty t^{a-1} e^{-t} dt, \quad \Gamma(a, z) = \int_z^\infty t^{a-1} e^{-t} dt.$$

Proof. Let \mathbf{X} be as in Definition 4.2.2. Then $X_n = \sum_{\lambda \in \Lambda_n} 2^{-s|\lambda|} X_\lambda^\alpha \psi_\lambda$ and its norm $\|X_n\|_{B_{p,p}^s(\mathbb{R}^d)}^p = \sum_{\lambda \in \Lambda_n} |X_\lambda^\alpha|^p$ reduces to a sum of n i.i.d. random variables with density (4.20). The resulting density can be calculated using the moment generating function (c.f., e.g. [27]) of the X_λ^α which is just the Laplace transform $\mathcal{L}(\cdot)$ of the probability density function. The moment generating function of a sum of random variables is given by the product of the single moment generating functions [27]. With $\pi_{|X_\lambda^\alpha|^p}$ from (4.20) we get

$$\mathcal{L}[\pi_{|X_\lambda^\alpha|^p}](s) = \left(1 + \frac{2s}{\alpha}\right)^{-1/p}$$

and obtain the probability density function of $\pi_{\sum_{\lambda \in \Lambda_n} |X_\lambda^\alpha|^p}(\xi)$, $\xi \geq 0$, via the inverse Laplace transform \mathcal{L}^{-1} ,

$$\pi_{\sum_{\lambda \in \Lambda_n} |X_\lambda^\alpha|^p}(\xi) = \mathcal{L}^{-1} \left[\left(1 + \frac{2s}{\alpha}\right)^{-n/p} \right] (\xi) = \frac{\xi^{\frac{n}{p}-1}}{\Gamma(\frac{n}{p})} \left(\frac{\alpha}{2}\right)^{\frac{n}{p}} e^{-\frac{\alpha}{2}\xi}. \quad (4.29)$$

Because $\sum_{\lambda \in \Lambda_n} |X_\lambda^\alpha|^p$ is non-negative, $\mathbb{P}(\|X_n\|_{B_{p,p}^s(\mathbb{R}^d)} > \varrho) = \mathbb{P}(\|X_n\|_{B_{p,p}^s(\mathbb{R}^d)}^p > \varrho^p)$. The claim follows by integrating (4.29) over ξ from ϱ^p to infinity. \square

Remark. Similar to the infinite dimensional setting, Chebyshev's inequality allows to estimate

$$\mathbb{P}(\|X_n\|_{B_{p,p}^s(\mathbb{R}^d)} > \varrho) \leq \frac{2n}{\alpha \varrho^p}. \quad (4.30)$$

This is indeed an upper bound for (4.28).

Remark. The relation between the two models is best seen by comparing (4.27) and (4.30). Both probabilities solely differ in a term describing the wavelet structure. In model (MI) the term $\ell_\phi + \ell_\psi \sum_{j=0}^\infty 2^{-j((r-s)p-d)}$ ensures that $\mathbb{E}(\|\mathbf{X}\|)$ is bounded independently of n , whereas in the finite model the expectation $\mathbb{E}(\|X_n\|)$ grows unbounded.

If ϱ is an a-priori estimate of $\|\mathbf{X}\|$ or $\|X_n\|$, respectively, it will have to be chosen differently for the two models, taking into account the different asymptotic behaviour of the respective random variables.

From the Bayesian point of view, Besov space priors are particularly interesting. Namely, it has been shown in [111, 112] that Besov priors are *discretization invariant*, i.e., fulfills the following three conditions.

- the solution to the Inverse Problem diverges as its discretization is refined
- the solution to the Inverse Problem diverges as the number of measured data point increase
- the representation of a-priori knowledge is incompatible with discretization

As a counterexample, it has been shown in [126] that discrete (non-Gaussian) total variation priors (see, e.g. [127]) converge to a smooth Gaussian prior the more the level of discretization is refined.

4.3 Convergence of maximum a posteriori solutions with Besov priors

4.3.1 Review of a deterministic result as basis for the lifting approach

Before we analyse convergence properties of Tikhonov regularization in the stochastic setting, i.e., of the maximum a posteriori solution (4.12), we want to review facts for the deterministic case. The following theorem has first been proven in the seminal paper [89]. A relaxed version has been shown in [91].

Theorem 4.3.1. *Assume that \mathbf{A} is a bounded operator from \mathcal{X} to \mathcal{Y} , that $1 \leq p \leq 2$, and that $c < \min_{\lambda} w_{\lambda}$, $\{w_{\lambda}\}_{\lambda \in \Lambda} = \mathbf{w}$ for some constant $c > 0$. Assume that (4.1) has a unique minimizer. Let $\mathbf{x}_{\hat{\alpha}}^*$ be this minimizer for given data \mathbf{y}^{δ} with $\|\mathbf{y} - \mathbf{y}^{\delta}\| \leq \delta$ and $\hat{\alpha} > 0$. If $\hat{\alpha} = \hat{\alpha}(\delta)$ satisfies the requirements*

$$\lim_{\delta \rightarrow 0} \hat{\alpha}(\delta) \rightarrow 0 \quad \text{and} \quad \lim_{\delta \rightarrow 0} \frac{\delta^2}{\hat{\alpha}(\delta)} = 0,$$

then we have, for any $\mathbf{x}_0 \in \mathcal{X}$,

$$\lim_{\delta \rightarrow 0} \left[\sup_{\|\mathbf{y} - \mathbf{y}^{\delta}\| \leq \delta} \|\mathbf{x}_{\hat{\alpha}}^* - \mathbf{x}^{\dagger}\| \right] = 0,$$

i.e., the regularized solutions converge to \mathbf{x}^{\dagger} , where \mathbf{x}^{\dagger} is the unique solution of the equation $\mathbf{A}\mathbf{x} = \mathbf{y}$ with minimal value of $\Phi(\cdot) = \|\cdot\|_{B_{p,p}^s(\mathbb{R}^d)}$.

Remark. Careful inspection of the proof in [89] reveals, that injectivity of the operator \mathbf{A} was only needed to guarantee the existence of a unique minimizer of (4.1) if $p = 1$. It has been shown in [91], that in this case injectivity on every finite dimensional subspace is already sufficient. For all $p > 1$, the functional (4.1) is convex and hence admits a unique minimizer.

Remark. For Besov spaces as defined in Section 4.2.2 the weights are given by $w_\lambda = 2^{s p |\lambda|} > 0$, $\varsigma = s + d(\frac{1}{2} - \frac{1}{p}) \geq 0$, compare (4.2) and (4.18).

It is well-known from deterministic theory (cf. [8]) that in general the convergence can be arbitrarily slow. In order to guarantee a certain decrease of the reconstruction error with respect to the noise parameter, i.e., to derive convergence rates, it is necessary to impose additional conditions on either the true solution, the operator, or both. We will require a smoothing property of the operator \mathbf{A} and an a priori bound of the norm of the solution. Since we are only interested in convergence with respect to the noise, we will consider the discretization levels m and n fixed. There are various results on convergence rates for sparsity promoting inversion, see for example [91, 89, 128, 129]. For our convergence rate analysis, we will follow the concept presented in [89]. Let us first summarize some results from [89, Section 4.2]. Assume $N(\mathbf{A}) = \{0\}$ for $p = 1$ and suppose that we know a priori a bound on the sparsity penalty of the exact solution, i.e., $\|\mathbf{x}^*\|_{B_{p,p}^s(\mathbb{R}^d)} \leq \varrho$ for some $\varrho > 0$. If we also know that \mathbf{y} lies within a distance δ of $\mathbf{A}\mathbf{x}^*$ in \mathcal{Y} , then the exact solution can be localized within the set

$$\mathcal{F}(\delta, \varrho) := \{\mathbf{x} \in \mathcal{X} : \|\mathbf{A}\mathbf{x} - \mathbf{y}\| \leq \delta, \|\mathbf{x}\|_{B_{p,p}^s(\mathbb{R}^d)} \leq \varrho\}.$$

The diameter of this set is a measure of the uncertainty of the solution for a given a priori constant ϱ and noise level δ . The maximum diameter of \mathcal{F} is bounded by $2M(\delta, \varrho)$ where $M(\delta, \varrho)$, defined by

$$M(\delta, \varrho) := \sup\{\|\mathbf{h}\| : \|\mathbf{A}\mathbf{h}\| \leq \delta, \|\mathbf{h}\|_{B_{p,p}^s(\mathbb{R}^d)} \leq \varrho\}, \quad (4.31)$$

is called the *modulus of continuity* of \mathbf{A}^{-1} under the a priori constraint. It can also be interpreted as the worst case error. An upper bound on the reconstruction error is given by the *modulus of convergence*

$$M_{\hat{\alpha}}(\delta, \varrho) := \sup\{\|\mathbf{x}_{\hat{\alpha}}^* - \mathbf{x}\| : \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y}, \|\mathbf{A}\mathbf{x} - \mathbf{y}\| \leq \delta, \|\mathbf{x}\|_{B_{p,p}^s(\mathbb{R}^d)} \leq \varrho\} \quad (4.32)$$

where $\mathbf{x}_{\hat{\alpha}}^*$ denotes the minimizer of the Tikhonov functional (4.1). The decay of this modulus of convergence as $\delta \rightarrow 0$ is governed by the decay of the modulus of continuity, as shown in the following proposition:

Proposition 4.3.2 ([89], Prop. 4.5). *The modulus of convergence (4.32) satisfies*

$$M(\delta, \varrho) \leq M_{\hat{\alpha}}(\delta, \varrho) \leq M(\delta + \delta', \varrho + \varrho'), \quad (4.33)$$

where

$$\delta' = (\delta^2 + \hat{\alpha}\varrho^p)^{\frac{1}{2}}, \quad \varrho' = (\varrho^p + \delta^2\hat{\alpha}^{-1})^{\frac{1}{p}} \quad (4.34)$$

and $M(\delta, \varrho)$ is defined in (4.31).

Thus it suffices to investigate the convergence behaviour of the modulus of continuity. As in [89], let us additionally assume that the operator \mathbf{A} is of smoothing order β , that is, we assume that for some $\beta > 0$ there exist constants A_l and A_u such that for all $\mathbf{h} \in L^2(\mathbb{R}^d)$

$$A_l^2 \sum_{\lambda \in \Lambda} 2^{-2|\lambda|\beta} |\langle \mathbf{h}, \psi_\lambda \rangle|^2 \leq \|\mathbf{A}\mathbf{h}\|^2 \leq A_u^2 \sum_{\lambda \in \Lambda} 2^{-2|\lambda|\beta} |\langle \mathbf{h}, \psi_\lambda \rangle|^2. \quad (4.35)$$

The decay of the modulus of continuity is then characterized as follows.

Proposition 4.3.3 ([89], Proposition 4.7). *If the operator \mathbf{A} satisfies the smoothing property (4.35), then the modulus of continuity $M(\delta, \varrho)$ satisfies*

$$c \left(\frac{\delta}{A_u} \right)^{\frac{\varsigma}{\varsigma+\beta}} \varrho^{\frac{\beta}{\beta+\varsigma}} \leq M(\delta, \varrho) \leq C \left(\frac{\delta}{A_l} \right)^{\frac{\varsigma}{\varsigma+\beta}} \varrho^{\frac{\beta}{\beta+\varsigma}},$$

where $\varsigma = s + d(\frac{1}{2} - \frac{1}{p}) \geq 0$ and c and C are constants depending on ς and β only.

4.3.2 Convergence in the stochastic setting

For the convergence analysis we mainly use a lifting argument from deterministic theory. In Theorem 4.1 of his PhD thesis [30], Hofinger proved how by means of the Ky Fan metric deterministic results can be lifted to the space of random variables. The same techniques can be used in our situation as well. Since the original source is not easily accessible, we include the proof for the convenience of the reader. Before the Theorem, we need the following Lemmata.

Lemma 4.3.4. ([130], see also [32]) *Let $(\Omega, \mathcal{F}, \mu)$ be a finite measure space. Let x_k and x be measurable functions from Ω into a metric space χ with metric d_χ . Suppose $x_k(\omega) \xrightarrow{d_\chi} x(\omega)$ for μ -almost all $\omega \in \Omega$. Then for any $\varepsilon > 0$ there is a set Ω_ε with $\mu(\Omega \setminus \Omega_\varepsilon) < \varepsilon$ such that $x_k \xrightarrow{d_\chi} x(\omega)$ uniformly on Ω_ε , that is*

$$\lim_{k \rightarrow \infty} \sup \{ d_\chi(x_k(\omega), x(\omega)) : \omega \in \Omega_\varepsilon \} = 0.$$

Lemma 4.3.5 ([30], Proposition 1.10). *Let $\{x_k\}_{k \in \mathbb{N}}$ be a sequence of random variables that converges to x in the Ky Fan metric. Then for any $\nu > 0$ and $\varepsilon > 0$ there exist $\Omega_\varepsilon \subset \Omega$, $\mathbb{P}(\Omega_\varepsilon) \geq 1 - \varepsilon$, and a subsequence x_{k_j} with*

$$\|x_{k_j}(\omega) - x(\omega)\| \leq (1 + \nu)\rho_K(x_{k_j}, x) \quad \forall \omega \in \Omega_\varepsilon.$$

Furthermore there exists a subsequence that converges to x almost surely.

Proof. We give a sketch of the proof for the first statement taken from [30].

Set $\sigma_k := (1 + \nu)\rho_K(x_k, x)$. By definition of the Ky Fan metric (1.24), for given σ_k , there exists a set Ω_{σ_k} with $\mathbb{P}(\Omega_{\sigma_k}) \geq 1 - \sigma_k$ and $\omega \in \Omega_{\sigma_k} \Rightarrow \|x(\omega) - x_k(\omega)\| \leq \sigma_k$. For arbitrary $\varepsilon > 0$ and $\sigma_k \rightarrow 0$ we pick a subsequence (σ_{k^j}) with $\sum_{j=1}^{\infty} \sigma_{k^j} \leq \varepsilon$ and introduce the set $\Omega_\varepsilon := \bigcap_{j=1}^{\infty} \Omega_{\sigma_{k^j}}$. One can check that $\mathbb{P}(\Omega_\varepsilon) \geq 1 - \varepsilon$. Since Ω_ε is a subset of every $\Omega_{\sigma_{k^j}}$ we have

$$\forall \omega \in \Omega_\varepsilon \subseteq \Omega_{\sigma_{k^j}} : \|x(\omega) - x_{k^j}(\omega)\| \leq \sigma_{k^j},$$

which proves the first statement. The second one follows since convergence in Ky Fan metric is equivalent to convergence in probability, which itself implies almost-sure convergence of a subsequence, cf [32]. \square

With this, we are ready for the convergence theorem.

Theorem 4.3.6. *Let $y = y(\omega)$ be the exact right hand side in (4.8) and $\{y^{\hat{\eta}_k}(\omega)\}_{k \in \mathbb{N}}$ be a sequence of noisy realizations of $y(\omega) + \epsilon(\omega)$ such that $\rho_K(y, y^{\hat{\eta}_k}) \leq \hat{\eta}_k$, $\hat{\eta}_k \rightarrow 0$ as $k \rightarrow \infty$. Let $\hat{\alpha}(\hat{\eta}_k)$ be a parameter choice rule such that $\hat{\alpha}(\hat{\eta}_k) \rightarrow 0$ and $\hat{\eta}^2/\hat{\alpha}(\hat{\eta}_k) \rightarrow 0$ as $\hat{\eta}_k \rightarrow 0$. Furthermore let the minimum norm solution x^\dagger , defined as in Theorem 4.3.1, be unique. Denote with $x_{\hat{\alpha}(\hat{\eta})}^*$ the minimizer of (4.1). Then*

$$\lim_{\hat{\eta} \rightarrow 0} \rho_K(x^\dagger, x_{\hat{\alpha}(\hat{\eta})}^*) = 0.$$

Proof of Theorem 4.3.6. Define $\theta := \limsup_{k \rightarrow \infty} \rho_K(x^\dagger, x_{\hat{\alpha}(\hat{\eta}_k)}^*)$. (Note that $0 \leq \theta \leq 1$ due to the properties of the Ky Fan metric). We show in the following that for arbitrary $\varepsilon > 0$ we have $\theta/2 \leq \varepsilon$ and hence $\limsup_{k \rightarrow \infty} \rho_K(x^\dagger, x_{\hat{\alpha}(\hat{\eta}_k)}^*) = \lim_{k \rightarrow \infty} \rho_K(x^\dagger, x_{\hat{\alpha}(\hat{\eta}_k)}^*) = 0$.

As a first step we pick a “worst case” subsequence $\{y^{\hat{\eta}_{k^j}}\}$ of $\{y^{\hat{\eta}_k}\}$, a subsequence for which the corresponding solutions satisfy $\rho_K(x^\dagger, x_{\hat{\alpha}(\hat{\eta}_{k^j})}^*) \geq \theta/2$. We now show that even from this “worst case” sequence we can pick a subsequence $\{y^{\hat{\eta}_{k_i^j}}\}$ for which we have $\limsup \rho_K(x^\dagger, x_{\hat{\alpha}(\hat{\eta}_{k_i^j})}^*) \leq \varepsilon$ for arbitrary $\varepsilon > 0$.

Let $\varepsilon > 0$. According to Lemma 4.3.5 we can pick a subsequence $\{y^{\hat{\eta}_{k_i^j}}\}$ and a set Ω_ε with $\mathbb{P}(\Omega_\varepsilon) \geq 1 - \frac{\varepsilon}{2}$ as well as $\|y(\omega) - y^{\hat{\eta}_{k_i^j}}(\omega)\| \leq 2\hat{\eta}_{k_i^j}$ on Ω_ε . For all $\omega \in \Omega_\varepsilon$, the noise tends to zero, we can therefore use the deterministic result and deduce via Theorem 4.3.1 that $x_{\hat{\alpha}(\hat{\eta}_{k_i^j})}^*(\omega)$ converges to the unique solution $x^\dagger(\omega)$ for $\hat{\eta}_{k_i^j} \rightarrow 0$, $\omega \in \Omega_\varepsilon$ if $\hat{\alpha}(\hat{\eta}) \rightarrow 0$ and $\hat{\eta}^2/\hat{\alpha}(\hat{\eta}) \rightarrow 0$ as $\hat{\eta} \rightarrow 0$. This convergence is not uniform in ω ; nevertheless, pointwise convergence implies uniform convergence except on sets of small measure according to Lemma 4.3.4. Therefore there exist $\Omega'_\varepsilon \subset \Omega_\varepsilon$, $\mathbb{P}(\Omega'_\varepsilon) < \frac{\varepsilon}{2}$ and

$j_0 \in \mathbb{N}$ such that $\|x_{\hat{\alpha}(\hat{\eta}_{k_l^j})}^*(\omega) - x^\dagger(\omega)\| < \varepsilon \forall \omega \in \Omega_\varepsilon \setminus \Omega'_\varepsilon$ and $j \geq j_0$. We thus have

$$\mathbb{P}\left(\left\{\omega \in \Omega_\varepsilon : \|x_{\hat{\alpha}(\hat{\eta}_{k_l^j})}^*(\omega) - x^\dagger(\omega)\| > \varepsilon\right\}\right) \leq \mathbb{P}(\Omega'_\varepsilon) \leq \varepsilon/2.$$

Since we split $\Omega = \Omega \setminus \Omega_\varepsilon \cup \Omega_\varepsilon \setminus \Omega'_\varepsilon \cup \Omega'_\varepsilon$ with $\mathbb{P}(\Omega \setminus \Omega_\varepsilon) < \frac{\varepsilon}{2}$, $\mathbb{P}(\Omega \setminus \Omega_\varepsilon) + \mathbb{P}(\Omega'_\varepsilon) \leq \varepsilon$ we have shown existence of a subsequence $\hat{\eta}_{k_l^j}$ such that

$$\mathbb{P}\left(\left\{\omega \in \Omega : \|x_{\hat{\alpha}(\hat{\eta}_{k_l^j})}^*(\omega) - x^\dagger(\omega)\| > \varepsilon\right\}\right) \leq \varepsilon$$

for $\hat{\eta}_{k_l^j}$ sufficiently small. This ε is by definition of the Ky Fan metric an upper bound for the distance between $x_{\hat{\alpha}(\hat{\eta}_{k_l^j})}^*$ and x^\dagger . Therefore we have $\limsup_{l \rightarrow \infty} \rho_K(x_{\hat{\alpha}(\hat{\eta}_{k_l^j})}^*, x^\dagger) \leq \varepsilon$.

On the other hand, the original sequence satisfied $\liminf_{j \rightarrow \infty} \rho_K(x^\dagger, x_{\hat{\alpha}(\hat{\eta}_{k_l^j})}^*) \geq \theta/2$. Since $\liminf_{j \rightarrow \infty} \rho_K(x^\dagger, x_{\hat{\alpha}(\hat{\eta}_{k_l^j})}^*) \leq \limsup_{l \rightarrow \infty} \rho_K(x_{\hat{\alpha}(\hat{\eta}_{k_l^j})}^*, x^\dagger)$ it follows $\theta/2 \leq \varepsilon$. Because $\varepsilon > 0$ was arbitrary, this implies $\theta = 0$, which concludes the proof. \square

In our case of Gaussian noise, we have the following corollary.

Corollary 4.3.7. *Let $\alpha, \eta > 0$, $1 \leq p \leq 2$ and x^\dagger be unique. Let $x_{\hat{\alpha}}^{\text{MAP}} = x_{\alpha, \eta}^{\text{MAP}}$ be the solution of (4.12). If $\alpha = \alpha(\eta)$ is chosen such that $\alpha \eta^2 \rightarrow 0$ and $\frac{|\ln \eta|}{\alpha} \rightarrow 0$ as $\eta \rightarrow 0$, then*

$$\lim_{\eta \rightarrow 0} \rho_K(x_{\hat{\alpha}}^{\text{MAP}}, x^\dagger) = 0.$$

Proof. From the definition of the Ky-Fan metric (1.24), with a $(\mathcal{X}, d_{\mathcal{X}}) = (\mathbb{R}^m, \|\cdot\|_2)$, one has

$$\begin{aligned} \rho_K(y^\eta, y) &= \inf_{\varepsilon > 0} \{\mathbb{P}(\{\omega \in \Omega : \|y(\omega) + \varepsilon(\omega) - y(\omega)\| > \varepsilon\}) < \varepsilon\} \\ &= \inf_{\varepsilon > 0} \{\mathbb{P}(\{\omega \in \Omega : \|\varepsilon(\omega) - 0\| > \varepsilon\}) < \varepsilon\} = \rho_K(\varepsilon, 0). \end{aligned} \quad (4.36)$$

Hence, by (1.30) we have

$$\rho_K(y, y^\eta) \leq \sqrt{2}\eta \sqrt{m - \ln^- \left(\eta^2 2\pi m^2 \left(\frac{\varepsilon}{2}\right)^m \right)} =: \hat{\eta}.$$

Recall the definition of $\hat{\alpha} = \alpha \eta^2$ and that the maximum a posteriori solution (4.12) coincides with the minimizer of the Tikhonov functional (4.1). Theorem 4.3.6 ensures convergence of $x_{\hat{\alpha}}^{\text{MAP}}$ to x^\dagger with respect to the Ky Fan metric if $\hat{\alpha} = \alpha \eta^2 \rightarrow 0$ and $\hat{\eta}^2 / \hat{\alpha}(\hat{\eta}) \rightarrow 0$. From the definitions of $\hat{\eta}$ and $\hat{\alpha}$ the condition

$$\frac{\left(\sqrt{2}\eta \sqrt{m - \ln^- \left(\eta^2 2\pi m^2 \left(\frac{\varepsilon}{2}\right)^m \right)}\right)^2}{\alpha \eta^2} \xrightarrow{\eta \rightarrow 0} 0$$

follows. This is fulfilled if α grows faster than $|\ln \eta|$ as $\eta \rightarrow 0$. \square

Reviewing the proof of Theorem 4.3.1, one notices that instead of $\lim_{\delta \rightarrow 0} \frac{\delta^2}{\hat{\alpha}(\delta)} = 0$, it is sufficient to require that $\frac{\delta^2}{\hat{\alpha}(\delta)} < C$ for an arbitrary positive constant C . Hence, as long as the logarithm is inactive, i.e., $m > m_{\min}$ from (1.31), it suffices to keep α constant with decreasing η . Although this is not important for theoretical purposes as the logarithm kicks in eventually, this explains why in numerical experiments α starts to grow only after η is small enough. In particular, this means that we are well situated in the Bayesian framework, where the prior is a fixed distribution, as long as $m > m_{\min}$. Now let us discuss the asymptotic setting. From a stochastic point of view, α can be interpreted as a measure for the variance of the prior. If $\alpha \rightarrow \infty$, i.e., this variance goes to zero, the coefficients $|X_\lambda^\alpha|^p$ are close to zero with high probability, which emphasizes the sparsity background. From the Bayesian perspective this is counter intuitive as has already been observed in [21]. If the prior variance is zero, that would mean to take the mean of the prior distribution as exact solution. However, we have to keep in mind that the product of α times the variance of the noise has to converge to zero, i.e., the noise variance decreases faster than the prior variance. Therefore, relative to the noise variance, the prior variance grows, i.e., the prior becomes non-informative. In other words, the smaller the noise is, the less (relative) regularization is needed. This seems to be a common interpretation of convergence results with respect to the noise covariance for maximum a-posteriori solutions in the Bayesian framework. As mentioned before, in [21] the authors ended up with a conclusion very similar to ours. Even in [23], where no regularization parameter is involved for the finite dimensional convergence analysis, one may conclude that, if the variance of the noise is sufficiently small, the prior does not influence the convergence behavior much. Hence, asymptotically, the prior distribution becomes non-informative. The requirement $\alpha \rightarrow \infty$ as observed in Corollary 4.3.7 can be interpreted as keeping the prior informative relative to the likelihood function as long as possible.

