
MASTER THESIS

DERIVATIVE-BASED REGULARIZATION OF INVERSE PROBLEMS IN ACOUSTIC HOLOGRAPHY

conducted at the
Institute of Electronic Music and Acoustics (IEM)
University of Music and Performing Arts Graz (KUG)

by
Manuel Pagavino

Supervisor:
Ass.Prof. DI Dr. Franz Zotter

Assessor:
O.Univ.Prof. Mag. DI Dr. Robert Höldrich



Graz, May 27, 2019

Institute of Electronic Music and Acoustics



Statutory Declaration

I declare that I have authored this thesis independently, that I have not used other than the declared sources/resources, and that I have explicitly marked all material which has been quoted either literally or by content from the used sources.

date

(signature)

Contents

1	Introduction	10
1.1	Introduction	10
1.2	Motivation and objective	11
2	Acoustic Background	13
2.1	The Euler equation	13
2.2	The homogeneous wave equation	14
2.3	The inhomogeneous wave equation	15
2.4	The free-space Green's function	16
2.5	The Kirchhoff Helmholtz integral equation	18
3	Sound Field Representation with the Equivalent Source Method	23
3.1	The Equivalent Source Method (ESM)	23
3.1.1	Formulation of ESM	24
3.1.2	Recovery of acoustic quantities of the sound field	27
3.2	Singular value decomposition	29
3.3	Propagating and evanescent waves	30
3.4	Ill-posedness of the inverse problem	31
3.5	Notes on measurement and retreat distance	32
4	Regularization in Convex Optimization Problems	35
4.1	Constrained and penalized optimization framework	35
4.2	Derivative-based regularizers	39
4.2.1	Introduction to derivatives in regularization	40
4.2.2	Gradient-based regularizer - The Total Variation	43
4.2.3	Laplacian-based regularizer - The Total Generalized Variation	43
4.2.4	Hessian-based regularizer - The Hessian Schatten-norm family	44
4.2.5	Fused sparse and sparse-differential regularizer	46
4.3	Comparison of first- and second-order derivative-based regularizers	47
4.4	Selection of regularization parameter	50
4.5	Discretization of derivatives via finite differences	54
5	The Algorithm - Proximal Splitting	59
5.1	Duality theory	60
5.1.1	Dual space and dual norm	60
5.1.2	Legendre-Fenchel conjugate and Moreau subdifferential	61
5.2	The projection operator	63
5.2.1	Projections onto vector- and matrix-norm balls	64
5.3	The proximity operator	66
5.3.1	Examples of proximal mappings	67
5.4	Linear operator and adjoint	67
5.5	Algorithm formulation	69
5.5.1	The composite problem	69
5.5.2	Primal-dual hybrid gradient (PDHG) algorithm	71
5.5.3	Practical application of the algorithm	72
6	Regularizer Characteristics - Simulation Study	75
6.1	Simulation setup	75
6.2	Figures of merit	76

6.3	Vibrating plate	77
6.4	Spatial response to a monopole	84
6.5	Rigid piston	88
6.6	Point source distribution	95
6.7	Sensitivity to noise and parameter choice	100
7	Regularizer Characteristics - Experimental Study	107
7.1	Experimental setup	107
7.2	Resonating plate	109
7.3	Loudspeaker driver layout	113
8	Summary and Conclusion	119

Abstract

The visualization of the sound field close to the source is often helpful to understand the vibro-acoustic origin. This gave rise to the development of several acoustic imaging techniques that can be used to model the measured sound field radiated by an arbitrary source. One of these models is the equivalent source method (ESM). It models the local sound field by superimposing distributed elementary sources of different strengths. From spatially discrete sound pressure measurements, the strength of these sources can be determined through solving a linear inverse problem.

Due to the underdetermined and ill-posed nature of the inverse problem, the introduction of some form of regularization is a prerequisite for obtaining a meaningful solution. Imposing additional constraints on the solution to enforce expected spatial structures can provide suitable regularization. Inverse problems with constraints typically minimize some norm functional acting on the spatial domain. Sparsity promotion through Compressive Sensing, based on L1-norm minimization, has received increasing attention in recent years due to its ability of providing solutions that are valid beyond the spatial sampling limit. However, typical vibro-acoustic source phenomena are not necessarily spatially sparse themselves, as they frequently contain spatially distributed patterns as well.

This thesis regards regularization methods that impose sparsity on first- and second-order spatial derivatives. This promotes piecewise constant or linear solutions with minimum curvature as a more probable spatial constraint. Such regularizers are heavily used in various fields of image processing. They were only recently introduced in acoustics, where they have consistently proven to effectively model common structures. In this thesis, I propose to adapt the Schatten-norms of the Hessian as regularizers, which to the best of my knowledge has not been considered for acoustic holography yet. What is more, a fused approach is considered where additional sparsity is imposed on the spatial domain, suitable for the characterization of sparse and extended sources. A proximal splitting algorithm is adopted to solve the minimization problem, which allows an efficient implementation of the proposed regularizers. This work provides the fundamental understanding of derivative-based regularization and reveals its characteristics and abilities. The proposed methods are investigated and verified by numerical simulations and by using measurements obtained from an experimental setup. The required theory behind the algorithm is examined and a detailed exposition of its use is provided.

Zusammenfassung

Die Visualisierung von Schallfeldern im Nahfeld einer Quelle ist oft hilfreich, um die vibroakustischen Eigenschaften des Strahlers zu verstehen. Aus dieser Notwendigkeit heraus wurden verschiedenste Methoden der akustischen Nahfeldholografie entwickelt, mit denen es möglich ist, das gemessene Schallfeld einer Quelle zu modellieren. Eine dieser Methoden ist die Äquivalenzquellen-Methode (ESM). Dabei wird das lokale Schallfeld als Überlagerung von unterschiedlich starken Elementarquellen betrachtet. Aus räumlich diskreten Schalldruckmessungen lassen sich die Quellstärken bestimmen, die dieses Schallfeld erzeugen. Dies erfordert die Lösung eines linearen inversen Problems.

Inverse Probleme sind meist schlecht gestellt und das resultierende Gleichungssystem häufig unterbesetzt, wodurch der Einsatz von Regularisierungsmaßnahmen zur Voraussetzung wird, um eine sinnvolle Lösung zu erhalten. Die Formulierung zusätzlicher Bedingungen an die Lösung um erwartete räumliche Strukturen zu forcieren stellt eine geeignete Regularisierungsform dar. An Bedingungen geknüpfte inverse Probleme erfordern üblicherweise die Minimierung bestimmter Normen der Lösung im Ortsbereich. Die Anwendung von Compressive Sensing durch Minimierung der L1-Norm erfreut sich immer größerer Beliebtheit, da sich damit gültige Lösungen auch abseits der räumlichen Abtastgrenze bestimmen lassen. Die implizite Forderung nach dünnbesetzten Lösungen beschränkt den Einsatz jedoch auf punktförmige und schwach ausgedehnte Quellen. Typische vibroakustische Quellenphänomene enthalten aber häufig auch räumlich ausgedehnte Muster.

Diese Arbeit untersucht Regularisierungsmethoden, die eine dünnbesetzte Struktur der räumlichen Ableitungen erster und zweiter Ordnung fordern. Dadurch ergeben sich im Ortsbereich abschnittsweise konstante oder lineare Lösungen mit minimaler Krümmung. Solche Regularisierungen sind gängige Praxis in vielen Bereichen der Bildverarbeitung. Sie kamen aber erstmals kürzlich im Bereich der Akustik zum Einsatz, wo sie ihre Fähigkeiten zur Rekonstruktion typischer Strukturen unter Beweis stellen konnten. In dieser Arbeit wird eine Regularisierung vorgeschlagen, die auf den Schatten-Normen der Hesse-Matrix basiert und nach meinem besten Wissen noch nicht für die akustische Nahfeldholografie eingesetzt wurde. Des Weiteren wird eine kombinierte Methode vorgeschlagen, die zusätzlich eine dünnbesetzte Struktur im Ortsbereich fordert. Mit ihr lassen sich sowohl räumlich kompakte als auch ausgedehnte Quellen charakterisieren. Für die Lösung des Minimierungsproblems wird ein Proximal Splitting Algorithmus verwendet. Dieser ermöglicht eine effiziente Implementierung der vorgeschlagenen Regularisierungsmethoden. Dem Leser dieser Arbeit wird das Konzept der ableitungsbasierten Regularisierungsverfahren nahegebracht und deren Charakteristika und Fähigkeiten aufgezeigt. Die Gültigkeit der vorgestellten Methoden wird sowohl durch numerische Simulation, als auch anhand realer Messdaten überprüft. Die Theorie hinter dem Algorithmus wird behandelt und dessen Gebrauch anhand der Probleme dieser Studie erläutert.

Acknowledgements

First and foremost, I want to express my deepest gratitude and appreciation to my supervisor Dr. Franz Zotter for giving me the opportunity to pursue my personal thesis interest in acoustic holography. His exceptional expertise and guidance, his deft editorial touch, along with his constant encouragement and enthusiasm for my research topic were invaluable for completion of this work. He has given me just the right amount of freedom and responsibility to explore my own ideas, while always being available for discussion regardless of time or day, showing a truly inspiring work ethic and dedication to his students.

I am also thankful to all the fellow students I had the pleasure to meet and work with during the course of my degree, who contributed their time and effort to various projects and assignments. I would also like to thank all my friends who provided welcome distractions from work to help me escape my tunnel vision and keep things in perspective. Thank you for always remaining faithful to me despite my limited time and unlimited excuses during the past few years of study. Your friendship means a lot to me and it will never be taken for granted.

My very special thanks go to my beloved family and relatives, who have always believed in me and my abilities and supported me in every aspect of my life. Words cannot express how grateful I am to my parents, Angelika and Hartwig, for their limitless love, support and encouragement to follow my ambitions, and for all the sacrifices they made for me to do so. They have always provided me with anything I ever needed and everything I could ever ask for, constantly trying to make my life easier. Thank you for all you've done for me and for always being there, no matter what!

Last, but by no means least, I am deeply indebted to my girlfriend Valentina for her unconditional love and endless support, which have sustained me throughout the last years. I am forever thankful for all she has done for me during this challenging time and for always showing immense patience and understanding in putting up with moody me and my relentless work schedule. I could not have made it without you and I am truly blessed to have you in my life!

1

Introduction

1.1 Introduction

Nearfield acoustic holography is an effective tool to retrieve the vibro-acoustic properties of complex sources from non-contact measurements of the sound field close to the source. A full description of the source facilitates the reconstruction of the emitted sound field at any given location in space. In industry and research, it has become an established method for localization and characterization of sound sources and the analysis of acoustic phenomena related to their structural vibrations.

With the introduction of the celebrated nearfield acoustic holography (NAH) by Williams and Maynard [DMWL85] in the early 1980's, it was first possible to reconstruct the sound field of planar radiators from array measurements by means of the spatial Fourier transform. Although it offers an elegant solution to acoustic radiation problems, it suffers from a number of inherent deficiencies and restrictions on source and measurement setup. This gave researchers the incentive for developing new holography methods based on numerical or hybrid approaches, all with different advantages and drawbacks. Some of the most prominent methods today are the statistically optimized NAH (SONAH) [Hal09], the inverse boundary element method (IBEM) [KI96, SHBRCH03, VW04], the Helmholtz equation least squares (HELs) [WW97] and the equivalent source method (ESM) [BFK91, Sar05].

Due to the typical ill-posedness of the reconstruction problem, regularization is required to stabilize the solution by imposing additional constraints on it. Traditionally, regularization is accomplished in a least-squares sense by minimizing the ℓ_2 -norm of the solution vector, formally known as ridge regression or Tikhonov regularization [GHO99]. It favors smooth estimates with minimum energy, as it tends to smear the solution over a large number of coefficients. This often yields a poor reconstruction with an over-smoothed result, significantly degrading the spatial resolution. Moreover, it fails at recovering spatially confined or flat regions with sharp edges. As a remedy, variational regularization methods have long time been considered in various fields of image processing [IROF92, BKP10, CCN⁺10, EAS16, Saw11], but they were only recently introduced in acoustics for the purpose of beamforming [XFGG16] and nearfield acoustic holography [FD16, FGD18]. They do not act directly on the solution in the spatial domain, but rather on its derivatives. Since the derivatives quantify the trajectory and curvature of a function's surface, it is possible to impose a certain structure on the solution by forcing its spatial derivatives to take on certain values. By sparsifying first- and second-order spatial derivatives, we respectively obtain piecewise-constant and piecewise-linear solutions that may better conform the physical structure of the source.

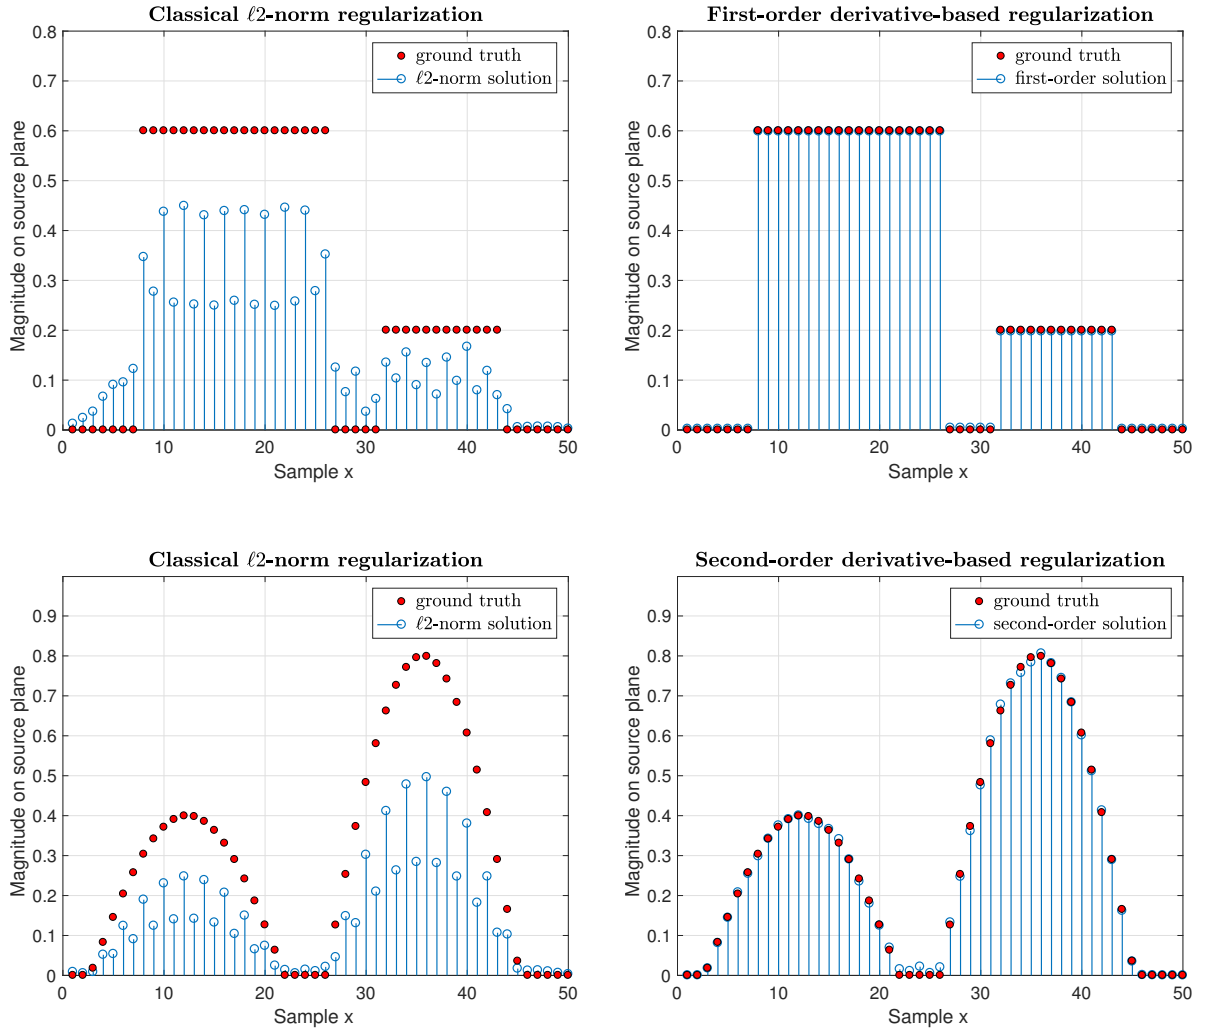


Figure 1.1: Comparison of classical ℓ_2 -norm and derivative-based regularization for a two-dimensional reconstruction scenario of a prescribed normal velocity distribution along a line; Piecewise-constant profile (top), Piecewise-linear profile (bottom).

1.2 Motivation and objective

The great potential of derivative-based regularization methods for sound source reconstruction and their ever-growing popularity and progress in the field of diagnostic imaging was motive to adapt some of their more recent advances to the task of acoustic holography. More precisely, a class of regularizers based on the Schatten-norms of the Hessian [LWU13] is investigated, as they have proven to consistently outperform the classical variational regularization schemes. Moreover, an extension of the regularizers by an additional sparsity constraint on the solution is proposed to further enhance the spatial resolution and to improve the source localization accuracy. In theory, these regularizers can be used for the retrieval of spatially confined and extended sources and their radiated sound fields.

In order to solve the associated optimization problem encountered in the reconstruction process, a suitable algorithm must be found, capable of dealing with the wide spectrum of regularizers examined in this work. In particular, the implementation of the proposed Hessian Schatten-norm regularizers is not straightforward and yet cannot be realized with standard convex optimization toolboxes like CVX [GB14]. Nevertheless, these regularizers can be implemented efficiently by making use of the so-called proximity operator, which has become an increasingly popular optimization tool. It builds the foundation of a relatively new class of optimization algorithms called proximal methods [PB14], which have gained great attention in recent years. The increasing popularity and ongoing developments of these methods in various fields of signal processing were motivation enough to also consider their use for the purpose of nearfield acoustic holography.

The main objective of this work is to review and analyze the most common derivative-based regularization methods and apply them to model the structures of typical acoustic sources. In particular, the superiority of the Hessian Schatten-norm regularizers shall be demonstrated. A comprehensive simulation study and experimental validation is conducted to prove the effectiveness of the proposed regularization techniques. Moreover, the implementation details are presented along with the basic theory behind the algorithm to provide the reader with the tools necessary for the practical application of derivative-based acoustic holography.

2

Acoustic Background

This chapter is dedicated to provide the basic theory relevant for acoustic holography. A derivation of the governing equations of acoustic fields is presented. They are then utilized to derive the famous Kirchhoff-Helmholtz integral, which serves as the fundamental equation for the holography method examined in this work. The following sections intend to offer a simple and compact formulation to provide a fundamental understanding of the acoustic basics. For a more detailed exposition of the governing theory of acoustics or for some more involved derivations, the interested reader is referred to the basic literature, e.g. [Ros07, Fah05]. More specialized treatises of the theory with the main focus on acoustic holography can be found in [BIB13] or [Wil99].

2.1 The Euler equation

The Euler equation describes the relation between the acoustic pressure, the particle velocity and the density of the medium.

Consider an infinitesimal volume element dV , fixed in space and exposed to external forces that are caused by the local pressure field. According to Newton's second law, the net force acting on the differential mass within the volume is given by

$$d\mathbf{F} = d m \mathbf{a} = \rho_0 dV \mathbf{a} = \rho_0 dV \frac{\partial \mathbf{v}}{\partial t}. \quad (2.1)$$

The magnitude of the vector component pointing in x direction can be interpreted as a pressure difference between both sides of the volume perpendicular to the corresponding force

$$\begin{aligned} \mathbf{F}_x &= [p(x, y, z) - p(x + dx, y, z)] dA_{yz} \mathbf{e}_x \\ &= \frac{[p(x, y, z) - p(x + dx, y, z)]}{dx} dx dA_{yz} \mathbf{e}_x = -\frac{\partial p}{\partial x} dV \mathbf{e}_x. \end{aligned} \quad (2.2)$$

By setting the above equations equal to each other, we obtain

$$\frac{\partial p}{\partial x} \mathbf{e}_x = -\rho_0 \frac{\partial \mathbf{v}_x}{\partial t}. \quad (2.3)$$

Further proceeding of the same calculation for the remaining directions leads to the Euler equation

$$\nabla p = -\rho_0 \frac{\partial \mathbf{v}}{\partial t}. \quad (2.4)$$

In a time-harmonic field, the equation transformed to the frequency domain becomes

$$\nabla p = -j\omega\rho_0\mathbf{v}. \quad (2.5)$$

2.2 The homogeneous wave equation

The homogeneous wave equation of an ideal gas can be derived by combining the Euler equation with the continuity equation and the adiabatic equation of state. It describes the propagation of sound waves through a source-free medium.

The equation of state relates the pressure variations p in a gas to a change of the density ρ from its equilibrium value ρ_0 . If this density change is induced by a travelling sound wave, propagated by adiabatic compression and expansion of the gas, the nonlinear adiabatic equation of state $p = p(\rho) = K\rho^\gamma$ holds, where $K = \frac{p_0}{\rho_0^\gamma}$ is a constant and γ denotes the specific heat ratio of the gas. For acoustic phenomena, the density fluctuations are small enough to be linearized using a Taylor series expansion around ρ_0 , neglecting higher-order terms. This leads to

$$p(\rho) = p_0 + \left(\frac{\partial p_0}{\partial \rho_0}\right)(\rho - \rho_0) + \dots \rightarrow dp = p - p_0 = \left(\frac{\partial p_0}{\partial \rho_0}\right)(\rho - \rho_0) = \left(\frac{\partial p_0}{\partial \rho_0}\right)d\rho = B\frac{d\rho}{\rho_0}, \quad (2.6)$$

where $B = \rho_0\left(\frac{\partial p_0}{\partial \rho_0}\right)$ is the bulk modulus of the gas. With the speed of sound $c = \sqrt{\frac{B}{\rho_0}}$ we can define the linearized, adiabatic equation of state as

$$dp = c^2 d\rho. \quad (2.7)$$

The continuity equation describes the transport of a conserved mass m with density ρ due to a velocity field \mathbf{v} . It expresses the conservation of mass inside a volume. The net flux of mass into or out of a fixed control volume V induces a change in the local fluid density. The temporal rate of mass change within the control volume is thus given by

$$\frac{dm}{dt} = \frac{d}{dt} \int_V \rho dV = \int_V \frac{\partial}{\partial t} \rho dV. \quad (2.8)$$

Utilizing Gauss's integral or divergence theorem, the net outflux of mass through the closed boundary surface S of the control volume V due to a velocity field \mathbf{v} can be expressed as

$$\oint_S \rho \mathbf{v} d\mathbf{S} = \oint_S \rho \mathbf{v} \mathbf{n} dS = \int_V \text{div}(\rho \mathbf{v}) dV = \int_V \nabla \cdot (\rho \mathbf{v}) dV, \quad (2.9)$$

where \mathbf{n} denotes the outward normal vector. A positive mass change over time always indicates a positive influx or negative outflux of mass through the volume, hence

$$\int_V \frac{\partial}{\partial t} \rho dV = - \int_V \nabla \cdot (\rho \mathbf{v}) dV \rightarrow \int_V \left[\frac{\partial}{\partial t} \rho dV + \nabla \cdot (\rho \mathbf{v}) dV \right] = 0. \quad (2.10)$$

Since the choice of control volume is arbitrary, the above expression can be generalized, leading

to the continuity equation

$$\frac{\partial}{\partial t}\rho + \nabla \cdot (\rho\mathbf{v}) = 0. \quad (2.11)$$

By taking the time derivative of the continuity equation

$$\frac{\partial^2}{\partial t^2}\rho + \frac{\partial}{\partial t}\nabla \cdot (\rho\mathbf{v}) = 0 \quad (2.12)$$

and the divergence of Euler's equation

$$\nabla \cdot \nabla p = \Delta p = -\rho \frac{\partial}{\partial t}\nabla \cdot \mathbf{v}, \quad (2.13)$$

the linearized homogeneous wave equation can be derived by substitution of these two equations together with the equation of state

$$\left(\Delta - \frac{1}{c^2} \frac{\partial^2}{\partial t^2}\right)p = 0. \quad (2.14)$$

For a time-harmonic dependence, the wave equation can be equivalently written in the frequency domain as

$$(\Delta + k^2)p = 0, \quad (2.15)$$

which is referred to as the homogeneous Helmholtz equation.