Again using a lifting technique, the deterministic convergence rate from Theorem 4.3.1 can be carried over into the stochastic setting. We will prove convergence rates using the finite model, then the infinite model. To simplify the notation we denote $L_m(\eta) := \ln^- (\eta^2 2\pi m^2 (\frac{\epsilon}{2})^m)$ from Proposition 1.4.5. Additionally we define

$$\mathcal{E}(\eta, m, \alpha) := \sqrt{2}\eta\sqrt{m - L_m(\eta)} + \sqrt{2}\eta\sqrt{m - L_m(\eta) + \frac{\alpha q^p}{2}}.$$

Theorem 4.3.8. *Let $X_n = T_n^* T_n \mathbf{X}$ be defined as $B_{p,p}^s(\mathbb{R}^d)$ -random variable according to Definition 4.2.2 for $1 \leq p \leq 2$ and take $s \in \mathbb{R}$ such that $\varsigma = s + d(\frac{1}{2} - \frac{1}{p}) \geq 0$. Let x_α^{MAP} be the maximum a-posteriori estimate (4.12) for the solution of (4.8) in the Bayesian framework with a $B_{p,p}^s$ -Besov space prior (4.11) according to model (MII). Assume we are given noisy data $y^\eta \in \mathbb{R}^m$ such that the error in each component of y^η is normally*

distributed with zero mean and variance η^2 . Assume additionally that the operator \mathbf{A} fulfils (4.35) with $\beta > 0$, $A_l > 0$. If the minimum norm solution \mathbf{x}^\dagger is unique and smooth enough, i.e., there is an a-priori estimate $\|\mathbf{x}^\dagger\|_{B_{p,p}^s(\mathbb{R}^d)} \leq \varrho$ for some $\varrho > 0$, then as $\eta \rightarrow 0$ the maximum a-posteriori solution $x_{\hat{\alpha}}^{\text{MAP}}$ converges to the solution with minimal norm $\|\cdot\|_{B_{p,p}^s(\mathbb{R}^d)}$ provided that the parameter $\alpha = \alpha(\eta, \varrho, \beta, \varsigma, p)$ is chosen such that

$$\begin{aligned} & \min \left(\left(\frac{\sqrt{2}}{A_l} \mathcal{E}(\sigma, m, \alpha) \right)^{\frac{\varsigma}{\beta+\varsigma}} \left(\varrho + \left(\varrho^p + \frac{2m - L_m(\eta)}{\alpha} \right)^{\frac{1}{p}} \right)^{\frac{\beta}{\beta+\varsigma}}, 1 \right) \\ &= \frac{\Gamma(\frac{m}{2}, m - \ln^-(\eta'(m)))}{\Gamma(\frac{m}{2})} + \frac{\Gamma(\frac{n}{p}, \frac{\alpha \varrho^p}{2})}{\Gamma(\frac{n}{p})} \end{aligned} \quad (4.37)$$

is fulfilled. Additionally, it is

$$\rho_K(x_{\hat{\alpha}}^{\text{MAP}}, x^\dagger) = \mathcal{O} \left(\left(\eta \sqrt{1 + |\ln(\eta)| + \frac{\alpha \varrho^p}{2}} \right)^{\frac{\varsigma}{\beta+\varsigma}} \varrho^{\frac{\beta}{\beta+\varsigma}} \right).$$

Proof. To further improve readability we define $\theta := \frac{\varsigma}{\beta+\varsigma}$, $\theta' := \frac{\beta}{\beta+\varsigma}$ and $\delta := \sqrt{2}\eta\sqrt{m - L_m(\eta)}$. Then according to Proposition 1.4.5, $\rho_K(y, y^\eta) \leq \delta$ holds. From (4.33) and Proposition 4.3.3 we know

$$\begin{aligned} & \sup\{\|x_{\hat{\alpha}}^{\text{MAP}} - x\| : x \in \mathcal{X}, y \in \mathcal{Y}, \|Ax - y\| \leq \delta, \|T_n^* x\|_{B_{p,p}^s(\mathbb{R}^d)} \leq \varrho\} \\ &= M_{\hat{\alpha}}(\delta, \varrho) < CA_l^{-\theta} (\delta + \delta')^\theta (\varrho + \varrho')^{\theta'}. \end{aligned}$$

$M_{\hat{\alpha}}$ is a deterministic quantity. In particular $\|x_{\hat{\alpha}}^{\text{MAP}} - x^\dagger\| \leq CA_l^{-\theta} (\delta + \delta')^\theta (\varrho + \varrho')^{\theta'}$ whenever $\|Ax^\dagger - y^\eta\| \leq \delta$ and $\|T_n^* x^\dagger\|_{B_{p,p}^s(\mathbb{R}^d)} \leq \varrho$. On the other hand, $\|x_{\hat{\alpha}}^{\text{MAP}} - x^\dagger\|$ may be larger than $CA_l^{-\theta} (\delta + \delta')^\theta (\varrho + \varrho')^{\theta'}$ if at least one of the conditions above is violated. Hence

$$\begin{aligned} & \mathbb{P}(\{\omega \in \Omega : \|x_{\hat{\alpha}}^{\text{MAP}}(\omega) - x^\dagger(\omega)\| > CA_l^{-\theta} (\delta + \delta')^\theta (\varrho + \varrho')^{\theta'}\}) \\ & \leq \mathbb{P}(\{\omega : \|Ax^\dagger(\omega) - y^\eta(\omega)\| > \delta \vee \|T_n^* x^\dagger(\omega)\|_{B_{p,p}^s(\mathbb{R}^d)} \geq \varrho\}) \\ & \leq \mathbb{P}(\{\omega : \|Ax^\dagger(\omega) - y^\eta(\omega)\| > \delta\}) + \mathbb{P}(\{\omega : \|T_n^* x^\dagger(\omega)\|_{B_{p,p}^s(\mathbb{R}^d)} \geq \varrho\}) \end{aligned} \quad (4.38)$$

because for $A, B \subset \Omega : \mathbb{P}(A \cup B) \leq \mathbb{P}(A) + \mathbb{P}(B)$. Please note that the set $\{\omega : \|Ax^\dagger(\omega) - y^\eta(\omega)\| > \delta\}$ corresponds to $\mathbb{P}(\Omega \setminus \Omega_\varepsilon)$ with Ω_ε from the proof of Theorem 4.3.6, i.e., the subset of Ω for which we do not have a worst-case error bound $\|y(\omega) - y^\eta(\omega)\| \leq \hat{\eta}$ (in the notation of Theorem 4.3.6). For the probability $\mathbb{P}(\{\omega : \|T_n^* x^\dagger(\omega)\|_{B_{p,p}^s(\mathbb{R}^d)} \geq \varrho\})$, $\varrho > 0$ we derived in Lemma 4.2.4

$$\mathbb{P}(\|T_n^* x^\dagger(\omega)\|_{B_{p,p}^s(\mathbb{R}^d)} \geq \varrho) = \frac{\Gamma(\frac{n}{p}, \frac{\alpha \varrho^p}{2})}{\Gamma(\frac{n}{p})}. \quad (4.39)$$

The probability $\mathbb{P}(\|Ax^\dagger - y\| \geq \delta)$ is given in Lemma 1.3.2 with $c = \delta$ and $e = Ax^\dagger - y$,

$$\mathbb{P}(\|Ax^\dagger - y\| \geq \delta) = \frac{\Gamma(\frac{m}{2}, \frac{\delta^2}{2\eta^2})}{\Gamma(\frac{m}{2})} = \frac{\Gamma(\frac{m}{2}, m - L_m(\eta))}{\Gamma(\frac{m}{2})}. \quad (4.40)$$

Inserting (4.39) and (4.40) into (4.38) we arrive at

$$\begin{aligned} \mathbb{P}(\{\omega \in \Omega : \|x_{\hat{\alpha}}^{\text{MAP}}(\omega) - x^\dagger(\omega)\| > CA_l^{-\theta}(\delta + \delta')^\theta(\varrho + \varrho')^{\theta'}\}) \\ \leq \frac{\Gamma(\frac{m}{2}, \frac{\delta^2}{2\eta^2})}{\Gamma(\frac{m}{2})} + \frac{\Gamma(\frac{n}{p}, \frac{\alpha\varrho^p}{2})}{\Gamma(\frac{n}{p})}. \end{aligned} \quad (4.41)$$

Comparing this with the definition of the Ky Fan metric (1.24), we get an upper bound for $\rho_K(x_{\hat{\alpha}}^{\text{MAP}}, x^\dagger)$ if we choose α such that

$$CA_l^{-\eta}(\delta + \delta')^\eta(\varrho + \varrho')^{\eta'} = \frac{\Gamma(\frac{m}{2}, \frac{\delta^2}{2\eta^2})}{\Gamma(\frac{m}{2})} + \frac{\Gamma(\frac{n}{p}, \frac{\alpha\varrho^p}{2})}{\Gamma(\frac{n}{p})}. \quad (4.42)$$

Before we can solve (4.42) we have to calculate $\delta + \delta'$ and $\varrho + \varrho'$. Resubstituting the error δ and $\hat{\alpha} = \eta^2\alpha$ into (4.34) we get

$$\delta' = \sqrt{2}\eta\sqrt{m - L_m(\eta) + \frac{\alpha\varrho^p}{2}}$$

and

$$\delta + \delta' = \sqrt{2}\eta\sqrt{m - L_m(\eta)} + \sqrt{2}\eta\sqrt{m - L_m(\eta) + \frac{\alpha\varrho^p}{2}} =: \mathcal{E}(\eta, m, \alpha).$$

Analogously we find

$$\varrho + \varrho' = \varrho + \left(\varrho^p + \frac{2m - L_m(\eta)}{\alpha}\right)^{1/p}$$

and (4.42) reads

$$\begin{aligned} C \left(\frac{\sqrt{2}}{A_l}\mathcal{E}(\eta, m, \alpha)\right)^\theta \left(\varrho + \left(\varrho^p + \frac{2m - L_m(\eta)}{\alpha}\right)^{1/p}\right)^{\theta'} \\ = \frac{\Gamma(\frac{m}{2}, m - L_m(\eta))}{\Gamma(\frac{m}{2})} + \frac{\Gamma(\frac{n}{p}, \frac{\alpha\varrho^p}{2})}{\Gamma(\frac{n}{p})}. \end{aligned} \quad (4.43)$$

The only unknown quantity in (4.43) is the constant C . Since we have no information about it, we neglect it and set it to one. Solving (4.43) for α immediately gives an upper bound for $\rho_K(x_{\hat{\alpha}}^{\text{MAP}}, x^\dagger)$ by definition of the Ky Fan metric. Although (4.43) does not have an analytical solution, the nonlinear equation can still be solved numerically. By construction the convergence rate is given by both the left hand side and the right hand side of (4.43). \square

We obtain a similar result for the infinite dimensional model. To this end, we only have to replace the probability $\mathbb{P}(\|T_n^* x^\dagger\|_{\cdot,p} \geq \varrho)$ from Lemma 4.2.4 by the one from Corollary 4.2.3. We obtain the following Corollary where, as before, we use the notation $\mathcal{E}(\eta, m, \alpha) := \sqrt{2}\eta\sqrt{m - L_m(\eta)} + \sqrt{2}\eta\sqrt{m - L_m(\eta) + \frac{\alpha\varrho^p}{2}}$.

Corollary 4.3.9. *Let $x_{\hat{\alpha}}^{\text{MAP}}$ be the maximum a-posteriori estimate (4.12) for the solution of (4.8) in the Bayesian framework with a $B_{p,p}^s$ -Besov space prior (4.11) with $s \in \mathbb{R}$ fulfilling $\varsigma = s + d(\frac{1}{2} - \frac{1}{p}) \geq 0$. Let \mathbf{X} be defined in $B_{p,p}^r(\mathbb{R}^d)$ according to model (MI) with $s < r - \frac{d}{p}$. Assume we are given noisy data $y^n \in \mathbb{R}^m$ such that the error in each component of y^n is normally distributed with zero mean and variance η^2 . Assume additionally that the operator \mathbf{A} fulfils (4.35) with $\beta > 0$, $A_l > 0$. If the minimum norm solution \mathbf{x}^\dagger is unique and smooth enough, i.e., there is an a-priori estimate $\|\mathbf{x}^\dagger\|_{B_{p,p}^s(\mathbb{R}^d)} \leq \varrho$ for some $\varrho > 0$, then as $\eta \rightarrow 0$ the maximum a-posteriori solution $x_{\hat{\alpha}}^{\text{MAP}}$ converges to the solution with minimal norm $\|\cdot\|_{B_{p,p}^s(\mathbb{R}^d)}$ provided that the parameter $\alpha = \alpha(\eta, \varrho, \beta, \varsigma, p)$ is chosen such that*

$$\begin{aligned} & \min \left(\left(\frac{\sqrt{2}}{A_u} \mathcal{E}(\eta, m, \alpha) \right)^{\frac{\varsigma}{\beta+\varsigma}} \left(\varrho + \left(\varrho^p + \frac{2m - L_m(\eta)}{\alpha} \right)^{\frac{1}{p}} \right)^{\frac{\beta}{\beta+\varsigma}}, 1 \right) \\ & = \frac{\Gamma(\frac{m}{2}, m - L_m(\eta))}{\Gamma(\frac{m}{2})} + \frac{1}{\varrho} \left(\frac{2}{\alpha p} \left(\ell_\phi + \ell_\psi \sum_{j=0}^{\infty} 2^{-j((s-r)p-d)} \right) \right)^{\frac{1}{p}} \end{aligned} \quad (4.44)$$

is fulfilled. Additionally, it is

$$\rho_K(x_{\hat{\alpha}}^{\text{MAP}}, x^\dagger) = \mathcal{O} \left(\left(\eta \sqrt{1 + |\ln(\eta)| + \frac{\alpha\varrho^p}{2}} \right)^{\frac{\varsigma}{\beta+\varsigma}} \varrho^{\frac{\beta}{\beta+\varsigma}} \right). \quad (4.45)$$

So far we only used the Ky Fan metric to measure the noise in the data. If instead one only knows the expectation of the noise, a construction analogously to (3.50) can be used.

4.4 Numerical examples for $p = 1$

In this section we want to illustrate our theoretical results with a specific example for the case $p = 1$, as it is well known that this choice produces sparse solutions. We will consider a deconvolution problem with a given kernel function. Convolution operators appear in many fields, e.g., in signal processing, where the output of a linear time-invariant system is given by the convolution of the input signal with the impulse response, a fixed function depending on the system. In image processing, blurring effects can often be modelled as

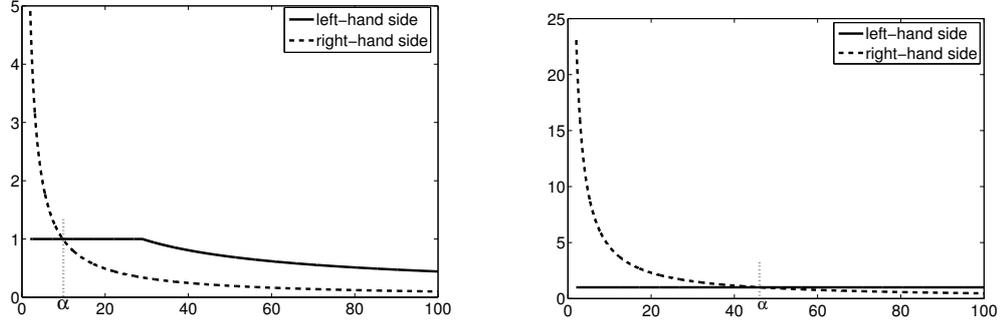


Figure 4.1: Plot of the left-hand side and right-hand side of (4.37) (MI, left) and (4.44) (MII, right) for $\eta = 0.01$, $m = 2500$, $\varsigma = 0.5$, $\beta = 1$, $\varrho = 2.16$. The optimal α is the one for which the intersection occurs.

convolution of an image with a smoothing kernel. Mathematically, we have an operator equation $\mathbf{A}\mathbf{x} = \mathbf{y}$ where $\mathbf{A} : L_2(\mathbb{R}^d) \rightarrow L_2(\mathbb{R}^d)$ is defined by

$$[\mathbf{A}\mathbf{x}](s) = [\mathbf{k} * \mathbf{x}](s) = \int_{\mathbb{R}^d} \mathbf{k}(s-t)\mathbf{x}(t)dt, \quad s \in \mathbb{R}^d \quad (4.46)$$

for some kernel function $\mathbf{k} \in L_2(\mathbb{R}^d)$. In order to use our theory, we have to require that \mathbf{A} fulfils (4.35) for some $\beta > 0$. Since the properties of \mathbf{A} are determined by its kernel \mathbf{k} , we just have to choose \mathbf{k} appropriately. Inequality (4.35) describes the equivalence of $\|\mathbf{A}\mathbf{h}\|_{L_2}$ with a norm of \mathbf{h} in a Sobolev space of negative order $H^{-\beta}$. Using Fourier analysis we have

$$\|\mathbf{h}\|_{H^{-\beta}} = \int_{\mathbb{R}^d} (1 + |\xi|^2)^{-\beta} |\widehat{\mathbf{h}}(\xi)|^2 d\xi, \quad (4.47)$$

where $\widehat{\mathbf{h}}$ denotes the Fourier transform of \mathbf{h} . Because of the Fourier-convolution theorem

$$\|\mathbf{A}\mathbf{h}\|_{L_2} = \|\widehat{\mathbf{k}} \cdot \widehat{\mathbf{h}}\|_{L_2} = \int_{\mathbb{R}^d} |\widehat{\mathbf{k}}(\xi) \cdot \widehat{\mathbf{h}}(\xi)|^2 d\xi. \quad (4.48)$$

Comparing (4.47) and (4.48) we may define $\widehat{\mathbf{k}}(\xi) := (1 + |\xi|^2)^{-\beta/2}$ and obtain equality in (4.35) with $A_u = A_l = 1$. To control the width of the convolution filter and its norm we introduce additional constants $\kappa > 0$, $c_{\kappa,\beta} > 0$ and define

$$\widehat{\mathbf{k}}(\xi) = \frac{c_{\kappa,\beta}}{(1 + \kappa|\xi|^2)^{\beta/2}}, \quad \xi \in \mathbb{R}^d. \quad (4.49)$$

Now (4.35) holds with $A_l = c_{\kappa,\beta}^2$ and $A_u = \frac{c_{\kappa,\beta}^2}{\kappa}$ for $\kappa \leq 1$ and vice versa for $\kappa > 1$. The maximum a-posteriori solution, i.e., the minimizer of the Tikhonov functional (4.12) was

calculated with the *iterative soft-shrinking algorithm* proposed in [89]. Here we use the formulation of [91] that admits slightly weaker restrictions. Starting from an initial guess x_0 , the iterates are given by

$$x_{k+1} = \mathcal{S}_{\mathbf{t},p}(x_k + \gamma A^*(y - Ax_k)), \quad k = 1, 2, \dots \quad (4.50)$$

with an appropriately chosen stepsize $\gamma < 2/\|A\|^2$, where the thresholding operator $\mathcal{S}_{\mathbf{t},p}(h) := \sum_{\lambda \in \Lambda} S_{\tau_\lambda,p}(\langle h, \psi_\lambda \rangle) \psi_\lambda$ is defined componentwise. For $p = 1$ we have

$$S_{\tau_\lambda,1}(x) := \begin{cases} x - \frac{\tau_\lambda}{2} & x \geq \frac{\tau_\lambda}{2} \\ 0 & |x| < \frac{\tau_\lambda}{2} \\ x + \frac{\tau_\lambda}{2} & x \leq -\frac{\tau_\lambda}{2} \end{cases}.$$

We chose $c_{\kappa,\beta}$ in (4.49) such that $\|\mathbf{k}\|_{L_1(\mathbb{R}^d)} < 1$ and hence (4.50) converges with $\gamma = 1$ to the minimizer of (4.12). The thresholding parameters τ_λ depend on the regularization parameter $\hat{\alpha}$ and the weights $2^{s p |\lambda|}$ from the definition of the Besov-space norm $\|\cdot\|_{B_{p,p}^s(\mathbb{R}^d)}$ in (4.18). Written in full dependence of all parameters, $\tau_\lambda = \alpha \eta^2 2^{s+d(\frac{1}{2}-\frac{1}{p})p|\lambda|}$ where $|\lambda|$ is the scale of the wavelet. Because the kernel is symmetric we have $A^* = A$. The application of A and A^* in (4.50) can easily be implemented using Fourier transformation and the convolution theorem. In order to calculate the regularization parameter $\hat{\alpha} = \alpha \eta^2$, we have to solve (4.37) or (4.44), respectively, for α . This can be done with Newton's method after obtaining a good initial guess, for example with the bisection method. A typical situation is shown in Figure 4.1.

Each of the two definitions of the Besov space random variable allows for a slightly different implementation of the parameter choice rule. To illustrate the behavior of both variants, we consider an academic example of a one dimensional signal \mathbf{x} that is sparse with respect to the Haar basis in $L_2(\mathbb{R})$ and its convolution with a kernel of type (4.49). A sample of signal, measurements and corresponding regularized solution is shown in Figure 4.2, where $m = n = 2500$ and $\beta = 1$. Next we want to compare the predicted convergence properties to the numerical results. The behaviour of our parameter choice rules (4.37) and (4.44) with respect to η is demonstrated in Figure 4.3. Both models (MI) and (MII) lead to parameters α and $\hat{\alpha}$ fulfilling the theoretical conditions. The numerically obtained errors follow the theoretically predicted convergence rates.

So far we did not address the question of convergence of the solutions if m and n are not fixed anymore but increasing. Although theoretical results are missing at this point, Figure 4.4 shows a comparison of the parameter choice rules for the models (MI) and (MII). In contrary to the convergence with respect to decreasing variance η , the two models show a distinct convergence behaviour with respect to m . It is future work to give detailed analysis on this.