2.3 The inhomogeneous wave equation

The inhomogeneous wave equation describes the propagation of acoustic waves in a free space excited by a concentrated point source. Such point source can be considered as a radially oscillating sphere with a radius much smaller compared to the emitted wavelength. The amount of mass released from this source per unit time can be quantified by the source strength or volume velocity

$$Q_S = \frac{d}{dt}V_S = \frac{d}{dt} \frac{m_s}{\rho}. \quad (2.16)$$

Assume a point source located at \mathbf{r}_0 emits mass m_s per time frame Δt into the observation Volume V enclosing the source. According to the law of mass conservation

$$\frac{dm}{dt} = \int_V \frac{\partial}{\partial t}\rho dV = - \int_V \nabla \cdot (\rho\mathbf{v}) dV + \frac{\partial m_S}{\partial t} \rightarrow \int_V \left[\frac{\partial}{\partial t}\rho + \nabla \cdot (\rho\mathbf{v}) - \frac{\partial m_S}{\partial t} \delta(\mathbf{r} - \mathbf{r}_0) \right] dV(\mathbf{r}) = 0 \quad (2.17)$$

the continuity equation for the resulting field gets

$$\frac{\partial}{\partial t}\rho + \nabla(\rho\mathbf{v}) = \frac{\partial m_S}{\partial t}\delta(\mathbf{r} - \mathbf{r}_0). \quad (2.18)$$

The substitution of the time derivative of the continuity equation for a point source excitation into the divergence of the Euler equation, while also considering the adiabatic equation of state, leads to the inhomogeneous wave equation

$$\left(\Delta - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right)p = -\frac{\partial^2}{\partial t^2}m_S\delta(\mathbf{r} - \mathbf{r}_0) = -\rho_0\frac{\partial}{\partial t}Q_S\delta(\mathbf{r} - \mathbf{r}_0). \quad (2.19)$$

Provided that we have time-harmonic dependence, the equation can be reduced to the Helmholtz equation in the frequency domain

$$(\Delta + k^2)p = -j\omega\rho_0Q_S\delta(\mathbf{r} - \mathbf{r}_0). \quad (2.20)$$

2.4 The free-space Green's function

The Green's function is the fundamental solution to the inhomogeneous Helmholtz equation for a concentrated point source. It can be regarded as the spatial response obtained at an arbitrary field point \mathbf{r} , induced by a monopole source located at the source point \mathbf{r}_0 . In the absence of external boundaries surrounding the source, it is referred to as the free-space Green's function. Various derivations can be found in literature, however, a simple and straightforward approach is offered by [RP56] and laid out below.

A prerequisite for the free-space Green's function is that it satisfies the inhomogeneous Helmholtz equation for a point source, as

$$(\Delta + k^2)G(\mathbf{r}, \mathbf{r}_0) = -\delta(\mathbf{r} - \mathbf{r}_0). \quad (2.21)$$

Due to the radially symmetric sound field produced by such point source, the Green's function only depends on the distance $r = \|\mathbf{r} - \mathbf{r}_0\|$ between field point and source point. This is referred to as the symmetry or reciprocity property. For such spherical symmetry the Laplacian Δ in the Helmholtz equation reduces to

$$\Delta G(\mathbf{r}, \mathbf{r}_0) = \frac{1}{r}\frac{d^2(rG(\mathbf{r}, \mathbf{r}_0))}{dr^2}. \quad (2.22)$$

Taking into account that the Green's function implicitly is a solution for the homogeneous Helmholtz equation for any $\mathbf{r} \neq \mathbf{r}_0$, i.e.

$$(\Delta + k^2)G(\mathbf{r}, \mathbf{r}_0) = \frac{1}{r}\frac{d(rG(\mathbf{r}, \mathbf{r}_0))}{dr^2} + k^2G(\mathbf{r}, \mathbf{r}_0) = \frac{d(rG(\mathbf{r}, \mathbf{r}_0))}{dr^2} + k^2(rG(\mathbf{r}, \mathbf{r}_0)) = 0, \quad (2.23)$$

it must have the form of a spherical wave solution that decreases with distance

$$G(\mathbf{r}, \mathbf{r}_0) = \frac{Ae^{jkr} + Be^{-jkr}}{r}. \quad (2.24)$$

Under the assumption of time-harmonic dependence of type $e^{+j\omega t}$, the negative sign in the exponent indicates diverging outgoing waves while the positive sign indicates converging incoming waves. Since for a single point source only diverging solutions are of relevance, we must have $A = 0$ in order for the first term to vanish. The remaining task is to find an expression for the complex amplitude B .

We set up a spherical coordinate system (r, φ, ϑ) with the point source located at its origin and with coordinates varying over the ranges of $-\infty \leq r \leq \infty$ and $0 \leq \varphi, \vartheta \leq \pi$. The Green's function can now be expressed as

$$G(\mathbf{r}, \mathbf{r}_0) = \frac{Be^{-jk|r|}}{|r|}. \quad (2.25)$$

The evaluation of the volume integral of the inhomogeneous Helmholtz equation over a concentric sphere of radius a leads to

$$\int_{-a}^a \int_0^\pi \int_0^\pi \left[\frac{1}{r} \frac{d}{dr^2} r \frac{Be^{-jk|r|}}{|r|} + k^2 \frac{Be^{-jk|r|}}{|r|} \right] r^2 \sin(\varphi) dr d\vartheta d\varphi = \int_V -\delta(\mathbf{r} - \mathbf{r}_0) dV(\mathbf{r}). \quad (2.26)$$

With the sifting property of the Dirac delta function

$$\int_V f(\mathbf{r}) \delta(\mathbf{r} - \mathbf{r}_0) dV(\mathbf{r}) = f(\mathbf{r}_0), \quad \mathbf{r}_0 \in V \quad (2.27)$$

we obtain

$$2\pi B \int_{-a}^a \left[\frac{1}{|r|^2} \frac{d}{dr^2} \left(-e^{-jk|r|} (1 + jk|r|) \frac{d|r|}{dr} \right) + k^2 \frac{e^{-jk|r|}}{|r|} \right] r^2 dr = -1. \quad (2.28)$$

If we consider a vanishingly small sphere of a radius $a \rightarrow 0$ which tends towards zero, we end up with $B = \frac{1}{4\pi}$ for the complex amplitude and consequently with the appropriate expression for the free-space Green's function

$$G(\mathbf{r}, \mathbf{r}_0) = \frac{e^{-jkr}}{4\pi r}. \quad (2.29)$$

2.5 The Kirchhoff Helmholtz integral equation

The Kirchhoff Helmholtz integral equation (KHIE) can be regarded as the mathematical formulation of the Huygen's principle, according to which any sound field can be represented by a superposition of elementary waves generated by a continuous distribution of secondary simple sources. It can be used to determine the sound pressure at any local point of the field produced by a given surface boundary value of either pressure or pressure gradient. The KHIE can be formulated for interior domain problems as well as for exterior domain or radiation problems. Since we are interested in solving the boundary value problem of a volume, where all sound sources and scattering objects are located within the domain of interest, we restrict ourselves to the derivation of the KHIE for exterior domain problems.




Figure 2.1: Illustration of the geometry used for the Kirchhoff-Helmholtz integral equation for exterior domain problems.

The exterior sound field must satisfy the wave equation for any given source distribution $f(\mathbf{r})$ inside the observation volume V . Thus the underlying boundary value problem can be formulated as

$$(\Delta + k^2)p(\mathbf{r}) = f(\mathbf{r}) \quad \text{subject to} \quad p(\mathbf{r}_0) \text{ or } \frac{\partial}{\partial n}p(\mathbf{r}_0), \mathbf{r}_0 \in \mathbf{S}_0, \quad (2.30)$$

where either the pressure $p(\mathbf{r}_0)$ (Dirichlet problem) or the normal pressure gradient $\frac{\partial}{\partial n}p(\mathbf{r}_0)$ (Neumann problem) on the surface boundary is known.

To solve this boundary value problem, we apply Green's second identity

$$\int_V [\phi L\psi - \psi L\phi] dV = \int_S \left[\phi \frac{\partial}{\partial n}\psi - \psi \frac{\partial}{\partial n}\phi \right] dS, \quad (2.31)$$

which is valid for any finite, continuous, twice differentiable scalar fields ϕ and ψ inside the volume V bounded by a surface S . L is a self-adjoint operator¹ and $\frac{\partial}{\partial n} = \nabla \mathbf{n}$ is the directional

¹ L is a self-adjoint operator if $L = L^*$, where L^* denotes the conjugate transpose or Hermitian transpose of L .

derivative with respect to the outward normal \mathbf{n} of the source surface. With the use of the Helmholtz operator as the self-adjoint operator $L = (\Delta + k^2)$ and under the assumption that the two scalar fields $\phi = p(\mathbf{r})$ and $\psi = G(\mathbf{r}, \mathbf{r}_0)$ are solutions of the homogeneous and inhomogeneous Helmholtz equation respectively, we get

$$\begin{aligned} \int_V [\phi L\psi - \psi L\phi] dV &= \int_V [p(\mathbf{r}) \underbrace{(\Delta + k^2)G(\mathbf{r}, \mathbf{r}_0)}_{=-\delta(\mathbf{r}-\mathbf{r}_0)} - G(\mathbf{r}, \mathbf{r}_0) \underbrace{(\Delta + k^2)p(\mathbf{r})}_{=0}] dV(\mathbf{r}) \\ &= - \int_V p(\mathbf{r}) \delta(\mathbf{r} - \mathbf{r}_0) dV(\mathbf{r}) = -p(\mathbf{r}_0) \end{aligned} \quad (2.32)$$

for the volume integral, and

$$\begin{aligned} \int_S \left[\phi \frac{\partial}{\partial n} \psi - \psi \frac{\partial}{\partial n} \phi \right] dS &= \int_{S_0} \left[p(\mathbf{r}) \frac{\partial}{\partial n} G(\mathbf{r}, \mathbf{r}_0) - G(\mathbf{r}, \mathbf{r}_0) \frac{\partial}{\partial n} p(\mathbf{r}) \right] dS(\mathbf{r}) \\ &+ \underbrace{\int_{S_\infty} \left[p(\mathbf{r}) \frac{\partial}{\partial n} G(\mathbf{r}, \mathbf{r}_0) - G(\mathbf{r}, \mathbf{r}_0) \frac{\partial}{\partial n} p(\mathbf{r}) \right] dS(\mathbf{r})}_{=0} \end{aligned} \quad (2.33)$$

for the surface integral of Green's second identity. Due to the Sommerfeld radiation condition, the impact of simple sources located at infinite distance on the local sound field is negligible, thus the surface integral of the outer enclosure S_∞ vanishes. By utilizing the fact that the Green's function is a radially symmetric function that only depends on distance $r = \|\mathbf{r} - \mathbf{r}_0\| = \|\mathbf{r}_0 - \mathbf{r}\|$, we can substitute field point \mathbf{r} and source point \mathbf{r}_0 to eventually end up with the Kirchhoff Helmholtz integral equation (KHIE) for exterior domain problems

$$\alpha p(\mathbf{r}) = \int_{S_0} \left[G(\mathbf{r}, \mathbf{r}_0) \frac{\partial}{\partial n} p(\mathbf{r}_0) - p(\mathbf{r}_0) \frac{\partial}{\partial n} G(\mathbf{r}, \mathbf{r}_0) \right] dS(\mathbf{r}_0), \quad (2.34)$$

with

$$\alpha = \alpha(\mathbf{r}) = \begin{cases} 1 & \text{for } \mathbf{r} \text{ outside } S_0 \\ \frac{1}{2} & \text{for } \mathbf{r} \in S_0 \\ 0 & \text{for } \mathbf{r} \text{ inside } S_0. \end{cases} \quad (2.35)$$

With the aid of the KHIE, the total acoustic pressure at any field point \mathbf{r} in the exterior domain or on the surface can be expressed as the superposition of field responses due to fictitious monopole $G(\mathbf{r}, \mathbf{r}_0)$ and dipole sources $\frac{\partial}{\partial n} G(\mathbf{r}, \mathbf{r}_0)$ distributed on the boundary of the surface. The first term of the integrand represents a monopole surface distribution, weighted by the normal pressure gradient $\frac{\partial}{\partial n} p(\mathbf{r}_0)$ and thus proportional to the normal acceleration of the surface. It represents the contribution of the sound generated by the normal motion of the body to the acoustic field. The second term represents a dipole surface distribution, weighted by the sound pressure $p(\mathbf{r}_0)$. The dipole sources are composed of two monopole sources of equal strength but opposite phase, located closely together and in line with the surface normal \mathbf{n} . It represents the interaction between the body as an acoustic obstacle and the sound field generated by itself and also by any other existing sources in space. The monopole and dipole distributions can be regarded as a single and double layer potential on the boundary, respectively.




Figure 2.2: Monopole and dipole representation of different boundary conditions in the presence of external sources.

In many real-world scenarios only one of the two boundary conditions is accessible. In this case, the underlying problem is called Neumann problem if the boundary pressure gradient $\frac{\partial}{\partial n}p(\mathbf{r}_0)$ is known, and Dirichlet problem if the boundary pressure $p(\mathbf{r}_0)$ is known. To cope with these types of problems, the free-space Green's function must be slightly modified to force the unknown quantity to become zero at the surface, while still satisfying the KHIE. For simple geometric structures we can apply a special type of Green's function being comprised of two terms that satisfy the inhomogeneous and homogeneous Helmholtz equation, respectively:

$$G_{BC}(\mathbf{r}, \mathbf{r}_0) = G(\mathbf{r}, \mathbf{r}_0) + H(\mathbf{r}, \mathbf{r}_0). \quad (2.36)$$

Rayleigh's first integral

For Neumann problems, where only the boundary pressure gradient $\frac{\partial}{\partial n}p(\mathbf{r}_0)$ is known, the Neumann boundary condition is imposed

$$\frac{\partial}{\partial n}G_N(\mathbf{r}, \mathbf{r}_0)|_{\mathbf{r}_0 \in \mathbf{S}} \stackrel{!}{=} 0. \quad (2.37)$$

As a consequence, the required boundary condition for $H(\mathbf{r}, \mathbf{r}_0)$ gets

$$\frac{\partial}{\partial n}H(\mathbf{r}, \mathbf{r}_0)|_{\mathbf{r}_0 \in \mathbf{S}} \stackrel{!}{=} -\frac{\partial}{\partial n}G(\mathbf{r}, \mathbf{r}_0)|_{\mathbf{r}_0 \in \mathbf{S}}. \quad (2.38)$$

In this case the boundary can be regarded as an acoustically rigid or sound hard surface. The dipole term of the integrand vanishes and the KHIE simplifies to Rayleigh's first integral

$$p(\mathbf{r}) = \int_{S_0} G_N(\mathbf{r}, \mathbf{r}_0) \frac{\partial}{\partial n}p(\mathbf{r}_0) dS(\mathbf{r}_0). \quad (2.39)$$

The remaining task is to find an expression for $H(\mathbf{r}, \mathbf{r}_0)$ that satisfies both the homogeneous Helmholtz equation and the required boundary condition. For simple geometries we can apply the mirror image source method for evaluation.

A monopole source is placed at position \mathbf{r}_0 above an infinite rigid surface, producing a sound field $G(\mathbf{r}, \mathbf{r}_0)$ in the free half-space $z > 0$ that is reflected by the boundary at $z = 0$. This reflection can be modeled by a sound field $H(\mathbf{r}, \mathbf{r}'_0)$ emitted by a mirror image source component located at position \mathbf{r}'_0 underneath the surface at $z < 0$. The local sound field $G(\mathbf{r}, \mathbf{r}_0)$ evaluated at an arbitrary field point \mathbf{r} above the plane is now comprised of the sound fields produced by the primary monopole source and the mirror image source. Thus the resulting Green's function for the problem becomes

$$G_N(\mathbf{r}, \mathbf{r}_0) = G(\mathbf{r}, \mathbf{r}_0) + H(\mathbf{r}, \mathbf{r}'_0) = \frac{e^{-jkr}}{4\pi r} + \frac{e^{-jkr'}}{4\pi r'}, \quad (2.40)$$

where

$$G_N(\mathbf{r}, \mathbf{r}_0)|_{\mathbf{r}_0 \in \mathbf{S}} = \frac{e^{-jkr}}{2\pi r}. \quad (2.41)$$

For this type of Green's function, the normal gradient becomes zero as we place the primal source, and consequently its image, at the boundary, such that $z_0 = 0$ and $r = r'$:

$$\begin{aligned} \frac{\partial}{\partial n} G_N(\mathbf{r}, \mathbf{r}_0)|_{\mathbf{r}_0 \in \mathbf{S}} &= \frac{\partial}{\partial z_0} \left(\frac{e^{-jkr}}{4\pi r} + \frac{e^{-jkr'}}{4\pi r'} \right) \Big|_{\mathbf{r}_0 \in \mathbf{S}} = \left[\frac{\partial r}{\partial z_0} \frac{\partial}{\partial r} \frac{e^{-jkr}}{4\pi r} + \frac{\partial r'}{\partial z_0} \frac{\partial}{\partial r'} \frac{e^{-jkr'}}{4\pi r'} \right] \Big|_{z_0=0} \\ &= \left[\frac{(z_0 - z)}{r} \frac{\partial}{\partial r} \frac{e^{-jkr}}{4\pi r} + \frac{(z_0 + z)}{r'} \frac{\partial}{\partial r'} \frac{e^{-jkr'}}{4\pi r'} \right] \Big|_{z_0=0} = \left(\frac{-z}{r} + \frac{z}{r} \right) \frac{\partial}{\partial r} \frac{e^{-jkr}}{4\pi r} = 0. \end{aligned} \quad (2.42)$$

Rayleigh's second integral

For Dirichlet problems, where only the pressure at the boundary $p(\mathbf{r}_0)$ is known, the Dirichlet boundary condition is imposed

$$G_D(\mathbf{r}, \mathbf{r}_0)|_{\mathbf{r}_0 \in \mathbf{S}} \stackrel{!}{=} 0. \quad (2.43)$$

Thus the required boundary condition for $H(\mathbf{r}, \mathbf{r}_0)$ gets

$$H(\mathbf{r}, \mathbf{r}_0)|_{\mathbf{r}_0 \in \mathbf{S}} \stackrel{!}{=} -G(\mathbf{r}, \mathbf{r}_0)|_{\mathbf{r}_0 \in \mathbf{S}}. \quad (2.44)$$

The Dirichlet boundary can be regarded as a sound soft surface with low acoustic impedance compared to the surrounding medium. In this case, the monopole term of the integrand vanishes and the KHIE simplifies to Rayleigh's second integral

$$p(\mathbf{r}) = \int_{S_0} \frac{\partial}{\partial n} G_D(\mathbf{r}, \mathbf{r}_0) p(\mathbf{r}_0) dS(\mathbf{r}_0). \quad (2.45)$$

Again, an appropriate expression for $\frac{\partial}{\partial n} H(\mathbf{r}, \mathbf{r}_0)$ must be found that fulfills the homogeneous Helmholtz equation and the required boundary condition. Once more, the image source method can be applied to find a suitable solution.

Following the same calculation path as above, but with a dipole primary and mirror source, the normal gradient of the Green's function for the combined local field becomes

$$\frac{\partial}{\partial n} G_D(\mathbf{r}, \mathbf{r}_0) = \frac{\partial}{\partial n} [g(\mathbf{r}, \mathbf{r}_0) + H(\mathbf{r}, \mathbf{r}'_0)] = \frac{\partial}{\partial n} \left[\frac{e^{-jkr}}{4\pi r} - \frac{e^{-jkr'}}{4\pi r'} \right] = \left[\frac{\partial r}{\partial z_0} \frac{\partial}{\partial r} \frac{e^{-jkr}}{4\pi r} - \frac{\partial r'}{\partial z_0} \frac{\partial}{\partial r'} \frac{e^{-jkr'}}{4\pi r'} \right]. \quad (2.46)$$

When taking the primal source, and consequently its mirrored image, to the boundary layer, the expression becomes

$$\begin{aligned} \frac{\partial}{\partial n} G_D(\mathbf{r}, \mathbf{r}_0)|_{\mathbf{r}_0 \in \mathbf{S}} &= \left[\frac{(z_0 + z)}{r} \frac{\partial}{\partial r} \frac{e^{-jkr}}{4\pi r} - \frac{(z_0 - z)}{r'} \frac{\partial}{\partial r'} \frac{e^{-jkr'}}{4\pi r'} \right] \Big|_{z_0=0} \\ &= 2 \frac{z}{r} \frac{\partial}{\partial r} \frac{e^{-jkr}}{4\pi r} = \frac{z}{2\pi r} \left[-\frac{1}{r^2} e^{-jkr} + \frac{jk}{r} e^{-jkr} \right] = -z \frac{e^{-jkr}}{2\pi r^3} (jkr + 1) = 2 \cos(\varphi) \frac{e^{-jkr}}{4\pi r} \left(-\frac{1}{r} - jk \right). \end{aligned} \quad (2.47)$$

Based on this acquired knowledge, we are now able to derive the Equivalent Source Method that is used throughout this theses to model the radiation of arbitrary sources.

3

Sound Field Representation with the Equivalent Source Method

3.1 The Equivalent Source Method (ESM)

The reconstruction and visualization of the spatial distribution of different acoustic quantities in close proximity to the source is often crucial for a better understanding of acoustic phenomena related to its radiation of sound. This gave rise to the development of several imaging techniques for the analysis of the sound field, one of them being the nearfield acoustic holography (NAH). It is based on discrete measurements of the acoustic pressure in the nearfield of the source under study and the back-propagation of the sound field to the requested source surface. Due to the evanescent acoustic field captured close to the source, the hologram can still be recovered with high spatial resolution. NAH provides a way to assess different quantities of the sound field at positions different from the actual measurements.

There are several NAH methods, including Fourier-based NAH, the inverse boundary element method (IBEM), methods based on basis function expansion, such as the Helmholtz equation least squares (HELs), and many others, all with their own advantages and drawbacks. The so-called Equivalent Source Method (ESM), also referred to as the wave superposition method or method of fundamental solutions [Lee16], is a common alternative to the classical IBEM. The basic principle of ESM is that the sound field radiated by any complex sound source can be substituted by the superposition of the sound fields generated by a distribution of auxiliary sources. The main task is to find a suitable estimation of the equivalent source strengths in terms of amplitude and phase. Once the source strengths are determined by solving an inverse problem, the sound field can be retrieved. The equivalent sources can be expressed as monopoles, dipoles or even higher-order multipoles and they are always placed within the interior domain of the sound source, usually retracted by some amount from the surface. This retraction circumvents the singularity at the boundary as the distance to the equivalent sources is never zero. Assuming a suitable source layout, ESM ensures great accuracy of the solution without involving complex numerical interpolation and integration methods. [Per13, BIB13] The simplicity of ESM permits low implementation effort and computational cost, and for the requirements in this thesis it is therefore the best choice.

The following section provides a compact formulation of ESM, as well as some useful insights on the underlying inverse problem. The impact of different distribution settings of the equivalent sources in terms of number and retreat distance is discussed and some suggestions are made. Moreover, the section provides an introduction of regularization methods to overcome the typical ill-conditioning of the ESM source strength retrieval.

3.1.1 Formulation of ESM




Figure 3.1: Typical single layer ESM setup for a planar source (left); Exemplary equivalent source configurations (right).