Until now we only considered the one dimensional case. However, Figure 4.5 shows that also for two dimensional problems our parameter choice rule works and leads to

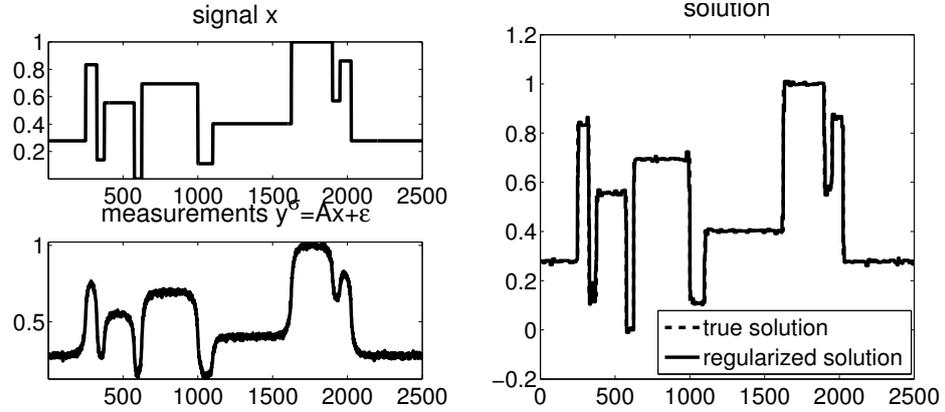


Figure 4.2: Signal, noisy measurements and regularized solution for $\eta = 0.01$, obtained with our parameter choice rule for the finite dimensional model (MII). We used the exact ϱ and chose $s = 1$ such that $\varsigma = 0.5$. Resulting from $\alpha = 45.85$ as solution of (4.43) we obtained the effective regularization parameter $\hat{\alpha} = \alpha \cdot \eta^2 = 0.0045$.

reasonable reconstructions. Figure 4.6 shows that in the 2-dimensional case the regularization parameter $\hat{\alpha} = \alpha\eta^2$ with α chosen according to (4.44) for the infinite dimensional model (MI) keeps the number of recovered non-zero coefficients nearly constant as $\eta \rightarrow 0$.

4.5 Application to Atmospheric Tomography

The theory of this chapter is also applicable to the problem of Atmospheric Tomography and, in particular, to the setting of MCAO as presented in Section 2.6. We will also adopt the notation from that section. Therefore the unknown turbulence profiles are denoted by Φ instead of x , which will be used as a spatial coordinate, and the measured wavefront data is φ instead of y . The van Karman model for the atmospheric turbulence (2.11) implies a certain smoothness of the turbulence in terms of the Sobolev scale. Let C_Φ be the covariance operator induced by the van Karman model. Then $\|C_\Phi^{-1/2}\Phi\|_{L_2}^2$ is a natural candidate for a prior in the Bayesian setting. In practice, one usually approximates the covariance via [131]

$$\|C_\Phi^{-1/2}\Phi\|_{L_2}^2 \approx \frac{1}{C_n^2(h)} \left(\tilde{L}_0^{-\frac{11}{3}} \|\Phi\|_{L_2} + \|(-\Delta)^{11}12\Phi\|_{L_2} \right) \quad (4.51)$$

where Δ is the Laplace operator. Since

$$\langle \Phi, \Psi \rangle_{H^s(\mathbb{R}^2)} = \int_{\mathbb{R}^2} ((I_\Delta)^{\frac{s}{2}}\Phi)(x)((I_\Delta)^{\frac{s}{2}}\Psi)(x) dx$$

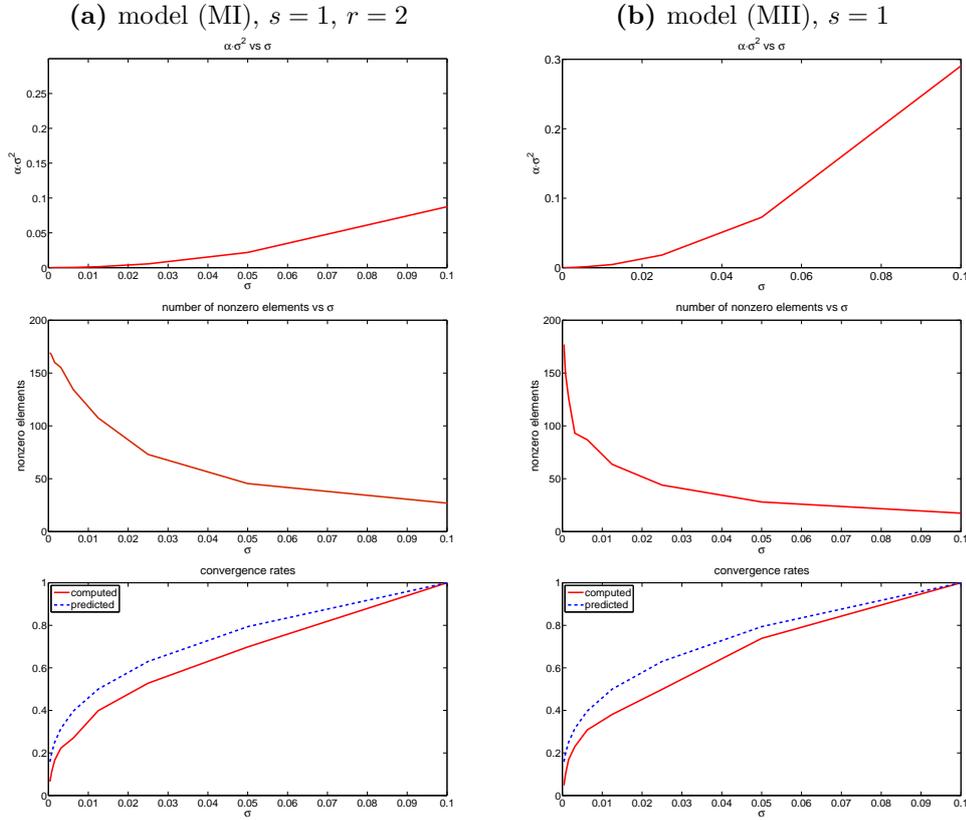


Figure 4.3: Numerical comparison of the two parameter choice rules (4.44) and (4.37), respectively, for the one dimensional deconvolution problem. The values for $\alpha \cdot \eta^2$, plotted against η , are shown in the first row. For sufficiently small η , α starts to grow, following the theory developed in Corollary 4.3.7. However, $\alpha \cdot \eta^2$ still goes to zero. The number of recovered non-zero coefficients in the solution is shown in the second row. Since $n = 2500$ we end up with a sparse solution for all η used in the simulations. In the last row we plotted the obtained reconstruction error $\|x^{MAP} - x^*\|$ and compare it with the convergence rates $O\left((\eta\sqrt{m})^{\frac{\zeta}{\zeta+\beta}}\right)$ predicted in Theorem 4.3.8 and Corollary 4.3.9, respectively.

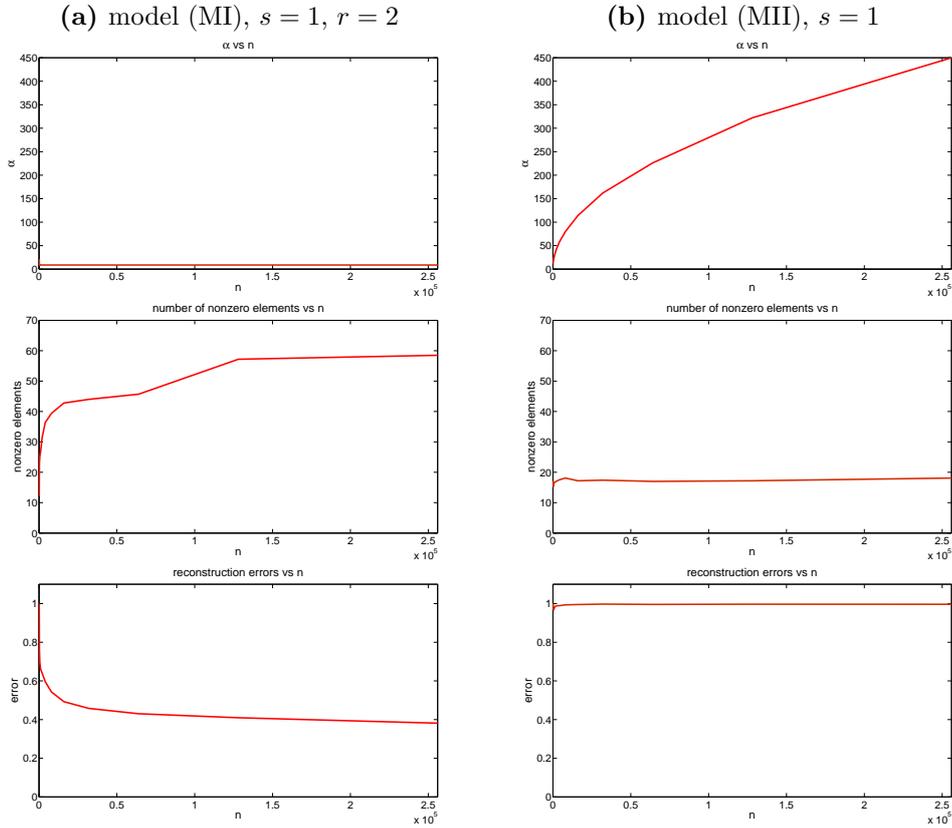


Figure 4.4: Comparison of the two implementations of the parameter choice rule with respect to varying level of discretization n . The variance $\eta = 0.01$ was held constant. As estimate for ϱ we used the exact value calculated from the true solution scaled according to the respective models. While the finite model leads to growing α for increasing n , the infinite model keeps the regularization parameter constant. The number of non-zero coefficients in the recovered solution behaves conversely. To plot the reconstruction errors we used the same scaling for both models. While in the finite model the error stays the same, it decreases for growing n in the infinite model.

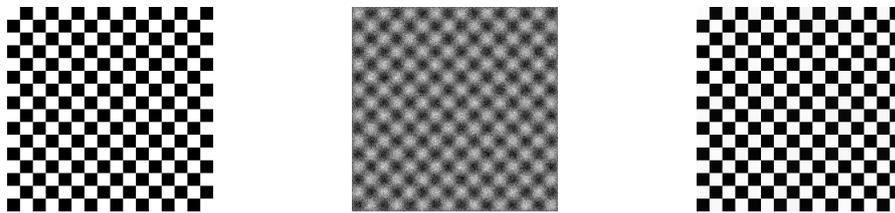


Figure 4.5: 2D example, left: signal x with 68 non-zero coefficients, middle: measurements $Ax + \epsilon$, $\eta = 0.1$, $\beta = 1$, right: solution with $\alpha = 130.5$, exactly the 68 true coefficients were recovered.

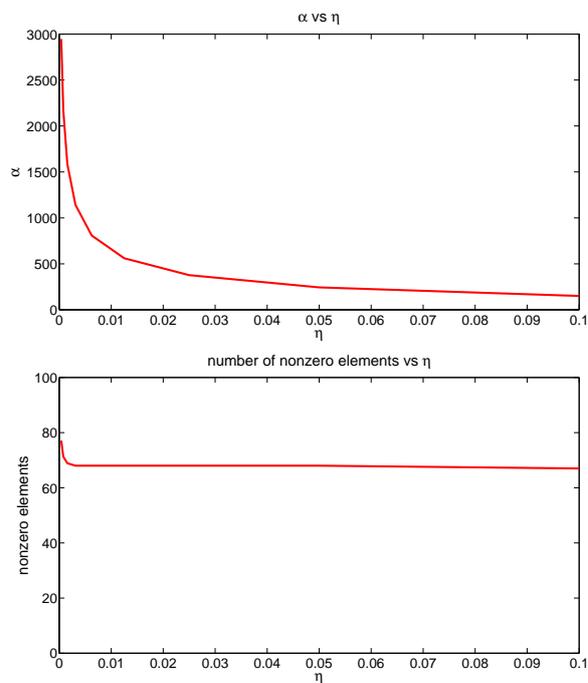


Figure 4.6: 2D example. Upper plot: regularization parameter α against η . α grows for decreasing η according to Corollary 4.3.7; lower plot: number of recovered non-zero elements against η . Out of $n = 65536$ coefficients of the original image 68 had non-zero values.

defines a scalar product in $H^s(\mathbb{R}^2)$ [112], an $H^{\frac{11}{6}}(\mathbb{R}^2)$ prior is equivalent to the characterization 4.51. Since $H^{\frac{11}{6}}(\mathbb{R}^2) = B_{2,2}^{\frac{11}{6}}(\mathbb{R}^2)$, one can use a Besov prior as introduced in Section 4.2.2. In the choice of the wavelet we follow [131] and use the *db3* wavelet. Since Φ is a vector valued function, we have to adjust the definition of the Besov space slightly. We define a norm tailored to the layer-model of the atmosphere via

$$\|\Phi\|_{\otimes B_{2,2}^{\frac{11}{6}}}^2 := \sum_{l=1}^L \frac{1}{c_l} \|\Phi^{(l)}\|_{B_{2,2}^{\frac{11}{6}}(\Omega_l)}. \quad (4.52)$$

With this we find Φ by minimizing

$$J(\Phi) = \|A\Phi - \varphi\|_{(L_2(\Omega_D))^G} + \hat{\alpha} \|\Phi\|_{\otimes B_{2,2}^{\frac{11}{6}}}^2 \quad (4.53)$$

with an appropriate regularization parameter $\hat{\alpha}$. The minimizer is calculated with a soft shrinkage algorithm analogously to (4.50). Starting from an initial guess Φ_0 , we iterate

$$\Phi_{k+1} = \mathcal{S}_{t,2}(\Phi_k + \gamma A^*(\varphi - A\Phi_k)), \quad k = 1, 2, \dots \quad (4.54)$$

where the shrinkage operator $\mathcal{S}_{t,2}(h) := \sum_{\lambda \in \Lambda} S_{\tau_\lambda,2}(\langle h, \psi_\lambda \rangle) \psi_\lambda$ is now for $p = 2$ defined via $S_{\tau_\lambda,2} := F_{\tau_\lambda,2}^{-1}$ with

$$F_{\tau_\lambda,p}^{-1}(t) = t + \frac{\tau_\lambda p}{2} \text{sign}(t) |t|^{p-1},$$

see [89]. In [132] Saxenhuber and Ramlau introduced a gradient based algorithm with the iterative procedure

$$\Phi_{k+1} = \Phi_k + \gamma A^*(\varphi - A\Phi_k), \quad k = 1, 2, \dots$$

where γ is a step size parameter. They provided the code for their algorithm and we just had to include the shrinkage operator to obtain 4.54. All that remains is to find an appropriate regularization parameter $\hat{\alpha}$. Since the problem fits into the framework of this chapter, we calculate α via Equation (4.44), i.e., we use the infinite dimensional model due to the given discretization. Namely, we have $m = 43350$ measured data point ans seek to reconstruct Φ on $n = 12243$ points. The smoothness parameter r in the infinite model was chosen to barely fulfill condition iii) of Lemma 4.2.2. We set $r = s + \frac{d}{p} + 0.0001$. Since we use $B_{2,2}^{\frac{11}{6}}(\mathbb{R}^2)$ priors, we have $\varsigma = \frac{11}{6}$ from (4.17). Because Atmospheric Tomography is a special case of limited angle tomography, we have $\beta = \frac{1}{2}$ in (4.35). The upper bound ϱ for the norm of the solution in (4.32) was estimated from solutions obtained with the gradient method. We typically obtained values of ϱ around 1.7 and chose $\varrho = 2$ for the calculation of α . Since the MOST simulation tool only allows to simulate the high flux case, we have a very small noise variance $\eta = 8.8 \cdot 10^{-19}$. Plugging all the constants into (4.44), we obtained $\hat{\alpha} = 2.27 \cdot 10^{-27}$ as the regularization

parameter. Although the value is small, we observed a clear improvement of about 3% to 4% between the Strehl ratio obtained with the soft-thresholding algorithm and the basic gradient algorithm, see Figure 4.7 for an example. We would expect to obtain even higher relative gain in Strehl ratio when the noise level increases.

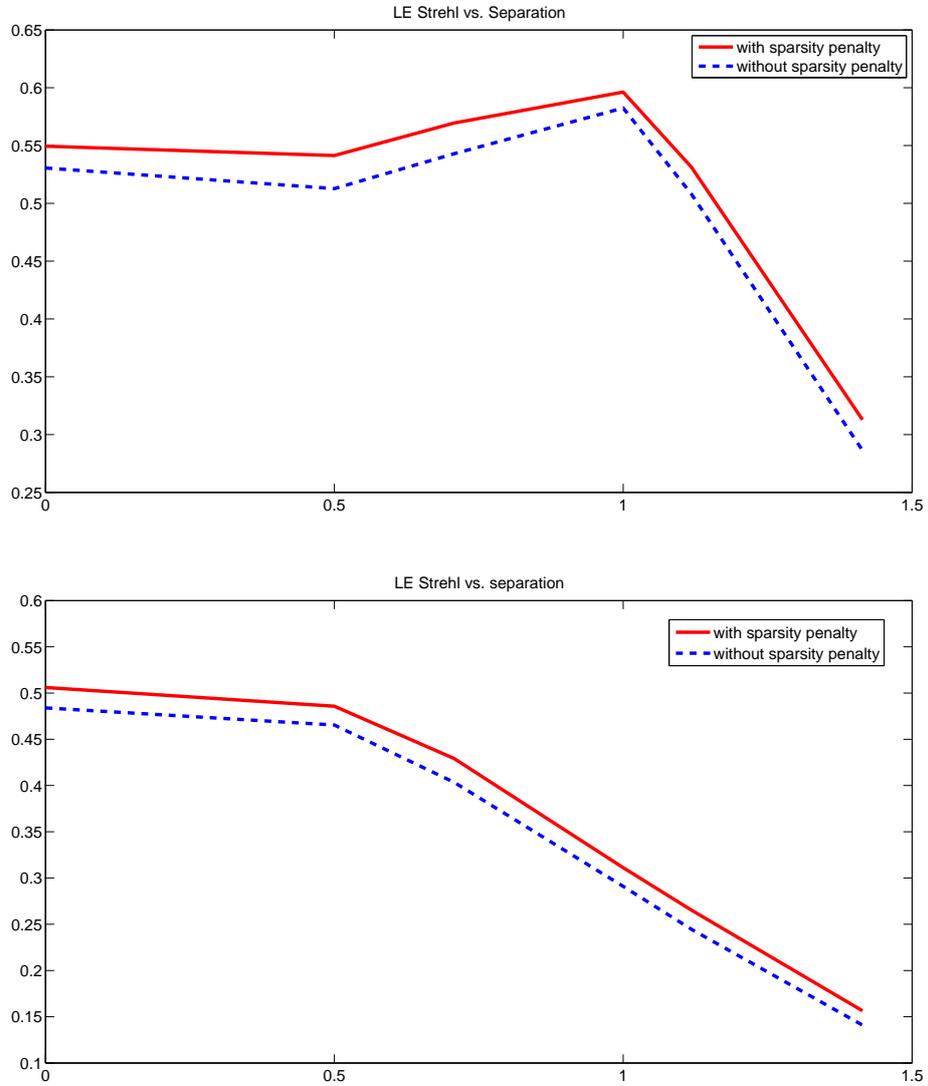


Figure 4.7: Gradient algorithm with and without thresholding. Top: $B_{2,2}^{11/6}$ -prior for 6 natural guide stars. Bottom: $B_{1,1}^1$ prior for 6 laser guide stars. In both experiments, the sparsity regularization improves the Strehl over the whole field of view.

Chapter 5

Summary and generalization of the lifting strategy

The lifting arguments from Chapter 3 and Chapter 4 are based on properties of the error and the fact that there is a deterministic regularization theory for the situation under investigation. However, neither the setting itself nor the actual regularization method influence the lifting analysis. In this chapter, we summarize and generalize the main results of the previous chapters.

5.1 Theoretical results

In particular, the lifting arguments are independent of the spaces involved and the structure of the Inverse Problem. We will therefore allow certain nonlinear ill-posed problems

$$F(x) = y \tag{5.1}$$

with an operator $F : \mathcal{X} \rightarrow \mathcal{Y}$ mapping between Banach or Hilbert spaces \mathcal{X} and \mathcal{Y} . Instead of the exact data, again only noisy data y^δ is available. We assume that in the deterministic setting a regularization method to solve (5.1) is available. That is, in generalization of Definition 1.2.2 an operator $R_\alpha : Y \rightarrow X$, $\mathbb{R}_\alpha(y^\delta) \mapsto x_\alpha^\delta$ is available that maps the given data to a regularized solution and fulfills, possibly under some assumptions \mathcal{A} ,

$$\lim_{\delta \rightarrow 0} \mathbb{R}_\alpha(y^\delta) = x^\dagger \tag{5.2}$$

given that the parameter α is chosen appropriately. As before, x^\dagger is a minimum norm solution to (5.1). Note that, in particular for nonlinear problems, x^\dagger does not need to be unique. Hofinger [30] pointed out that this is problematic for the lifting arguments. One often proves convergence of subsequences to the desired solution, and then deduces convergence of the whole series of regularized solutions if possible. In the stochastic setting, this is not possible in general since subsequences for different ω do not have to be related. A constructed example for this behavior can be found in [30]. In order to lift general deterministic regularization methods into the stochastic setting we must therefore require that x^\dagger is unique. We obtain the following Theorem which is a straightforward generalization of Theorem 4.3.6. The main ingredient of the proof is that the Ky Fan metric allows a natural decomposition of Ω into two subsets. On the first, a deterministic error bound is available. On the second subset, which has low probability, this is not the case. If one uses the expectation to describe the properties of the noise, the same technique as in Section 3.3.1 can be applied. In order to achieve convergence, we have to “inflate” the expectation which we realize again by adding a factor $\tau = \tau(\eta)$ to the expectation.

Theorem 5.1.1. *Let R_α be a regularization method for the solution of (5.1) in the deterministic setting under a suitable choice of the regularization parameter. Let now $y^\eta = y + \epsilon(\eta)$ where $\epsilon(\eta)$ is a stochastic error such that*

a) $\rho_K(y, y^\eta) \rightarrow 0$ or

b) $\mathbb{E}(\|\epsilon\|) = f(\eta) \rightarrow 0$

as $\eta \rightarrow 0$. Then, assuming (5.1) has a unique solution x^\dagger and all necessary assumptions for the deterministic theory (except the bound on the noise) hold with probability one, the regularization method R_α fulfills

$$\lim_{\eta \rightarrow 0} \rho_K(x^\dagger, R_\alpha(y^\eta)) = 0$$

under the same parameter choice rule as in the deterministic setting with δ replaced by $\rho_K(y, y^\eta)$ (case a)) or $\tau(\eta)\mathbb{E}(\|\epsilon\|)$ where $\tau(\eta)$ fulfills (3.50) (case b)). If the regularized solutions are defined by (1.26) if they fail to fulfill Assumption 1.4.2, then in both cases

$$\lim_{\eta \rightarrow 0} \mathbb{E}(\|x^\dagger - R_\alpha(y^\eta)\|) = 0.$$

Proof. In case a) where the Ky Fan distance between noisy and true data is given, the proof coincides with the one from Theorem 4.3.6. Only the parameter choice has to be adjusted such that $\rho_k(y, y^\eta)$ is used instead of the deterministic δ . In case b), where we assume to know the expectation of the error, a small adjustment has to be made.

Set $\varepsilon := \frac{1}{\tau(\eta)}$. Then $\varepsilon \rightarrow 0$ as $\eta \rightarrow 0$ and

$$\mathbb{P}(\|y - y^\eta\| \leq 2\tau(\eta)\mathbb{E}(\|\epsilon\|)) \leq \frac{1}{2\tau(\eta)} = \frac{\varepsilon}{2}.$$

Then we again have found Ω_ε with $\mathbb{P}(\Omega_\varepsilon) \geq 1 - \frac{\varepsilon}{2}$. The remainder of the proof remains the same. The parameter choice rule now has to be carried out with $\tau(\eta)\mathbb{E}(\|\epsilon\|)$ instead of δ . \square

Hence it is possible to use rather general deterministic regularization methods in a stochastic setting. However, the expectation of the noise alone decreases too fast; it has to be slowed down to find an appropriate regularization parameter or replaced by the Ky Fan distance which is slower than the expectation. A numerical example will be given at the end of this chapter.

As usual, convergence rates are obtained under additional assumptions. Sometimes they are already such that (local) uniqueness of the true solution is ensured. If not, we have to require such a property for the same reason as previously.