Assume that the acoustic pressure of a sound field is measured at M array microphones on the measurement surface S_M . According to ESM [BIB13, Per13], the sound pressure at a specified microphone location \mathbf{r} can be considered as the superposition of sound waves radiated by a continuous distribution of elementary sources at the equivalent source surface S_0 , which may be retracted from the actual source surface S . As a major simplification only a single layer potential is considered, reducing the KHIE to

$$p(\mathbf{r}) = j\omega\rho \int_{S_0} q(\mathbf{r}_0)G(\mathbf{r}, \mathbf{r}_0)dS_0(\mathbf{r}_0), \quad (3.1)$$

where $q(\mathbf{r}_0)$ represents the unknown source strength of the equivalent source located at \mathbf{r}_0 . In practice, equation (3.1) is discretized as

$$p(\mathbf{r}) \approx j\omega\rho \sum_{n=1}^N [q(\mathbf{r}_{0,n})\Delta S_0(\mathbf{r}_{0,n})]G(\mathbf{r}, \mathbf{r}_{0,n}) = \sum_{n=1}^N q_n G(\mathbf{r}, \mathbf{r}_{0,n}), \quad (3.2)$$

where $q_n = Q_n = v_n dS_0$ is the source strength of the n -th equivalent source, and N is the number of sources. For a condensed notation, the frequency dependency term $j\omega\rho$ can be incorporated either in the source strength or in the Green's function. The latter is more convenient as the term is already known and need not to be estimated together with the source strengths. In matrix notation, the resulting system of linear equations can be expressed as

$$\mathbf{p} = \mathbf{G}\mathbf{q} + \mathbf{n}, \quad (3.3)$$

where $\mathbf{p} = [p(\mathbf{r}_1), p(\mathbf{r}_2), \dots, p(\mathbf{r}_M)]^T \in \mathbb{C}^M$ is a vector containing the M discrete-space sound pressure samples, $\mathbf{q} = [q(\mathbf{r}_{0,1}), q(\mathbf{r}_{0,2}), \dots, q(\mathbf{r}_{0,N})]^T \in \mathbb{C}^N$ is the coefficient vector consisting of the unknown strengths of the N equivalent sources and $\mathbf{G} \in \mathbb{C}^{M \times N}$ is the propagation matrix

whose elements are given by the free-field Green's functions associating the source strengths with the sound pressure samples, such that

$$\mathbf{G} = \begin{bmatrix} G(\mathbf{r}_1, \mathbf{r}_{0,1}) & G(\mathbf{r}_1, \mathbf{r}_{0,2}) & \dots & G(\mathbf{r}_1, \mathbf{r}_{0,N}) \\ G(\mathbf{r}_2, \mathbf{r}_{0,1}) & G(\mathbf{r}_2, \mathbf{r}_{0,2}) & \dots & G(\mathbf{r}_2, \mathbf{r}_{0,N}) \\ \vdots & \vdots & \ddots & \vdots \\ G(\mathbf{r}_M, \mathbf{r}_{0,1}) & G(\mathbf{r}_M, \mathbf{r}_{0,2}) & \dots & G(\mathbf{r}_M, \mathbf{r}_{0,N}) \end{bmatrix}. \quad (3.4)$$

In addition, the corruption of the measurements with additive white Gaussian noise is considered in form of the vector $\mathbf{n} \in \mathbb{C}^M$. The remaining task is to retrieve the equivalent source strengths (inverse problem), leading to a solution vector $\hat{\mathbf{q}}$ that can be used to reconstruct different quantities of the sound field at any desired point in space (forward problem). However, the solution of equation (3.3) critically depends on the extent of both the measurement and model layout.

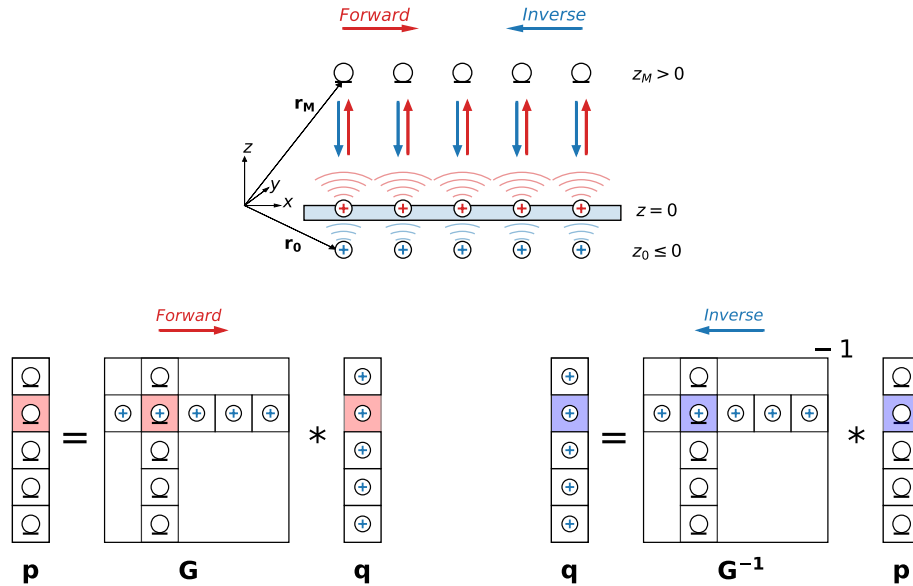
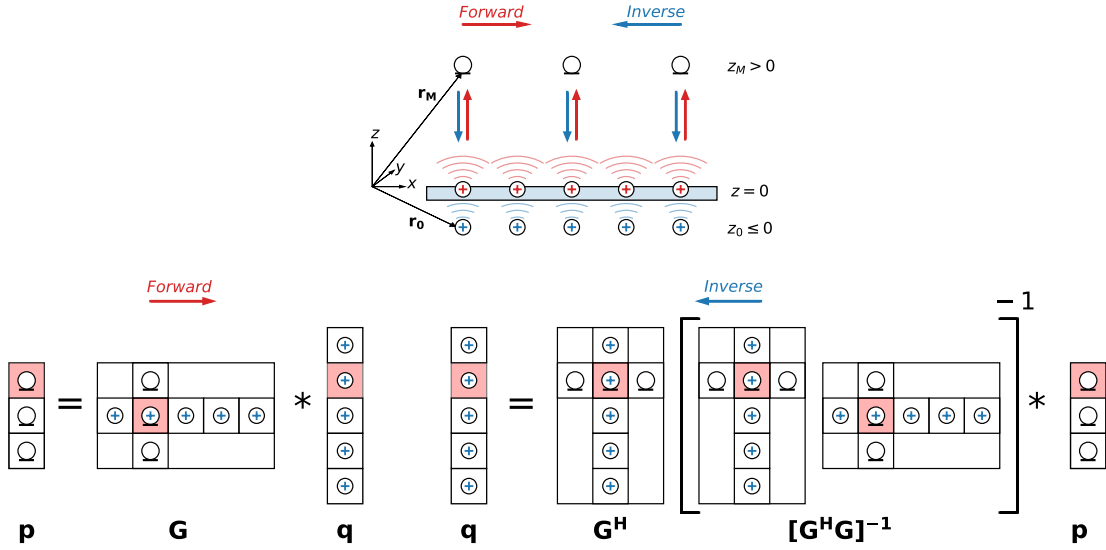


Figure 3.2: Exactly determined system ($M=N$).

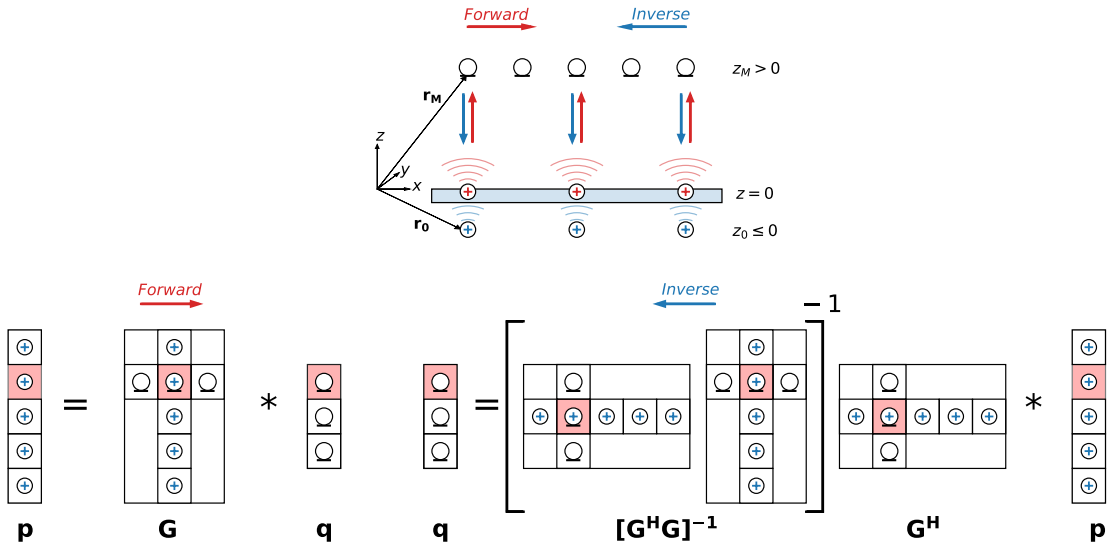
Every linear equation system either has a unique solution, an infinite number of solutions or no solution at all. For an *exactly determined* system ($M=N$), where the number of microphone positions matches the number of equivalent sources, \mathbf{G} becomes a square matrix and the inverse \mathbf{G}^{-1} can be computed, provided that the matrix is regular or non-singular. In this case, a unique solution is obtained by

$$\hat{\mathbf{q}} = \mathbf{G}^{-1}\mathbf{p}. \quad (3.5)$$


 Figure 3.3: Overdetermined system ($M > N$).

When the number of measurements exceeds the number of auxiliary sources, the system becomes *overdetermined* ($M > N$). Since a linear system with more equations than unknowns usually does not provide a solution, the best way to proceed is to search for an estimate of the solution associated with the minimum residual error in a least-squares sense. This requires solving the following optimization problem to minimize the sum of squares of the residual error where the solution is obtained by the left inverse of \mathbf{G} as

$$\hat{\mathbf{q}} = \arg \min_{\mathbf{q}} \|\mathbf{p} - \mathbf{G}\mathbf{q}\|_2^2 = (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{G}^H \mathbf{p}. \quad (3.6)$$


 Figure 3.4: Underdetermined system ($M < N$).

However, the systems encountered in practical implementation of ESM are typically *underdetermined* as the number of measurement points is usually far less than the number of equivalent sources ($M < N$). In this case, an infinite number of possible solutions can be discovered, hence such system provides no unique solution. It commonly requires to seek the solution subspace for the one solution that best fits an additional condition, usually given as some minimum norm. The most preferred choice is to find the solution that exhibits minimum energy which is accompanied by minimizing the squared ℓ_2 -norm. This can be cast into the following equality constrained minimization problem

$$\hat{\mathbf{q}} = \arg \min_{\mathbf{q}} \|\mathbf{q}\|_2^2 \quad \text{subject to} \quad \mathbf{p} = \mathbf{G}\mathbf{q}. \quad (3.7)$$

One way of solving (3.7) is to rephrase it into an unconstrained optimization problem by introducing the so-called Lagrangian

$$L(\mathbf{q}, \lambda) = \mathbf{q}^H \mathbf{q} + \lambda^H (\mathbf{G}\mathbf{q} - \mathbf{p}) \quad (3.8)$$

as the new cost function to minimize, where λ^H denotes the Lagrange multiplier. The necessary condition for $L(\hat{\mathbf{q}}, \lambda)$ to be a minimum is a vanishing complex gradient with respect to \mathbf{q} , hence it must hold that

$$\nabla_{\mathbf{q}} L(\mathbf{q}, \lambda) = 2\mathbf{q} + \mathbf{G}^H \lambda = \mathbf{0}. \quad (3.9)$$

It remains to determine the unique Lagrange multiplier λ that satisfies the optimality equation. Expanding the equation by the matrix \mathbf{G} followed by a substitution with the equality constraint $\mathbf{p} = \mathbf{G}\mathbf{q}$ results in

$$\begin{aligned} 2\mathbf{G}\mathbf{q} + \mathbf{G}\mathbf{G}^H \lambda &= 2\mathbf{p} + \mathbf{G}\mathbf{G}^H \lambda = \mathbf{0} \\ \lambda &= -2(\mathbf{G}\mathbf{G}^H)^{-1} \mathbf{p}. \end{aligned} \quad (3.10)$$

Solving the optimality equation (3.9) with the unique Lagrange multiplier of equation (3.10) leads to the classical least-norm solution of the underdetermined system

$$\hat{\mathbf{q}} = \mathbf{G}^H (\mathbf{G}\mathbf{G}^H)^{-1} \mathbf{p} \quad (3.11)$$

involving the right inverse of \mathbf{G} .

3.1.2 Recovery of acoustic quantities of the sound field

Based on the estimated source strengths $\hat{\mathbf{q}}$ of the equivalent sources, the entire sound field can be reconstructed at any given location in the exterior of the source [FGD18]. The recovery of the sound field requires the use of the reconstruction matrix $\mathbf{G}_{\mathbf{S}}$, whose elements are the free-space Green's functions $G(\mathbf{r}_{\mathbf{S}}, \mathbf{r}_{\mathbf{0}})$. They describe the propagation of sound waves from the source point $\mathbf{r}_{\mathbf{0}}$ to an arbitrary field point $\mathbf{r}_{\mathbf{S}}$ at the recovery surface. The reconstruction matrix can be

defined as

$$\mathbf{G}_S = \begin{bmatrix} G(\mathbf{r}_{S,1}, \mathbf{r}_{0,1}) & G(\mathbf{r}_{S,1}, \mathbf{r}_{0,2}) & \dots & G(\mathbf{r}_{S,1}, \mathbf{r}_{0,N}) \\ G(\mathbf{r}_{S,2}, \mathbf{r}_{0,1}) & G(\mathbf{r}_{S,2}, \mathbf{r}_{0,2}) & \dots & G(\mathbf{r}_{S,2}, \mathbf{r}_{0,N}) \\ \vdots & \vdots & \ddots & \vdots \\ G(\mathbf{r}_{S,M}, \mathbf{r}_{0,1}) & G(\mathbf{r}_{S,M}, \mathbf{r}_{0,2}) & \dots & G(\mathbf{r}_{S,M}, \mathbf{r}_{0,N}) \end{bmatrix}. \quad (3.12)$$

Note that there is no restriction on the number of reconstruction points M , allowing a recovery surface with enhanced spatial resolution. The reconstructed acoustic pressure on the recovery surface becomes

$$\mathbf{p}_S = \mathbf{G}_S \hat{\mathbf{q}}. \quad (3.13)$$

Several other physical quantities of the sound field can be deduced from the source strength, including the particle velocity \mathbf{v}_n normal to the surface, the normal sound intensity \mathbf{I}_n and from that, the total radiated power \mathbf{P} of the source. The normal velocity can be derived from Euler's equation (2.5) as

$$\mathbf{v}_{S,n} = -\frac{1}{j\omega\rho} \nabla_n \mathbf{p}_S = -\frac{1}{j\omega\rho} \frac{\partial \mathbf{G}_S}{\partial n} \hat{\mathbf{q}}, \quad (3.14)$$

where the normal gradient $\nabla_n = \nabla \cdot \mathbf{n} = \frac{\partial}{\partial n}$ is applied elementwise on the free-space Green's functions, whose normal derivative is given by

$$\frac{\partial G(\mathbf{r}_S, \mathbf{r}_0)}{\partial n} = 2 \cos(\varphi) \frac{e^{-jk r}}{4\pi r} \left(-\frac{1}{r} - jk\right), \quad (3.15)$$

where $r = \|\mathbf{r}_S - \mathbf{r}_0\|$ denotes the distance between source and reconstruction point.

The normal component of the sound intensity depends on the real part of the product of sound pressure and particle velocity at the hologram surface. It is always oriented along the particle velocity vector, pointing in positive or negative direction depending on the resulting sign. It can be obtained by

$$\mathbf{I}_{S,n} = \frac{1}{2} \text{Re} \left\{ \mathbf{p}_S \odot \mathbf{v}_{S,n}^* \right\}, \quad (3.16)$$

where $\text{Re}\{\cdot\}$ gives the real component of its argument, the raised asterisk $*$ marks the complex conjugate and \odot denotes a pointwise multiplication (Hadamard/Schur product) of the two vectors.

The radiated sound power is calculated by taking the integral of the estimated normal intensity over a closed, imaginary surface S_∞ surrounding the source, under the assumption of a source-free environment

$$\mathbf{P} = \int_{S_\infty} \mathbf{I}_{S,n}(\mathbf{r}_S) dS_\infty. \quad (3.17)$$

3.2 Singular value decomposition

One of the most powerful tools from linear algebra is the Singular Value Decomposition (SVD) of a matrix as it helps characterizing the underlying problem. Any square or rectangular matrix $\mathbf{G} \in \mathbb{C}^{M \times N}$ can be factorized into the product of three matrices

$$\mathbf{G} = \mathbf{U}\mathbf{S}\mathbf{V}^H, \quad (3.18)$$

where $\mathbf{U} \in \mathbb{C}^{M \times M}$ and $\mathbf{V} \in \mathbb{C}^{N \times N}$ are unitary matrices

$$\mathbf{U}\mathbf{U}^H = \mathbf{U}^H\mathbf{U} = \mathbf{I}, \quad \mathbf{V}\mathbf{V}^H = \mathbf{V}^H\mathbf{V} = \mathbf{I}, \quad (3.19)$$

whose columns are comprised of the left and right singular vectors of \mathbf{G} , respectively. The matrix $\mathbf{S} \in \mathbb{R}^{M \times N}$ is a diagonal matrix where the diagonal elements contain the positive, real-valued singular values σ of \mathbf{G} in descending order

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_R \geq 0, \quad \sigma_{R+1} = \dots = \sigma_s = 0, \quad (3.20)$$

where $s = \min(M, N)$ and $R = \text{rank}(\mathbf{G})$ is defined as the number of linearly independent rows or columns in \mathbf{G} .

To provide a physical interpretation of the decomposition, the left and right singular vectors can be understood as sets of discrete orthogonal basis functions describing the sound field [NK01]. The left singular values \mathbf{u}_i can be regarded as a set of individual field modes characterizing the spatial structure or shape of the emitted acoustic pressure field at the observation surface. The right singular values \mathbf{v}_i are linked to the corresponding set of source modes specifying the spatial variation of the equivalent source distribution. Due to orthogonality, each pair of left and right singular vectors independently and uniquely propagate a specific source mode to a specific field mode. The propagation from a specific source to its field mode is determined by their singular value σ_i . From this perspective, the propagation matrix \mathbf{G} can be formulated as the linear superposition of decoupled weighted rank-one matrices consisting of a pair of field and source modes weighted by their propagation strength

$$\mathbf{G} = \mathbf{U}\mathbf{S}\mathbf{V}^H = \begin{bmatrix} \mathbf{u}_1 & \mathbf{u}_2 & \dots & \mathbf{u}_s \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & \dots & 0 \\ 0 & \sigma_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_s \end{bmatrix} \begin{bmatrix} \mathbf{v}_1^H \\ \mathbf{v}_2^H \\ \vdots \\ \mathbf{v}_s^H \end{bmatrix} = \sum_{i=1}^s \sigma_i \mathbf{u}_i \mathbf{v}_i^H. \quad (3.21)$$

The decomposed matrix can be used to express the governing equation system in (3.3) as

$$\mathbf{p} = \mathbf{U}\mathbf{S}\mathbf{V}^H \mathbf{q}. \quad (3.22)$$

Left-multiplying by \mathbf{U}^H and using the orthogonality property of equation (3.19) structurally shows the relation above, the source strength decomposed into source modes $\mathbf{V}^H \mathbf{q}$ and the sound pressure decomposed into field modes $\mathbf{U}^H \mathbf{p}$ are linked by the singular values (propagation strengths) in the diagonal matrix \mathbf{S} as

$$\mathbf{U}^H \mathbf{p} = \underbrace{\mathbf{U}^H \mathbf{U}}_{\mathbf{I}} \mathbf{S} \mathbf{V}^H \mathbf{q} = \mathbf{S} \mathbf{V}^H \mathbf{q}. \quad (3.23)$$

The application of SVD naturally yields the solution with least-squared error in case of an exactly or overdetermined system ($M \geq N$) and provides the least-norm solution for an underdetermined system ($M < N$), where the corresponding (pseudo-)inverse of \mathbf{G} is defined as

$$\mathbf{G}^{-1} = \mathbf{V}\mathbf{S}^{-1}\mathbf{U}^H = \sum_{i=1}^R \frac{\mathbf{v}_i\mathbf{u}_i^H}{\sigma_i}. \quad (3.24)$$

3.3 Propagating and evanescent waves

The sound field near a vibro-acoustic source consists of two types of waves: propagating waves and evanescent waves. Both types carry information about the surface vibrations, thus both need to be captured for the reconstruction process in order to end up with a truthful representation of the source. While propagating waves radiate away from their origin and transmit energy to the farfield, evanescent waves only exist in the immediate vicinity of the source, propagating in parallel to its surface and exponentially decay with distance. These waves can be best described by utilizing the concept of the traditional Fourier acoustic [Wil99], which decomposes the sound field into a distribution of traveling plane waves and evanescent waves with different complex amplitudes and directions. This requires a reformulation of the problem in the wavenumber domain or *k-space*, which represents the spatial frequency domain. Each point in *k-space* represents a specific wave component in the spatial domain, traveling at a certain frequency and amplitude in a certain direction to a reference plane.

Consider a single, monofrequent plane wave traveling in *x*-direction (no variation along the *y*-direction) on an infinite plate located in the reference plane (*xy*-plane), causing radiation in *z*-direction. Depending on the wavenumber $k_x = \frac{\omega}{c_x} = \frac{2\pi}{\lambda_x}$ of the plane wave on the plate, a diverging wave of fixed wavenumber $k = \frac{\omega}{c} = \frac{2\pi}{\lambda}$ is either propagating or evanescent. At wavenumbers $k_x < k$, the wavelength λ_x on the plate is bigger than the acoustic wavelength λ in the surrounding medium. In this case, a propagating plane wave is emitted at a certain elevation angle, determined by their wavenumber relation. This type of wave is also referred to as *supersonic wave*, since its causing wave on the plate travels at higher speed than the speed of sound in the medium. When $k_x > k$, the wavelength on the plate becomes smaller than the one in the medium and the propagating wave turns into an evanescent wave. The amplitude of this wave fades out exponentially with increasing distance to the plate. This type of wave is also called *subsonic wave*, as its inducing plane wave on the surface moves slower than the acoustic wave. The wavenumber at which surface vibration and acoustic wave are of equal wavelength and travel in phase at same speed and in same direction ($k_x = k$) is called the *coincidence frequency* (or wavenumber). It marks the transition point between propagating and evanescent waves, where the acoustic wave travels in parallel to the source plane. Figure 3.5 illustrates the sound field of a vibrating plate for these three particular cases.




Figure 3.5: Illustration of propagating and evanescent waves.

The physical explanation for the existence of evanescent waves is the fact that all solid media are dispersive, hence the speed at which a wave travels through the medium depends on its frequency. Bending waves are most responsible for sound field excitation as they lead to a normal deflection of the surface. The speed of the induced sound wave in air most likely differs from the speed of the bending wave, yielding different wavelengths at the same frequency. The amount of deviation and the location of the coincidence frequency depends on the material properties and geometric dimensions of the vibrating structure. As already stated, wave propagation to the farfield only occurs when the wavelength of the surface wave exceeds the one in air. There is no sound radiation due to the bending waves, when their wavelengths become smaller than that in air, resulting in a hydrodynamic short circuit due to an increasing impedance mismatch. The motion of air particles is limited to the nearfield of the source, as they flow back and forth between vibrational regions of opposite phase, while particles in the farfield do not move at all. These nearfield waves are representing the real evanescent waves that contain high spatial frequency information about the surface vibration.

3.4 Ill-posedness of the inverse problem

The holographic problem addressed in this thesis deals with the inversion of an underdetermined system and its regularization.

For any inverse problem, Hadamard defines the criteria for a well-posed problem: (i) a solution exists, (ii) the solution is unique, (iii) the solution changes continuously with continuous changes of the initial conditions.

Any system of equations may contain conditions contradicting the equations, which would turn the problem into an ill-posed one, as no solution can exist.

Assuming the problem concerns the minimization of a cost function, as in our case, then the problem can be underdetermined as soon as not all degrees of freedom of the solution affect the cost function minimized. Then there exist many solutions, and the problem is ill-posed. In this case, regularization involving all degrees of freedom is helpful.

Moreover, non-convex minimization problems can have multiple minima, i.e. solutions, which would make them ill-posed; we deal with convex formulations here.

Even if a unique solution exists (i)+(ii), the problem can be (numerically) discontinuous with regard to the initial conditions. In our scenario, this would be the case if the problem is numerically ill-conditioned. Given the model matrix \mathbf{G} and its SVD in (3.21), the spectral decomposition theorem states that the norm of its output vector $\mathbf{p} = \mathbf{G}\mathbf{q}$ lies between the largest and smallest

singular value for a normalized input \mathbf{q} , such that

$$\sigma_{min} \leq \left\| \frac{\mathbf{G}\mathbf{q}}{\|\mathbf{q}\|_2} \right\| \leq \sigma_{max}. \quad (3.25)$$

The ill-conditioning of a system therefore typically uses the condition number

$$\text{cond}(\mathbf{G}) = \frac{\sigma_{max}}{\sigma_{min}} \quad (3.26)$$

that measures which numerical range is covered by the norm of the possible output vectors. It hereby measures the continuity of output vectors depending on changes of the input vector.

It becomes apparent that with an ill-conditioned matrix \mathbf{G} with a large condition number, even the presence of small perturbations in the acquired data due to measurement noise or numerical error may yield an excessive boost of errors in the retrieved data. The condition number therefore provides an upper bound on the amplification of measurement noise. A physical explanation about the inherent ill-conditioning of the inverse problem is the absence of evanescent waves in the measurements. These waves only appear within the nearfield of the source, as they exponentially decay with increasing distance. Actually, all of the evanescent source modes are required to retrieve detailed structural information about the source. To avoid excessive loss in the propagation of evanescent components, measurements are required to be taken in close proximity of the source, with the geometric accuracy, gain matching and signal-to-noise ratio (SNR) of the sensors as the limiting factors. Ill-conditioned measurement layouts cause the associated propagators' singular values close to zero.

To overcome numerical stability problems, regularization is a powerful tool. Typically, regularization introduces additional constraints on the solution, such that the modified problem no longer suffers from ill-posedness. In fact, regularization hereby characterizes criteria that are considered to yield a valid solution. Regarding regularization as a rather open concept, it can essentially prescribe a criterion for preferable solutions not only in norm but also in shape. This powerful concept deserves thorough investigation, and yet regularization is still only a penalty term in the cost function of our underdetermined holographic system, in addition to its linear constraint that describes the wave propagation.

3.5 Notes on measurement and retreat distance

As already mentioned, probably the most crucial physical influence factor on well-conditioning is proximity to the source. The capture of evanescent waves due to the slender structural vibrations of higher wavenumbers is essential to gather detailed information about the source. The amount of evanescent waves captured in the measured data also determines the size of the singular values of the propagation matrix \mathbf{G} . The more so, proximity of the measurement surface to the source is of the essence, as spatially fine-structured evanescent waves decay exponentially with distance. Conversely, when the measurement array is positioned further away from the source, the most fine-structured evanescent waves are drastically attenuated with distance. If the measurement distance is large compared to the emitted wavelength, the high wavenumber components of the propagation matrix, related to the evanescent waves, almost vanish completely, resulting in singularity of the matrix. Figure 3.6 illustrates the impact of different measurement distances on the evanescent wave components of the propagator.

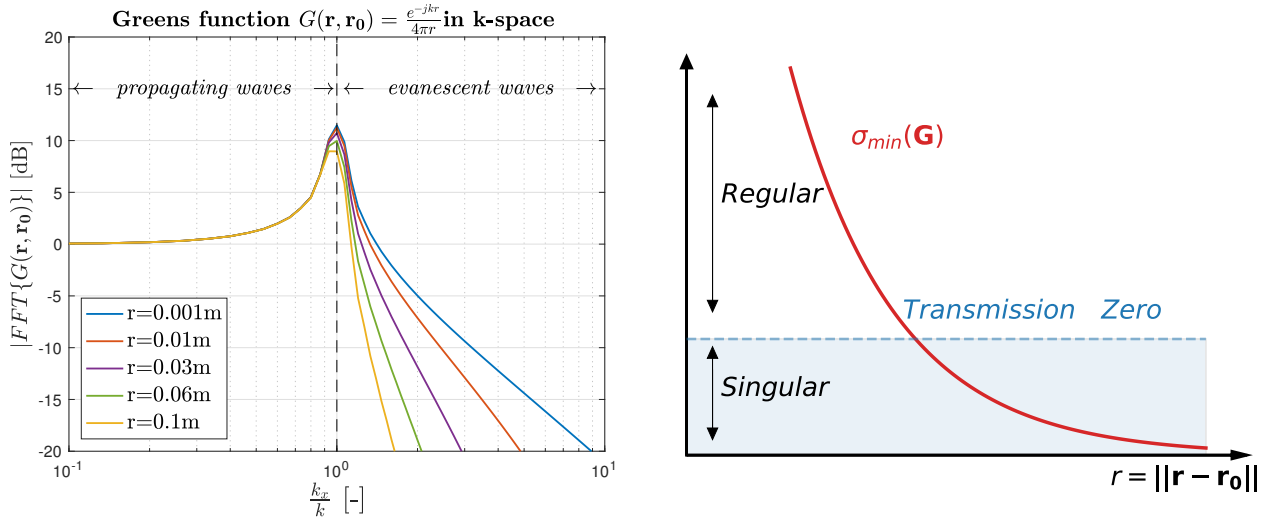


Figure 3.6: Free-space Green's function in the wavenumber domain (k -space) (left); High-order singular value (evanescent wave mode) vs. distance (right).

The transfer function, given by the radially symmetric free-field Green's function, exhibits a typical low pass characteristic with resonance at the coincidence frequency, that marks the transition between propagating and evanescent waves. An increase of the hologram distance leads to a steeper descent of the slope, controlling the cut-off applied to the higher wavenumber components. The farther we move to the farfield, the more negligible becomes the contribution of the evanescent waves to the measurements. An increase in distance obviously results in a spatial smoothing. At infinite distance, the contribution of evanescent waves, as with any exponentially decaying function, vanishes.

A high-order singular value, associated to a source mode of higher wavenumber components, exponentially decreases in magnitude with increasing distance to the source. As it gradually approaches zero, there comes a point at which the associated propagation matrix turns singular. This is referred to as the transmission zero [BIB13], already occurring at very small singular values, see Figure 3.6.

These issues above must also be taken into account when setting up the equivalent sources to model the radiated sound field. The formulation of ESM typically requires a suitable retreat distance between the equivalent source grid and the source surface to prevent potential singularities of point sources. For the recovery of sound field quantities on the source surface without retreat distance, the distance between equivalent-source and field points may become zero which would entail a division by zero in some Green's functions. In particular, this would yield infinite values in each source column of the propagation matrix \mathbf{G} . A suitable retreat distance helps. Finding an optimal distance is not trivial as it depends on many different aspects regarding source, model and measurement setup. While a large retreat distance diminishes infinite values in \mathbf{G} , too much increase promotes the ill-posedness of the inverse problem by attenuating the evanescent source modes, as laid out above. A practical choice must consider a trade-off between either large values in the columns of \mathbf{G} or strongly attenuated evanescent source modes in \mathbf{G} . A rule of thumb often applied suggests the use of one-to-two times the microphone spacing. A detailed study on the effects of the retreat distance in ESM can be found in [BCJH11] and the literature referred therein.

4

Regularization in Convex Optimization Problems

4.1 Constrained and penalized optimization framework

The remaining task is to find an appropriate way of solving the linear inverse problem, i.e. solving the linear observation model

$$\mathbf{p} = \mathbf{G}\mathbf{q} + \mathbf{n} \tag{4.1}$$

for the unknown source strength vector $\mathbf{q} \in \mathbb{C}^N$, given the propagation matrix $\mathbf{G} \in \mathbb{C}^{M \times N}$ and the measured pressure $\mathbf{p} \in \mathbb{C}^M$, which is corrupted by additive Gaussian noise $\mathbf{n} \in \mathbb{C}^M$ of variance ϵ . Since the number of microphones is typically smaller than the number of equivalent sources ($M \ll N$), the underlying problem is underdetermined. In this case, the problem can either not be solved or an infinite number of possible solutions exists. To obtain a unique solution, additional criteria may be imposed in order to promote certain characteristics of the source strength distribution. The use of an appropriate assumption on the source strength crucially restricts the set of possible solutions to a satisfactory one. From an optimization perspective, the problem above can be regarded as a constrained minimization problem [FGXG17,FGD18,Hal16]. The most generic constrained formulation of the problem is

$$\min_{\mathbf{q}} \mathcal{R}(\mathbf{q}) \quad \text{subject to} \quad \mathbf{G}\mathbf{q} = \mathbf{p}, \tag{4.2}$$

where $\mathcal{R}(\mathbf{q})$ denotes the so-called regularizer and $\mathbf{G}\mathbf{q} = \mathbf{p}$ is the acoustical constraint imposed on the solution by the equivalent source model. Because of the inevitable perturbations in the measurements due to the presence of noise or numerical error \mathbf{n} , a relaxed formulation of the optimization is often preferred. It permits a small deviation $\epsilon \leq \|\mathbf{n}\| \leq 0$ between model $\mathbf{G}\mathbf{q}$ and measurement \mathbf{p}

$$\min_{\mathbf{q}} \mathcal{R}(\mathbf{q}) \quad \text{subject to} \quad \|\mathbf{p} - \mathbf{G}\mathbf{q}\|_2 \leq \epsilon. \tag{4.3}$$

The term $\|\mathbf{p} - \mathbf{G}\mathbf{q}\|_2$ can be called model mismatch. From all possible solutions \mathbf{q} of limited mismatch, we look for a unique solution that minimizes the proposed regularizer $\mathcal{R}(\mathbf{q})$.

Alternatively, based on Bayesian estimation theory, an optimization task may be defined that incorporates desired prior knowledge in form of a penalty function that replaces the constraint. Under the assumption of a Gaussian distribution of the measurement noise \mathbf{n} , the maximum a

posteriori probability estimation yields the following regularized least-squares problem

$$\min_{\mathbf{q}} \left\{ \|\mathbf{p} - \mathbf{G}\mathbf{q}\|_2^2 + \lambda \mathcal{R}(\mathbf{q}) \right\}. \quad (4.4)$$

In a similar manner to ϵ , the parameter $\lambda \geq 0$ controls the tradeoff between the mismatch and regularizer.

Certain values of the parameters ϵ and λ lead to equality of the constrained and penalized variant of the optimization problem.

For a better understanding of the concept of norm-based regularizers, a geometric interpretation of their minimization in the context of a simple two-dimensional linear inverse problem is provided. [Ela10] All conclusions drawn from the following analysis can be validly extended to higher dimensions. Lets consider the general constrained minimization problem of the form (4.3) where the linear system $\mathbf{A}\mathbf{x} = \mathbf{y}$ is underdetermined and thus provides multiple solutions, and the regularizer $\mathcal{R}(\mathbf{x})$ is given by an arbitrary ℓ_p -norm, such that

$$\min_{\mathbf{x}} \|\mathbf{x}\|_p \quad \text{subject to} \quad \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2 \leq \epsilon. \quad (4.5)$$

The set of all possible solutions to the linear system is a single line (or a hyperplane for an N-dimensional problem), which is equivalent to the solution space to the constraint for $\epsilon = 0$. Due to the permissible deviation $\epsilon > 0$ this solution space is extended to a continuous set of lines parallel to the existing one, surrounding it within the tolerated distance in both directions. The graphical solution to this problem makes use of the concept of ℓ_p -norm balls, which describe a set of vectors satisfying a certain ℓ_p -norm. Figure 4.1 shows some examples of scaled ℓ_p -norm balls that cover the ones most frequently used in optimization.




Figure 4.1: Examples for two-dimensional ℓ_p -norm balls.

In case of minimization problems, the goal is to determine the solution which provides the smallest possible ℓ_p -norm, while still solving the given linear constraint. This can be achieved

by successively minimizing the associated ℓ_p -norm ball until a unique intersection point between its contour and any of the lines marking the solution space is found.

If the regularization term involves the ℓ_1 -norm, the intersection occurs at either one of the axes, implying that one component of the solution vector will always be zero. Accordingly, it can be deduced that the minimization of the ℓ_1 -norm results in a sparse representation of the solution, and in general, as with the other $0 \leq p < 1$ norms, it will promote the fewest-possible non-zero entries. Although very unlikely, if the model hyperplane fits on an entire side of the ℓ_1 -norm ball, this will yield a non-unique solution, e.g. in 2D for $\mathbf{A} = [\pm\lambda, \pm\lambda]$. From this example it becomes obvious that only norms with $0 \leq p < 1$ could still induce strict sparsity by potentially countable, non-unique and non-convex solutions, while the ℓ_1 -norm is the most sparsity-enforcing norm that is still convex but would become non-unique in the special case.




Figure 4.2: Geometric interpretation of the ℓ_1 -norm minimization.

When considering the minimization of the ℓ_2 -norm instead, the intersection point defines a vector, whose components most likely differ from zero, thus providing a non-sparse solution. It corresponds to the solution that is closest to the origin, which therefore exhibits minimum energy. Moreover, while the ℓ_1 -norm minimization generates a solution where the energy is concentrated to a small number of coefficients, the ℓ_2 -norm tends to distribute or smear the energy equally amongst all dimensions. However, if the line of solutions is perpendicular to one of the axes, it is still possible to obtain a sparse solution that lies on an axis, but this is very unlikely to happen.




Figure 4.3: Geometric interpretation of the ℓ_2 -norm minimization.

The conventional approach of solving an underdetermined inverse problem is in the sense of least squares (LS) obtained by the regularized right inverse $\hat{\mathbf{q}}_{LS} = \mathbf{G}^H(\mathbf{G}\mathbf{G}^H + \lambda\mathbf{I})^{-1}\mathbf{p}$. This corresponds to the solution of a squared ℓ_2 -norm minimization problem with the quadratic, sum of squares regularizer $\mathcal{R}(\mathbf{q}) = \|\mathbf{q}\|_2^2$, seeking for a smooth estimate with minimum energy $\hat{\mathbf{q}}_{LS}$. Using the above equivalence between ϵ and λ , this is either

$$\hat{\mathbf{q}}_{LS} = \arg \min_{\mathbf{q}} \|\mathbf{q}\|_2^2 \quad \text{subject to} \quad \|\mathbf{p} - \mathbf{G}\mathbf{q}\|_2 \leq \epsilon, \quad \text{or} \quad (4.6)$$

$$\hat{\mathbf{q}}_{LS} = \arg \min_{\mathbf{q}} \left\{ \|\mathbf{p} - \mathbf{G}\mathbf{q}\|_2^2 + \lambda \|\mathbf{q}\|_2^2 \right\}.$$

In general the LS approach yields robust reconstruction in the presence of noise and it is computational efficient. One of its drawbacks is that it is mainly applicable for low to mid frequencies as it is known to underestimate the sound power at high frequencies. Moreover, the frequency range of unique solutions is limited by the used spatial sampling of the sound pressure at higher frequencies. Since the reconstruction error almost evenly spreads throughout the solution subspace, the estimate also provides a blurred or over-smoothed version of the underlying source distribution. This can affect the solution in terms of sound source localization, especially when the underlying source strength distribution is sparse, i.e. comprised of only a few separated sources.

The common approach to cope with these problems is based on the compressive sensing (CS) methodology, which favors and promotes sparsity in the solution. It produces a source vector that is composed of only a few entries that are non-zero. This is ideally achieved by minimizing the ℓ_0 -pseudo-norm, since it is a measure of the cardinality or the number of non-zero terms of a vector. The resulting optimization task is an NP-hard problem that typically requires an intractable non-convex combinatorial search. The ℓ_1 -norm can be considered as a close approximation of the ℓ_0 -pseudo-norm leading to a convex relaxation of the problem. Despite the lack of a closed form solution for this minimization problem, it can be efficiently computed by using state of the art algorithms for convex optimization. The ℓ_1 -norm minimization problem can be formulated with the aid of the non-quadratic, sum of absolutes regularizer $\mathcal{R}(\mathbf{q}) = \|\mathbf{q}\|_1$, seeking for a sparse estimate with only few non-zero entries $\hat{\mathbf{q}}_{CS}$. With the above equivalence between ϵ and λ , this

either becomes

$$\begin{aligned} \hat{\mathbf{q}}_{CS} &= \arg \min_{\mathbf{q}} \|\mathbf{q}\|_1 \quad \text{subject to} \quad \|\mathbf{p} - \mathbf{G}\mathbf{q}\|_2 \leq \epsilon, \\ & \hspace{15em} \text{or} \\ \hat{\mathbf{q}}_{CS} &= \arg \min_{\mathbf{q}} \left\{ \|\mathbf{p} - \mathbf{G}\mathbf{q}\|_2^2 + \lambda \|\mathbf{q}\|_1 \right\}. \end{aligned} \tag{4.7}$$

One distinct advantage of promoting sparse solutions is the suppression of side lobes. However, when exploring non-sparse source strength distributions, the promotion of sparsity may yield unrealistic representations of the actual source. Spatially extended sources may be decomposed into a sparse distribution of scattered point sources.

4.2 Derivative-based regularizers

In Bayesian estimation theory, both the LS and CS solution can be considered as synthesis priors as the desired structure of the solution is imposed directly in the spatial domain. An alternative approach to incorporate prior knowledge is to impose a specific structure on some linear basis or transformed domain of the solution. In this case, the solution falls under the category of analysis priors. [FGD18]

Most of the regularizers typically used in the context of optimization problems can be generally described by [LWU13]

$$\mathcal{R}(q(\mathbf{r})) = \int_{\Omega} \Phi(\mathbf{D}q(\mathbf{r}))d\mathbf{r} \quad \rightarrow \quad \mathcal{R}(\mathbf{q}) = \sum_{i=1}^N \Phi(\mathbf{D}\mathbf{q}_i) = \|\Phi(\mathbf{D}\mathbf{q})\|_1, \tag{4.8}$$

where \mathbf{D} denotes a linear operator applied on the solution vector \mathbf{q} and $\Phi(\cdot)$ is a potential function, usually involving some norm distance $\|\cdot\|_p^q$. Most common regularizers impose a direct spatial structure on the solution rather than on a transformed domain. In this case, the linear operator simplifies to the identity operator $\mathbf{D} = \mathbf{I}$.

Another popular choice for the linear mapping is some type of differential operator (i.e. Gradient, Laplacian, Hessian) to penalize the spatial derivatives. The use of a differential operator in combination with a quadratic regularizer promotes derivatives with minimum-energy, hence leading to over-smoothed solutions. For this reason, non-quadratic derivative-based regularizers may be the preferred choice. The application of Compressive Sensing, i.e. the ℓ_1 -norm, on the differential domain promotes solutions with vanishing spatial derivatives except of a few non-zero occurrences. The ℓ_1 -norm penalization of first-order derivatives promotes piecewise-constant solutions. They contain regions of steady state magnitude where the gradient vanishes, and a few magnitude alterations in case of a non-zero gradient. This approach introduces the appearance of blocking artifacts, better known as the staircase effect [Jal14]. Since it yields a solution with constant regions and sharp edges, this regularizer may not be adequate for the representation of a smooth distribution of acoustic sources. However, sparsifying second-order derivatives promotes piecewise-linear solutions with a small, piecewise-constant gradient and thus minimum curvature. Therefore, it seems to be more suitable for recovering the spatial amplitude variations of acoustic sources. In general, the promotion of sparse n-th order spatial derivatives forces the solution to fit a n-1 degree polynomial trendline (i.e. constant, linear, quadratic, etc.) [XFGG16].

A possible improvement could be achieved by simultaneously forcing sparsity in the spatial and differential domain. This is accompanied by the indirect promotion of block- or group-sparse

solutions in the spatial domain. In this case, the solution vector is not compulsory sparse, but all the occurrences of non-zero entries are clustered or grouped together. In theory, such regularizer can be applied for the reconstruction of compact sources as well as extended sources, which makes it a preferable choice for the general case of sound field reconstruction.

In addition, a distinction can be made between isotropic and anisotropic regularization penalties. In case of isotropic smoothing, the presence of a derivative with high energy along one direction prohibits the attenuation of a smaller derivative along the other directions at a specific source point. This leads to an enhancement of point-like features in the resulting representation of the source distribution. In case of anisotropic smoothing, strong derivatives along one direction have no effect on the minimization of the smaller derivatives along other directions. This is known to encourage line-like or elliptic features and well-preserved edges and contours. [HJ12]

The following section gives an intuitive introduction to the use of derivatives in regularization. It illustrates that different choices of the differential operator as regularization penalty have their specific abilities and shortcomings. The discussion focuses on the most common derivative-based regularizers of first and second order. They share favorable properties such as convexity, rotation-, translation-, contrast- and scale-invariance.

4.2.1 Introduction to derivatives in regularization

To provide a deep understanding of common derivative-based regularizers, this section is dedicated to the analysis of first- and second-order derivatives and how they can be used to impose a certain structure on the underlying function. This already involves all differential operators needed for the implementation of subsequently proposed regularizers and points out their interrelations. For a compact formulation of the basic differential-geometric principles, the reader is referred to [Pan14].

Let us first consider a one-dimensional function that is assumed to be continuously differentiable in its whole range of definition. The first-order derivative at a particular point is formally defined as the instantaneous rate of change as it provides the amount the function value changes due to a small change in its evaluation position. The second-order derivative is defined as the derivative of the derivative, thus it describes the rate of change of the slope and hence the curvature of the function. Figure 4.4 demonstrates the behavior of the derivatives around an inflection point or edge indicated by a local change of intensity.




Figure 4.4: Local behavior of first- and second-order derivatives around an edge.

Points lying on an edge are associated with local minima or maxima of the first-order derivative and zero-crossings of the second-order derivative. A positive first-order derivative denotes an increase while a negative one indicates a decrease of function value. A positive second-order derivative relates to a convex function while a negative one implies a concave curvature of the function profile. Furthermore, regions of constant function value are characterized by zero first-order derivatives and consistently changing regions that exhibit equal consecutive slopes feature zero second-order derivatives.

First-order differentiation extends to two-dimensional functions by the gradient that defines the vector of partial derivatives along the two orthogonal directions of the coordinate system. Hence, for a two-dimensional function $f(x, y)$, the gradient is defined as

$$\nabla f(x, y) = \begin{bmatrix} \frac{\partial f(x, y)}{\partial x} \\ \frac{\partial f(x, y)}{\partial y} \end{bmatrix}. \quad (4.9)$$

The gradient is a vector pointing towards the direction of the steepest ascent whereby its magnitude quantifies the amount of that increase. In this way, the gradient provides the slope of the tangent plane to the graph at the evaluation point.




Figure 4.5: Gradient field of a mode shape and its associated divergence field.

To observe the curvature of a function, second-order derivatives need to be considered. The Hessian is a matrix composed by all second-order derivatives (univariate and mixed) that exist in an N-dimensional space. For the two-dimensional case, the Hessian is defined as

$$\mathcal{H}f(x, y) = \begin{bmatrix} \frac{\partial^2 f(x, y)}{\partial x^2} & \frac{\partial^2 f(x, y)}{\partial x \partial y} \\ \frac{\partial^2 f(x, y)}{\partial y \partial x} & \frac{\partial^2 f(x, y)}{\partial y^2} \end{bmatrix}. \quad (4.10)$$

The mixed partial derivatives in the off-diagonal elements are equal for continuous functions, regardless of the order of differentiation, which makes the Hessian symmetric. The eigenvalues and eigenvectors of the Hessian classify the orientation and isotropy/anisotropy of the curvature. The eigenvectors are always orthogonal and point into the directions where the function exhibits minimum and maximum curvature, respectively. They are often referred to as the principal directions of curvature. The curvature in those directions is determined by the magnitude and sign of the corresponding eigenvalues. Figure 4.6 illustrates the geometric meaning of the Hessian eigenvectors that define the axes of ellipses that characterize the local principal curvature.