Theorem 5.1.2. *Let R_α be a regularization method for the solution of (5.1) in the deterministic setting such that, under a set of assumptions on the operator F and the solutions x^\dagger and a suitable choice of the regularization parameter,*

$$d_{\mathcal{X}}(x^\dagger, R_\alpha(y^\delta)) \leq \Phi(d_{\mathcal{Y}}(y, y^\delta))$$

with a monotonically increasing right-continuous function Φ .

Let now $y^\eta = y + \epsilon(\eta)$ where $\epsilon(\eta)$ is a stochastic error such that

a) $\rho_K(y, y^n) \rightarrow 0$ or

b) $\mathbb{E}(\|\epsilon\|) = f(\eta) \rightarrow 0$

as $\eta \rightarrow 0$. Then, assuming (5.1) has a (locally) unique solution x^\dagger and all necessary assumptions for the deterministic theory (except the bound on the noise) hold with probability one, the regularization method R_α fulfills

$$\rho_K(x^\dagger, R_\alpha(y^n)) = \mathcal{O}(\max\{\Phi(\rho_K(y, y^n)), \rho_K(y, y^n)\})$$

in case a) or, respectively, in case b),

$$\rho_K(x^\dagger, R_\alpha(y^n)) = \mathcal{O}(\max\{\Phi(\tau(\eta)\mathbb{E}(\|\epsilon\|)), \mathbb{P}(\|\epsilon\| \geq \tau(\eta)\mathbb{E}(\|\epsilon\|))\})$$

under the same parameter choice rule as in the deterministic setting with δ replaced by $\rho_K(y, y^n)$ (case a)) or $\tau(\eta)\mathbb{E}(\|\epsilon\|)$ where $\tau(\eta)$ fulfills (3.50) (case b)).

Proof. Convergence rates in the Ky Fan metric follow directly from Theorem 1.4.3. If the expectation is used as measure for the data error, we have

$$\mathbb{P}(\|\epsilon\| \geq \tau(\eta)\mathbb{E}(\|\epsilon\|)) \leq \frac{1}{\tau(\eta)}$$

as in (3.48). Hence, with probability $1 - \frac{1}{\tau(\eta)}$ we are in the deterministic setting with $\delta = \tau(\eta)\mathbb{E}(\|\epsilon\|)$. The convergence rate follows by the definition of the Ky Fan metric. \square

For many nonlinear Inverse Problems the requirement of a unique solution is too strong. Often one has several solutions of the same quality, in particular there exists more than one minimum norm solution. In this case, Theorem 5.1.1 is not applicable. In Hofingers example [30, Example 4.3 and 4.5] with two minimum norm solutions the noise was constructed such that, while the error in the data converges to zero, for each fixed $\omega \in \Omega$ the regularized solutions jump between both solutions such that no converging subsequence can be found. The main problem there is that the Ky Fan distance cannot incorporate the concept that all minimum norm solutions are equally acceptable. We will now define a pseudo metric that resolves this issue.

Definition 5.1.1. Let $(\mathcal{X}, d_{\mathcal{X}})$ be a metric space. Denote with \mathcal{L} the set of minimum-norm solutions to (5.1). Then

$$\rho_K^{\mathcal{L}}(x) := \inf_{\varepsilon > 0} \left\{ \mathbb{P} \left(\inf_{x^\dagger \in \mathcal{L}} d_{\mathcal{X}}(x, x^\dagger) > \varepsilon \right) \leq \varepsilon \right\} \quad (5.3)$$

measures the distance between an element $x \in \mathcal{X}$ and the set \mathcal{L} , in particular it is

$$\rho_K^{\mathcal{L}}(x) = 0 \quad \Leftrightarrow \quad x \in \mathcal{L} \quad \text{a.s.}$$

With this, one can define a pseudometric on $(\Omega, \mathcal{F}, \mathbb{P})$ via

$$\rho_K^{\mathcal{L}}(x_1, x_2) =: \max\{\rho_K^{\mathcal{L}}(x_1), \rho_K^{\mathcal{L}}(x_2)\}. \quad (5.4)$$

Obviously (5.4) is positive, symmetric and fulfills the triangle inequality. However, $\rho_K^{\mathcal{L}}(x_1, x_2) = 0$ does not imply $x_1 = x_2$ a.e. but instead $x_1 \wedge x_2 \in \mathcal{L}$ which fixes the aforementioned issue of the Ky Fan metric and allows the following theorems.

Theorem 5.1.3. *Let R_α be a regularization method for the solution of (5.1) in the deterministic setting under a suitable choice of the regularization parameter. Let now $y^\eta = y + \epsilon(\eta)$ where $\epsilon(\eta)$ is a stochastic error such that*

- a) $\rho_K(y, y^\eta) \rightarrow 0$ or
- b) $\mathbb{E}(\|\epsilon\|) = f(\eta) \rightarrow 0$

as $\eta \rightarrow 0$. Then, assuming all necessary assumptions for the deterministic theory (except the bound on the noise) hold with probability one, the regularization method R_α fulfills

$$\lim_{\eta \rightarrow 0} \rho_K^{\mathcal{L}}(R_\alpha(y^\eta)) = 0$$

under the same parameter choice rule as in the deterministic setting with δ replaced by $\rho_K(y, y^\eta)$ (case a)) or $\tau(\eta)\mathbb{E}(\|\epsilon\|)$ where $\tau(\eta)$ fulfills (3.50) (case b)). In particular, the series of regularized solutions fulfills

$$\lim_{\eta_1, \eta_2 \rightarrow 0} \rho_K^{\mathcal{L}}(R_\alpha(y^{\eta_1}), R_\alpha(y^{\eta_2})) = 0$$

Proof. The proof follows the lines of the one of Theorem 4.3.6 with $\rho_K(\cdot, x^\dagger)$ replaced by $\rho_K^{\mathcal{L}}(\cdot)$. Also Lemma 4.3.4 is easily adjusted to incorporate multiple solutions. \square

5.2 An example in theory and numerical results

We return to sparsity promoting inversion similar to the strategy in the previous chapter. Instead of a linear Inverse Problem and a Besov-space penalty term we now consider a nonlinear problem (5.1) with an ℓ_1 -penalty. The regularization parameter will be determined with the discrepancy principle. The deterministic results are taken from [11]. In particular, we will adopt the numerical example and present the theory already tailored to this problem.

We consider an autoconvolution equation

$$[F(x)](s) = \int_0^s x(s-t)x(t) dt, \quad 0 \leq s \leq 1 \quad (5.5)$$

between Hilbert spaces $\mathcal{X} = L_2[0, 1]$ and $\mathcal{Y} = L_2[0, 1]$ where $x \in \mathcal{D}(F)$. Such an equation is of great interest in, for example, stochastics or spectroscopy and has been analyzed in detail in [133]. Recently, a more complicated autoconvolution problem has emerged from a novel method to characterize ultra-short laser pulses [134, 135].

Using the Haar-wavelet basis (see, e.g., [120] and Section 4.4), the authors of [11] reformulate (5.5) as an equation from ℓ_2 to ℓ_2 by switching to the space of coefficients in the Haar basis. In order to stabilize the inversion, an ℓ_1 penalty term is used such that the task is to minimize the functional

$$J_\alpha(x) = \|F(x) - y^\delta\|^2 + \alpha \|x\|_1. \quad (5.6)$$

For simplicity we use the same notation as before although the objects are slightly different now. The regularization parameter α in (5.6) is chosen according to the discrepancy principle. In [11], the following formulation is used: For $1 < \tau_1 \leq \tau_2$ choose $\alpha = \alpha(\delta, y^\delta)$ such that

$$\tau_1 \delta \leq \|F(x_\alpha^\delta) - y^\delta\| \leq \tau_2 \delta \quad (5.7)$$

holds. The authors show that this leads to a convergence of the regularized solutions against a solution of (5.5) with minimal ℓ_1 -norm of its coefficients. It was also shown that the regularization parameter fulfills

$$\alpha(\delta, y^\delta) \rightarrow 0, \quad \frac{\delta^2}{\alpha(\delta, y^\delta)} \rightarrow 0 \quad \text{as } \delta \rightarrow 0. \quad (5.8)$$

By courtesy of S. Anzengruber we were allowed to use the original code for the numerical simulation in [11]. We only changed the parts directly connected to the data noise. Namely, we replaced the deterministic error $\|y - y^\delta\| \leq \delta$ with the i.i.d Gaussian noise from the previous chapters,

$$y^\eta = y + \epsilon$$

with $\epsilon \sim \mathcal{N}(0, \eta^2 I)$. The discretization is due to the truncation of the expansion of the functions in the Haar-basis after m elements. Therefore

$$\mathbb{E}(\|\epsilon\|) = \frac{\eta}{\sqrt{2}} \frac{\Gamma(\frac{m+1}{2})}{\Gamma(\frac{m}{2})},$$

see Lemma 1.3.2. In practice, we used the upper bound

$$\mathbb{E}(\|\epsilon\|) \leq \eta \sqrt{m}$$

since, as shown in this chapter, the expectation has to be “blown up” anyway. The parameter choice (5.7) was realized with δ replaced by $\tau(\eta) \mathbb{E}(\|\epsilon\|)$ in accord with Theorem 5.1.1. In a first experiment we let $\tau(\epsilon) = 1.3 = \text{const}$. In this case, the numerical results suggest that the regularization parameter decreases too fast, i.e., $\frac{(\tau(\eta) \mathbb{E}(\|\epsilon\|))^2}{\alpha}$ does not converge to zero as in (5.8), see Figure 5.1. For comparison, in a second run we chose $\tau(\eta) = \sqrt{1 - \log(\eta^2 2\pi m^2 (\frac{\epsilon}{2})^m)}$ where m is the amount of data points. This way, $\tau(\eta) \mathbb{E}(\|\epsilon\|) \propto \rho_K(y, y^\eta)$. Now $\frac{(\tau(\eta) \mathbb{E}(\|\epsilon\|))^2}{\alpha}$ converges to zero as it should be the case.

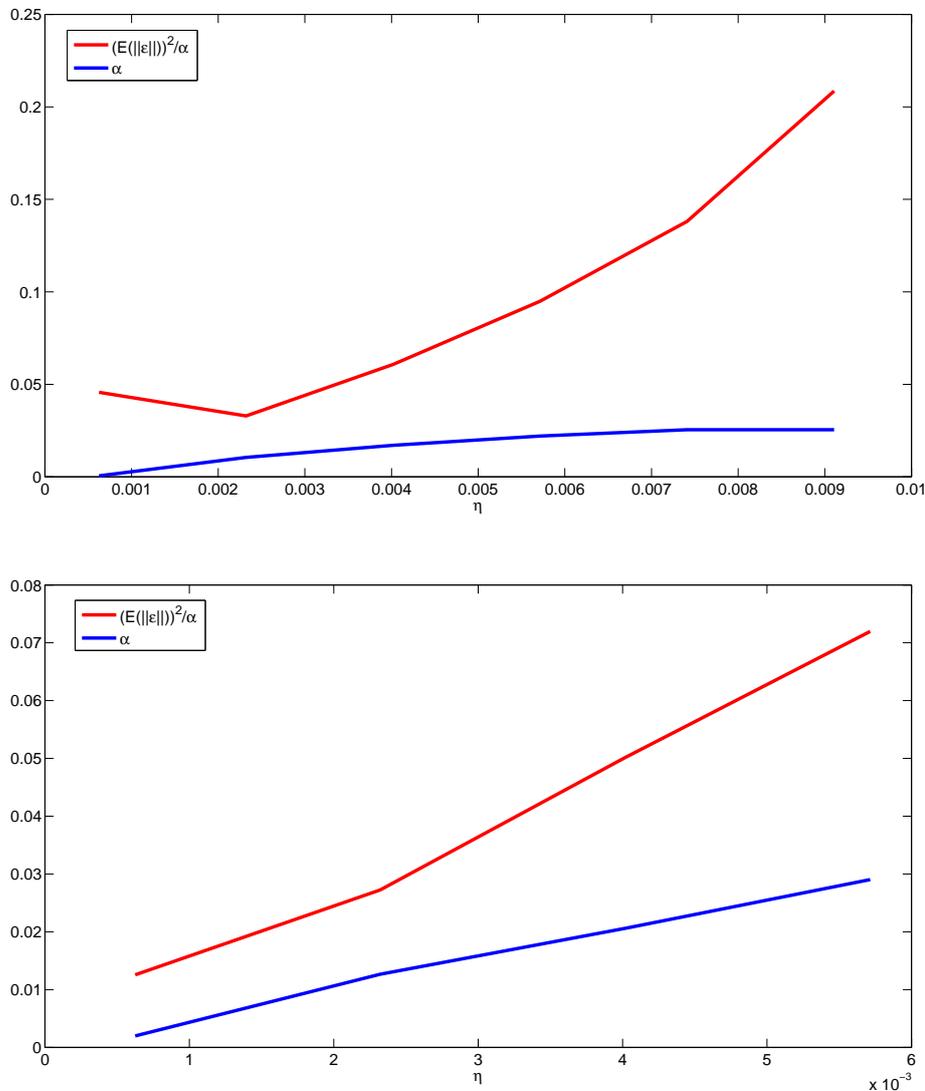


Figure 5.1: Convergence behavior for sparsity regularization with a nonlinear operator. Top: α and $\frac{(\tau(\eta)\mathbb{E}(\|\epsilon\|))^2}{\alpha}$ for constant $\tau(\eta) = 1.3$. The regularization parameter becomes too small too fast. Bottom: Same experiment with $\tau = \sqrt{1 - \log(\eta^2 2\pi m^2 (\frac{\epsilon}{2})^m)}$. Adjusting τ resolves the problem of α being too small.

Chapter 6

The method of the Approximate Inverse for Adaptive Optics

In this chapter we move from rather theoretical aspects of regularization theory to the design of a novel reconstruction algorithm for atmospheric tomography. Namely, we employ the method of the approximate inverse which was proposed by Louis and Maaß [136] for integral equations of the first kind. First we introduce this strategy for scalar functions and summarize important results. Then we apply the approach to the problem of atmospheric tomography. While at first we stick to the notation of the previous chapters, we will later use the notation of the Adaptive Optics community which we introduced in Chapter 2. A large part of this chapter is taken from the paper [137].

6.1 Method of the Approximate Inverse

6.1.1 Reconstruction of scalar functions

We present the theory of the approximate inverse for the standard setting employed in several of the previous chapters. The approximate inverse is a regularization scheme to solve linear inverse problems (1.1),

$$Ax = y \tag{6.1}$$

and an operator A between Hilbert spaces [138, 136]. An extension to Banach spaces is possible, see for example [139, 140].

Let us start in the classical setting, more precisely let $A : L_2(\Omega_X) \rightarrow L_2(\Omega_Y)$ be a linear operator where the subsets $\Omega_X \subset \mathbb{R}^n$ and $\Omega_Y \subset \mathbb{R}^m$ are bounded and assume that $x \in L_2(\Omega_X)$ and $y \in L_2(\Omega_Y)$ are scalar valued functions. In order to obtain a stable approximation of the function x from the data $y = Ax$, we consider a smoothed version

$$x^\gamma(t) := \langle x, e_t^\gamma \rangle_{L_2(\Omega_X)} \tag{6.2}$$

with a chosen *mollifier* $e_t^\gamma \in L_2(\Omega_X)$ with some $\gamma > 0$. The mollifier e_t^γ can be seen as an approximation to the delta distribution δ_t . A precise definition of a mollifier is given in the following, see [141].

Definition 6.1.1. For all $t \in \Omega_X$, $\gamma > 0$ let $e_t^\gamma \in L_2(\Omega_X)$ with

$$\int_{\Omega_X} e_t^\gamma(z) dz = 1.$$

Let further

$$x^\gamma(t) = \int_{\Omega_X} x(z) e_t^\gamma(z) dz$$

converge to x in $L_2(\Omega_X)$ as $\gamma \rightarrow 0$. Then e_t^γ is called a mollifier.

Now, instead of the original equation (6.1), one solves the auxiliary problem

$$A^* \psi_t^\gamma = e_t^\gamma \quad (6.3)$$

where A^* denotes as usual the adjoint of A . With this, the mollified version x^γ can be computed from the measured data via

$$x^\gamma(t) = \langle x, e_t^\gamma \rangle_{L_2(\Omega_X)} = \langle x, A^* \psi_t^\gamma \rangle_{L_2(\Omega_X)} = \langle Ax, \psi_t^\gamma \rangle_{L_2(\Omega_Y)} = \langle y, \psi_t^\gamma \rangle_{L_2(\Omega_Y)}.$$

The auxiliary problem (6.3) is only solvable if $e_t^\gamma \in \mathcal{R}(A^*)$. Otherwise, if at least $e_t^\gamma \in \mathcal{R}(A^*) \oplus \mathcal{R}(A^*)^\perp$, the function ψ_t^γ can be computed by minimizing the defect $\|A^* \psi_t^\gamma - e_t^\gamma\|^2$, which is equivalent to the solution of the normal equation

$$AA^* \psi_t^\gamma = Ae_t^\gamma. \quad (6.4)$$

In both cases (6.3) and (6.4), the functional $\langle y, \psi_t^\gamma \rangle_{L_2(\Omega_Y)}$ reconstructs $(\mathcal{P}_{\mathcal{N}(A)^\perp} x)^\gamma(t)$ with the orthogonal projection \mathcal{P} on $\mathcal{N}(A)^\perp$. We will show in Section 6.2 that in our application $e_x^\gamma \notin \mathcal{R}(A^*)$, i.e., we will determine the reconstruction kernels by solving the normal equation (6.4). Hence, from now on we will focus on the theory for this case.

Definition 6.1.2. Let e_t^γ be a mollifier according to Definition 6.1.1. The operator $S^\gamma : L_2(\Omega_Y) \rightarrow L_2(\Omega_X)$ with

$$S^\gamma y(t) := \langle y, \psi_t^\gamma \rangle_{L_2(\Omega_Y)}, \quad (6.5)$$

where ψ_t^γ solves (6.4), is called the *approximate inverse* of A to compute an approximation of x . The function ψ_t^γ is called *reconstruction kernel*.

Since the auxiliary problem (6.4) is independent of the data, the reconstruction kernels ψ_t^γ can be precomputed. However, the dependence on the reconstruction point $t \in \Omega_X$ requires the solution of a possibly different auxiliary problem for each of these points. Using suitable invariances of the operator A , this computational effort can be dramatically reduced, see e.g. [138, 141].

Theorem 6.1.1. Let the operators $T_1^t : L_2(\Omega_X) \rightarrow L_2(\Omega_X)$, $T_2^t : L_2(\Omega_Y) \rightarrow L_2(\Omega_Y)$ and $T_3^t : L_2(\Omega_Y) \rightarrow L_2(\Omega_Y)$ be linear and bounded for $t \in \Omega_X$, satisfying

$$AT_1^t = T_2^t A, \quad T_2^t AA^* = AA^* T_3^t. \quad (6.6)$$

Further assume that the mollifier e_t^γ is generated by T_1^t , i.e. $e_t^\gamma := T_1^t e^\gamma$ with $e^\gamma \in \mathcal{D}((A^*)^\dagger)$. If ψ^γ solves $AA^* \psi^\gamma = Ae^\gamma$, then the special reconstruction kernels ψ_t^γ are given by

$$\psi_x^\gamma = T_3^t \psi^\gamma.$$

Proof Using the intertwining properties (6.6), it holds

$$Ae_t^\gamma = AT_1^t e^\gamma = T_2^t A e^\gamma = T_2^t A A^* \psi^\gamma = A A^* T_3^t \psi^\gamma = A A^* \psi_t^\gamma.$$

□

Theorem 6.1.1 supposes a global invariance property on A , namely $AT_1^t x = T_2^t Ax$ for all $x \in L_2(\Omega_X)$. However, according to the proof, the statement of the theorem is also true if this invariance holds at least for the prescribed mollifier e^γ . This fact will be used in Section 6.2, since there is no global invariance in case of the atmospheric tomography operator.

Using such invariances, only one single auxiliary problem has to be solved, and the special reconstruction kernels ψ_t^γ are generated by this solution ψ^γ and the operator T_3^t , leading to the reconstruction

$$x^\gamma(t) = \langle y, T_3^t \psi^\gamma \rangle_{L_2(\Omega_Y)}. \quad (6.7)$$

Note that no artificial discretization of the unknown x is introduced. The reconstruction point only influences the translation operator T_1^t and can be chosen freely.

To close this section we would like to mention that Louis [138] showed, that in special cases, e.g. for quadratic operators, the application of the method of the approximate inverse to nonlinear problems is possible.

6.1.2 The approximate inverse and filter-based methods

The approximate inverse as defined in the previous section can be interpreted as evaluation of a linear functional of the data y . It is known, that this leads to regularization methods, see the theoretical result of Likht [142].

For the convergence analysis we make use of the connection between the method of the approximate inverse and filter based regularization methods as discussed in Chapter 3. Let therefore again $\{\sigma_n, u_n, v_n\}$ be the singular system of the linear compact operator A . The following theorems are due to Louis [138].

Theorem 6.1.2. *Let the regularization method R_α in (3.4) be given with a filter F_α . Then this method can be written as an approximate inverse with mollifier*

$$e_t^\alpha(s) = \sum_n F_\alpha(\sigma_n) v_n(s) v_n(t).$$

Theorem 6.1.3. *Let e^α be decomposed into*

$$e_t^\alpha(s) = \sum_{m,n} e_{m,n}^\alpha v_n(s) v_m(t).$$

Then the approximate inverse can be represented as regularization method based on a filter F_α if and only if the filter satisfies

$$e_{m,n}^\alpha = F_\alpha(\sigma_n)\delta_{mn}$$

with the Kronecker symbol δ_{mn} .

The first Theorem states that every filter method can, at least in theory, be expressed as an approximate inverse. Hence, from this point of view, the approximate inverse is more general than the filter based methods. On the other hand, under the condition of the second Theorem, one may find a filter such that the approximate inverse coincides with a filter-based method. If this is the case, regularization properties and convergence rates follow directly from Theorem 3.1.2 and Theorem 3.1.3, respectively, assuming that the conditions on the filter F_α are satisfied. Therefore, in this case also the theory of Section 3.3.1 can be carried over directly.

In another paper, Louis shows that both approximate inverse and filter-based regularization methods can be treated as subclasses of a larger class of regularization methods for which, under certain conditions similarly to the ones in Theorem 3.1.2 and Theorem 3.1.3, one can show regularization properties and optimal convergence rates under an a priori parameter choice.

6.1.3 Reconstruction of vector-valued functions

Previously, the regularization scheme was presented in the case of scalar-valued functions. However, the mathematical model of atmospheric tomography, see Section 2.6.2, is described by an operator A relating vector-valued functions. We will therefore and throughout the remaining part of the chapter switch to the notation common in astronomy and previously introduced in Chapter 2. The task is to solve the equation

$$A\Phi = \varphi$$

where $\Phi \in \bigotimes_{l=1}^L L_2(\Omega_l)$ comprises the layers of the atmosphere and the available data is $\varphi \in (L_2(\Omega_D))^G$. In particular, x from now on denotes a spatial coordinate rather than a function. The method of the approximate inverse has been extended to the reconstruction of vector fields already in the case of the three-dimensional Doppler transform [144]. In the following, we apply the method to the setting of atmospheric tomography by choosing appropriate mollifiers for the different layers. This is justified since the individual layers can be reconstructed independently of each other. Contrariwise, we will later show that it is not possible to formulate a mollifier for the whole product space given the available data.