Figure 4.6: Geometric meaning of Hessian eigenvectors.

Note that the major axis of the ellipses is always related to the smaller eigenvalue as it indicates a lower principal curvature along its direction, and vice versa. In general, flat regions of constant value yield smaller ellipses while regions of high local variations result in larger ellipses. Furthermore, very narrow and elongated ellipses indicate an anisotropic curvature evolution, whereas an isotropic curvature creates more compressed ellipses that tends to become circles. The direction and magnitude of the principal curvature are both intrinsic properties of the function as they do not depend on the orientation of a coordinate system.

Another important tool for the characterization of any function is the divergence of its gradient field, which is a measure of the net flux of the vector field into and out of the evaluation point. This scalar is formally known as the Laplacian which corresponds to the trace of the Hessian

$$\Delta f(x, y) = \nabla^2 f(x, y) = \frac{\partial^2 f(x, y)}{\partial x^2} + \frac{\partial^2 f(x, y)}{\partial y^2} = \text{trace}(\mathcal{H}f(x, y)). \quad (4.11)$$

The Laplacian summarizes the curvature as twice the mean curvature which equals the sum of the principal curvatures at a point

$$\text{trace}(\mathcal{H}f(x, y)) = \text{trace}(\mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{-1}) = \text{trace}(\mathbf{Q}^{-1}\mathbf{Q}\mathbf{\Lambda}) = \text{trace}(\mathbf{\Lambda}) = \lambda_1 + \lambda_2. \quad (4.12)$$

In summary, the evaluation of the Gradient, Laplacian or Hessian operator delivers valuable structural characterization about the evolution of a function. It is possible to make use of those measures to promote certain shapes or characters in a function by enforcing either of these operators to take on certain values. Since we consider two-dimensional source strength distributions in our holographic task, these operators are highly qualified for regularization purposes.

4.2.2 Gradient-based regularizer - The Total Variation

One of the most popular choice of derivative-based regularizers is the Total Variation (∇ TV) functional [IROF92], where the linear operator in equation (4.8) is given by the gradient $\mathbf{D} = \nabla$ in equation (4.9) and the potential function involves the magnitude or Euclidean norm $\Phi(\cdot) = \|\cdot\|_2$

$$\begin{aligned} \mathcal{R}_{\nabla TV}(q(\mathbf{r})) &= \int_{\Omega} |\nabla q(\mathbf{r})| d\mathbf{r} = \int_{\Omega} \sqrt{[\partial_x q(\mathbf{r})]^2 + [\partial_y q(\mathbf{r})]^2} d\mathbf{r} \\ &\rightarrow \mathcal{R}_{\nabla TV}(\mathbf{q}) = \sum_{i=1}^N \|\nabla \mathbf{q}_i\|_2 = \|\nabla \mathbf{q}\|_{2,1}. \end{aligned} \quad (4.13)$$

By summation of the gradient vector ℓ_2 -norms across all locations, this regularizer incorporates a mixture of the ℓ_1 - and ℓ_2 -norm. The mixed $\ell_{1,2}$ -norm is known to globally promote joint-sparsity of the directional derivatives over all positions in \mathbf{q} . The minimization results not only in a sparse gradient representation but also in a grouping of the solution coefficients. As a consequence, this regularizer encourages piecewise-constant solutions with regions of homogeneous source strengths with vanishing gradient and only few variations in amplitude with non-zero gradient. Despite the isotropic formulation, Total Variation regularization leads to anisotropic smoothing along the edges, orthogonal to the direction of the gradient [HJ12, HORJ14]. The attenuation of smoothing across the edges in gradient direction induces block- or patch-like artifacts with sharp edges, better known as the staircase effect. The anisotropic smoothing may be a suitable property for characterizing clearly confined sources like a radiating piston, but is highly undesirable for the representation of generic acoustic sources with continuously varying amplitudes.

4.2.3 Laplacian-based regularizer - The Total Generalized Variation

Another common approach to prevent the staircase effect and to preserve smooth transitions in the solution is the penalization of higher-order spatial derivatives. They are known to offer an improved approximation of the signal with piecewise polynomials of higher degree. In most cases, it is sufficient to consider the penalization of second-order derivatives for satisfactory results, as the use of higher-order derivatives often provides no further improvements. A typical choice for the second-order differential operator in equation (4.8) is the scalar Laplacian $\mathbf{D} = \nabla^2 = \partial_{xx} + \partial_{yy}$. One of the major drawbacks of this regularizer is the large size of its null space. For any homogeneous solution of the Laplacian $\nabla^2 q(\mathbf{r}) = 0$, this regularizer will become $\mathcal{R}_{\nabla^2}(q(\mathbf{r})) = 0$. For this reason, the scalar Laplacian penalty may not be an ideal choice for regularization as the problem may still remain ill-posed [HJ12, HORJ14]. One way to circumvent the large null space of the regularizer is to choose the vectorized Laplacian $\mathbf{D} = \Delta$ as the differential operator, whose entries are the second-order derivatives along the coordinate directions. Considering the magnitude or Euclidean norm as the potential function $\Phi(\cdot) = \|\cdot\|_2$, we end up with the Laplacian-based regularizer as a second-order extension of Total Variation (Δ TV2)

$$\begin{aligned} \mathcal{R}_{\Delta TV2}(q(\mathbf{r})) &= \int_{\Omega} |\Delta q(\mathbf{r})| d\mathbf{r} = \int_{\Omega} \sqrt{[\partial_{xx} q(\mathbf{r})]^2 + [\partial_{yy} q(\mathbf{r})]^2} d\mathbf{r} \\ &\rightarrow \mathcal{R}_{\Delta TV2}(\mathbf{q}) = \sum_{i=1}^N \|\Delta \mathbf{q}_i\|_2 = \|\Delta \mathbf{q}\|_{2,1}. \end{aligned} \quad (4.14)$$

This penalty imposes global sparsity on the second-order derivatives, which leads to piecewise-constant gradients with small magnitude and thus smooth changes in the underlying spatial dis-

tribution. When compared to the ∇TV , ΔTV2 encourages piecewise-linear or -affine solutions as opposed to piecewise-constant solutions, leading to an enhancement of the spatial resolution and a reduction of blocking artifacts. This seems more suitable for the localization and representation of acoustic sources with continuously varying amplitudes and phases. Moreover, the Laplacian-based regularizer provides isotropic, circularly symmetric smoothing and thus also encourages point-like features of the source distribution. This can be a desirable property for the localization or identification of point sources, but may yield poor reconstruction of arbitrarily shaped extended sources.

4.2.4 Hessian-based regularizer - The Hessian Schatten-norm family

To overcome the shortcomings of the Laplacian, another family of regularizers based on the Schatten-norms of the Hessian matrix proposed in [LWU13] are valuable. They can also be considered as second-order extensions of TV, as they inherit the same favourable convexity and invariance properties while penalizing second-degree polynomials instead. Their ability to avoid the staircase effect and to significantly reduce other artifacts (ringing, ripple, etc.) introduced by the reconstruction process endorse their use as a regularizer. In order to define the Hessian-based regularizers from the general viewpoint of equation (4.8), the differential operator is chosen as the Hessian $\mathbf{D} = \mathcal{H}$ while the potential function involves the Schatten p -norm $\Phi(\cdot) = \|\cdot\|_{S_p}$ of the matrix.

The Schatten p -norm of the Hessian in equation (4.10) is defined as the ℓ_p -norm of a vector $\boldsymbol{\sigma}$ whose entries contain the two singular values σ_k of the Hessian

$$\|\mathcal{H}q(\mathbf{r})\|_{S_p} = \|\boldsymbol{\sigma}(\mathbf{r})\|_p = \left(\sum_{k=1}^2 \sigma_k^p(\mathcal{H}q(\mathbf{r})) \right)^{\frac{1}{p}}, \quad p \geq 1. \quad (4.15)$$

This leads to the following definition of the general Hessian Schatten-norm regularizer

$$\begin{aligned} \mathcal{R}_{\mathcal{H}S_p}(q(\mathbf{r})) &= \int_{\Omega} \left(\sum_{k=1}^2 \sigma_k^p(\mathcal{H}q(\mathbf{r})) \right)^{\frac{1}{p}} d\mathbf{r} \\ \rightarrow \mathcal{R}_{\mathcal{H}S_p}(\mathbf{q}) &= \sum_{i=1}^N \|\mathcal{H}\mathbf{q}_i\|_{S_p} = \|\mathcal{H}\mathbf{q}\|_{S_p,1}. \end{aligned} \quad (4.16)$$

The Schatten p -norm in the regularizer yields three special cases of matrix norms, including the nuclear or trace norm ($p = 1$)

$$\|\mathcal{H}q(\mathbf{r})\|_{S_1} = \|\boldsymbol{\sigma}(\mathbf{r})\|_1 = \sum_{k=1}^2 \sigma_k(\mathcal{H}q(\mathbf{r})), \quad (4.17)$$

the Frobenius norm ($p = 2$)

$$\begin{aligned} \|\mathcal{H}q(\mathbf{r})\|_{S_2} &= \|\boldsymbol{\sigma}(\mathbf{r})\|_2 = \sqrt{\sum_{k=1}^2 \sigma_k^2(\mathcal{H}q(\mathbf{r}))} \\ &\equiv \|\mathcal{H}q(\mathbf{r})\|_F = \sqrt{[\partial_{xx}q(\mathbf{r})]^2 + [\partial_{yy}q(\mathbf{r})]^2 + 2[\partial_{xy}q(\mathbf{r})]^2} \end{aligned} \quad (4.18)$$

and the spectral or operator norm ($p = \infty$)

$$\|\mathcal{H}q(\mathbf{r})\|_{\mathcal{S}_\infty} = \|\boldsymbol{\sigma}(\mathbf{r})\|_\infty = \max_{k=1,2} \sigma_k(\mathcal{H}q(\mathbf{r})). \quad (4.19)$$

With these types of regularizers, global sparsity is not directly imposed on the second-order derivatives but rather on measures depending on the local singular values of the Hessian matrix. In case of a 2×2 Hessian, the singular values correspond to the absolute eigenvalues and the involved SVD can be reduced to an eigendecomposition. Since the eigenvalues of the Hessian describe the respective amounts of surface diffraction in the direction of the corresponding eigenvectors, they can be considered as a measure of the local principal curvature at each evaluation point. Accordingly, the enforcement of sparse eigenvalues leads to piecewise-linear solutions with smooth, minimum curvature. The indirect way of penalizing second-order derivatives makes the Hessian-based regularizer superior to the scalar Laplacian penalty as it overcomes the issue of a potentially remaining ill-posedness of the problem due to the large null space of the regularizer. Furthermore, the Hessian-Schatten regularizers provide both cases of isotropic and anisotropic smoothing, depending on the choice of p . For $p = 1$, the regularizer favors smoothing along edges over smoothing across them, thus it locally promotes sharp-edged, one-sided curvature. It is known to encourage line-like or elliptic features in the retrieval. For $p = 2$, the regularizer behaves similar to the Laplacian-based regularizer, as it only differs by also penalizing the cross-derivative. Hence it applies additional diagonal smoothing and provides improved enhancement of point-like features. For $p = \infty$, the regularizer favors smoothing across edges and promotes ridges and valleys with isotropic curvature.

4.2.5 Fused sparse and sparse-differential regularizer

To further extend the applicability of the proposed regularizers to the various spatial structures encountered in the reconstruction of acoustic sources, the enforcement of sparsity on spatial derivatives may be combined with direct spatial sparsity by using an additional ℓ_1 -norm penalty on the solution. This promotes block-sparsity in the solution, meaning that the solution is not necessarily sparse in the spatial domain but rather exhibits isolated, active regions where non-zero patches are grouped together. The promotion of block- or group-sparse solutions leads to further enhancement of the spatial resolution and is applicable to the characterization of sound fields involving spatially compact and extended sources.

There are two possible ways to additionally enforce sparsity in space. One being the use of a regularizer comprised of the sum of some form of TV and a ℓ_1 -norm penalty on the solution

$$\mathcal{R}_{\mathbf{D}-\ell_1}(\mathbf{q}) = \mathcal{R}_{\mathbf{D}}(\mathbf{q}) + \mu\|\mathbf{q}\|_1 = \sum_{i=1}^N \|\mathbf{D}\mathbf{q}_i\|_p + \mu\|\mathbf{q}\|_1 = \|\mathbf{D}\mathbf{q}\|_{p,1} + \mu\|\mathbf{q}\|_1. \quad (4.20)$$

Here \mathbf{D} is the spatially discrete differential operator associated with the regularization method of choice, p defines the regularizers' inner p-norm (or Schatten p-norm of the Hessian) and μ denotes a hyperparameter controlling the balance between sparsity of the spatial derivatives and direct spatial sparsity. In this case, the enforcement of sparse derivatives and coefficients is fully separable. Inactive coefficients can exhibit non-zero spatial derivatives and vice versa. This may lead to the appearance of occasional artifacts in inactive regions and affects the amount of suppression of noise components.

Another way of promoting block-sparsity is to incorporate the spatial coefficients directly in the mixed-norm of the regularizer by embedding an additional identity operator within the differential operator

$$\mathcal{R}_{\mathbf{D}-SV}(\mathbf{q}) = \sum_{i=1}^N \|\mathbf{K}\mathbf{q}_i\|_2 = \|\mathbf{K}\mathbf{q}\|_{2,1} \quad \text{with} \quad \mathbf{K} = \begin{bmatrix} \mathbf{d} \\ \mu\mathbf{I} \end{bmatrix}. \quad (4.21)$$

This generalization includes the $\mathcal{R}_{\nabla TV-SV}$, $\mathcal{R}_{\Delta TV2-SV}$ and $\mathcal{R}_{\mathcal{H}_{S_2}-SV}$ regularizers, where \mathbf{d} is the vectorized discrete differential operator of choice, whose entries are the associated spatial derivatives. The operator \mathbf{K} can be considered as an expansion of this linear mapping to an additional dimension containing direct spatial information. Since sparsity is imposed simultaneously on the elements grouped by the inner norm, this penalty can lead to a joint-sparse spatial and differential distribution when considering the non-separable norm with $p = 2$. Either an estimated component is non-zero in space with non-zero, penalized derivatives, or it vanishes concurrently in both domains. This enforces sparsity on adjacent elements of large inactive regions that may otherwise contain erroneous artifacts while allowing smooth spatial variation within non-zero regions, suitable for the characterization of acoustic sources. In this case, even the use of the gradient operator results in piecewise-linear solutions within active regions since a zero gradient is always attached to a zero spatial component. Thus, this type of regularizer is denoted as Sparse Variation (SV) [EDTV15, NPCK17].

For the regularizers $\mathcal{R}_{\mathcal{H}_{S_1}-SV}$ ($p = 1$) and $\mathcal{R}_{\mathcal{H}_{S_\infty}-SV}$ ($p = \infty$), the spatial information is nested in the associated inner norm acting on the singular value vector $\boldsymbol{\sigma}$ of the Hessian, with

$$\mathcal{R}_{\mathbf{D}-SV}(\mathbf{q}) = \sum_{i=1}^N \left\| \begin{bmatrix} \|\mathcal{H}\mathbf{q}_i\|_{S_p} \\ \mu\mathbf{q}_i \end{bmatrix} \right\|_p. \quad (4.22)$$

4.3 Comparison of first- and second-order derivative-based regularizers

For a better understanding of the certain characteristics of derivative-based regularizers, let's consider a simple two-dimensional radiation scenario. The source distribution model on the boundary is prescribed by three synthetic signals of N samples, each representing a particular type of spatial structure commonly found in acoustic imaging problems. The first signal consists of only a few non-zero components and represents a series of sparsely distributed monopole sources. The second signal contains piecewise-constant regions with sharp edges that can be associated with a vibrating rectangular aperture. The last signal is comprised of piecewise-linear parts that represent the continuous variation of mode patterns. The minimization problem is solved numerically using the unconstrained formulation of (4.4). The entries of the convolution operator \mathbf{G} are the two-dimensional free-space Green's functions $g(\mathbf{y}, \mathbf{y}_0) = \frac{j}{4} H_0^{(1)}(kr)$ where $H_0^{(1)}$ is the Hankel function of the first kind and zeroth order and $r = \|\mathbf{y} - \mathbf{y}_0\|$ is the distance between field and source point. The resulting sound field is observed at $M < N$ discrete observation points at a distance of $y=5$ cm away from the source plane and under ideal conditions in the absence of noise. In order to display the peculiarities of the individual regularizations, each scenario is solved using four different types of regularizers $\mathcal{R}(\cdot)$. They include the ℓ_1 - and ℓ_2 -norm acting on the spatial domain, as well as two composite $\ell_{1,2}$ -norms including the first- and second-order differential operators in form of the gradient ∇_x and Laplacian Δ_x , penalizing a structural domain of the source distribution.

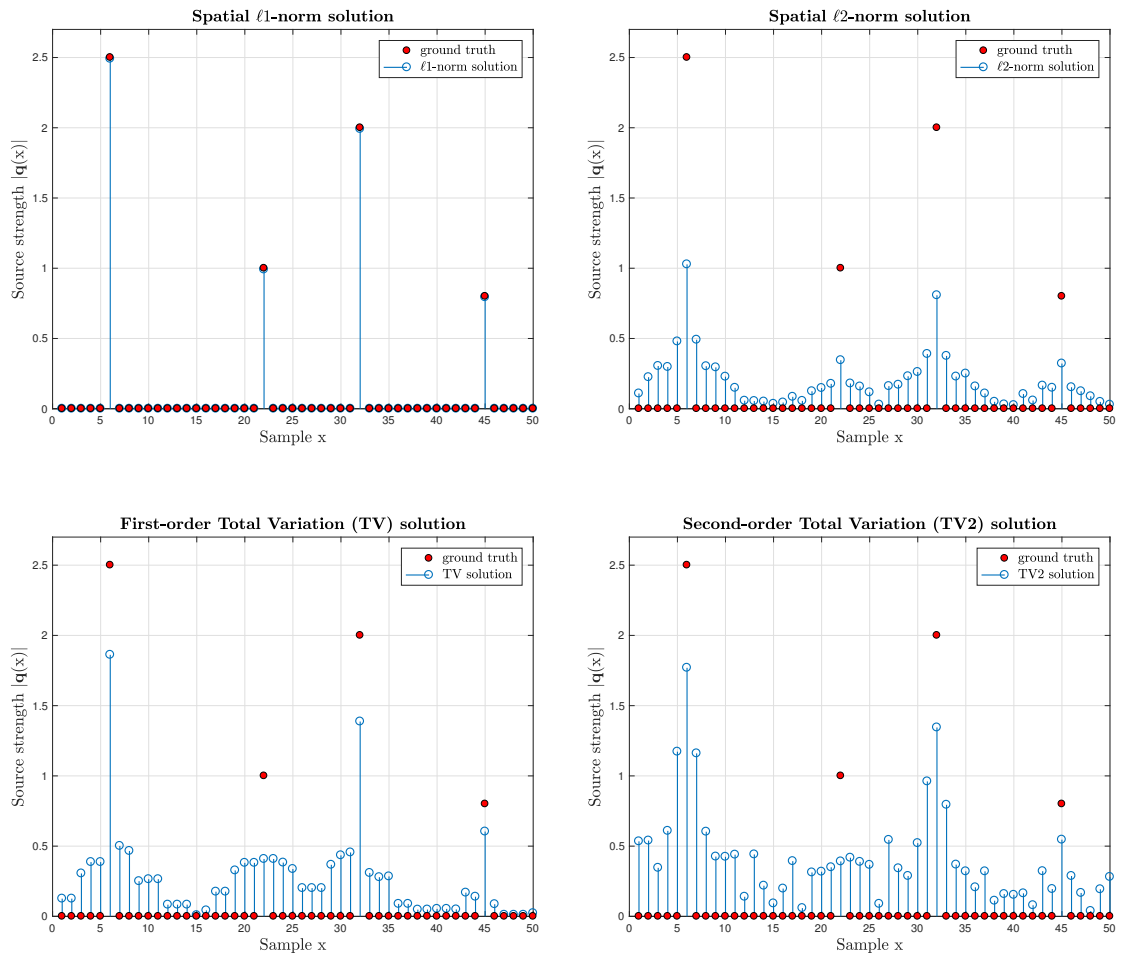


Figure 4.7: 2D Holographic reconstruction example of a known sparse source strength \mathbf{q} based on ℓ_2 , ℓ_1 , TV, TV2.

Figure 4.7 provides the obtained solutions to the case of the isolated monopoles. The application of the ℓ_1 -norm results in a sparse representation of the source field that exactly matches the target distribution. This makes it the preferred choice for reconstructing sound fields comprised of only few point sources. On the contrary, the least squares solution obtained by the ℓ_2 -norm cannot recover the sparse data as it equally spreads the energy of the solution vector across all its components. This leads to a spatially smooth source distribution of the least squares. It suffers from a high contamination of sidelobes, which comes with a significant degradation of localization accuracy for the considered point sources. The gradient-based solution exhibits distinct peaks at the locations of the monopoles, while showing confined regions of equal magnitude in the intermediate spaces. Small deviations from the piecewise-constant profile only occur at the transition points between different plateaus of magnitude. Similar to the ℓ_2 -norm, it does not concentrate the energy in a few equivalent sources but rather distribute it over multiple ones surrounding the locations of the actual point sources. Although it does not promote sparsity in the spatial domain, it provides an improved localization reliability. The Laplacian-based solution TV2 forms a sequence of consecutive linear segments that peak at the positions of the true sources and alternate in ascending and descending order in between. As it also does not impose direct sparsity on the solution, a strong activity among all components can be observed. In fact, it suffers a level of sidelobe contamination that is worse in comparison to the ℓ_2 -norm regularizer.




Figure 4.8: 2D Holographic reconstruction example of a known piecewise-constant source strength \mathbf{q} based on ℓ_2 , ℓ_1 , TV, TV2.

Figure 4.8 illustrates three line-like pistons vibrating with different magnitudes at separate locations. The ℓ_1 -norm yields a sparse estimate where the pistons are represented by a composition of a few monopoles. Since a high number of coefficients is forced to zero, the remaining ones partially exceed the maximum value of the respective piston magnitudes to be able to produce the same sound field. As the individual pistons cannot even be identified as such anymore, this approach seems fairly inadequate for the retrieval of extended sources. In case of the ℓ_2 -norm, the main part of the sources is distributed over the segments covered by the pistons. As this approach is known to smear the source strength distribution, it does not preserve the sharp edges of the sources and thus tends to overestimate their spatial extent. In addition, we can observe its tendency to underestimate the true magnitude of the sources. Nevertheless, this regularizer is an appropriate choice for the reconstruction of spatially extended sources as it yields a good, but smoothed estimate of the sources which accurately reflects their true position. The gradient-based TV regularization achieves a nearly perfect retrieval in terms of magnitude and shape of the sources. It preserves the well-defined, sharp edges of the pistons and shows a good suppression of sidelobes. Note that, although the coefficients tend towards zero in the source-free regions, they always remain active which causes the small deviation from the true magnitude of the sources. This stems from the fact that no spatial sparsity is imposed on the solution. At last, the solution obtained by the Laplacian-based regularizer TV2 exhibits its characteristic piecewise-linear profile which does not conform with the piecewise-constant one describing the source distribution. The regions of constant value are represented by a sawtooth-like structure, partially exceeding the true maximum.

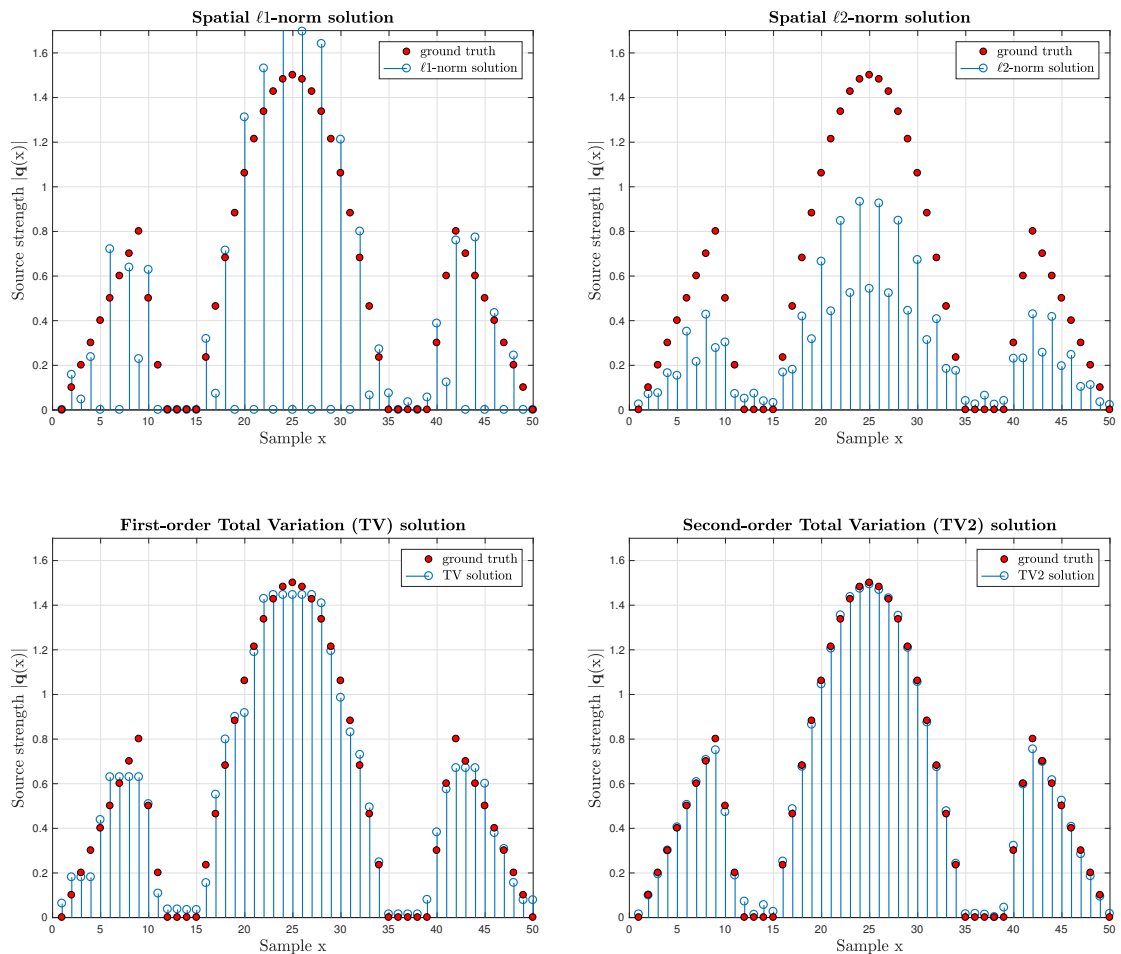


Figure 4.9: 2D Holographic reconstruction example of a known piecewise constantly-curved source strength \mathbf{q} based on ℓ_2 , ℓ_1 , TV, TV2.

Figure 4.9 presents a partly continuous variation of magnitude describing different vibrational shapes. With the ℓ_1 -norm, the continuous variations are again decomposed into a collection of a few point sources, providing an incorrect estimation of the source extent. The ℓ_2 -norm successfully retrieves the general shape of the source, even though it induces energy leaps between consecutive coefficients. According to the minimum energy property, it underestimates the true magnitude distribution. The estimate obtained by the gradient-based regularizer approximates the true function profile quite well, despite the introduction of the above-mentioned staircase effect. In order to avoid the staircase effect, the Laplacian can be utilized as a regularizer. It provides a smooth solution with minimum curvature, which allows a nearly perfect reconstruction of the original source distribution. The only deficiency may lie in the introduction of small variations between active regions that actually should be zero. Hence, the penalization of second-order derivatives seems the most suitable approach for dealing with continuously varying source distributions.

4.4 Selection of regularization parameter

A crucial aspect of regularization is the appropriate choice of the regularization parameter λ for the use of the unconstrained minimization problem (4.4). It determines the quality of reconstruction as well as the characteristic properties of the solution, as it controls the trade-off between matching and regularization. Finding the optimal amount of regularization is a non-trivial task, as it mainly depends on the relative noise level, which is usually unknown in practice. For the ℓ_2 -norm regularized approach (4.6) the regularization parameter controls the amount of smoothing of the solution. There exist a variety of methods for the automated determination of the optimal parameter value that proceed without any prior knowledge of the noise floor. This makes these methods very attractive and feasible in use, such that they have turned into state of the art techniques, including the Generalized Cross Validation (GCV) and the L-Curve method [Gom08]. When the problem involves a sparsifying regularizer, characterized by the ℓ_1 -norm, the absence of such automatic methods makes the optimal parameter choice a separated issue that deserves dedicated investigation. By successively changing the parameter λ , the evolution of the trajectory of the optimal solution takes the path of a piecewise linear function, denoted as the solution path. The points of discontinuity in this path indicate a change in sparsity level and thus can be utilized for proper parameter selection for a desired level of sparsity [GXM15]. While a majority of the methods available are based on such approaches, the decision fell on another approach proposed in [Hal18] that chooses the parameter according to the applied dynamic range D expressed in decibels. The reason for this is twofold: Firstly, a method based on the solution path is not practical for a comparative study since the involved regularization terms are composed of different linear operators that respectively modify the solution path in a different way. Hence the same sparsity level might be associated with distinct regularization parameters making a direct comparison difficult. Secondly, for the examples provided in this thesis, the noise floor is either known a priori or can at least be reliably estimated. Therefore, employing ad hoc parameter selection based on the given signal-to-noise ratio seems the most appropriate approach. The adopted method aims at rejecting singular values that are much smaller than the largest singular value for reconstruction. With the SVD in (3.21), we define a data vector $\mathbf{p} = \mathbf{c}\mathbf{u}_j$ consisting of a single field mode component associated with the singular value σ_j , which is supposed to lie D decibels beneath the maximum singular value σ_{max} . It is determined as the cut-off singular value below which truncation occurs, satisfying the relation

$$D = 20\log\left(\frac{\sigma_{max}}{\sigma_j}\right) \quad \rightarrow \quad \sigma_j = \sigma_{max}10^{-\frac{D}{20}}. \quad (4.23)$$

For an initial estimate $\mathbf{q}_0 = \mathbf{0}$, the model-mismatch $\|\mathbf{G}\mathbf{q} - \mathbf{p}\|_2^2$ in (4.4) takes on the value $\|\mathbf{p}\|_2^2 = |c|^2$. In order to drive this error for the given data vector to zero, the solution estimate $\hat{\mathbf{q}}$ must take the form

$$\begin{aligned}\|\mathbf{G}\hat{\mathbf{q}} - \mathbf{p}\|_2^2 &= \|\sigma_j \mathbf{u}_j \mathbf{v}_j^H \hat{\mathbf{q}} - c \mathbf{u}_j\|_2^2 = 0 \\ \sigma_j \mathbf{u}_j \mathbf{v}_j^H \hat{\mathbf{q}} &= c \mathbf{u}_j \\ \hat{\mathbf{q}} &= \left(\frac{c}{\sigma_j}\right) \mathbf{v}_j.\end{aligned}\tag{4.24}$$

For the truncation to take place around the singular value σ_j , the additional regularization term $\lambda \|\hat{\mathbf{q}}\|_1$ induced by this solution must compensate the resulting decrease $\|\mathbf{p}\|_2^2$ of the mismatch. By this demand, we obtain a suitable expression for λ , as it follows that

$$\begin{aligned}\lambda_{\text{opt}} \|\hat{\mathbf{q}}\|_1 &= \lambda_{\text{opt}} \left(\frac{c}{\sigma_j}\right) \|\mathbf{v}_j\|_1 = \|\mathbf{p}\|_2^2 = c \|\mathbf{p}\|_2 \\ \lambda_{\text{opt}} &= c \left(\frac{\sigma_j}{c}\right) \frac{\|\mathbf{p}\|_2}{\|\mathbf{v}_j\|_1} = \sigma_j \frac{\|\mathbf{p}\|_2}{\|\mathbf{v}_j\|_1}.\end{aligned}\tag{4.25}$$

Substitution of equation (4.23) into (4.25) then leads to the regularization parameter as a function of the applied dynamic range D

$$\lambda_{\text{opt}} = \sigma_{\text{max}} 10^{-\frac{D}{20}} \frac{\|\mathbf{p}\|_2}{\|\mathbf{v}_j\|_1}.\tag{4.26}$$

The value range of $\|\mathbf{v}_j\|_1$, which is normalized as $\|\mathbf{v}_j\|_2 = 1$, lies between $1 \leq \|\mathbf{v}_j\|_1 \leq \sqrt{N}$, i.e. 1 if the unit vector is either sparse $\mathbf{v}_j = [\dots 0 0 1 0 \dots]$ or \sqrt{N} when it is dense $\mathbf{v}_j = [1 1 \dots 1 1]/\sqrt{N}$, and we may approximate as minimum parameter

$$\lambda_{\text{opt}} = \sigma_{\text{max}} 10^{-\frac{D}{20}} \frac{\|\mathbf{p}\|_2}{\sqrt{N}}.\tag{4.27}$$

In a more realistic scenario, the data vector \mathbf{p} may consist of more than one field mode component \mathbf{u}_i , which causes an ℓ_1 -norm of the solution \mathbf{q} that is no longer additive over its elements along the source mode components \mathbf{v}_i . Anyway, for the examples provided, the lower-limit expression is a reasonably good ad-hoc choice from which further tuned adjustments can be made.

When intending to investigate the influence of the regularization parameter on the solution, the solution path delivers valuable insight. The following exposition of the various solution paths of the sparsifying regularizers is inspired by the study in [XFGG16]. Figure 4.10 illustrates the solution path for the ℓ_1 -norm-regularized estimate, providing the magnitude of its coefficients as a function of the regularization parameter λ . For $\lambda = 0$ we end up with the unregularized least-squares solution. It is apparent that an increase of λ fades any solution component to zero, hence the number of active coefficients gradually decreases. For this reason, the method is also referred to as the *least absolute shrinkage and selection operator* or *LASSO*.

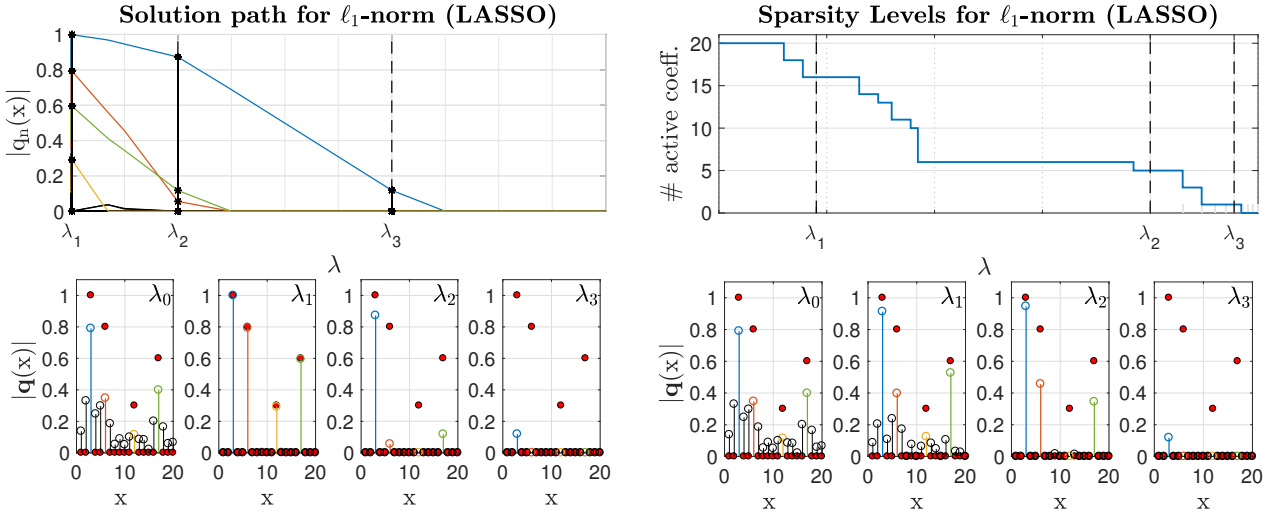


Figure 4.10: Solution path for the ℓ_1 -norm (left); Evolving sparsity levels (right).

With the penalty composed of a different linear operator, the characteristics of the solution and consequently also its evolution due to a change in the regularization parameter will differ. In case the linear operator is the gradient, block-sparsity is imposed on the solution vector where its components are gradually grouped together as λ increases. This comes along with their increasingly sparse first-order derivatives that induce a piecewise-constant shape. Smaller values of λ permit shapes that split into an increasing numbers of smaller groups with different constant amplitudes. For $\lambda = 0$, we obtain the fully ungrouped coefficients of the unregularized solution with no mismatch. This behaviour is depicted in Figure 4.11.

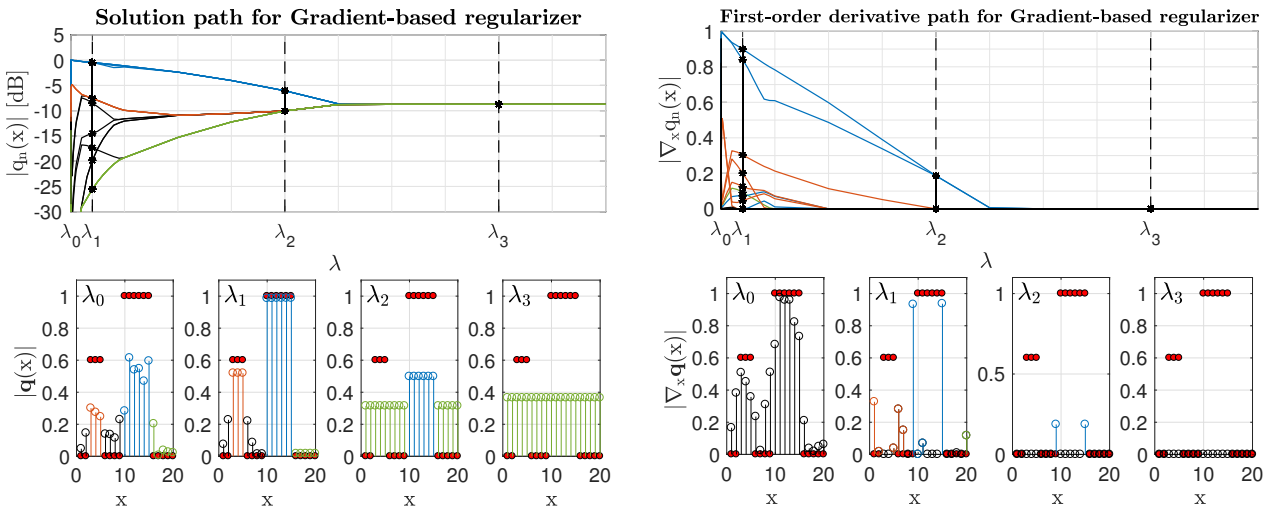


Figure 4.11: Solution path for the gradient-based TV regularizer (left); First-order derivative path (right).

A similar observation can be made when considering the Laplacian as the linear operator applied on the solution. An increase of λ again yields a gradual grouping of the coefficients, which, in contrast, now exhibit a linear profile. As increasing sparsity is enforced on the second-order derivatives, the first-order derivatives are fused into fewer groups of constant value. As a result, the number of piecewise-affine segments in the solution vector decreases until we end up with a single straight line for very large choices of λ . Conversely, for a decrease of λ the solution breaks up into an increasing number of small linear segments, causing the solution to exhibit a zig-zag

structure. For $\lambda = 0$, no regularization is applied, recovering the least-squares solution.

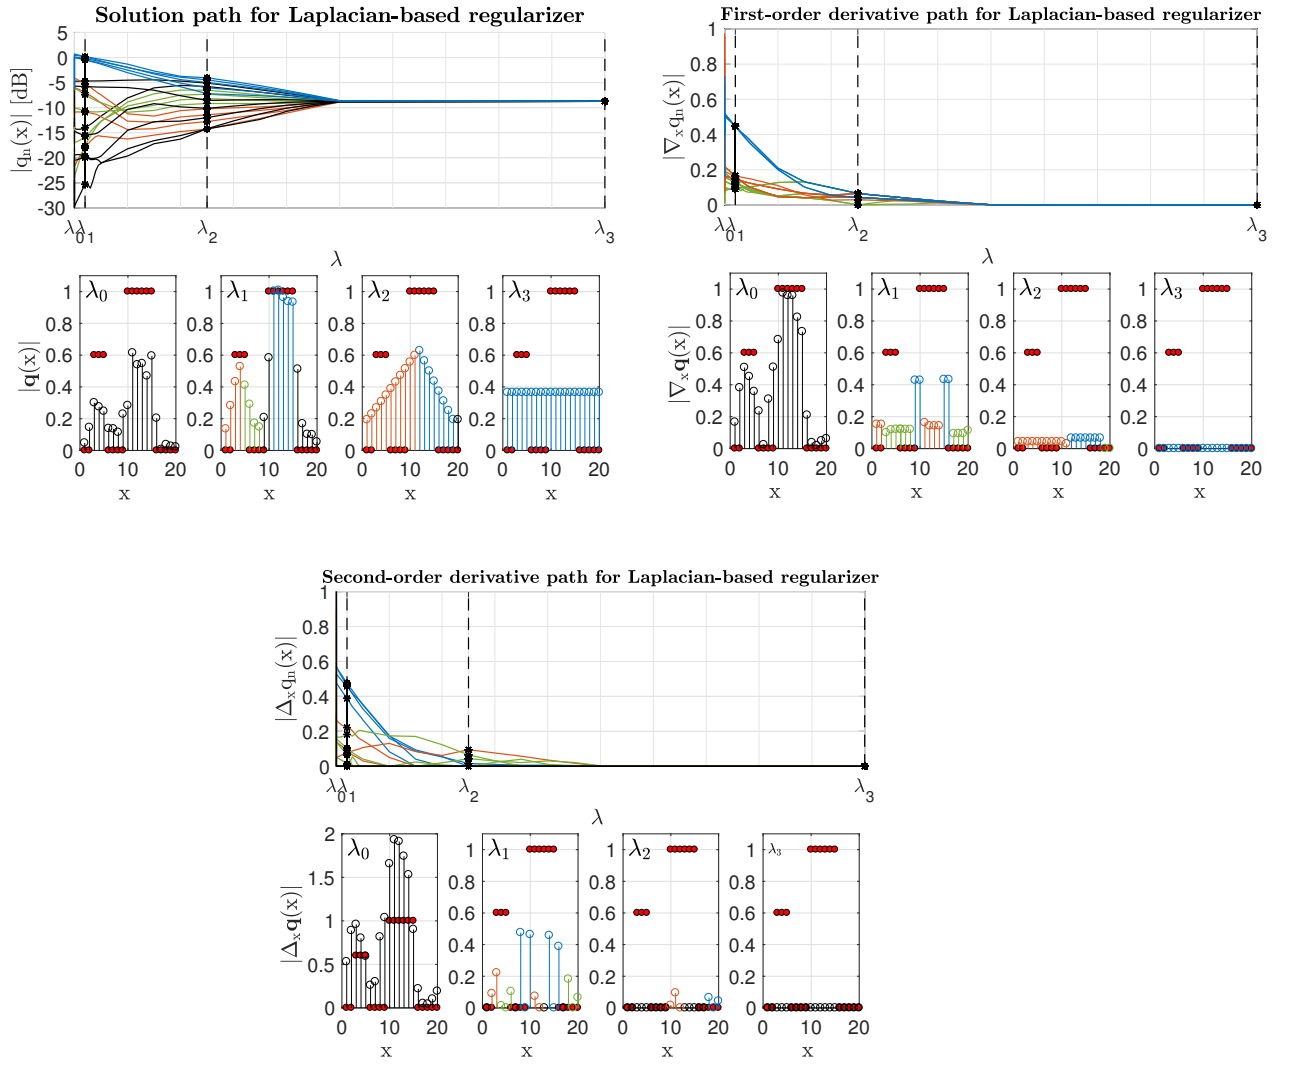


Figure 4.12: Solution path for the Laplacian-based TV2 regularizer (top-left); First-order derivative path (top-right); Second-order derivative path (bottom).

As a concluding remark, note that in neither of the derivative-based regularization methods sparsity is induced on the solution coefficients, but only on their consecutive differences. Hence, although they might be very small, the coefficients are always non-zero along the whole solution path. To moreover enforce sparsity on the spatial domain, an additional regularization term involving the ℓ_1 -norm must be applied directly on the solution.

4.5 Discretization of derivatives via finite differences




Figure 4.13: The three different types of finite difference (FD) approximations and boundary conditions (BC) for finite domains in x .

The discrete formulation of the objective function also requires the discretization of the linear operator involved in the regularization term. Without exception, the operators under study are all composed of first- or second-order partial derivatives of the source distribution. Since the source distribution consists of discrete point sources, the derivatives of the distribution function must be approximated with a numerical scheme. A common approach to approximate the partial derivatives within the operator is the use of finite differences [LeV07]. They rely on discarding the limit in the standard definition of the derivative, which is the limit of the difference quotient as the observed interval approaches zero. Instead, the infinitesimal interval is replaced by a sufficiently small but finite displacement that leads to a good approximation of the slope at the evaluation point. In general it holds that the higher the spatial resolution and thus the smaller the displacement, the more accurate the approximation. There are three typical slope estimators by first-order finite difference, the forward \vec{D}_x , central $\leftrightarrow D_x$ and backward \overleftarrow{D}_x difference

$$\vec{D}_x f(x) = \frac{f(x + \Delta x) - f(x)}{\Delta x}, \quad (4.28)$$

$$\leftrightarrow D_x f(x) = \frac{f(x + \Delta x) - f(x - \Delta x)}{2\Delta x}, \quad (4.29)$$

$$\overleftarrow{D}_x f(x) = \frac{f(x) - f(x - \Delta x)}{\Delta x}. \quad (4.30)$$

They all provide slopes of different secant lines that approximate the slope of the tangent to the graph at the evaluation point (see Figure 4.13). The estimation error can be specified by applying a Taylor series expansion, which approximates a differentiable function close to a given point in terms of its derivatives. This analysis shall now be exemplified by means of the forward finite difference operator. The Taylor series around a point x , describing the function in its

neighbourhood $x + \Delta x$, is given by

$$\begin{aligned}
 f(x + \Delta x) &= \sum_{n=0}^{\infty} \frac{f^{(n)}(x)}{n!} \Delta x^n \\
 &= f(x) + \Delta x f'(x) + \frac{1}{2} \Delta x^2 f''(x) + \frac{1}{6} \Delta x^3 f'''(x) + \underbrace{\mathcal{O}(\Delta x^4)}_{\text{higher order terms}}.
 \end{aligned} \tag{4.31}$$

Here $f^{(n)}(x)$ denotes the n -th derivative of f at the point x and $n!$ denotes the factorial of n . The substitution of (4.31) into (4.28) leads to

$$\begin{aligned}
 \vec{D}_x f(x) &= \frac{f(x + \Delta x) - f(x)}{\Delta x} = f'(x) + \frac{1}{2} \Delta x f''(x) + \frac{1}{6} \Delta x^2 f'''(x) + \mathcal{O}(\Delta x^3) \\
 \vec{D}_x f(x) - f'(x) &= \frac{1}{2} \Delta x f''(x) + \frac{1}{6} \Delta x^2 f'''(x) + \mathcal{O}(\Delta x^3) \stackrel{\Delta x \ll 1}{\approx} \frac{1}{2} \Delta x f''(x).
 \end{aligned} \tag{4.32}$$

Provided that Δx is sufficiently small, the truncation error, defined as the difference between the true derivative and its approximation, is mainly determined by the first residual $\frac{1}{2} \Delta x f''(x)$ while all remaining higher-order residuals are typically negligible. In the case of forward finite difference, the error is proportional to the interval Δx . A similar analysis can be carried out for backward finite difference, where the truncation error behaves the same, and for central difference, it is proportional to Δx^2 , thus yields a more accurate approximation, provided that Δx is small.