Let δ_x denote the delta distribution with $x \in \mathbb{R}^2$. The first layer of the atmosphere, i.e. the first component of the vector $\Phi = (\Phi^{(1)}, \dots, \Phi^{(L)})^T$, is given by

$$\begin{aligned}\Phi^{(1)}(x) &= \langle \Phi^{(1)}, \delta_x \rangle_{L_2(\Omega_1)} \\ &= \left(\frac{1}{c_1} \langle \Phi^{(1)}, c_1 \delta_x \rangle_{L_2(\Omega_1)} + \sum_{l=2}^L \frac{1}{c_l} \langle \Phi^{(l)}, 0 \rangle_{L_2(\Omega_l)} \right) \\ &= \langle \Phi, \bar{\delta}_{x,1} \rangle_{\otimes_{l=1}^L L_2(\Omega_l)}.\end{aligned}$$

Here $\bar{\delta}_{x,1} := (c_1 \delta_x, 0, \dots, 0)^T$ denotes an L -dimensional vector, where the first component is $c_1 \delta_x$ and the remaining components correspond to the zero functions in $L_2(\Omega_l)$, $l = 2, \dots, L$. Replacing the delta-distribution δ_x by a scalar-valued mollifier $e_x^\gamma \in L_2(\Omega_1)$ fulfilling Definition 6.1.1 leads to the mollified version

$$\Phi_\gamma^{(1)}(x) := \langle \Phi, \bar{e}_{x,1}^\gamma \rangle_{\otimes_{l=1}^L L_2(\Omega_l)}$$

with $\bar{e}_{x,1}^\gamma := (c_1 e_x^\gamma, 0, \dots, 0)^T \in \otimes_{l=1}^L L_2(\Omega_l)$. Analogously, we obtain a mollified version $\Phi_\gamma^{(l)}$ of the l -th layer by

$$\Phi_\gamma^{(l)}(x) = \left\langle \Phi, \bar{e}_{x,l}^\gamma \right\rangle_{\otimes_{i=1}^L L_2(\Omega_i)} \quad \text{for } x \in \Omega_l \quad (6.8)$$

where

$$\bar{e}_{x,l}^\gamma := (c_l e_x^\gamma \delta_{i,l})_{i=1, \dots, L} \quad (6.9)$$

with the Kronecker symbol $\delta_{i,l}$. This notation abbreviates the fact that only the l -th component of $\bar{e}_{x,l}^\gamma$ is nonzero with a scalar-valued mollifier $e_x^\gamma \in L_2(\Omega_l)$. It can be verified in a straight forward calculation that $\bar{e}_{x,l}^\gamma$ is in fact a mollifier to approximate $\Phi^{(l)}$ in accordance to Definition 6.1.1.

Now, let $\Psi_l^\gamma \in L_2(\Omega_D)^G$ solve the normal equation

$$AA^* \Psi_l^\gamma = A \bar{e}_l^\gamma, \quad (6.10)$$

where $\bar{e}_l^\gamma := \bar{e}_{0,l}^\gamma \in \mathcal{R}(A^*) \oplus \mathcal{R}(A^*)^\perp$ is, without loss of generality, the mollifier centred at zero and let $T_1^x : \otimes_{l=1}^L L_2(\Omega_l) \rightarrow \otimes_{l=1}^L L_2(\Omega_l)$, $T_2^x : (L_2(\Omega_D))^G \rightarrow (L_2(\Omega_D))^G$ and $T_3^x : (L_2(\Omega_D))^G \rightarrow (L_2(\Omega_D))^G$ denote suitable invariance operators for A in accordance to Theorem 6.1.1. Then, the mollified version of the layer $\Phi^{(l)}$ can be computed from the measured data via

$$\Phi_\gamma^{(l)}(x) = \langle \varphi, T_3^x \Psi_l^\gamma \rangle_{(L_2(\Omega_D))^G} = \sum_{g=1}^G \langle \varphi_{\beta_g}, (T_3^x \Psi_l^\gamma)_g \rangle_{L_2(\Omega_D)}.$$

Remark. The mollifier proposed here is applied to each layer separately. It would also be possible to define a mollifier on the whole product space $\bigotimes_{l=1}^L L_2(\Omega_l)$. However, in order to use the method of the approximate inverse with this mollifier of the product space, a different kind of data would be required which is not available in practice. Namely, let $\tilde{e}_x^\gamma := (e_{x,1}^\gamma, e_{x,2}^\gamma, \dots, e_{x,L}^\gamma) \in \bigotimes_{l=1}^L L_2(\Omega_l)$ be a mollifier according to Definition 6.1.1 and assume we found $\Psi = (\Psi_g)_{g=1, \dots, G} \in L_2^G(\mathbb{R}^2)$ such that $A^*\Psi = \tilde{e}_x^\gamma$. Then it is (with the calculus of the adjoint of Theorem 6.2.1)

$$\begin{aligned} \Phi_\gamma(x) &= \left(\frac{1}{c_l} \langle \Phi^{(l)}, (\tilde{e}_x^\gamma)^{(l)} \rangle \right)_{l=1, \dots, L} = \left(\frac{1}{c_l} \langle \Phi^{(l)}, (A^*\Psi)^{(l)} \rangle \right)_{l=1, \dots, L} \\ &= \left(\sum_{g=1}^G \langle \mathcal{P}_{\Phi^{(l)}, g}, \Psi_g \rangle \right)_{l=1, \dots, L}, \end{aligned}$$

where $\mathcal{P}_{\Phi^{(l)}} = A((\delta_{i,l} \Phi^{(i)})_{i=1, \dots, L}) \in (L_2(\Omega_D))^G$ is the projection of the single layer $\Phi^{(l)}$ onto the telescope aperture in direction β_g . This data is not available as the wavefronts $\varphi_{\beta_g} = A\Phi$ can only be reconstructed from projections of all layers.

6.2 Application to Atmospheric Tomography

Using the mollifiers defined in the previous section we now proceed in applying the method of the approximate inverse to the problem of atmospheric tomography. That is, we need to solve an adjoint system analogously to (6.4) and define proper invariance operators in accordance with Theorem 6.1.1.

For each layer $\Phi^{(l)}$, the operator A only takes the respective set Ω_l into account, i.e., the part of the sky covered by the telescope. Thus, an intertwining property (6.6) holds only for reconstruction points $x \in \bigcap_{g=1}^G \Omega_D^{\mu_l, \beta_g}(h_l \beta_g)$, i.e. only on the part of the layer which is seen from all directions β_g , $g = 1, \dots, G$. To overcome this restriction, a continuation of A to the whole spatial domain \mathbb{R}^2 is presented in the following. For this new operator, an intertwining property can be found which holds for all reconstruction points $x \in \Omega_l$ if either only natural guide stars or only laser guide stars are considered. Despite the slightly changed setting, the searched-for layers will still be adequately reconstructed, i.e. the induced error is small, as discussed in Section 6.2.3. At the end of that section, we will briefly discuss an alternative approach to deal with the problem of finding invariances. Instead of extending the operator, one could simply introduce a cut-off function ensuring only the domains visible by the telescope are used. However, this leads to further restrictions on the invariances and does not provide numerical benefits.

6.2.1 Continuation of the operator A

Each layer $\Phi^{(l)} \in L_2(\Omega_l)$ can be extended to \mathbb{R}^2 via

$$\tilde{\Phi}^{(l)}(\varrho) := \begin{cases} \Phi^{(l)}, & \varrho \in \Omega_l \\ 0, & \varrho \notin \Omega_l \end{cases}. \quad (6.11)$$

We now consider the mapping

$$\begin{aligned} \mathcal{A} : \bigotimes_{l=1}^L L_2(\mathbb{R}^2) &\longrightarrow (L_2(\mathbb{R}^2))^G \\ \mathcal{A}\tilde{\Phi}(r) &:= \left(\sum_{l=1}^L \tilde{\Phi}^{(l)}(\mu_{l,\beta_g} r + h_l \beta_g) \right)_{g=1,\dots,G}. \end{aligned} \quad (6.12)$$

This new operator considers the whole spatial domain \mathbb{R}^2 instead of only the subsets Ω_l , $l = 1, \dots, L$. With the correlation of $\tilde{\Phi}$ and Φ and $r \in \Omega_D$, it is for $g \in \{1, \dots, G\}$

$$\begin{aligned} (\mathcal{A}\tilde{\Phi})_g(r) &= \sum_{l=1}^L \tilde{\Phi}^{(l)}(\mu_{l,\beta_g} r + \beta_g h_l) \\ &= \sum_{l=1}^L \Phi^{(l)}(\mu_{l,\beta_g} r + \beta_g h_l) = (A\Phi)_g(r) = \varphi_{\beta_g}(r). \end{aligned}$$

Hence, on the telescope pupil, $\mathcal{A}\tilde{\Phi}$ corresponds to the measured data. However, the extended operator \mathcal{A} leads to artificial, unknown data for $r \notin \Omega_D$. Denote

$$\tilde{\Omega} := \bigcup_{g=1}^G \{r \in \mathbb{R}^2 \setminus \Omega_D : \exists l \text{ with } \mu_{l,\beta_g} r + \beta_g h_l \in \Omega_l\}. \quad (6.13)$$

This set $\tilde{\Omega}$ comprises all points r outside the telescope aperture which are still so close that $\mu_{l,\beta_g} r + \beta_g h_l \in \Omega_l$ for at least one height h_l and one unit vector β_g , see Section 6.2.3 for numerical examples. For $r \in \mathbb{R}^2 \setminus (\tilde{\Omega} \cup \Omega_D)$, it holds

$$\mathcal{A}\tilde{\Phi}^{(l)}(r) = 0,$$

and for $r \in \tilde{\Omega}$, it is

$$\mathcal{A}\tilde{\Phi}^{(l)}(r) = \tilde{\varphi}(r)$$

with an unknown function $\tilde{\varphi} \neq 0$. However, the effect of this deviation on the reconstruction of $\Phi^{(l)}$ is very weak, as discussed in Section 6.2.3.

Analogously to (2.19), the inner product on $\bigotimes_{l=1}^L L_2(\mathbb{R}^2)$ is given by

$$\langle \Phi, \Psi \rangle_{\bigotimes_{l=1}^L L_2(\mathbb{R}^2)} := \sum_{l=1}^L \frac{1}{c_l} \langle \Phi^{(l)}, \Psi^{(l)} \rangle_{L_2(\mathbb{R}^2)}.$$

In order to apply the method of the approximate inverse, we first have to compute the adjoint \mathcal{A}^* of \mathcal{A} . A physical interpretation is given afterwards.

Theorem 6.2.1. *The adjoint of \mathcal{A} is given by*

$$\begin{aligned} \mathcal{A}^* : (L_2(\mathbb{R}^2))^G &\longrightarrow \bigotimes_{l=1}^L L_2(\mathbb{R}^2) \\ (\mathcal{A}^* \varphi)^{(l)}(\varrho) &= c_l \sum_{g=1}^G \mu_{l,\beta_g}^{-2} \varphi_{\beta_g} \left(\mu_{l,\beta_g}^{-1} (\varrho - \beta_g h_l) \right), \quad l = 1, \dots, L. \end{aligned} \quad (6.14)$$

Proof Using the definition of \mathcal{A} , we obtain

$$\begin{aligned} \langle \mathcal{A}\Phi, \varphi \rangle_{(L_2(\mathbb{R}^2))^G} &= \sum_{g=1}^G \langle (\mathcal{A}\Phi)_g, \varphi_{\beta_g} \rangle_{L_2(\mathbb{R}^2)} \\ &= \sum_{g=1}^G \int_{\mathbb{R}^2} \sum_{l=1}^L \Phi^{(l)}(\mu_{l,\beta_g} r + h_l \beta_g) \varphi_{\beta_g}(r) \, dr. \end{aligned}$$

With the substitution $\varrho := \mu_{l,\beta_g} r + h_l \beta_g$ it holds

$$\begin{aligned} \langle \mathcal{A}\Phi, \varphi \rangle_{(L_2(\mathbb{R}^2))^G} &= \sum_{g=1}^G \sum_{l=1}^L \int_{\mathbb{R}^2} \mu_{l,\beta_g}^{-2} \Phi^{(l)}(\varrho) \varphi_{\beta_g}(\mu_{l,\beta_g}^{-1} (\varrho - h_l \beta_g)) \, d\varrho \\ &= \sum_{l=1}^L \int_{\mathbb{R}^2} \Phi^{(l)}(\varrho) \sum_{g=1}^G \mu_{l,\beta_g}^{-2} \varphi_{\beta_g}(\mu_{l,\beta_g}^{-1} (\varrho - h_l \beta_g)) \, d\varrho \\ &= \sum_{l=1}^L \frac{1}{c_l} \left\langle \Phi^{(l)}, c_l \sum_{g=1}^G \mu_{l,\beta_g}^{-2} \varphi_{\beta_g}(\mu_{l,\beta_g}^{-1} (\cdot - h_l \beta_g)) \right\rangle_{L_2(\mathbb{R}^2)} \\ &= \langle \Phi, \mathcal{A}^* \varphi \rangle_{\bigotimes_{l=1}^L L_2(\mathbb{R}^2)}. \end{aligned}$$

□

Recall that the operator \mathcal{A} sums up the wavefronts of all layers for each direction of a guide star. The adjoint takes the wavefronts φ and projects them onto the layers in the atmosphere, in each layer summing up the contribution of each guide star. Additionally, each layer is weighted such that the “energy” is preserved.

In order to solve the normal equation (6.10) and define proper intertwining operators, we will use the following result on the composition of \mathcal{A} and \mathcal{A}^* .

Theorem 6.2.2. *It holds*

$$\mathcal{A}\mathcal{A}^* = B + \lambda I$$

with the identity I , $\lambda := (\lambda_g)_{g=1,\dots,G}$ where $\lambda_g := \sum_{l=1}^L c_l \mu_{l,\beta_g}^{-2}$ and $B\Psi := (B_g\Psi)_{g=1,\dots,G}$,

$$B_g\Psi(r) := \sum_{l=1}^L c_l \sum_{i=1, i \neq g}^G \mu_{l,\beta_i}^{-2} \Psi_i \left(\frac{\mu_{l,\beta_g}}{\mu_{l,\beta_i}} r + \mu_{l,\beta_i}^{-1} h_l(\beta_g - \beta_i) \right). \quad (6.15)$$

Proof. With the definition of the operators \mathcal{A} (6.12) and \mathcal{A}^* (6.14), we obtain for $g = 1, \dots, G$

$$\begin{aligned} (\mathcal{A}\mathcal{A}^*\Psi)_g(r) &= \sum_{l=1}^L (\mathcal{A}^*\Psi)^{(l)}(\mu_{l,\beta_g} r + h_l \beta_g) \\ &= \sum_{l=1}^L c_l \sum_{i=1}^G \mu_{l,\beta_i}^{-2} \Psi_i \left(\mu_{l,\beta_i}^{-1} ((\mu_{l,\beta_g} r + h_l \beta_g) - h_l \beta_i) \right) \\ &= \sum_{l=1}^L c_l \left(\sum_{i=1, i \neq g}^G \mu_{l,\beta_i}^{-2} \Psi_i \left(\frac{\mu_{l,\beta_g}}{\mu_{l,\beta_i}} r + \mu_{l,\beta_i}^{-1} h_l(\beta_g - \beta_i) \right) + \mu_{l,\beta_g}^{-2} \Psi_g(r) \right) \\ &= \sum_{l=1}^L c_l \sum_{i=1, i \neq g}^G \mu_{l,\beta_i}^{-2} \Psi_i \left(\frac{\mu_{l,\beta_g}}{\mu_{l,\beta_i}} r + \mu_{l,\beta_i}^{-1} h_l(\beta_g - \beta_i) \right) + \sum_{l=1}^L c_l \mu_{l,\beta_g}^{-2} \Psi_g(r) \\ &= B_g\Psi(r) + \lambda_g \Psi_g \end{aligned}$$

with B and λ defined above. \square

When a single reconstructed wavefront φ_g , $g \in 1, \dots, G$ is projected onto the layers and then back down onto the wavefront sensors, one obtains the (rescaled) φ_g again. However, additionally φ_g is influenced by a difficult structure of cross-projections of the other reconstructed wavefronts φ_i , $i \neq g$, collected in the operator B .

Using the representation of Theorem 6.2.2, we obtain the following invariances.

Theorem 6.2.3. *Let $\bar{e}_l^\gamma(\varrho) = (c_l e^\gamma(\varrho) \delta_{i,l})_{i=1,\dots,L}$ be the mollifier from (6.9) centred at zero, with $l \in \{1, \dots, L\}$ arbitrary but fixed. Define for $j = 1, \dots, L$ the linear operators*

$$\begin{aligned} T_{1,l}^x &: \bigotimes_{j=1}^L L_2(\mathbb{R}^2) \longrightarrow \bigotimes_{j=1}^L L_2(\mathbb{R}^2), \quad (T_{1,l}^x \bar{e}_l^\gamma)^{(j)}(\varrho) := (\bar{e}_l^\gamma)^{(j)}(\varrho - x) \delta_{j,l} \quad \text{and} \\ T_{2,l}^x &: (L_2(\mathbb{R}^2))^G \longrightarrow (L_2(\mathbb{R}^2))^G, \quad (T_{2,l}^x \Psi)_g(r) := \Psi_g \left(r - \frac{x}{\mu_{l,\beta_g}} \right) \end{aligned}$$

for $g = 1, \dots, G$. If for all $l = 1, \dots, L$, $\mu_{l,\beta_g} = \mu_l$ independent of $g \in \{1, \dots, G\}$, then it holds

$$\mathcal{A}T_{1,l}^x \bar{e}_l^\gamma = T_{2,l}^x \mathcal{A} \bar{e}_l^\gamma, \quad T_{2,l}^x \mathcal{A}\mathcal{A}^* = \mathcal{A}\mathcal{A}^* T_{2,l}^x. \quad (6.16)$$

Proof Using the definition of the operators $T_{1,l}^x$ and $T_{2,l}^x$, we obtain for $g = 1, \dots, G$

$$\begin{aligned}
(\mathcal{A}T_1^x \bar{e}_l^\gamma)_g(r) &= \sum_{j=1}^L (T_1^x \bar{e}_l^\gamma)^{(j)} (\mu_{j,\beta_g} r + h_j \beta_g) \\
&= c_l e^\gamma (\mu_{l,\beta_g} r + h_l \beta_g - x) \\
&= c_l e^\gamma \left(\mu_{l,\beta_g} \left(r - \frac{x}{\mu_{l,\beta_g}} \right) + h_l \beta_g \right) \\
&= \sum_{j=1}^L (\bar{e}_l^\gamma)^{(j)} \left(\mu_{j,\beta_g} \left(r - \frac{x}{\mu_{l,\beta_g}} \right) + h_j \beta_g \right) \\
&= (\mathcal{A} \bar{e}_l^\gamma)_g \left(r - \frac{x}{\mu_{l,\beta_g}} \right) \\
&= (T_{2,l}^x \mathcal{A} \bar{e}_l^\gamma)_g(r).
\end{aligned}$$

For the second invariance, we have the following situation.

$$\begin{aligned}
(T_{2,l}^x \mathcal{A} \mathcal{A}^* \Psi)_g(r) &= (\mathcal{A} \mathcal{A}^* \Psi)_g \left(r - \frac{x}{\mu_{l,\beta_g}} \right) \\
&= \sum_{j=1}^L c_l \sum_{i=1, i \neq g}^G \mu_{j,\beta_i}^{-2} \Psi_i \left(\frac{\mu_{j,\beta_g}}{\mu_{j,\beta_i}} \left(r - \frac{x}{\mu_{l,\beta_g}} \right) + \mu_{j,\beta_i}^{-1} h_j (\beta_g - \beta_i) \right) \\
&\quad + \lambda_g \Psi_g \left(r - \frac{x}{\mu_{l,\beta_g}} \right). \tag{6.17}
\end{aligned}$$

On the other hand, we obtain

$$\begin{aligned}
(\mathcal{A} \mathcal{A}^* T_{2,l}^x \Psi)_g(r) &= \sum_{j=1}^L c_l \sum_{i=1, i \neq g}^G \mu_{j,\beta_i}^{-2} \Psi_i \left(\frac{\mu_{j,\beta_g}}{\mu_{j,\beta_i}} \left(r - \frac{\mu_{j,\beta_i} x}{\mu_{j,\beta_g} \mu_{l,\beta_i}} \right) + \mu_{j,\beta_i}^{-1} h_j (\beta_g - \beta_i) \right) \\
&\quad + \lambda_g \Psi_g \left(r - \frac{x}{\mu_{l,\beta_g}} \right). \tag{6.18}
\end{aligned}$$

Hence, with $\mu_{j,\beta_i} = \mu_{j,\beta_g}$ and $\mu_{l,\beta_i} = \mu_{l,\beta_g}$, it is $T_{2,l}^x \mathcal{A} \mathcal{A}^* = \mathcal{A} \mathcal{A}^* T_{2,l}^x$. \square

The theorem, in principle, states that it does not matter if one first translates the mollifier on the layer by x and then projects it onto the wavefront sensors or first projects it onto the wavefronts sensors and then translates the measurements by $\frac{x}{\mu_{l,\beta_g}}$. However, it reveals an issue when laser guide stars and natural guide stars are combined. Comparing

the action of $T_{2,l}^x$ in (6.17) and (6.18), respectively, one can spot a difference in the “cross-projections”, i.e., when projecting a component Ψ_i on a layer and then projecting it back down in direction β_g , $i \neq g$. The difference lies in the scaling factor of the shift, $r - \frac{x}{\mu_{l,\beta_g}}$ compared to $r - \frac{\mu_{j,\beta_i}x}{\mu_{j,\beta_g}\mu_{l,\beta_i}}$. In case we only consider either natural guide stars or laser guide stars, it is $\mu_{j,\beta_i} = \mu_{j,\beta_g}$ for all i, g and hence the invariance holds true. If both types of guide stars are mixed, however, this factor adds a systematic error in the invariances. Therefore we cannot expect the approximate inverse to perform as well for mixed guide stars as for a single type of guide stars. Numerically, however, the proposed invariances still lead to good results, see Section 6.3.2. The problem may be resolved when tip/tilt indetermination of laser guide stars is part of the model. In this case, the natural guide stars can be used only to correct the reconstructions of the layers performed with laser guide stars. This so called separate tip/tilt reconstruction technique or split tomography is described in [68, 52]. Using this, one may use a setup with only laser guide stars and reconstruct a turbulence profile from this. The tip/tilt information is provided from additional natural guide star measurements via a separate approach. Since we did not include tip/tilt indetermination in this paper, we do not know if additional problems might appear. However, from the current situation this should avoid the problem of mixing the guide star types for the method of the approximate inverse. If we only consider laser guide stars, the following set of invariance operators leads to a faster reconstruction algorithm.

Let \bar{e}_l^γ be again the mollifier from (6.9) centered at zero, with $l \in \{1, \dots, L\}$ arbitrary but fixed. If for all $l = 1, \dots, L$, $\mu_{l,\beta_g} = \mu_l$ independent of $g \in \{1, \dots, G\}$ the linear operators T_{1,μ_l}^x, T_2^x are defined via

$$\begin{aligned} T_{1,\mu_l}^x : \bigotimes_{j=1}^L L_2(\mathbb{R}^2) &\longrightarrow \bigotimes_{j=1}^L L_2(\mathbb{R}^2), & (T_{1,\mu_l}^x \bar{e}_l^\gamma)^{(j)}(\varrho) &= (\bar{e}_l^\gamma)^{(j)}(\varrho - \mu_l x) \delta_{j,l}, \\ T_2^x : L_2(\mathbb{R}^2)^G &\longrightarrow L_2(\mathbb{R}^2)^G, & (T_2^x \Psi)_g(r) &:= \Psi_g(r - x), \end{aligned}$$

then it holds

$$\mathcal{A} T_{1,\mu_l}^x \bar{e}_l^\gamma = T_2^x \mathcal{A} \bar{e}_l^\gamma, \quad T_2^x \mathcal{A} \mathcal{A}^* = \mathcal{A} \mathcal{A}^* T_2^x.$$

Thus, it is

$$\bar{e}_{x,l}^\gamma(\varrho) = T_{1,\mu_l}^x \bar{e}_l^\gamma(\varrho) = (c_l e^\gamma(\varrho_l - \mu_l x) \delta_{i,l})_{i=1\dots L},$$

leading to

$$\langle \Phi, \bar{e}_{x,l}^\gamma \rangle_{\bigotimes_{i=1}^L L_2(\mathbb{R}^2)} = \langle \Phi^{(l)}, e_l^\gamma(\cdot - \mu_l x) \rangle_{L_2(\mathbb{R}^2)} = \Phi_\gamma^{(l)}(\mu_l x).$$

Hence, the mollified version $\langle \Phi, \bar{e}_{x,l}^\gamma \rangle_{\bigotimes_{i=1}^L L_2(\mathbb{R}^2)}$ approximates a scaled version of the layer $\Phi^{(l)}$, instead of the layer itself. The corresponding reconstruction kernels represent then unscaled translated versions of Ψ_l^γ ,

$$\Psi_{x,l}^\gamma(r) = T_2^x \Psi_l^\gamma(r) = \Psi_l^\gamma(r - x).$$

Thus, the functional $\langle \varphi, \Psi_{x,l}^\gamma \rangle_{L_2(\mathbb{R}^2)^G}$ can be computed without interpolation of the data vector. Instead, this interpolation step is transferred to an interpolation on the layer. For a discussion on the numerical effort we refer to Section 6.3.3. The image quality obtained with these invariances in numerical simulations was of the same quality as with the previously introduced invariances as long as the guide star types were not mixed. However, for the mixed guide star setting this set of invariance operators gave significantly worse results.