A possible way to also derive a finite difference representation of the second-order derivative is the repeated application of first-order differences. This leads to the following second-order expressions of forward \vec{D}_{xx} , central $\overleftrightarrow{D}_{xx}$ and backward \overleftarrow{D}_{xx} finite difference

$$\vec{D}_{xx} f(x) = \vec{D}_x \vec{D}_x f(x) = \frac{1}{\Delta x^2} \left[f(x + 2\Delta x) - 2f(x + \Delta x) + f(x) \right], \tag{4.33}$$

$$\overleftrightarrow{D}_{xx} f(x) = \vec{D}_x \overleftarrow{D}_x f(x) = \overleftarrow{D}_x \vec{D}_x f(x) = \frac{1}{\Delta x^2} \left[f(x + \Delta x) - 2f(x) + f(x - \Delta x) \right], \tag{4.34}$$

$$\overleftarrow{D}_{xx} f(x) = \overleftarrow{D}_x \overleftarrow{D}_x f(x) = \frac{1}{\Delta x^2} \left[f(x - 2\Delta x) - 2f(x - \Delta x) + f(x) \right]. \tag{4.35}$$

Another issue that is crucial for the quality of reconstruction is the choice of boundary conditions [Cap04] for the limits of the domain in x . Since the domain is finite, an appropriate boundary condition should be imposed to enable the computation of derivatives at its limits. A poor choice may introduce discontinuity at the domain boundaries. The source distribution map retrieved may contain ringing artifacts near its limits or spurious ghost sources.

There are three types of boundary conditions commonly used for numerical simulation: If *Dirichlet (zero) boundary conditions* are imposed, the outside boundary is assumed to be zero everywhere. If the actual source distribution does not tend to zero towards the boundaries, this type of boundary condition will introduce discontinuity in the spatial and differential domain that result in a poor reconstruction. The associated derivative operator can be implemented as a convolution matrix with block-Toeplitz structure. A slight improvement can be achieved when considering *periodic/circular boundary conditions* instead. They assume the periodic repetition of the source distribution in all directions, thus synthetically extending it in a cyclic manner.

These boundary conditions are particularly suitable if the underlying distribution is periodic by nature, which, for example, is the case for vibrational mode shapes. If, on the other hand, there is no relation between the solution values on the opposite boundaries, the periodic extension can also cause strong discontinuity at the borders. In case of periodic boundary conditions, the discrete derivative operator can be realized as a block-circulant matrix. The one type remaining is the *Neumann (reflective) boundary condition*, where the values outside the domain are treated as a reflection of the actual data inside, with the boundary serving as the mirror. These boundary conditions are capable of significantly reducing the ringing artifacts as the spatial discontinuities at the edges are annihilated. However, the differential distribution around the boundary may still contain discontinuities that could have a negative impact on the reconstruction. The corresponding derivative operator can be composed of four different types of block matrices. In general, the preferred choice of boundary conditions strongly depends on the nature of the signal to be differentiated as well as the level of measurement noise. The effect of the different boundary conditions on the reconstruction is demonstrated in Figure 4.14.




Figure 4.14: Effect of the boundary conditions (BC) - Dirichlet/zero (left); periodic/circular (center); Neumann/reflective (right).

When extending the concept of finite differences to the two-dimensional case, the task is greatly facilitated by considering the matrix representation of the source field rather than the vectorized one. Let the source distribution $q(x, y)$ be defined on a uniform grid of size $N_x \times N_y$ with the same spacing $\Delta x = \Delta y = h$ along both directions. For the definition of the finite difference operators, it is convenient in image processing to set the grid spacing to $h = 1$ such that each vertex depicts the size of the reference unit. According to their implementation, we specify the differential operators as a vectorized composition of n first- or second-order differences acting as the linear map $q \in \mathbb{C}^{N_x \times N_y} \mapsto \mathbf{u} \in \mathbb{C}^{(N_x \times N_y)^n}$. The operators of interest include the gradient $\nabla q = [D_x q, D_y q]$, the vectorized Laplacian $\Delta q = [D_{xx} q, D_{yy} q]$ and the vectorized Hessian $\mathcal{H} q = [D_{xx} q, D_{yy} q, D_{xy} q]$, where $D_{xy} = D_x(D_y) = D_y(D_x) = D_{yx}$.

Normally, one would use central finite differences due to the higher accuracy of approximation, but in this work the use of forward differences is considered instead. The reason for this lies in the fact that we need not only to discretize the differential operators presented above. For the further course of this study, there is an additional need of defining the discrete representation of the so called adjoint of the differential operator.

Adjoint differential operator

There exists an adjoint relation between the gradient and the divergence operator [Pap15,PSS13], according to which the adjoint of the gradient is the negative divergence of its generated vector field, and vice versa. This implies that they satisfy the following adjoint equality

$$-\operatorname{div}\mathbf{u}(x, y) \cdot f(x, y) = \mathbf{u}(x, y) \cdot \nabla f(x, y). \quad (4.36)$$

Here the divergence operator is equivalent to applying the Nabla operator ∇ on a vector field, which gives a scalar field in return and thus denotes the mapping $\mathbb{C}^{(N_x \times N_y)^2} \mapsto \mathbb{C}^{N_x \times N_y}$. In order to preserve the above relation, the divergence operator can be computed using backward finite differences when considering a forward finite difference implementation of the gradient, and vice versa. In doing so, the boundary conditions must be kept the same. In this particular case, the adjoint of the gradient operator can therefore be defined as

$$\nabla^* q(x, y) = -\operatorname{div}\mathbf{u}(x, y) = -\overleftarrow{D}_x u_1(x, y) - \overleftarrow{D}_y u_2(x, y). \quad (4.37)$$

In the same way, the discrete adjoint operator of a generalized vector ∇^2 comprised of second-order derivatives can be constructed, using a second-order extension of the adjoint equality in (4.36)

$$\operatorname{div}^2 \mathbf{u}(x, y) \cdot f(x, y) = \mathbf{u}(x, y) \cdot \nabla^2 f(x, y). \quad (4.38)$$

Similar to the first-order case, when the original operator is implemented using forward finite differences, the calculation of the squared divergence operator must be performed with second-order backward finite differences to satisfy the above relation, and vice versa. This leads to the following expressions for the adjoint of the vectorized Laplacian

$$\Delta^* q(x, y) = \operatorname{div}^2 \mathbf{u}(x, y) = \overleftarrow{D}_{xx} u_1(x, y) + \overleftarrow{D}_{yy} u_2(x, y), \quad (4.39)$$

which denotes the mapping $\mathbb{C}^{(N_x \times N_y)^2} \mapsto \mathbb{C}^{N_x \times N_y}$, as well as the adjoint of the vectorized Hessian

$$\mathcal{H}^* q(x, y) = \operatorname{div}^2 \mathbf{u}(x, y) = \overleftarrow{D}_{xx} u_1(x, y) + \overleftarrow{D}_{yy} u_2(x, y) + 2\overleftarrow{D}_{xy} u_3(x, y), \quad (4.40)$$

which denotes the mapping $\mathbb{C}^{(N_x \times N_y)^3} \mapsto \mathbb{C}^{N_x \times N_y}$. When considering an additional sparsity constraint in the spatial domain, the identity operator \mathbf{I} may be incorporated into the regularizer. As the identity operator is known to be self-adjoint, its adjoint operator is given by $\mathbf{I}^* = \mathbf{I}$.

5

The Algorithm - Proximal Splitting

Despite there exist toolboxes for convex optimization (CVX), the targeted optimization tasks (Hessian Schatten-norm, etc.) and the required algorithms, extensions, and combinations were reason enough in this work to implement all optimization from scratch in MATLAB. To gather all required techniques, implementation requires profound theoretical knowledge that is laid out in this chapter, with all the necessary explanations and references for further reading. The variety of considered regularizers and fuzzy approaches can be dealt with efficiently by utilizing the concept of the proximity operator, which is the main ingredient of a class of optimization algorithms called proximal methods. They have gained great popularity in recent years, mostly due to their capability of minimizing possibly non-smooth cost functions and their great efficiency when solving large-scale problems. In return, they may exhibit a slower convergence rate compared to other state-of-the-art methods. Moreover, most of these algorithms are initially designed for optimizing objective functions of real-valued variables. Although a complex-valued vector quantity can always be expressed as a real-valued vector with twice the dimensions by stacking of the real and imaginary components, this involves a consistent remapping of the variables leading to unnecessary computation steps. To cope with this issue, the extension of various proximity operators to the complex-valued case became subject of many recent studies. In this thesis, we adopt a full splitting based primal-dual algorithm that only exploits first-order information of the objective function. A detailed exposition of its use for the cases under study is presented and closed-form expressions for the proximity operator of the involved norm functionals are given. For simplicity, the theory is laid out for the real-valued case and is only extended to the complex-valued cases of interest when discussing the practical application of the algorithm. For more detailed information on the topics of convex optimization, duality and proximal algorithms, the interested reader is referred to the standard literature [BV09, Bec17, PB14, Roc74]. A more compact description of the topic and a survey of different proximal methodologies can be found in [KP14, CP09].

5.1 Duality theory

5.1.1 Dual space and dual norm

For a general understanding of the proposed algorithm, it is necessary to provide the reader with the definition of the dual space and its immanent norms [Bec17]. Let \mathbb{X} be a vector space equipped with an inner product $\langle \cdot, \cdot \rangle$ and a norm $\|\cdot\|$. Any linear functional f on that vector space is a linear transformation $f : \mathbb{X} \mapsto \mathbb{R}$ that maps a point $\mathbf{x} \in \mathbb{X}$ to the scalar field of real numbers \mathbb{R} . The set of all possible linear functionals on \mathbb{X} is considered as the dual space \mathbb{X}^* such that $f \in \mathbb{X}^*$. Any linear functional on an inner product space \mathbb{X} can be alternatively expressed by

$$f(\mathbf{x}) = \langle \mathbf{v}, \mathbf{x} \rangle = \mathbf{v}^T \mathbf{x} \quad (5.1)$$

as there is always a unique vector $\mathbf{v} \in \mathbb{X}$ that satisfies this equality. This relationship is a reasonable justification to also define the dual space \mathbb{X}^* as a vector space that inherits the same structure and operations as \mathbb{X} . The set of linear functionals f is now represented by the collection of vectors $\mathbf{y} \in \mathbb{X}^*$.

However, the major difference between the two spaces lies in the use of norms. In dual space, the vector norm of \mathbf{y} must be a quantification of a linear functional f in terms of size which of course is always relative to the size of its input vector \mathbf{x} . To provide a meaningful measure, it is considered the largest possible value of the functional applied to \mathbf{x} where \mathbf{x} ranges over all elements in the closed unit ball of a normed vector space. Hence the norm of the dual space, referred to as the dual norm, is defined as

$$\|\mathbf{y}\|_* = \sup_{\mathbf{x}} \{\mathbf{y}^T \mathbf{x} : \|\mathbf{x}\| \leq 1\} = \sup_{\mathbf{x} \in \mathcal{B}} \{\mathbf{y}^T \mathbf{x}\} = S_{\mathcal{B}}, \quad (5.2)$$

which can be considered as the support function $S_{\mathcal{B}}$ of the unit norm ball $\mathcal{B} = \{\mathbf{x} \in \mathbb{X} : \|\mathbf{x}\| \leq 1\}$. In a finite-dimensional vector space, the dual of the dual norm always yields the original norm, thus $\|\cdot\|_{**} = \|\cdot\|$.

In general, any vector space \mathbb{R}^N equipped with an arbitrary ℓ_p -norm admits a dual space, where the corresponding dual norm is the ℓ_q -norm. Here $p, q \geq 1$ are conjugate exponents such that $\frac{1}{p} + \frac{1}{q} = 1$. This leads to the following dual representations of the three types of ℓ_p -norms frequently encountered in this thesis, where $p \in \{1, 2, \infty\}$ [BV09].

The ℓ_2 -norm or Euclidean norm $\|\cdot\|_2$ is said to be self-dual as it is equivalent to its dual norm, since

$$\|\mathbf{y}\|_{2,*} = \sup_{\mathbf{x}} \{\mathbf{y}^T \mathbf{x} : \|\mathbf{x}\|_2 \leq 1\} = \|\mathbf{y}\|_2. \quad (5.3)$$

The dual norm to the ℓ_∞ -norm $\|\cdot\|_\infty$ is the ℓ_1 -norm $\|\cdot\|_1$, as

$$\|\mathbf{y}\|_{\infty,*} = \sup_{\mathbf{x}} \{\mathbf{y}^T \mathbf{x} : \|\mathbf{x}\|_\infty \leq 1\} = \sum_i y_i = \|\mathbf{y}\|_1. \quad (5.4)$$

Since the dual of the dual is the primal, the dual of the ℓ_1 -norm yields the ℓ_∞ -norm, hence they are mutually dual

$$\|\mathbf{y}\|_{1,*} = \sup_{\mathbf{x}} \{\mathbf{y}^T \mathbf{x} : \|\mathbf{x}\|_1 \leq 1\} = \|\mathbf{y}\|_\infty. \quad (5.5)$$

Note that minimizing the dual norm is not equivalent to minimizing its corresponding primal norm, but it often can be a helpful alternative optimization task as it minimizes its upper bound.

5.1.2 Legendre-Fenchel conjugate and Moreau subdifferential

In many optimization problems the solution process may be simplified considerably by utilizing the concept of duality. The fundamental idea is to derive an abstract representation of the original problem in another domain that can be solved alternatively to the primal problem whenever it improves convergence. This is especially useful for optimization tasks involving non-smooth functions that are not uniformly differentiable. A central point of duality is the construction of the Legendre-Fenchel conjugate [KP14, BV09], defined as

$$f^*(\mathbf{u}) = \sup_{\mathbf{x} \in \mathbb{R}^N} \{\mathbf{x}^T \mathbf{u} - f(\mathbf{x})\}. \quad (5.6)$$

This transformation naturally leads to the corresponding dual representation $f^*(\mathbf{u})$ of the primal problem $f(\mathbf{x})$ at hand. It is a map between the two spaces of the primal variable \mathbf{x} and the dual variable \mathbf{u} , where \mathbf{u} represents the slope of $f(\mathbf{x})$ evaluated at point \mathbf{x} . Every slope \mathbf{u} defines a linear function $\mathbf{x}^T \mathbf{u}$ for which we target the point \mathbf{x} of $f(\mathbf{x})$ where the gap between the functions attains its maximum value. At this point the slope \mathbf{u} provides a tangent to $f(\mathbf{x})$ that intersects the ordinate at the negative of this maximum. Hence the Legendre-Fenchel conjugate $f^*(\mathbf{u})$ can be interpreted as the negative y-intercept of the tangent to the graph of $f(\mathbf{x})$ at point \mathbf{x} . Figure 5.1 provides two graphical examples of this idea. As the supremum of affine functions, the Legendre-Fenchel conjugate is always convex, no matter of the type of primal function. For this reason it is also often referred to as the convex conjugate.




Figure 5.1: Graphical interpretation of the Legendre-Fenchel conjugate.

When dealing with non-smooth functions, there are several ways to assign a tangent to points of discontinuity as the derivatives at those points, giving the tangential slopes, are not well-defined. In this case, we have to utilize the concept of the subgradient as a generalized notation of the gradient for non-differentiable convex functions. Any subgradient defines an affine tangent function that globally minorizes $f(\mathbf{x})$ outside the point of contact \mathbf{x} . The set of all possible subgradients \mathbf{u} of $f(\mathbf{x})$ at point \mathbf{x} is called the Moreau subdifferential defined as

$$\partial f(\mathbf{x}) = \{\mathbf{u} \in \mathbb{R}^N \mid (\forall \mathbf{y} \in \mathbb{R}^N) \quad f(\mathbf{y}) \geq f(\mathbf{x}) + \mathbf{u}^T (\mathbf{y} - \mathbf{x})\}. \quad (5.7)$$

Note that if $f(\mathbf{x})$ is a convex, differentiable function, its subgradient is uniquely defined everywhere as the gradient and the subdifferential reduces to $\partial f(\mathbf{x}) = \{\nabla f(\mathbf{x})\}$. A geometrical interpretation of this concept is illustrated in Figure 5.2.




Figure 5.2: Graphical interpretation of the Moreau subdifferential.

The dual variable \mathbf{u} in equation (5.6) can take on any of the subgradients out of the subdifferential $\partial f(\mathbf{x})$ at point \mathbf{x} such that the supremum is obtained.

At last we shall define the biconjugate, considered as the conjugate of the conjugate, which under a few assumptions² again yields the initial function

$$f^{**}(\mathbf{x}) = f(\mathbf{x}) = \sup_{\mathbf{u} \in \mathbb{R}^N} \{\mathbf{x}^T \mathbf{u} - f^*(\mathbf{u})\}. \quad (5.8)$$

To close on the issue of the Legendre-Fenchel conjugate, its definition for a generic norm $\|\cdot\|$ is provided which counterintuitively does not coincide with the dual norm. This becomes more obvious when we recall the definition of a norm in equation (5.2) as the support function of the unit ball of its dual norm. The convex conjugate of a support function of a closed, convex set \mathcal{C} is the indicator function of the set [PB14, Bec17], hence $S_{\mathcal{C}}^* = i_{\mathcal{C}}$, which is defined as

$$i_{\mathcal{C}}(\mathbf{x}) = \begin{cases} 0 & \mathbf{x} \in \mathcal{C} \\ \infty & \mathbf{x} \notin \mathcal{C} \end{cases}. \quad (5.9)$$

For this reason, the conjugate of any norm yields the indicator function of the unit ball of its dual norm

$$\|\mathbf{x}\|^* = \left(\sup_{\mathbf{x} \in \mathcal{B}_*} \{\mathbf{v}^T \mathbf{x}\} \right)^* = S_{\mathcal{B}_*}^* = i_{\mathcal{B}_*}, \quad (5.10)$$

where $\mathcal{B}_* = \{\mathbf{x} \in \mathbb{R}^N : \|\mathbf{x}\|_* \leq 1\}$.

This turns out to be an important property for the computation of the so called proximity operator of a norm, subsequently introduced in this chapter as a crucial part of the proposed algorithm.

² proper, lower-semicontinuous, convex function

5.2 The projection operator

In advance of providing the definition of the projection operator, let us first consider a simple minimization problem of an arbitrary objective function $f(\mathbf{x}) : \mathbb{R}^N \mapsto \mathbb{R}$ with respect to the vector $\mathbf{x} \in \mathbb{R}^N$. Let \mathcal{C} be a closed, nonempty and convex subset of \mathbb{R}^N . If we want to restrict the solution to lie in the domain of \mathcal{C} , this can be achieved either by using a constrained formulation or by introducing an auxiliary function to the unconstrained formulation of the problem, such that

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathcal{C}} f(\mathbf{x}) = \arg \min_{\mathbf{x}} \left\{ f(\mathbf{x}) + i_{\mathcal{C}}(\mathbf{x}) \right\}, \quad (5.11)$$

where $i_{\mathcal{C}}(\cdot)$ denotes the indicator function of the corresponding set \mathcal{C} .

Now let $f(\mathbf{x})$ be a measure of the Euclidean distance between $\mathbf{x} \in \mathbb{R}^N$ and an arbitrary point $\mathbf{u} \in \mathcal{C}$ inside the subset. The solution to the resulting minimization problem is given by the evaluation of the projection operator [CP09, Bec17]

$$P_{\mathcal{C}}(\mathbf{x}) = \hat{\mathbf{x}} = \arg \min_{\mathbf{u} \in \mathcal{C}} \frac{1}{2} \|\mathbf{u} - \mathbf{x}\|_2^2 = \arg \min_{\mathbf{u}} \left\{ \frac{1}{2} \|\mathbf{u} - \mathbf{x}\|_2^2 + i_{\mathcal{C}}(\mathbf{u}) \right\}. \quad (5.12)$$

It can be regarded as the orthogonal projection of \mathbf{x} onto \mathcal{C} as it provides the point $\hat{\mathbf{x}} \in \mathcal{C}$ that is closest to the initial vector. This point is always the unique minimum solution of (5.12) provided that \mathcal{C} is a closed, convex set. Note that this property is satisfied by any ℓ_p -norm ball for $p \geq 1$. Figure 5.3 provides a geometric interpretation of the projection operator by means of a projection of a vector onto the scaled ℓ_2 -norm ball.




Figure 5.3: Illustration of the orthogonal projection onto the convex set $\{\mathcal{B}_2 : \|\mathbf{x}\|_2 \leq \lambda, \mathbf{x} \in \mathbb{R}^2\}$ describing the λ -scaled unit ℓ_2 -norm ball.

5.2.1 Projections onto vector- and matrix-norm balls

This section is dedicated to providing examples of the projection operator for the orthogonal projection of vectors onto ℓ_p -norm balls for the cases $p \in \{1, 2, \infty\}$ [PB14]. In addition, the close connection between the projection of vectors onto ℓ_p -norm balls and matrices onto \mathcal{S}_p -norm (Schatten-norm) balls is explained [LWU13]. Note that in both cases the norm balls can be of arbitrary size determined by the norm scaling factor λ .

Scaling projects a vector onto the ℓ_2 -norm or Euclidean norm ball $\mathcal{B}_2 = \{\mathbf{x} \in \mathbb{R}^N : \|\mathbf{x}\|_2 \leq \lambda\}$ by a simple normalization to unit length and a subsequent scaling by the radius λ if the vector is located outside the ball. No further change must be applied on a vector within the ball. This leads to the following expression for the projection operator

$$P_{\mathcal{B}_2}(\mathbf{x}) = \begin{cases} \lambda \frac{\mathbf{x}}{\|\mathbf{x}\|_2} & \|\mathbf{x}\|_2 > \lambda \\ \mathbf{x} & \|\mathbf{x}\|_2 \leq \lambda \end{cases} = \frac{\mathbf{x}}{\max(1, \frac{\|\mathbf{x}\|_2}{\lambda})}. \quad (5.13)$$

Limitation of its entries projects a vector onto the ℓ_∞ -norm ball $\mathcal{B}_\infty = \{\mathbf{x} \in \mathbb{R}^N : \|\mathbf{x}\|_\infty \leq \lambda\}$. The associated projection operator is the component-wise operation

$$\left(P_{\mathcal{B}_\infty}(\mathbf{x})\right)_i = \begin{cases} \lambda & x_i > \lambda \\ x_i & |x_i| \leq \lambda \\ -\lambda & x_i < -\lambda \end{cases} \quad (5.14)$$

which can be written in a more compact notation as

$$P_{\mathcal{B}_\infty}(\mathbf{x}) = \min(\lambda, \max(-\lambda, \mathbf{x})). \quad (5.15)$$

An iterative soft-thresholding operator projects a vector onto the ℓ_1 -norm ball $\mathcal{B}_1 = \{\mathbf{x} \in \mathbb{R}^N : \|\mathbf{x}\|_1 \leq \lambda\}$ and it involves a threshold that is not known in advance. Its estimation is the task of a separate operation which makes the application of the projection operator more expensive. Literature provide alternative, fast and efficient computation of the projection onto ℓ_1 -norm balls. The projection is the soft-thresholding operator

$$\left(P_{\mathcal{B}_1}(\mathbf{x})\right)_i = \begin{cases} x_i - \gamma & x_i > \gamma \\ x_i & |x_i| \leq \gamma \\ x_i + \gamma & x_i < -\gamma \end{cases} \quad (5.16)$$

which is applied component-wise and can be equivalently expressed in condensed notation as

$$P_{\mathcal{B}_1}(\mathbf{x}) = \frac{\mathbf{x}}{\|\mathbf{x}\|_1} \max(\|\mathbf{x}\|_1 - \gamma, 0) = \text{sgn}(\mathbf{x}) \max(\|\mathbf{x}\|_1 - \gamma, 0). \quad (5.17)$$

If the vector lies in the interior of the norm ball, the threshold γ becomes zero. Otherwise, it can be obtained by solving the following equality

$$\sum_i \max(|x_i| - \gamma, 0) = \lambda. \quad (5.18)$$

Apart from the projection of vectors onto vector ℓ_p -norm balls, the targeted application also

considers the projection of matrices onto Schatten \mathcal{S}_p -norm balls. The projection of a matrix onto the \mathcal{S}_p -norm ball is directly associated with the projection of its singular values onto the corresponding ℓ_p -norm ball. In particular, let $\mathbf{X} \in \mathbb{C}^{M \times N}$ be an arbitrary matrix expressed by the SVD $\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^H$ and the set $\mathcal{B}_{\mathcal{S}_p} = \{\mathbf{X} \in \mathbb{C}^{M \times N} : \|\mathbf{X}\|_{\mathcal{S}_p} \leq \lambda\}$ be the norm ball for a general \mathcal{S}_p -norm. Then the projection of \mathbf{X} onto $\mathcal{B}_{\mathcal{S}_p}$ is defined as

$$P_{\mathcal{B}_{\mathcal{S}_p}}(\mathbf{X}) = \mathbf{U} \text{diag}(P_{\mathcal{B}_p}(\boldsymbol{\sigma}(\mathbf{X}))) \mathbf{V}^H, \quad (5.19)$$

where $\boldsymbol{\sigma}(\mathbf{X})$ is a vector of the singular values of \mathbf{X} , $P_{\mathcal{B}_p}(\cdot)$ denotes the projection of a vector onto the ℓ_p -norm ball and $\text{diag}(\cdot)$ creates a matrix with diagonal elements equal to the entries of the vector the function is applied on.