6.2.2 Calculation of the reconstruction kernel

To compute the reconstruction kernels, we have to solve the auxiliary problem

$$\mathcal{A}^* \Psi_{x,l}^\gamma = \bar{e}_{x,l}^\gamma \quad (6.19)$$

for each point x we wish to reconstruct. The following Lemma shows that it is not possible to find an exact solution for a mollifier as defined in (6.9), i.e., an individual mollifier for each layer.

Lemma 6.2.4. *Let $\bar{e}_{x,l}^\gamma$ as in (6.9). Then $\bar{e}_{x,l}^\gamma \notin \mathcal{R}(\mathcal{A}^*)$.*

Proof Assume $\bar{e}_{x,l}^\gamma \in \mathcal{R}(\mathcal{A}^*)$. Then $\bar{e}_{x,l}^\gamma \in \overline{\mathcal{R}(\mathcal{A}^*)} = \mathcal{N}(\mathcal{A})^\perp$, i.e., $\langle \bar{e}_{x,l}^\gamma, \Phi \rangle_{\otimes_{l=1}^L L_2(\mathbb{R}^2)} = 0$ for all $\Phi \in \mathcal{N}(\mathcal{A})$. Let $\tilde{\Phi} := (C_l \chi_{\Omega_l})_{l=1,\dots,L}$ be constant on each layer with $C_l \in \mathbb{R}$, $\sum_{l=1}^L C_l = 0$. Then $\mathcal{A}\tilde{\Phi} = 0$, i.e., $\tilde{\Phi} \in \mathcal{N}(\mathcal{A})$ but

$$\langle \bar{e}_{x,l}^\gamma, \tilde{\Phi} \rangle_{\otimes_{l=1}^L L_2(\mathbb{R}^2)} = \langle e_x^\gamma, \tilde{\Phi}^{(l)} \rangle_{L_2(\Omega_l)} = C_l \langle e_x^\gamma, 1 \rangle_{L_2(\Omega_l)} \neq 0$$

because of Definition 6.1.1. Hence $\bar{e}_{x,l}^\gamma \notin \mathcal{N}(\mathcal{A})^\perp = \overline{\mathcal{R}(\mathcal{A}^*)}$, contradicting the original assumption. \square

Since $\bar{e}_{x,l}^\gamma \notin \mathcal{R}(\mathcal{A}^*)$, we solve the normal equation

$$\mathcal{A}\mathcal{A}^* \Psi_{x,l}^\gamma = \mathcal{A}\bar{e}_{x,l}^\gamma \quad (6.20)$$

instead of (6.19). In order to circumvent the dependency of this equation on the point x , we use the invariance properties (6.16) of \mathcal{A} and $\mathcal{A}\mathcal{A}^*$. Hence the mollifiers $\bar{e}_{x,l}^\gamma$ are generated by

$$\bar{e}_{x,l}^\gamma = T_{1,l}^x \bar{e}_l^\gamma \quad (6.21)$$

with $\bar{e}_l^\gamma := (c_l e^\gamma \delta_{i,l})_{i=1\dots L}$ analogously to (6.9). Again, $e^\gamma \in L_2(\mathbb{R}^2)$ is a prescribed mollifier and $\delta_{i,l}$ represents the Kronecker symbol. According to Theorem 6.1.1 and Theorem 6.2.3, the corresponding special reconstruction kernels $\Psi_{x,l}^\gamma$ are given by

$$\left(\Psi_{x,l}^\gamma \right)_g (r) = \left(T_{2,l}^x \Psi_l^\gamma \right)_g (r) = \left(\Psi_l^\gamma \right)_g \left(r - \frac{x}{\mu_l, \beta_g} \right),$$

where Ψ_l^γ solves

$$\mathcal{A}\mathcal{A}^*\Psi_l^\gamma = \mathcal{A}\bar{e}_l^\gamma \quad (6.22)$$

i.e., Ψ_l^γ minimizes $\|\mathcal{A}^*\Psi_l^\gamma - \bar{e}_l^\gamma\|^2$. It remains to solve this problem once for each layer $l = 1, \dots, L$.

A common method for solving equations of type (6.22) are iterative methods, in particular gradient methods. Starting from an initial guess, one iterates

$$\Psi_l^\gamma \leftarrow \Psi_l^\gamma - \tau(\mathcal{A}\mathcal{A}^*\Psi_l^\gamma - \mathcal{A}\bar{e}_l^\gamma) \quad (6.23)$$

until a stopping criterion is met, e.g., a maximum number of iterations. The steplength τ is used to ensure convergence and has to be chosen appropriately, see below. With $\mathcal{A}\mathcal{A}^* = B + \lambda I$ from Theorem 6.2.2, (6.23) can be written

$$\Psi_l^\gamma \leftarrow (1 - \tau\lambda)\Psi_l^\gamma - \tau(B\Psi_l^\gamma - \mathcal{A}\bar{e}_l^\gamma).$$

This procedure updates all G components of Ψ_l^γ at once. However, the numerical results we obtained with this approach were of low quality, i.e., the obtained Strehl ratios were close to zero. Alternatively, one may update each component separately and immediately use the updated information in the next step. This approach was introduced by Kaczmarz for the solution of linear systems of equations. It can be extended to more general settings, in particular to Adaptive Optics, see, e.g., [41, 60]. In each step, we now minimize $\|\mathcal{A}^*\Psi_l^\gamma - \bar{e}_l^\gamma\|^2$ in direction of component $(\Psi_l^\gamma)_g$ in a cyclical way. The gradient w.r.t. to the g -th component, i.e., the Fréchet derivative in the direction of the g -th component, is given by

$$\nabla_g \|\mathcal{A}^*\Psi_l^\gamma - \bar{e}_l^\gamma\|^2 = (\mathcal{A}\mathcal{A}^*\Psi_l^\gamma - \mathcal{A}\bar{e}_l^\gamma)_g = \lambda_g(\Psi_l^\gamma)_g + B_g\Psi_l^\gamma - (\mathcal{A}\bar{e}_l^\gamma)_g.$$

With this as descent direction we construct an iterative Kaczmarz-type algorithm. Let $\bar{k} := \text{mod}(k, G) + 1$. Then, again starting from an initial guess which in practice we simply chose zero, we update

$$(\Psi_l^\gamma)_{\bar{k}} \leftarrow (1 - \tau_k\lambda_{\bar{k}})(\Psi_l^\gamma)_{\bar{k}} - \tau_k(B_{\bar{k}}\Psi_l^\gamma - (\mathcal{A}\bar{e}_l^\gamma)_{\bar{k}}) \quad (6.24)$$

for $k = 1, 2, \dots$ until a stopping criterion is met. In our numerical experiments we let

$$\tau_k = \frac{\|\lambda_{\bar{k}}(\Psi_l^\gamma)_{\bar{k}} + B_{\bar{k}}\Psi_l^\gamma - (\mathcal{A}\bar{e}_l^\gamma)_{\bar{k}}\|_{(L^2(\mathbb{R}^2))}}{\|\mathcal{A}^*(\lambda_{\bar{k}}(\Psi_l^\gamma)_{\bar{k}} + B_{\bar{k}}\Psi_l^\gamma - (\mathcal{A}\bar{e}_l^\gamma)_{\bar{k}})\|_{\otimes_{l=1}^L L_2(\mathbb{R}^2)}}. \quad (6.25)$$

This choice of τ_k minimizes the residual in the search direction, c.f. [145],

$$\tau_k = \min_{\tau} \|\mathcal{A}^*(\Psi_l^\gamma + \tau(\lambda_g(\Psi_l^\gamma)_g + B_g\Psi_l^\gamma - (\mathcal{A}\bar{e}_l^\gamma)_g)) - \bar{e}_l^\gamma\|_{\otimes_{l=1}^L L_2(\mathbb{R}^2)}^2.$$

With the calculated Ψ_l^γ we can then reconstruct the layer l in any point $x \in \Omega_l$ via

$$\Phi_\gamma^{(l)}(x) = \langle \mathcal{A}\Phi, T_{2,l}^x \Psi_l^\gamma \rangle_{(L_2(\mathbb{R}^2))^G}. \quad (6.26)$$

6.2.3 The extension error

Our main goal is the reconstruction of the layers $\Phi^{(l)}$, $l = 1 \dots L$, from the measured data $A\Phi = \varphi$. Using the extension $\tilde{\Phi}$ of Φ to the whole spatial domain \mathbb{R}^2 , (6.11), the relation between $\mathcal{A}\tilde{\Phi}$ and $A\Phi$ is given by

$$\mathcal{A}\tilde{\Phi}(r) = \begin{cases} A\Phi(r) & r \in \Omega_D \\ 0 & r \in \mathbb{R}^2 \setminus (\Omega_D \cup \tilde{\Omega}) \\ \tilde{\varphi}(r) & r \in \tilde{\Omega} \end{cases}$$

with an unknown function $\tilde{\varphi} \neq 0$ and $\tilde{\Omega}$ defined in (6.13). Now, for $x \in \Omega_l$, it is

$$\begin{aligned} \Phi_\gamma^{(l)}(x) &= \tilde{\Phi}_\gamma^{(l)}(x) = \langle \mathcal{A}\tilde{\Phi}, T_{2,l}^x \Psi_l^\gamma \rangle_{(L_2(\mathbb{R}^2))^G} \\ &= \langle A\Phi, T_{2,l}^x \Psi_l^\gamma \rangle_{(L_2(\Omega_D))^G} + \langle \tilde{\varphi}, T_{2,l}^x \Psi_l^\gamma \rangle_{L_2(\tilde{\Omega})} \\ &= \langle A\Phi, T_{2,l}^x \Psi_l^\gamma \rangle_{(L_2(\Omega_D))^G} + \varepsilon. \end{aligned} \quad (6.27)$$

The angles β_g cover just a small range, so the domain $\tilde{\Omega}$ will be relatively small compared to Ω_D . Besides, with the distance of r to Ω_D getting larger, less layers give a contribution to $\sum_{l=1}^L \Phi^{(l)}(\mu_l r + h_l \beta_g)$, i.e. the value of $\tilde{\varphi}(r)$ is getting smaller. In order to verify this numerically, we conduct the following experiment using the guide star setup corresponding to the one from Section 6.3.2 with 6 natural guide stars.

For the 42m diameter telescope used in the simulations, $\tilde{\Omega}$ corresponds to a ring of width 4m around the telescope mirror, see the left image in Figure 6.1. There, the inner circle corresponds to the mirror of the telescope Ω_D . We define $\bar{\Phi}(\varrho) := (\chi_{\Omega_l}(\varrho))_{l=1,\dots,L}$ where χ_{Ω_l} is the characteristic function of Ω_l , i.e., has the value one for $\varrho \in \Omega_l$ and zero otherwise. From this we calculate $\varphi := (\mathcal{A}\bar{\Phi}) \cdot \chi_{\Omega_D}$ and $\tilde{\varphi} := (\mathcal{A}\bar{\Phi}) \cdot (1 - \chi_{\Omega_D})$. Hence, φ corresponds to the available data and $\tilde{\varphi}$ to the artificial data. In the right picture of Figure 6.1 we show that the value of $\tilde{\varphi}$ indeed becomes smaller the further away the position from the mirror is. From both data sets, we can numerically evaluate the ratio R of the integrals over φ and $\tilde{\varphi}$,

$$R = \frac{\int_{\tilde{\Omega}} \tilde{\varphi}(r) dr}{\int_{\Omega_D} \varphi(r) dr} = 0.01,$$

i.e., the ‘‘magnitude’’ of $\tilde{\varphi}$ is about one percent of that of φ . Additionally, the reconstruction kernels Ψ_l^γ decay away from a few spikes situated around the center of the telescope pupil (compare Figure 6.2). Hence both the contribution of the reconstruction kernels and of $\tilde{\varphi}$ to $\varepsilon = \langle \tilde{\varphi}, T_{2,l}^x \Psi_l^\gamma \rangle_{L_2(\tilde{\Omega})}$ is expected to be small. So we approximate

$$\Phi_\gamma^{(l)}(x) = \langle A\Phi, T_{2,l}^x \Psi_l^\gamma \rangle_{(L_2(\Omega_D))^G}. \quad (6.28)$$

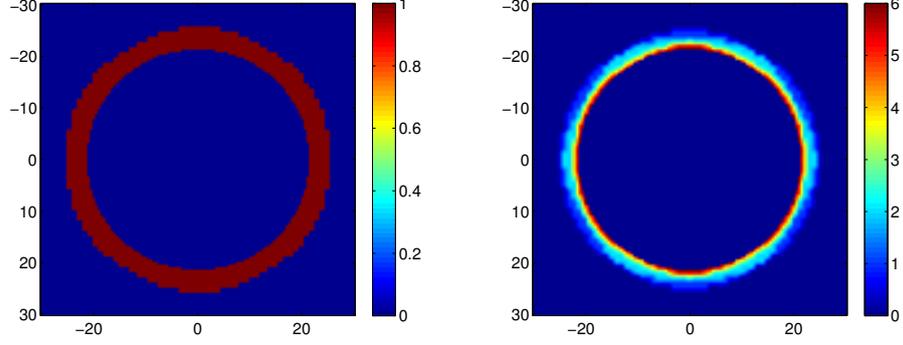


Figure 6.1: $\tilde{\Omega}$ for a telescope mirror of 42m diameter, 6 natural guide stars in a regular hexagon, each NGS 5 arcmin from zenith. On the left side, the domain qualitative shape of $\tilde{\Omega}$ is shown. On the right side, we give a more quantitative image showing the “magnitude” of the artificially created data. This figure suggests that the contribution of $\tilde{\Omega}$ can be neglected making only a small approximation error.

As mentioned in the introduction of this section, one could also, instead of extending the layer as in (6.12), introduce a rigorous cut-off of the operator, i.e., define

$$\begin{aligned} \tilde{\mathcal{A}} &: \bigotimes_{l=1}^L L_2(\mathbb{R}^2) \longrightarrow (L_2(\mathbb{R}^2))^G \\ \tilde{\mathcal{A}}\Phi(r) &:= \left(\chi_{\Omega_D}(r) \sum_{l=1}^L \Phi^{(l)}(\mu_{l,\beta_g} r + h_l \alpha_g) \right)_{g=1,\dots,G}. \end{aligned} \quad (6.29)$$

The characteristic function χ_{Ω_D} then appears in the adjoint operator as well as in the operator B when translating Theorem 6.2.1 and Theorem 6.15, respectively. With this operator, however the invariances of Theorem 6.2.3 only hold for $r \in \Omega_D(\cdot) \cap \Omega_D^{x,\mu}$ where $\Omega_D^{x,\mu} = \{r \in \mathbb{R}^2 : r - \frac{x}{\mu_{l,\beta_g}} \in \Omega_D\}$. Otherwise, an additional error is introduced. One can show that, with $T_{1,l}^x$ and $T_{2,l}^x$ from Theorem 6.2.3, it holds

$$\tilde{\mathcal{A}}T_{1,l}^x \bar{e}_l^\gamma = T_{2,l}^x \tilde{\mathcal{A}} \bar{e}_l^\gamma + \varepsilon_1, \quad T_{2,l}^x \tilde{\mathcal{A}} \tilde{\mathcal{A}}^* = \tilde{\mathcal{A}} \tilde{\mathcal{A}}^* T_{2,l}^x + \varepsilon_2$$

with $\varepsilon_1, \varepsilon_2 \in L_2(\Omega_D)$. Consequently it is, analogously to (6.27)

$$\begin{aligned} \Phi_\gamma^{(l)}(x) &\approx \langle \tilde{\mathcal{A}}\Phi, T_{2,l}^x \Psi_l^\gamma \rangle_{(L_2(\Omega_D))^G} + \langle \tilde{\mathcal{A}}\Phi, \varepsilon_1 + \varepsilon_2 \rangle_{L_2(\tilde{\Omega})} \\ &= \langle \varphi, T_{2,l}^x \Psi_l^\gamma \rangle_{(L_2(\Omega_D))^G} + \tilde{\varepsilon} \end{aligned}$$

with $\tilde{\varepsilon} \in \mathbb{R}$. Numerically, both methods are almost identical. In several simulations, the maximal improvement in Strehl ratio when using the cut-off instead of the extension of the layer was 0.086%. We preferred the continued layer model as it leads to simpler notation as well as less computational effort.

6.3 Numerical Results

In this section we discuss and present the numerical implementation of the method of the approximate inverse for MCAO. The setting was presented in Section 2.7.

6.3.1 Reconstruction Method

Previously, all functions were considered in infinite dimensional spaces. For the numerical treatment of the problem, however, a discretization is necessary. The discretization grid for this problem is given by the wavefront sensors and deformable mirrors used in the setup of the telescope. The shape of each deformable mirror is controlled by a fixed number of actuators, see Section 2.5.2. The grid of these actuator positions is, in our case, quadratic with the spacing between two neighbouring actuators being Δ_l , that is, each mirror might have a different but still uniform grid. These grids are used directly as the discretization of the layers of the atmosphere. This is possible since each DM is conjugated to a certain height in the atmosphere as discussed in Section 2.6. Thus we reconstruct for $l = 1, \dots, L$ the discretized layer $\Phi^{(l)}(x_{a_l})$ where $x_{a_l} \in X_{a_l}$ and X_{a_l} denotes the set of the positions of all actuators of the l -th deformable mirror. Similarly, the discretization of the measured wavefronts is given by the grid of the subapertures of the wavefront sensors (Section 2.5.1). Hence, the available data is, for $g = 1, \dots, G$, given by $\varphi_{\beta_g}(x_{sa_g})$ where $x_{sa_g} \in X_{sa_g}$ and X_{sa_g} denotes the set of coordinates of the subapertures of the g -th wavefront sensor.

The reconstruction of each layer $\Phi^{(l)}$ requires three basic steps. First, the vectors \tilde{e}_l^γ and hence the mollifier $e^\gamma \in L_2(\mathbb{R}^2)$ has to be set up. Since we are free to choose the parameter γ , we might make a different choice for each layer l . Hence, γ_l denotes the mollification parameter used to reconstruct layer l . In our first simulations we chose the Gaussian

$$e^{\gamma_l}(\varrho) := \frac{1}{2\pi\gamma_l^2} \exp\left(-\frac{\|\varrho\|^2}{2\gamma_l^2}\right)$$

with $\gamma_l > 0$. For this mollifier, it holds $\|e^{\gamma_l}\|_{L_2(\mathbb{R}^2)} = 1$. However, the given discretization seemed to be too coarse and a high quality solution could not be found since either γ was so large that the mollifier was blurring too much or it was so small that its discrete norm was far away from 1. Instead, we followed the idea that the mollifier is supposed

to approximate the delta distribution. We let

$$e^\gamma(\varrho) = \begin{cases} \gamma_l \Delta_l^{-2}, & \varrho = 0 \\ 0, & \varrho \neq 0 \end{cases}, \quad (6.30)$$

with γ_l being a free parameter and Δ_l again the actuator spacing. This definition ensures $\|e^\gamma\| \approx 1$ on the actuator grid of each mirror. For the results presented subsequently we let $\gamma_1 = \mu_1^{-1}$, $\gamma_2 = \mu_2^{-1}$ and $\gamma_3 = \mu_3^{-1}$. Using this mollifier e^γ , the l -th mollifying vector is given by $\bar{e}_l^\gamma = (c_l e^\gamma \delta_{i,l})_{i=1,\dots,L}$.

The second step for the reconstruction of the l -th layer is to solve equation (6.22) via the iteration procedure (6.24) to obtain the reconstruction kernels. The procedure is described in Algorithm 1. An example of such a reconstruction kernel is given in Figure 6.2. In order to get an idea about the reconstruction quality, we can check the mollifier vector obtained from the reconstruction kernels. In the optimal case we would get $\mathcal{A}^* \psi_l^\gamma = \bar{e}_l^\gamma$ for $l = 1, 2, 3$. However, since we can only solve the normal equation, i.e., look for an approximate solution, layers on which the mollifier should be zero (as would be the case if we were able to solve (6.19)) show some artefacts, when the mollifier is reconstructed from the computed reconstruction kernels. An example is given in Figure 6.3. In Lemma 6.2.4 we showed that a function that is nonzero on only one layer (as we would desire for the mollifier) is not in the range of \mathcal{A}^* . Hence, if the reconstruction kernels ψ_l^γ are obtained by solving the normal equation (6.22), the reconstructed mollifier $\tilde{e}_l^\gamma := \mathcal{A}^* \psi_l^\gamma$ will not fulfil this condition. In a last step we reconstruct the layers via the method of the approximate inverse, see Algorithm 2. Having available the precomputed reconstruction kernels, the function value of each point $x_{a_l} \in X_{a_l}$ on a layer is given by $\Phi_\gamma^{(l)}(x_{a_l}) = \langle \varphi, T_{2,l}^{x_{a_l}} \Psi_l^\gamma \rangle_{(L_2(\mathbb{R}^2))^G}$, i.e., the sum of G single L_2 -scalar products $\langle \varphi_{\beta_g}, (\Psi_l^\gamma)_g(\cdot - \frac{x_{a_l}}{\mu_{l,\beta_g}}) \rangle_{L_2(X_{s_{a_l}})}$. In order to evaluate this in the discrete setting, each shifted reconstruction kernel (according to the invariance operator $T_{2,l}^{x_{a_l}}$ from Theorem 6.2.3) has to be interpolated on the grid of the data first. This procedure is repeated for each point $x_{a_l} \in X_{a_l}$ and for all layers $l = 1, \dots, L$. If we would reconstruct more layers than the number of deformable mirrors, an additional fitting step would have to be introduced as mentioned in Section 2.5.2. Algorithm 2 is repeated for every time step. The reconstruction kernels are not changed between different time steps.

6.3.2 Simulation results

We compare the performance of the method of the approximate inverse (AI) with a gradient-based algorithm [132] and a Kaczmarz-type method [60, 41]. Both are known to produce solutions of the same quality as the methods currently used in practice for smaller telescopes. As evaluation criterion we use the long exposure (LE) Strehl ratio in K-band (for a wavelength of 2200 nm) after 100 time steps of 2 ms, according to the

Algorithm 1 Algorithm for the calculation of the reconstruction kernels. We stop the iteration after a fixed number of steps K has been performed. In the experiments here, we used $K = 1800$

Input: telescope geometry, guide star setting

for $l = 1, \dots, L$ **do**

Initialise discretized mollifiers $\bar{e}_l^\gamma(x_{a_l})$, $l = 1, \dots, L$ from (6.30)

Initialise discretized reconstruction kernels $(\Psi_l^\gamma)_g(x_{sa_g}) = 0$, $g = 1, \dots, G$, set $k = 0$

while $k < K$ **do**

$\bar{k} = \text{mod}(k, G) + 1$

 calculate τ_k according to (6.25)

 update $(\Psi_l^\gamma)_{\bar{k}}$ according to (6.24)

$k = k + 1$

end while

Output: reconstruction kernels $(\Psi_l^\gamma)_g(x_{sa_g})$ for $g = 1, \dots, G$

end for

Algorithm 2 The method of the approximate inverse for MCAO. Reconstruction of the layers in one time step.

Input: telescope, geometry, guide star setting, reconstructed wavefronts φ_{β_g} , $g = 1, \dots, G$

for $l = 1, \dots, L$ **do**

Input: precomputed reconstruction kernels $(\Psi_l^\gamma)_g(x_{sa_g})$ for $g = 1, \dots, G$

for each $x_{a_l} \in X_{a_l}$ **do**

 evaluate (6.26): Initialise $\Phi^{(l)}(x_{a_l}) = 0$

for $g = 1, \dots, G$ **do**

 interpolate $(\Psi_l^\gamma)_g(\cdot - \frac{x_{a_l}}{\mu_{l,\beta_g}})$ on the grid X_{sa_g}

$\Phi^{(l)}(x_{a_l}) = \Phi^{(l)}(x_{a_l}) + \langle \varphi_{\beta_g}, (\Psi_l^\gamma)_g(\cdot - \frac{x_{a_l}}{\mu_{l,\beta_g}}) \rangle_{L_2(X_{sa_l})}$

end for

end for

Output: reconstructed layer $\Phi^{(l)}(x_{a_l})$, $x_{a_l} \in X_{a_l}$

end for

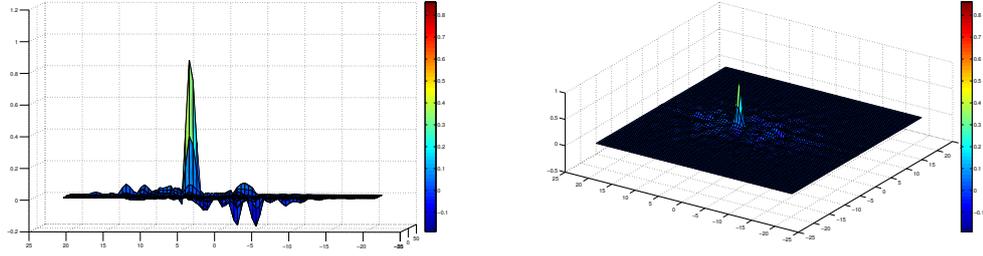


Figure 6.2: Reconstruction kernel Ψ_3^γ for layer 3 and a laser guide star

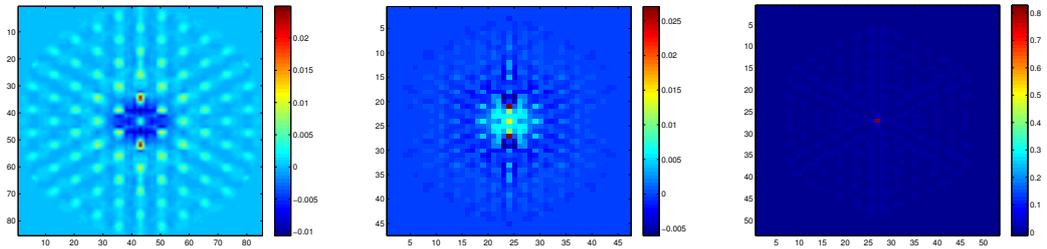


Figure 6.3: Mollification vector \bar{e}_l^γ obtained from reconstruction kernels for $\mathcal{A}^*\Psi_3^\gamma$ for the reconstruction of layer 3. From left to right: $(\mathcal{A}^*\Psi_3^\gamma)^{(l)}$ for $l = 1, 2, 3$. Although the third layer is reconstructed adequately, the other layers show some artefacts instead of being plain zero as would be the optimal case.

frequency of the AO system. For each setup we provide LE Strehl plots of the field of view of the telescope from reconstructions obtained with the method of the approximate inverse and of the two reference methods. Below these figures, we show a comparison of LE Strehl versus separation. Here, separation is the deviation of the viewing angle from zenith. The radial Strehl has been azimuthally averaged. Running the simulation against the gradient method and Kaczmarz method, we observe that the approximate inverse performs comparatively to the other ones. In Figure 6.4 we show the result for natural guide stars, where the approximate inverse produces a slightly worse result than the reference methods. For laser guide stars, however, the approximate inverse gave the best result, although by a small margin only; see Figure 6.5. “Better” results here correspond to higher Strehl value which would lead to the observed images showing more details than for lower Strehl ratios. Since the invariances of Theorem 6.2.3 only hold for a single type of guide stars used, we expect a slight drop in reconstruction quality when both guide star types are mixed. Numerical results confirm this, see Figure 6.6. There, we used 6 laser guide stars and 3 natural guide stars. However, the Strehl obtained with the approximate inverse is still very close to the reference methods. In this experiment,

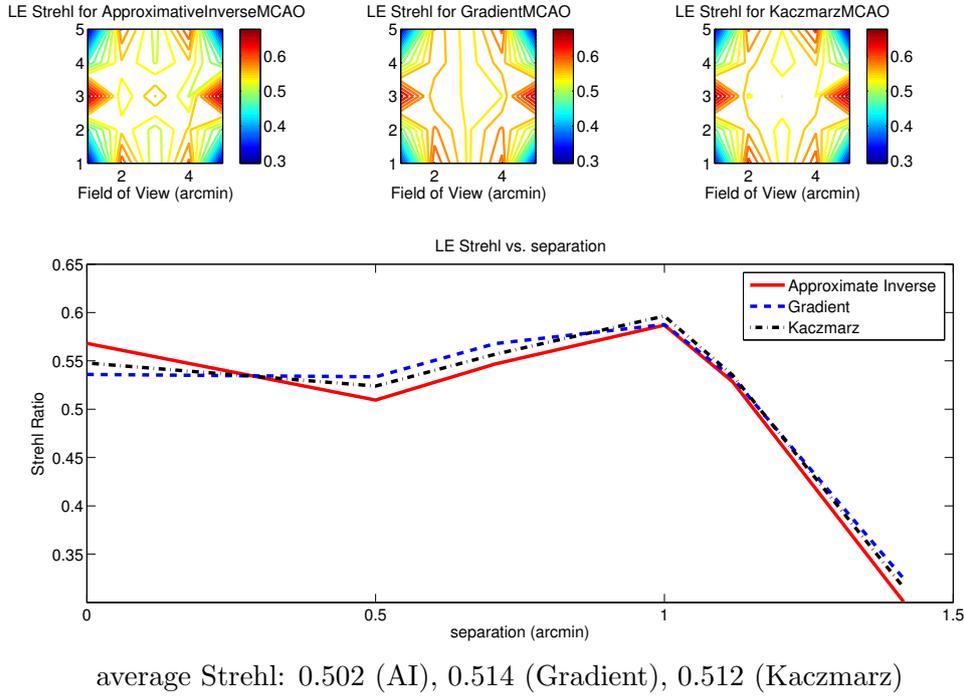


Figure 6.4: Strehl for 6 natural guide stars. The Strehl ratio obtained with the method of the approximate inverse is comparable with the reference methods. In the center of the field of view, it is even larger, corresponding to better image quality for astronomical objects positioned there.

our methods drops in Strehl in particular for large separations since the difference in the invariance operators matters most for larger angles. In order to stabilize the updates of the mirror shape, which have to be performed every 2 ms, we use a pseudo-open loop control (POLC) for all three methods. In Adaptive Optics, POLC has been introduced in, e.g., [64].

6.3.3 Computational complexity

In Adaptive Optics, speed is one of the most critical criteria for acceptance of a method. For the approximate inverse, the situation is as follows. In a first step, the wavefronts have to be reconstructed for each guide star from the Shack-Hartmann wavefront sensors with CuReD. The computational complexity is $20n_{sa}$ where n_{sa} is the total number of subapertures in the Shack-Hartmann wavefront sensor [48]. This has to be done for each of the G guide stars. CuReD is parallelizable. The reconstructed wavefronts are then

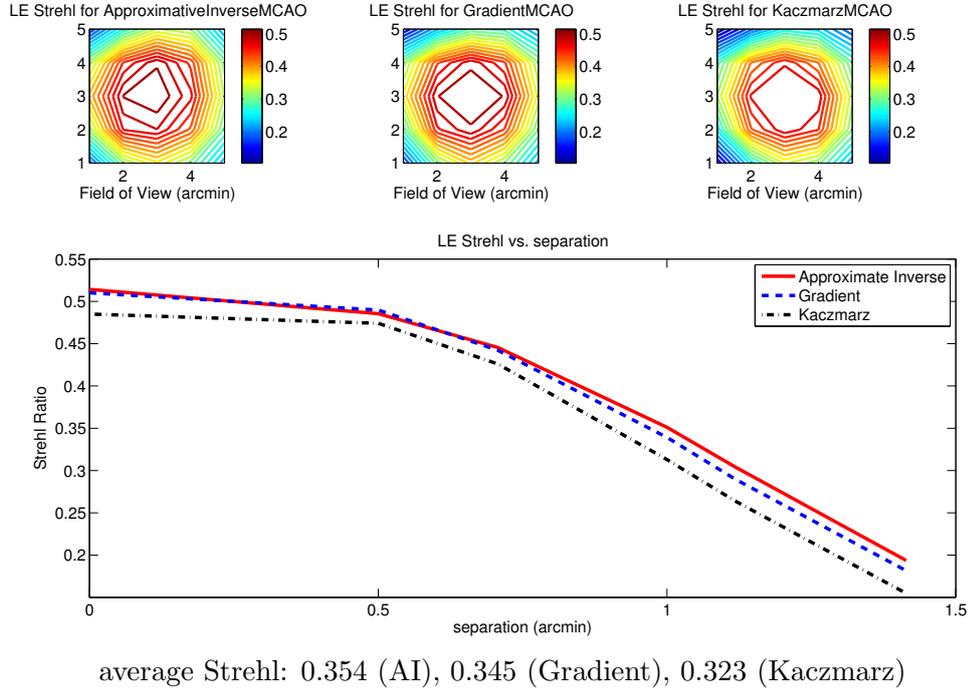


Figure 6.5: Strehl for 6 laser guide stars. The Strehl ratio obtained with the method of the approximate inverse is higher than the one of the reference methods, i.e, the observed image would be of higher quality.

used to calculate the shape of the deformable mirrors using the reconstruction kernels. The kernels are precomputed, see Algorithm 1. Hence this rather time consuming step gives no contribution during the actual computations. In particular, no runtime evaluation of the forward or adjoint operator is needed. The remaining effort lies in the evaluation of the scalar products, where the reconstruction kernels are merely an input, see Algorithm 2. For the reconstruction of a single point of one layer, G inner products are required, each consisting of $2 \cdot n_{sa}$ operations for the evaluation. If the invariance operators from Theorem 6.2.3 are used, further $6n_{sa}$ operations are required for the evaluation of each inner product to interpolate the shifted reconstruction kernels on the grid given by the wavefront sensors. Since we can reconstruct the artificial layers directly on this grid, we have to do this for n_{act} points, where n_{act} denotes the combined number of actuators for all mirrors. Assuming that for each mirror we have approximately n_{sa} actuators, the overall complexity for the reconstruction of a whole atmosphere in one time step is $G \cdot 20 \cdot n_{sa} + n_{act} \cdot G \cdot 2n_{sa} \cdot 6 \cdot n_{sa} \approx G \cdot 20 \cdot n_{sa} + 12L \cdot G \cdot n_{sa}^3$. Using the alternative invariance operators from the end of Section 6.2.1, we can avoid the interpolation of the

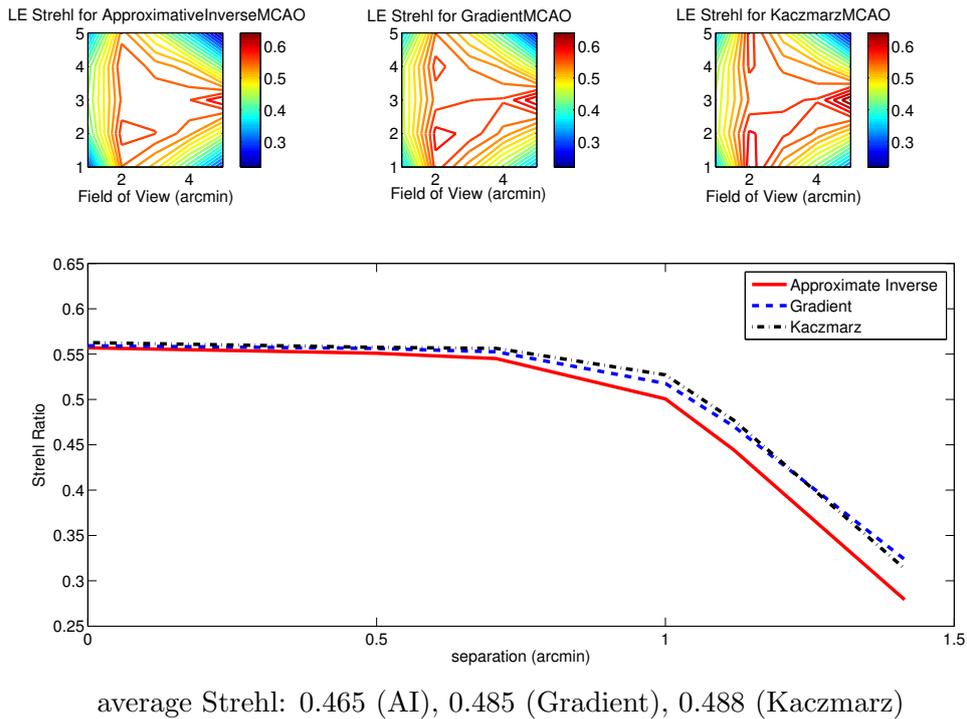


Figure 6.6: Strehl for 6 laser guide stars and 3 (full) natural guide stars. Although the invariances do not hold exactly for mixed guide stars, the Strehl ratio obtained with the method of the approximate inverse is still close to the reference methods. We explained at the end of Section 4.1 that this is not a setting relevant in practice.

reconstruction kernels, but require an interpolation of the reconstructed mirror shape to its actual domain. However, since there are only very few layers compared to the amount of actuators, the computational cost is significantly reduced. The overall complexity can then be estimated as $G \cdot 20 \cdot n_{sa} + n_{act} \cdot G \cdot 2n_{sa} + 6 \cdot n_{act} \approx G \cdot 26 \cdot n_{sa} + 12L \cdot G \cdot n_{sa}^2$. Since each inner product is independent of the other ones, the procedure is parallelizable. Note that, due to the nature of the method, an extension of the model, e.g., an inclusion of additional effects of laser guide stars, does not increase the relevant computational effort as all operations involving the model are solely performed in the calculation of the reconstruction kernels. Finally, we would like to briefly discuss the option of estimating more layers, i.e., more turbulence profiles of the atmosphere, than deformable mirrors which usually leads to improved Strehl ratio. Each additional layer to be reconstructed adds about $12G \cdot n_{sa}^3$ (or $12G \cdot n_{sa}^2$, respectively) operations. Then, one needs an optimization routine which finds the configuration of the available deformable mirrors that corrects best for the computed atmospheric turbulence profile. This can be done in linear complexity, see for example [146]. Hence, adding more layers is very costly. Although we do not have experimental results, we expect the possible gain in estimating more layers to be comparable to the moderate gain achieved with other methods, for example [75]. In summary, we would expect that estimating more layers than deformable mirrors is not advisable as the costs are too high.

Chapter 7

Conclusions

In Chapter 1 we introduced the deterministic and stochastic setting of Inverse Problems. We discussed advantages and disadvantages of both approaches and showed connections between the two approaches of noise modeling. We moved on to the introduction Adaptive Optics in Atmospheric Tomography as a particular application of an Inverse Problem. In Chapter 3 we reviewed filter-based regularization methods with focus on two regularization methods called fractional Tikhonov regularization. We showed both in theory and numerical experiments that, although both are of optimal convergence order with a-priori parameter choice as well as with the discrepancy principle, their superiority to Tikhonov regularization in standard form is restricted to the case when the underlying problem is severely ill-posed or the noise is concentrated in low frequencies. We continued to show that a-priori convergence results for filter-based regularization methods carry over directly to the stochastic setting. The transition is less straight forward for the discrepancy principle for which we presented two stochastic formulations. The first one, based on the modification of the residual as a stochastic quantity, inherits the deterministic convergence rates in expectation while for the second one, assuming only one realization of the noisy data is available, we showed convergence rates in the slightly weaker Ky Fan metric. Again, the theoretical results were exemplified numerically. We moved on to sparsity-promoting regularization. Starting from the Bayesian approach in the stochastic setting we derived a Tikhonov-type functional for the maximum a-priori solution and showed convergence and convergence rates in the Ky Fan metric by lifting deterministic results to the stochastic setting. This lifting approach allowed the construction of a novel parameter choice rule based on the stochastic modeling of noise and unknown required by the Bayesian approach. We applied the theory to the problem of Atmospheric tomography, confirming the viability of the new parameter choice rule. We then summarized the results of lifting deterministic convergence rates to the stochastic setting. Given a deterministic convergence theory, convergence in the Ky Fan metric is directly implied when instead of the worst case estimate δ the Ky Fan distance between true and noisy data or the slightly slowed down expectation of the error is used. Finally, we apply the method of the approximate inverse to Atmospheric Tomography. We show how the traditional formulation for scalar functions has to be adapted for vector valued functions. Finding special invariance operators allows to reduce the computational effort by having to solve only one adjoint problem to obtain the so called reconstruction kernels. This is done with a Kaczmarz-type algorithm. In numerical simulations, the method of the approximate inverse leads to results which can compete with state of the art methods in several astronomical configurations.

Bibliography

- [1] A. Tikhonov. On the stability of inverse problems. *Dokl. Akad. Nauk SSSR*, 39:195–198, 1943.
- [2] A. Tikhonov. On solution of ill-posed problems and the regularization method. *Dokl. Akad. Nauk SSSR*, 151:501–504, 1963.
- [3] F. Natterer. *The Mathematics of Computerized Tomography*. Wiley, 1986.
- [4] J. Hadamard. Sur les problèmes aux dérivées partielles et leur signification physique. *Princeton University Bulletin*, 1902.
- [5] J. Kaipio and E. Somersalo. Statistical inverse problems: Discretization, model reduction and inverse crimes. *J. Comput. Appl. Math.*, 198:493504, 2007.
- [6] J. Kaipio and E. Somersalo. *Statistical and computational inverse problems*, volume 160 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 2005.
- [7] A. K. Louis. *Inverse und schlecht gestellte Probleme*. B. G. Teubner, Stuttgart, 1989.
- [8] H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of inverse problems*, volume 375 of *Mathematics and its Applications*. Kluwer Academic Publishers Group, Dordrecht, 1996.
- [9] B. Hofmann. *Regularization for applied inverse and ill-posed problems*. Teubner, 1986.
- [10] Andreas Rieder. *Keine Probleme mit inversen Problemen*. Friedr. Vieweg & Sohn, Braunschweig, 2003.
- [11] S. W. Anzengruber and R. Ramlau. Morozovs discrepancy principle for tikhonov-type functionals with nonlinear operators. *Inverse Problems*, 26:025001, 2010.
- [12] T. Hohage and F. Werner. Convergence rates for inverse problems with impulsive noise. *SIAM J. Numer. Anal.*, 52(3):1203–1221, 2014.
- [13] T. Kösters C. Brune, A. Sawatzky F. Wübbeling and M. Burger. Forward backward em-tv methods for inverse problems with poisson noise. *Preprint*, 2010.

- [14] A. Munk N. Bissantz, T. Hohage and F. Ruymgaart. Convergence rates of general regularization methods for statistical inverse problems and applications. *SIAM J. Numer. Anal.*, 45(6):2610–2636, 2007.
- [15] M. Lassas H. Kekkonen and S. Siltanen. Analysis of regularized inversion of data corrupted by white Gaussian noise. *Inverse Problems*, 30(4):045009, 18, 2014.
- [16] A. M. Stuart. Inverse problems: A bayesian perspective. *Acta Numerica*, 19:451–559, 2010.
- [17] K. Mosegaard and M. Sambridge. Monte carlo analysis of inverse problems. *Inverse Problems*, 18, 2002.
- [18] S. N. Evans and P. B. Stark. Inverse problems as statistics. *Inverse Problems*, 18, 2002.
- [19] A. Tarantola. *Inverse problem theory and methods for model parameter estimation*. SIAM, 2005.
- [20] D. Calvetti and E. Somersalo. *An introduction to Bayesian scientific computing Ten lectures on subjective computing*. Springer, 2007.
- [21] A. Hofinger and H. K. Pikkarainen. Convergence rate for the Bayesian approach to linear inverse problems. *Inverse Problems*, 23(6):2469–2484, 2007.
- [22] A. Hofinger and H. K. Pikkarainen. Convergence rates for linear inverse problems in the presence of an additive normal noise. *Stochastic Analysis and Application*, 27(2):240–257, 2009.
- [23] A. Neubauer and H. K. Pikkarainen. Convergence results for the Bayesian inversion theory. *J. Inverse Ill-Posed Probl.*, 16(6):601–613, 2008.
- [24] Map estimators and their consistency in bayesian nonparametric inverse problems. *Inverse Problems*, 29(9):095017, 2013.
- [25] Maximum a posteriori probability estimates in infinite-dimensional bayesian inverse problems. preprint, arXiv:1412.5816v3.
- [26] S. Lasanen. Non-gaussian statistical inverse problems. part i: Posterior distributions. *Inv. Probl. Imaging*, 6:215–266, 2012.
- [27] Ian F. Blake. *An introduction to applied probability*. John Wiley & Sons, New York-Chichester-Brisbane, 1979.
- [28] G. Blanchard and P. Mathé. Discrepancy principle for statistical inverse problems with application to conjugate gradient iteration. *Inverse Problems*, 28(11):115011, 2012.

- [29] T. Hohage and F. Werner. Convergence rates in expectation for tikhonov-type regularization of inverse problems with poisson data. *Inverse Problems*, 28:104004, 2012.
- [30] A. Hofinger. *Ill-posed problems: Extending the Deterministic Theory to a Stochastic Setup*. PhD thesis, Johannes Kepler University Linz, 2006.
- [31] Ky Fan. Entfernung zweier zufälligen Grössen und die Konvergenz nach Wahrscheinlichkeit. *Math. Z.*, 49:681–683, 1944.
- [32] R. M. Dudley. *Real analysis and probability*. The Wadsworth & Brooks/Cole Mathematics Series. Wadsworth & Brooks/Cole Advanced Books & Software, Pacific Grove, CA, 1989.
- [33] B. Eisenberg and B.K. Ghosh. A generalization of markov’s inequality. *Statistics & Probability Letters*, 53(1):59 – 65, 2001.
- [34] V. I. Bogachev. *Gaussian measures*, volume 62 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1998.
- [35] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth. On the Lambert W function. *Adv. Comput. Math.*, 5(4):329–359, 1996.
- [36] F. Roddier. *Adaptive Optics in Astronomy*. Cambridge, U.K. ; New York : Cambridge University Press, Cambridge, 1999.
- [37] M. C. Roggemann and B. Welsh. *Imaging through turbulence*. CRC Press laser and optical science and technology series. CRC Press, 1996.
- [38] F. Roddier. The effects of atmospheric turbulence in optical astronomy. *Progress in optics*, 19:281–376, 1981.
- [39] A. Tokovinin. Adaptive optics tutorial at CTIO, 2001. <http://www.ctio.noao.edu/~atokovin/tutorial/intro.html>.
- [40] NASA. James webb space telescope (JWST) independent comprehensive review panel (ICRP), 2012. http://www.nasa.gov/pdf/499224main_JWST-ICRP_Report-FINAL.pdf.
- [41] M. Rosensteiner and R. Ramlau. The Kaczmarz algorithm for multi-conjugate adaptive optics with laser guide stars. *J. Opt. Soc. Am. A*, 30(8):1680–1686, 2013.
- [42] A. Quirrenbach. The effects of atmospheric turbulence on astronomical observations. In *Center for Adaptive Optics Summer School*. ASP Conference Series, 2002.

- [43] A. N. Kolmogorov. Local structure of turbulence in incompressible viscous fluid for very large reynolds numbers. *DAN SSSR*, 30, 1941. in Russian.
- [44] V. I. Tatarski. *wave Propagation in a turbulent medium*. McGraw Hill, New York, 1961.
- [45] D. L. Fried. Statistics of a geometric representation of wavefront distortion. *J. Opt. Soc. Am.*, pages 1427–1431, 1965.
- [46] M. Zhariy, A. Neubauer, M. Rosensteiner, and R. Ramlau. Cumulative wavefront reconstructor for the Shack-Hartman sensor. *Inverse Problems and Imaging*, 5(4):893–913, Nov 2011.
- [47] M. Rosensteiner. Cumulative reconstructor: fast wavefront reconstruction algorithm for Extremely Large Telescopes. *J. Opt. Soc. Am. A*, 28(10):2132–2138, Oct 2011.
- [48] M. Rosensteiner. Wavefront reconstruction for extremely large telescopes via CuRe with domain decomposition. *J. Opt. Soc. Am. A*, 29(11):2328–2336, Nov 2012.
- [49] B.L. Ellerbroek and C.R. Vogel. Inverse problems in astronomical adaptive optics. *Inverse Problems*, 25:1–37, 2009.
- [50] B. Ellerbroek, L. Gilles, and C.R. Vogel. Numerical simulations of multiconjugate adaptive optics wavefront reconstruction on giant telescopes. *Applied Optics*, 42:4811–4818, 2003.
- [51] N. Ageorges and C. Dainty, editors. *Laser Guide Star Adaptive Optics for Astronomy*. Series C: Mathematical and Physical Sciences - Vol. 551. Kluwer Academic Publishers, 2000.
- [52] R. Ramlau, A. Obereder, M. Rosensteiner, and D. Saxenhuber. Efficient iterative tip/tilt reconstruction for atmospheric tomography. *submitted*, 2013.
- [53] Béchet, C., Le Louarn, M., Clare, R., Tallon, M., Tallon-Bosc, I., and Thiébaud, É. Closed-loop ground layer adaptive optics simulations with elongated spots : impact of modeling noise correlations. *1st AO4ELT conference - Adaptive Optics for Extremely Large Telescopes*, page 03004, 2010.
- [54] Richard M. Clare, Miska Le Louarn, and Clementine Béchet. Optimal noise-weighted reconstruction with elongated shack–hartmann wavefront sensor images for laser tomography adaptive optics. *Appl. Opt.*, 49(31):G27–G36, Nov 2010.

- [55] Michel Tallon, Isabelle Tallon-Bosc, Clémentine Béchet, and Eric Thiébaud. Shack-hartmann wavefront reconstruction with elongated sodium laser guide stars: improvements with priors and noise correlations. volume 7015, page 70151N. SPIE, 2008.
- [56] M.A. Davison. The ill-conditioned nature of the limited angle tomography problem. *SIAM J. Appl. Math.*, 43:428–448, 1983.
- [57] D. Gavel. Tomography for multiconjugate adaptive optics systems using laser guide stars. *SPIE Astronomical Telescopes and Instrumentation*, 5490:1356–1373, Jun 2004.
- [58] B. Ellerbroek, L. Gilles, and C.R. Vogel. A computationally efficient wavefront reconstructor for simulation or multi-conjugate adaptive optics on giant telescopes. *Proc. SPIE*, 4839, 2002.
- [59] T. Fusco, J.-M. Conan, G. Rousset, L.M. Mugnier, and V. Michau. Optimal wavefront reconstruction strategies for multi conjugate adaptive optics. *J. Opt. Soc. Am. A*, 18(10):2527–2538, 2001.
- [60] R. Ramlau and M. Rosensteiner. An efficient solution to the atmospheric turbulence tomography problem using Kaczmarz iteration. *Inverse Problems*, 28(9):095004, 2012.
- [61] G. Auzinger. *New Reconstruction Approaches in Adaptive Optics for Extremely Large Telescopes*. PhD thesis, Johannes Kepler University Linz, 2014.
- [62] E. Marchetti, R. Brast, and B. Delabre et al. On-sky testing of the multi-conjugate adaptive optics demonstrator. *The Messenger*, 189:8–13, 2007.
- [63] L. Wang and E. Ellerbroek. Computer simulations and real-time control of elt ao systems using graphic processing units. In *Adaptive Optics systems III*, volume 8447 of *Proc. Soc. Photo-Opt. Instrum. Eng.*, 2012.
- [64] B. Ellerbroek and C.R. Vogel. Simulations of closed-loop wavefront reconstruction for multiconjugate adaptive optics on giant telescopes. *Proc. SPIE*, 5169:206–217, 2003.
- [65] L. Gilles, B. Ellerbroek, and C.R. Vogel. Layer-oriented multigrid wavefront reconstruction algorithms for multi-conjugate adaptive optics. *Proc. SPIE*, 4839, 2002.
- [66] L. Gilles, B.L. Ellebroek, and C.R. Vogel. Preconditioned conjugate gradient wave-front reconstructors for multiconjugate adaptive optics. *Applied Optics*, 42(26):5233–5250, 2003.

- [67] L. Gilles, B. Ellerbroek, and C. A. comparison of Multigrid V-cycle versus Fourier Domain Preconditioning for Laser Guide Star Atmospheric Tomography. In *Adaptive Optics: Analysis and Methods/Computational Optical Sensing and Imaging/Information Photonics/Signal Recovery and Synthesis Topical Meetings on CD-ROM, OSA Technical Digest (CD)*. Optical Society of America, 2007.
- [68] L. Gilles and B. Ellerbroek. Split atmospheric tomography using laser and natural guide stars. *J. Opt. Soc. Am.*, 25(10):2427–35, 2008.
- [69] C.R. and Q. Yang. Fast optimal wavefront reconstruction for multi-conjugate adaptive optics using the Fourier domain preconditioned conjugate gradient algorithm. *Optics Express*, 14(17), 2006.
- [70] Q. Yang, C.R., and B.L. Ellerbroek. Fourier domain preconditioned conjugate gradient algorithm for atmospheric tomography. *Applied Optics*, 45(21):5281–5293, 2006.
- [71] A. Tokovinin and M. Le Louarn. Isoplanatism in a multiconjugate adaptive optics system. *J. Opt. Soc. Am. A*, 17(10):1819–1827, Oct 2000.
- [72] A. Tokovinin and E. Viard. Limiting precision of tomographic phase estimation. *J. Opt. Soc. Am. A*, 18(4):873–882, Apr 2001.
- [73] L. Gilles, P. Massioni, C. Kulcsár, H.-F. Raynaud, and B. Ellerbroeck. Distributed kalman filtering compared to fourier domain preconditioned conjugate gradient for laser guide star tomography on extremely large telescopes. *JOSA A*, 30:898–909, 2013.
- [74] M. Tallon, I. Tallon-Bosc, C. Béchet, F. Momey, M. Fradin, and E. Thiébaud. Fractal iterative method for fast atmospheric tomography on extremely large telescopes. *Proc. SPIE 7736*, page 77360X, 2010.
- [75] T. Helin and M. Yudytskiy. Wavelet methods in multi-conjugate adaptive optics. *Inverse Problems*, 29(8):085003, 2013.
- [76] M. Yudytskiy, T. Helin, and R. Ramlau. A frequency dependent preconditioned wavelet method for atmospheric tomography. In *Third AO4ELT Conference - Adaptive Optics for Extremely Large Telescopes*, May 2013.
- [77] M. Rosensteiner and R. Ramlau. Efficient iterative atmospheric tomography reconstruction from LGS and additional tip/tilt measurements. In *SPIE 8447, Adaptive Optics Systems III*, pages 84475S–84475S–6, 2012.
- [78] Austrian Adaptive Optics Team. Mathematical Algorithms for E-ELT Adaptive Optics. <http://eso-ao.indmath.uni-linz.ac.at/>.

- [79] E. Diolaiti, A. Baruffolo, M. Bellazzini, V. Biliotti, G. Bregoli, C. Butler, P. Ciliegi, J.-M. Conan, G. Cosentino, S. D’Odorico, B. Delabre, H. Foppiani, T. Fusco, N. Hubin, M. Lombini, E. Marchetti, S. Meimon, C. Petit, C. Robert, P. Rossettini, L. Schreiber, and R. Tomelleri. MAORY: A Multi-conjugate Adaptive Optics RelaY for the E-ELT. *Messenger*, pages 28–9, June 2010.
- [80] Daniel Gerth, Esther Klann, Ronny Ramlau, and Lothar Reichel. On fractional Tikhonov regularization. *J. Inv. Ill-posed Probl.*, 30(5):245–266, 2015.
- [81] E. Klann and R. Ramlau. Regularization by fractional filter methods and data smoothing. *Inverse Problems*, 24(2):025018, 26, 2008.
- [82] Michiel E. Hochstenbach and Lothar Reichel. Fractional Tikhonov regularization for linear discrete ill-posed problems. *BIT Numerical Mathematics*, 51:197–215, 2011.
- [83] Peter Mathé and Ulrich Tautenhahn. Enhancing linear regularization to treat large noise. *J. Inverse Ill-Posed Probl.*, 19(6):859–879, 2011.
- [84] S. Noschese M. E. Hochstenbach and L. Reichel. Fractional regularization matrices for linear discrete ill-posed problems. *J. Eng. Math.*, in press.
- [85] V.A. Morozov. *Methods for solving incorrectly posed problems*. New York etc.: Springer-Verlag. XVIII, 257 p., 1984.
- [86] Christopher T. H. Baker. *The numerical treatment of integral equations*. Clarendon Press, Oxford, 1977. Monographs on Numerical Analysis.
- [87] Per Christian Hansen. Regularization Tools version 4.0 for Matlab 7.3. *Numer. Algorithms*, 46(2):189–194, 2007.
- [88] Daniel Gerth and Ronny Ramlau. A stochastic convergence analysis for Tikhonov regularization with sparsity constraints. *Inverse Problems*, 30(5):055009, 24, 2014.
- [89] I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Comm. Pure Appl. Math.*, 57(11):1413–1457, 2004.
- [90] M. Haltmeier M. Grasmair and O. Scherzer. Sparse regularization with l^q penalty term. *Inverse Problems*, 24:055020, 2008.
- [91] K. Bredies and D. Lorenz. Linear convergence of iterative soft-thresholding. *J Fourier Anal Appl*, 14:813837, 2008.
- [92] D. Lorenz. Convergence rates and source conditions for tikhonov regularization with sparsity constraints. *J. Inv. Ill-posed Probl.*, 16(5):463478, 2008.

- [93] K. Bredies and D. Lorenz. Iterated hard shrinkage for minimization problems with sparsity constraints. *SIAM J. Sci. Comput.*, 30(2):657683, 2008.
- [94] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sciences*, 2(1):183–202, 2009.
- [95] M. Fornasier I. Daubechies and I. Loris. Accelerated projected gradient method for linear inverse problems with sparsity constraints. *Journal of Fourier Analysis and Applications*, 14(5):764–792, 2008.
- [96] R. Griesse and D. Lorenz. A semismooth newton method for tikhonov functionals with sparsity constraints. *Inverse Problems*, 24:035007, 2008.
- [97] M. Fornasier. Domain decomposition methods for linear inverse problems with sparsity constraints. *Inverse Problems*, 23:2505–2526, 2007.
- [98] R. Ramlau. Regularization properties of Tikhonov regularization with sparsity constraints. *Electron. Trans. Numer. Anal.*, 30:54–74, 2008.
- [99] G. Teschke and R. Ramlau. An iterative algorithm for nonlinear inverse problems with joint sparsity constraints in vector valued regimes and an application to color image inpainting. *Inverse Problems*, 23:1851–1870, 2007.
- [100] G. Teschke and C. Borries. Accelerated projected steepest descent method for nonlinear inverse problems with sparsity constraints. *Inverse Problems*, 26(2):025007, 2010.
- [101] K. Bredies. An iterative thresholding-like algorithm for inverse problems with sparsity constraints in Banach space. *J. Inverse Ill-Posed Probl.*, 17(1):19–26, 2009.
- [102] M. Grasmair. Well-posedness convergence rates for sparse regularization with sublinear l^q penalty term. *Inv. Probl. Imaging*, 3:383–387, 2009.
- [103] M. Grasmair. Non-convex sparse regularisation. *Journal of Mathematical Analysis and Applications*, 365:19–28, 2010.
- [104] C. A. Zarzer. On tikhonov regularization with non-convex sparsity constraints. *Inverse Problems*, 25:025006, 2009.
- [105] B. Hofmann M. Burger and J. Flemming. Convergence rates in l^1 -regularization if the sparsity assumption fails. *Inverse Problems*, 29:025013, 2013.
- [106] B. Hofmann S. Anzengruber and R. Ramlau. On the interplay of basis smoothness and specific range conditions occurring in sparsity regularization. *Inverse Problems*, 29(12):125002, 2013.

- [107] S. Lu and S. Pereverzev. Sparse recovery by the standard tikhonov method. *Numer. Math*, 112:403–424, 2009.
- [108] R. Ramlau M. Fornasier and G. Teschke. The application of joint sparsity and total variation minimization algorithms to a real-life art restoration problem. *Adv. Comput. Math*, 31:157–184, 2009.
- [109] S. Kindermann and R. Ramlau. Surrogate functionals and thresholding for inverse interface problems. *J. Inv. Ill-posed Probl.*, 15:387401, 2007.
- [110] B. Jin and P. Maass. Sparsity regularization for parameter identification problems. *Inverse Problems*, 28:123001, 2012.
- [111] V. Kolehmainen, M. Lassas, K. Niinimäki, and S. Siltanen. Sparsity-promoting Bayesian inversion. *Inverse Problems*, 28(2), 2012.
- [112] M. Lassas, E. Saksman, and S. Siltanen. Discretization-invariant Bayesian inversion and Besov space priors. *Inverse Probl. Imaging*, 3(1):87–122, 2009.
- [113] S. Vänskä, M. Lassas, and S. Siltanen. Statistical X-ray tomography using empirical Besov priors. *Int. J. Tomogr. Stat.*, 11(S09):3–32, 2009.
- [114] S. Harris M. Dashti and A. Stuart. Besov priors for bayesian inverse problems. *Inverse Probl. Imaging*, 6:183200, 2012.
- [115] D. Lorenz and D. Trede. Optimal convergence rates for tikhonov regularization in besov scales. *Inverse Problems*, 24:05010, 2008.
- [116] S. Vänskä S. Järvenpää M. Kalke M. Lassas J. Moberg M. Rantala and S. Siltanen. Wavelet-based reconstruction for limited-angle x-ray tomography. *IEEE Tr. Medical Imaging*, 25:210217, 2006.
- [117] S. Siltanen K. Niinimäki and V. Kolehmainen. Bayesian multiresolution method for local tomography in dental x-ray imaging. *Phys. Med. Biol.*, 52:6663–6678, 2007.
- [118] S. Vää M. Lassas M. Kalke, S. Siltanen and M. Rantala. Method and arrangement for multiresolutive reconstruction for medical x-ray imaging. United States patent 7215730, 2007.
- [119] D. Leporini and J.-C. Pesquet. Bayesian wavelet denoising: Besov priors and non-gaussian noises. *Signal Processing*, 81:55–67, 2001.
- [120] I. Daubechies. *Ten lectures on wavelets*, volume 61 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1992.

- [121] Y. Meyer. *Wavelets and operators*, volume 37 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1992.
- [122] A. K. Louis, P. Maaß, and A. Rieder. *Wavelets*. B. G. Teubner, Stuttgart, 1998.
- [123] R. A. Adams and J. F. J. Fournier. *Sobolev spaces*, volume 140 of *Pure and Applied Mathematics (Amsterdam)*. Elsevier/Academic Press, Amsterdam, second edition, 2003.
- [124] M. Lindemann. *Approximation Properties of Non-separable Wavelet Bases with Isotropic Scaling Matrices*. PhD thesis, University Bremen, 2005.
- [125] S. Dahlke. *Wavelets: Construction Principles and Applications to the Numerical Treatment of Operator Equations*. PhD thesis, RWTH Aachen, 1996. Habilitation thesis.
- [126] M. Lassas and S. Siltanen. Can one use total variation prior for edge-preserving Bayesian inversion? *Inverse Problems*, 20(5):1537–1563, 2004.
- [127] S. Osher L. Rudin and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
- [128] Markus Grasmair, Markus Haltmeier, and Otmar Scherzer. Sparse regularization with l^q penalty term. *Inverse Problems*, 24(5):055020, 13, 2008.
- [129] D. A. Lorenz. Convergence rates and source conditions for Tikhonov regularization with sparsity constraints. *J. Inverse Ill-Posed Probl.*, 16(5):463–478, 2008.
- [130] D. Th. Egoroff. Sur les suites de fonctions mesurables. *C. R.*, 152:244–246, 1911.
- [131] M. Yudtyskiy, T. Helin, and R. Ramlau. Finite element-wavelet hybrid algorithm for atmospheric tomography. *J. Opt. Soc. Am. A*, 31(3):550–560, Mar 2014.
- [132] D. Saxenhuber and R. Ramlau. A gradient-based method for atmospheric tomography. Technical Report E-TRE-AAO-528-0043, Austrian In-Kind Contribution - AO, 2013.
- [133] R. Gorenflo and B. Hofmann. On autoconvolution and regularization. *Inverse Problems*, 10(2):353–373, 1994.
- [134] Daniel Gerth, Bernd Hofmann, Simon Birkholz, Sebastian Koke, and Günter Steinmeyer. Regularization of an autoconvolution problem in ultrashort laser pulse characterization. *Inverse Probl. Sci. Eng.*, 22(2):245–266, 2014.
- [135] Simon Birkholz, Günter Steinmeyer, Sebastian Koke, Daniel Gerth, Steven Bürger, and Bernd Hofmann. Phase retrieval via regularization in self-diffraction-based spectral interferometry. *JOSA B*, 32(5):983–992, 2015.

- [136] A. K. Louis and P. Maass. A mollifier method for linear operator equations of the first kind. *Inverse Problems*, 6(3):427–440, 1990.
- [137] Daniel Gerth, Bernadette Hahn, and Ronny Ramlau. The method of the approximate inverse for atmospheric tomography. *Inverse Problems*, 30(5):245–266, to appear.
- [138] A. K. Louis. Approximate inverse for linear and some nonlinear problems'. *Inverse Problems*, 12(2):175–190, 1996.
- [139] A. K. Louis. Feature reconstruction in inverse problems. *Inverse Problems*, 27:065010(21pp), 2011.
- [140] T. Schuster and F. Schöpfer. Solving linear operator equations in banach spaces non-iteratively by the method of approximate inverse. *Inverse Problems*, 26:085006(19pp), 2010.
- [141] T. Schuster. *The method of approximate inverse: theory and applications*, volume 1906 of *Lecture Notes in Mathematics*. Springer, Berlin, 2007.
- [142] M. K. Likht. On the calculation of functionals in the solution of linear equations of the first kind. *USSR Comput. Math. and Math. Phys.*, 7:271–8, 1967.
- [143] A. K. Louis. A unified approach to regularization methods for linear ill-posed problems. *Inverse Problems*, 15(2):489–498, 1999.
- [144] T. Schuster. An efficient mollifier method for three-dimensional vector tomography: convergence analysis and implementation. *Inverse Problems*, 17(4):739–766, 2001.
- [145] A. De Cezaro, M. Haltmeier, A. Leitão, and O. Scherzer. On steepest-descent-kaczmarz methods for regularizing systems of nonlinear ill-posed equations. *App. Math. Comput.*, 202(2):596–607, 2008.
- [146] R. Ramlau, D. Saxenhuber, and M. Yudytiskiy. Iterative reconstruction methods in atmospheric tomography: Fewha, kaczmarz and gradient-based algorithm. In *Proc. SPIE 9148, Adaptive Optics Systems IV*, 2014.

Curriculum Vitae

- **Personal details:**

Full Name: Daniel Gerth
Title: Dipl.-Math. techn., Diploma in Industrial Mathematics
Date of Birth: November 20, 1986
Place of Birth: Karl-Marx-Stadt, GDR (now Chemnitz, Germany)
Nationality: German

- **Education:**

2005 High School Diploma (Abitur) in Limbach-Oberfrohna, Germany
2005 – 2011 Studies in “Technomathematik” (Industrial Mathematics) at Chemnitz University of Technology, Chemnitz, Germany
Diploma thesis: *Regularization of an autoconvolution problem occurring in measurements of ultra-short laser pulses*
since Nov 2011 PhD student in the Doctoral Program “Computational Mathematics”, project DK08, Johannes Kepler University Linz, Austria

- **Teaching:**

SS 2012 Instructor: Mathematics for Chemistry Students II, JKU Linz, Austria
WS 2013 Instructor: Mathematics for Chemistry Students I, JKU Linz, Austria
SS 2013 Instructor: Mathematics for Chemistry Students II, JKU Linz, Austria
SS 2014 Lecturer: Mathematics for Chemistry Students II, JKU Linz, Austria
WS 2014 Lecturer: Mathematics for Chemistry Students I, JKU Linz, Austria
SS 2015 Lecturer: Mathematics for Chemistry Students II, JKU Linz, Austria

- **Research stays, with:**

Oct 2012 Prof. Shuai Lu at Fudan University, Shanghai, China
Dec 2012 Prof. Samuli Siltanen at University of Helsinki, Helsinki, Finland
Oct 2013 – Dec 2013 Prof. Lothar Reichel at Kent State University, Kent, Ohio, USA
Jan 2014 – Feb 2014 Prof. Chein-Shan Liu at National Taiwan University, Taipei, Taiwan (RoC)
Feb 2015 Prof. Samuli Siltanen at University of Helsinki, Helsinki, Finland
Jul 2015– Aug 2015 Prof. Alfred K. Louis at Saarland University, Saarbrücken, Germany

- **Participation in International Conferences**

- Sep 21–23 2011 Chemnitz Symposium on Inverse Problems 2011, Chemnitz, Germany (Talk)
- Apr 23–27 2012 ESI Workshop “Computational Inverse Problems” 2012, Vienna, Austria (without Talk)
- May 21–26 2012 6th International Conference “Inverse Problems: Modeling and Simulation”, Antalya, Turkey (Talk)
- Oct 19–20 2012 International Workshop on Inverse Problems: New trends in theory and applications, Shanghai, China (Talk)
- Oct 21–26 2012 International Conference on Inverse Problems and Related Topics 2012, Nanjing, China (Talk)
- Dec 17–19 2012 Inverse Days 2012, Jyväskylä, Finland (Talk)
- Jul 01–05 2013 Applied Inverse Problems Conference, Daejeon, South Korea (Talk)
- Sep 08–13 2013 26th IFIP TC7 Conference, Klagenfurt, Austria (Talk)
- Mar 26–28 2014 Sparse Tomo Days, Technical University of Denmark, Lyngby, Denmark (Talk)
- May 26–31 2014 7th International Conference “Inverse Problems: Modeling and Simulation”, Fethiye, Turkey (Talk)
- Aug 26–28 2014 Conference “Inverse Problems - from Theory to Application”, Bristol, UK (Poster presentation)
- May 24–28 2015 Applied Inverse Problems Conference, Helsinki, Finland (Talk)

- **Publications:**

- [1] D. Gerth, B. Hofmann, S. Birkholz, S. Koke and G. Steinmeyer, *Regularization of an autoconvolution problem in ultrashort laser pulse characterization*, Inverse Problems in Science and Engineering 22(2): 245–266, 2014
- [2] D. Gerth and R. Ramlau, *A stochastic convergence analysis for Tikhonov regularization with sparsity constraints*, Inverse Problems 30(5):055009, 2014
- [3] D. Gerth, E. Klann, R. Ramlau and L. Reichel, *On fractional Tikhonov regularization*, accepted to Journal of Inverse and Ill-posed problems
- [4] D. Gerth, B. Hahn and R. Ramlau, *The method of the approximate inverse for atmospheric tomography*, Inverse Problems 31:065002, 2015
- [5] S. Birkholz, G. Steinmeyer, S. Koke, D. Gerth, S. Bürger and B. Hofmann, *Phase retrieval via regularization in self-diffraction based spectral interferometry*, JOSA B 32(5) 983–992, 2015

Eidesstattliche Erklärung

Ich erkläre an Eides statt, dass ich die vorliegende Dissertation selbstständig und ohne fremde Hilfe verfasst, andere als die angegebenen Quellen und Hilfsmittel nicht benutzt bzw. die wörtlich oder sinngemäß entnommenen Stellen als solche kenntlich gemacht habe. Die vorliegende Dissertation ist mit dem elektronisch übermittelten Textdokument identisch.

Linz, am 17. September 2015

Daniel Gerth