We may as well provide closed-form expressions for projecting onto \mathcal{S}_p -norm balls for the particular cases where $p \in \{1, 2, \infty\}$ [LWU13, LU13]. The orthogonal projection onto the \mathcal{S}_2 -norm ball $\mathcal{B}_{\mathcal{S}_2} = \{\mathbf{X} \in \mathbb{C}^{M \times N} : \|\mathbf{X}\|_{\mathcal{S}_2} \leq \lambda\}$ can be achieved without the necessity of SVD by simply normalizing all matrix elements to a certain Frobenius norm, as the Frobenius (norm of all matrix elements) equals the ℓ_2 -norm of $\boldsymbol{\sigma}$,

$$P_{\mathcal{B}_{\mathcal{S}_2}}(\mathbf{X}) = \begin{cases} \lambda \frac{\mathbf{X}}{\|\mathbf{X}\|_F} & \|\mathbf{X}\|_F > \lambda \\ \mathbf{X} & \|\mathbf{X}\|_F \leq \lambda \end{cases} = \frac{\mathbf{X}}{\max(1, \frac{\|\mathbf{X}\|_F}{\lambda})}. \quad (5.20)$$

In this context $\|\cdot\|_F = \|\cdot\|_{\mathcal{S}_2}$ denotes the Frobenius norm of a matrix that coincides with its Schatten \mathcal{S}_2 -norm.

To project a matrix onto the \mathcal{S}_∞ -norm ball $\mathcal{B}_{\mathcal{S}_\infty} = \{\mathbf{X} \in \mathbb{C}^{M \times N} : \|\mathbf{X}\|_{\mathcal{S}_\infty} \leq \lambda\}$ its singular values are projected onto the ℓ_∞ -norm ball \mathcal{B}_∞ . Since the singular values of any matrix are always non-negative, real numbers, limitation only addresses their positive boundary value λ

$$P_{\mathcal{B}_{\mathcal{S}_\infty}}(\mathbf{X}) = \mathbf{U} \text{diag}(\min(\lambda, \boldsymbol{\sigma}(\mathbf{X}))) \mathbf{V}^H. \quad (5.21)$$

At last, the matrix projection onto the \mathcal{S}_1 -norm ball $\mathcal{B}_{\mathcal{S}_1} = \{\mathbf{X} \in \mathbb{C}^{M \times N} : \|\mathbf{X}\|_{\mathcal{S}_1} \leq \lambda\}$ is considered which does not universally admit a closed form solution as it involves the orthogonal projection of singular values onto the ℓ_1 -norm ball \mathcal{B}_1 . This is usually achieved through soft-thresholding by applying the associated operator given in equation (5.17). This yields the following general solution

$$P_{\mathcal{B}_{\mathcal{S}_1}}(\mathbf{X}) = \mathbf{U} \text{diag}(\max(\boldsymbol{\sigma}(\mathbf{X}) - \gamma, 0)) \mathbf{V}^H, \quad (5.22)$$

which requires a numerical computation of the suitable threshold level γ .

However, for matrices of small size with a modest number of singular values, as it is the case for the discrete Hessian frequently used throughout this thesis, we are still able to derive a solution for the threshold level in compact form. With that said, the particular case of projecting a square matrix $\mathbf{X} \in \mathbb{C}^{2 \times 2}$ onto the \mathcal{S}_1 -norm ball involves soft-thresholding with the threshold defined as

$$\gamma = \begin{cases} 0 & \sigma_1(\mathbf{X}) \leq \lambda - \sigma_2(\mathbf{X}) \\ \frac{\sigma_1(\mathbf{X}) + \sigma_2(\mathbf{X}) - \lambda}{2} & \lambda - \sigma_2(\mathbf{X}) < \sigma_1(\mathbf{X}) \leq \lambda + \sigma_2(\mathbf{X}) \\ \sigma_1(\mathbf{X}) - \lambda & \sigma_1(\mathbf{X}) > \lambda + \sigma_2(\mathbf{X}) \end{cases} \quad (5.23)$$

where $\sigma_1(\mathbf{X}) \geq \sigma_2(\mathbf{X})$.

5.3 The proximity operator

Utilizing the concept of the projection operator, Moreau [Mor65] suggested the use of an arbitrary function $g(\cdot)$ replacing the indicator function $i_C(\cdot)$ to expand its applicability to a larger class of problems. This generalization is referred to as the Moreau proximity operator or proximal map defined as

$$\text{prox}_{\tau g}(\mathbf{x}) = \hat{\mathbf{x}} = \arg \min_{\mathbf{u}} \left\{ \frac{1}{2} \|\mathbf{u} - \mathbf{x}\|_2^2 + \tau g(\mathbf{u}) \right\}. \quad (5.24)$$

It can be viewed as the optimum solution of a regularized minimization problem which is unique as long as $g(\cdot) : \mathbb{R}^N \mapsto \mathbb{R}$ represents a closed, convex function. This solution $\hat{\mathbf{x}}$ minimizes $g(\cdot)$ to some extent while simultaneously being located in the vicinity of \mathbf{x} . The optional parameter τ regulates the amount of coefficient movement towards the minimum of $g(\cdot)$ thus serving as some sort of step size. The operating principle of the proximal mapping is illustrated in Figure 5.4.




Figure 5.4: Illustration of the concept of proximal mapping.

The evaluation of the proximity operator does not always require numerical computation as simple closed-form expressions are available for a variety of functions.

A key property is the so called Moreau decomposition or Moreau's identity which denotes the link between the proximity operator of a function $g(\cdot)$ and the proximity operator of its convex conjugate $g(\cdot)^*$

$$\text{prox}_{\tau g}(\mathbf{x}) = \mathbf{x} - \tau \text{prox}_{\frac{g^*}{\tau}}\left(\frac{\mathbf{x}}{\tau}\right). \quad (5.25)$$

In many cases it is easier to obtain the proximity operator using this relation, as the conjugate function might provide a simplified proximal mapping. Since this thesis is mainly concerned with the minimization of norms, Moreau's identity for the associated proximity operators is subject to further investigation. From duality theory we know that the convex conjugate of a generic vector norm $g(\cdot) = \|\cdot\|_p$ is the indicator function $g(\cdot)^* = i_{\mathcal{B}_{p,*}}$ of the dual norm unit ball defined by the set $\{\mathcal{B}_{p,*} : \|\mathbf{x}\|_{p,*} \leq 1, \mathbf{x} \in \mathbb{R}^N\}$. In this case, the proximity operator reduces to a simple projection onto this set, as stated in the previous section. According to the Moreau decomposition, it follows that

$$\text{prox}_{\tau \|\cdot\|_p}(\mathbf{x}) = \mathbf{x} - \tau P_{\mathcal{B}_{p,*}}\left(\frac{\mathbf{x}}{\tau}\right) = \mathbf{x} - P_{\mathcal{B}_{p,*} \leq \tau}(\mathbf{x}), \quad (5.26)$$

where τ defines the radius of the dual norm ball. It reveals that the proximity operator associated with a norm can always be interpreted as the residual of a projection onto the corresponding dual norm ball.

5.3.1 Examples of proximal mappings

Equipped with the knowledge that the proximity operator of a norm is closely tied to the projection onto the dual norm ball, we can readily derive analytic expressions for the cases of interest [PB14].

As the ℓ_2 -norm is self-dual, its proximity operator involves the projection onto the Euclidean norm ball $\mathcal{B}_{2,*} = \mathcal{B}_2$ and leads to the occasionally called block-soft-thresholding operator

$$\text{prox}_{\lambda\|\cdot\|_2}(\mathbf{x}) = \mathbf{x} - \lambda P_{\mathcal{B}_2}\left(\frac{\mathbf{x}}{\lambda}\right) = \mathbf{x} - P_{\mathcal{B}_{2,\leq\lambda}}(\mathbf{x}) = \max\left(1 - \frac{\lambda}{\|\mathbf{x}\|_2}, 0\right)\mathbf{x}. \quad (5.27)$$

The proximity operator of the ℓ_1 -norm is associated with the projection onto the ℓ_∞ -norm ball $\mathcal{B}_{1,*} = \mathcal{B}_\infty$ and can be evaluated by the soft-thresholding operator of equation (5.17), therefore we get

$$\text{prox}_{\lambda\|\cdot\|_1}(\mathbf{x}) = \mathbf{x} - \lambda P_{\mathcal{B}_\infty}\left(\frac{\mathbf{x}}{\lambda}\right) = \mathbf{x} - P_{\mathcal{B}_{\infty,\leq\lambda}}(\mathbf{x}) = \text{sgn}(\mathbf{x}) \max(|\mathbf{x}| - \lambda, 0). \quad (5.28)$$

Since the proximity operator of the ℓ_∞ -norm entails a projection onto the ℓ_1 -norm ball $\mathcal{B}_{\infty,*} = \mathcal{B}_1$ its computation is more involved as it requires a subroutine for the estimation of the threshold parameter. Nevertheless, the operator exhibits the following generic form

$$\text{prox}_{\lambda\|\cdot\|_\infty}(\mathbf{x}) = \mathbf{x} - \lambda P_{\mathcal{B}_1}\left(\frac{\mathbf{x}}{\lambda}\right) = \mathbf{x} - P_{\mathcal{B}_{1,\leq\lambda}}(\mathbf{x}). \quad (5.29)$$

5.4 Linear operator and adjoint

Maintaining the definition of the inner product spaces \mathbb{X} and \mathbb{Y} , we now further specify the concept of linear transformations between two vector spaces [BB98, Bec17]. Let $\mathcal{A} : \mathbb{X} \mapsto \mathbb{Y}$ define an arbitrary transformation from \mathbb{X} to \mathbb{Y} , whose application on a vector $\mathbf{x} \in \mathbb{X}$ results in its mapping onto a certain vector $\mathbf{y} \in \mathbb{Y}$. The transformation \mathcal{A} is said to be *linear*, if it satisfies the following property for any pair of vectors in its domain and the arbitrary constants $\alpha, \beta \in \mathbb{R}$:

$$\mathcal{A}(\alpha\mathbf{x}_1 + \beta\mathbf{x}_2) = \alpha\mathcal{A}(\mathbf{x}_1) + \beta\mathcal{A}(\mathbf{x}_2). \quad (5.30)$$

Every such linear transformation can be expressed in condensed notation using a linear operator or matrix, such that

$$\mathcal{A}(\mathbf{x}) = \mathbf{A}\mathbf{x}. \quad (5.31)$$

Note that through linear transformation it is also possible to assign a vector to a matrix or even to a multidimensional array when applied elementwise, as it is the case for the Hessian operator implemented during this thesis.

A linear operator is said to be *bounded*, if there exists a real constant M for which the following inequality holds for every $\mathbf{x} \in \mathbb{X}$

$$\|\mathbf{A}\mathbf{x}\|_{\mathbb{Y}} \leq M\|\mathbf{x}\|_{\mathbb{X}}. \quad (5.32)$$

The smallest number M for which equation (5.32) is fulfilled is determined by the induced operator norm $\|\cdot\|$ of \mathbf{A} , hence M provides an upper bound on $\|\mathbf{A}\|$. In this case, equality holds and the bound M can be calculated as

$$M = \|\mathbf{A}\| = \sup_{\mathbf{x} \in \mathbb{X} \setminus \{0\}} \frac{\|\mathbf{Ax}\|_{\mathbb{Y}}}{\|\mathbf{x}\|_{\mathbb{X}}} = \sup_{\|\mathbf{x}\|_{\mathbb{X}}=1} \|\mathbf{Ax}\|_{\mathbb{Y}} = \sqrt{\lambda_{\max}(\mathbf{A}^T \mathbf{A})} = \sigma_{\max}(\mathbf{A}), \quad (5.33)$$

where $\lambda_{\max}(\cdot)$ denotes the largest eigenvalue and $\sigma_{\max}(\cdot)$ the largest singular value of the operator in their respective argument.

The transpose of the linear transformation \mathcal{A} describes a mapping $\mathcal{A}^T : \mathbb{Y}^* \mapsto \mathbb{X}^*$ between the dual spaces that acts in the opposite direction. Since the dual spaces encompass the set of all linear functionals $f_x \in \mathbb{X}^*$ and $f_y \in \mathbb{Y}^*$ on the respective primal spaces, the application of the map \mathcal{A}^T on such linear functional f_y can be regarded as a linear functional $f_x = \mathcal{A}^T(f_y)$ itself, living in the dual space \mathbb{X}^* .

A concept that generalizes the transpose is the adjoint of a linear transformation, denoted as \mathcal{A}^* . It is defined as the backward mapping $\mathcal{A}^* : \mathbb{Y} \mapsto \mathbb{X}$ that yields the unique vector in \mathbb{X} associated with the functional $f_x = \mathcal{A}^T(f_y)$. It essentially converts a vector $\mathbf{y} \in \mathbb{Y}$ into its functional equivalent $f_y \in \mathbb{Y}^*$, uses the transposed transformation \mathcal{A}^T to obtain the associated functional $f_x = \mathcal{A}^T(f_y) \in \mathbb{X}^*$ and eventually yields its corresponding unique vector expression in \mathbb{X} . Accordingly, the adjoint \mathcal{A}^* is defined as the transformation that satisfies the equality

$$\langle \mathbf{y}, \mathbf{Ax} \rangle_{\mathbb{Y}} = \langle \mathbf{A}^* \mathbf{y}, \mathbf{x} \rangle_{\mathbb{X}}. \quad (5.34)$$

For any linear, bounded operator $\mathbf{A} : \mathbb{X} \mapsto \mathbb{Y}$ there always exists a unique adjoint operator $\mathbf{A}^* : \mathbb{Y} \mapsto \mathbb{X}$ with equal operator norm $\|\mathbf{A}\| = \|\mathbf{A}^*\|$. If $\mathbf{A} = \mathbf{A}^*$, the operator is said to be self-adjoint, which, for instance, is the case for the identity operator \mathbf{I} .

Figure 5.5 provides a graphical illustration of the linear mappings between the primal and dual spaces described above.




Figure 5.5: Commutative diagram of linear transformations between two vector spaces and their dual spaces.

In this thesis we encounter the particular case of a linear mapping $\mathbf{A} : \mathbb{X} \mapsto \mathbb{M}$, where $\mathbb{X} = \mathbb{R}^N$ and $\mathbb{M} = \mathbb{R}^{N \times M \times K}$, which transforms a vector $\mathbf{x} \in \mathbb{X}$ into a vector $\mathbf{X} \in \mathbb{M}$ whose elements are

matrices \mathbf{X}_i of dimension $M \times K$. Let the space \mathbb{M} be endowed with the inner product

$$\langle \mathbf{X}, \mathbf{Y} \rangle_{\mathbb{M}} = \sum_{i=1}^N \text{trace}(\mathbf{Y}_i^T \mathbf{X}_i), \quad (5.35)$$

where $\mathbf{X}, \mathbf{Y} \in \mathbb{M}$ are two arbitrary vectors comprised of matrices living in that space. Then the adjoint $\mathbf{A}^* : \mathbb{M} \mapsto \mathbb{X}$ can be defined as the linear mapping for which the following equality holds [LWU13]

$$\langle \mathbf{Y}, \mathbf{A}\mathbf{x} \rangle_{\mathbb{M}} = \langle \mathbf{A}^*\mathbf{Y}, \mathbf{x} \rangle_{\mathbb{X}}. \quad (5.36)$$

5.5 Algorithm formulation

5.5.1 The composite problem

A wide variety of convex optimization problems that arise in the context of linear inverse problems can be covered by the following generic framework [KP14, Con13, PG18]

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathbb{R}^N} \left\{ f(\mathbf{x}) + g(\mathbf{x}) + h(\mathbf{L}\mathbf{x}) \right\} = \arg \min_{\mathbf{x} \in \mathbb{R}^N} \mathcal{P}(\mathbf{x}). \quad (5.37)$$

The objective function to minimize is given by the sum of three convex functions, which may be differentiable or not and possibly involve an additional linear operator. Each of these functions can either represent a data fidelity term describing the model error or a regularization term imposing certain constraints on the solution or its linear transformation. Hereafter, problem (5.37) is referred to as the *primal problem* in the light of duality theory. To narrow down the choice of possible functions, suppose that $f(\cdot) \in \mathbb{R}^{\mathbb{N}}$ is a smooth, differentiable function with Lipschitz continuous gradient $\nabla f(\cdot)$ for some Lipschitz constant $\beta = [0, \infty[$, such that

$$\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{x}')\|_2 \leq \beta \|\mathbf{x} - \mathbf{x}'\|_2 \quad \forall (\mathbf{x}, \mathbf{x}') \in (\mathbb{R}^{\mathbb{N}})^2, \quad (5.38)$$

which in general holds for least squares type cost functions. Moreover, suppose that $g(\cdot) \in \mathbb{R}^{\mathbb{N}}$ and $h(\cdot) \in \mathbb{R}^{\mathbb{K}}$ are potentially non-smooth, proximable functions that are simple enough to allow the computation of their proximity operators in closed-form or at least with high speed and accuracy. The operator $\mathbf{L} : \mathbb{R}^{\mathbb{N}} \mapsto \mathbb{R}^{\mathbb{K}}$ is bounded and represents a linear mapping to an arbitrary domain. Note that the composite problem offers flexibility by allowing some of these functions to become zero and the linear operator \mathbf{L} to denote the identity operator \mathbf{I} .

Sometimes it seems more convenient to tackle the problem in the dual space. Solving the dual problem can bring useful information about the primal problem as it provides a lower bound on its optimal solution. The corresponding dual formulation of the original problem is given by

$$\hat{\mathbf{u}} = \arg \min_{\mathbf{u} \in \mathbb{R}^{\mathbb{K}}} \left\{ (f + g)^*(-\mathbf{L}^*\mathbf{u}) + h^*(\mathbf{u}) \right\} = \arg \min_{\mathbf{u} \in \mathbb{R}^{\mathbb{K}}} \mathcal{D}(\mathbf{u}), \quad (5.39)$$

where the asterisk as superscript $*$ indicates the Legendre-Fenchel conjugate in the case of functions and the adjoint in the case of linear operators. Furthermore, the expression $(f + g)^*(-\mathbf{L}^*\mathbf{u}) = \min_{\mathbf{x}' \in \mathbb{R}^{\mathbb{N}}} \{ f^*(-\mathbf{L}^*\mathbf{u} - \mathbf{x}') + g^*(\mathbf{x}') \}$ denotes a so-called infimal convolution.

The difference between the optimal solution of the primal and dual problem is referred to as the duality gap. It provides information about the precision with which the primal problem can be solved with implicitly by using the dual one. The duality gap satisfies the following inequality

$$\mathcal{G}(\hat{\mathbf{x}}, \hat{\mathbf{u}}) = \mathcal{P}(\hat{\mathbf{x}}) - \mathcal{D}(\hat{\mathbf{u}}) \geq 0. \quad (5.40)$$

If the values of the primal and dual objectives at the optimum solution coincide, then the duality gap vanishes which means that strong duality holds for the problem. In this case, solving the dual equals solving the primal problem. Whenever there is a remaining duality gap, the problem exhibits weak duality and the dual solution only serves as an approximation of the primal one.

An equivalent formulation of the primal as well as the dual problem can be obtained by using the definition of the Legendre-Fenchel conjugate

$$\begin{aligned} & \min_{\mathbf{x} \in \mathbb{R}^N} \left\{ f(\mathbf{x}) + g(\mathbf{x}) + h(\mathbf{L}\mathbf{x}) \right\} \\ &= \min_{\mathbf{x} \in \mathbb{R}^N} \left\{ f(\mathbf{x}) + g(\mathbf{x}) + \max_{\mathbf{u} \in \mathbb{R}^K} [\langle \mathbf{L}\mathbf{x}, \mathbf{u} \rangle - h^*(\mathbf{u})] \right\} \\ &= \min_{\mathbf{x} \in \mathbb{R}^N} \max_{\mathbf{u} \in \mathbb{R}^K} \left\{ f(\mathbf{x}) + g(\mathbf{x}) + \langle \mathbf{L}\mathbf{x}, \mathbf{u} \rangle - h^*(\mathbf{u}) \right\} = \min_{\mathbf{x} \in \mathbb{R}^N} \max_{\mathbf{u} \in \mathbb{R}^K} \mathcal{L}(\mathbf{x}, \mathbf{u}). \end{aligned} \quad (5.41)$$

Now the equation takes the form of a minimax- or saddle-point search which incorporates both the primal and dual problem, hence it solves both problems simultaneously. Its solution is the pair $(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \mathbb{R}^N \times \mathbb{R}^K$ which represents a so-called Kuhn-Tucker point satisfying the conditions

$$-\mathbf{L}^* \hat{\mathbf{u}} - \nabla f(\hat{\mathbf{x}}) \in \partial g(\hat{\mathbf{x}}), \quad \mathbf{L} \hat{\mathbf{x}} \in \partial h^*(\hat{\mathbf{u}}). \quad (5.42)$$

Furthermore, the following inequality holds for a given saddle-point

$$\min_{\mathbf{x} \in \mathbb{R}^N} \mathcal{L}(\mathbf{x}, \mathbf{u}) = \mathcal{L}(\hat{\mathbf{x}}, \mathbf{u}) \leq \mathcal{L}(\hat{\mathbf{x}}, \hat{\mathbf{u}}) = \max_{\mathbf{u} \in \mathbb{R}^K} \mathcal{L}(\hat{\mathbf{x}}, \mathbf{u}). \quad (5.43)$$

By solving the saddle-point problem for the pair $(\hat{\mathbf{x}}, \hat{\mathbf{u}})$, we jointly obtain the optimal solution $\hat{\mathbf{x}}$ to the primal problem and the optimal solution $\hat{\mathbf{u}}$ to the dual problem. The following section introduces a proximal algorithm that aims at solving this alternative formulation of the original problem by a full splitting approach.

5.5.2 Primal-dual hybrid gradient (PDHG) algorithm

An efficient strategy for solving convex optimization problems of the generalized form (5.37) is to utilize its primal-dual formulation (5.41) to achieve what is known as full splitting. By this approach, the problem is decomposed into an iterative sequence of simpler subproblems, each of which is dedicated to a specific function of the original problem. The smooth functions usually involve the computation of the gradient, whereas non-smooth functions require the evaluation of their respective proximity operators. Thus only first-order information of the functions is exploited, making these splitting schemes suitable for solving large-scale problems with high computational efficiency. In addition, also the linear operator involved in the problem is treated separately which excludes the necessity of its inversion. This is a desirable feature as the computation of the inverse is an expensive task particularly for huge matrices or it may not even exist at all. The algorithm applied in this thesis is a forward-backward splitting based technique referred to as the primal-dual hybrid gradient (PDHG), often also called the Condat-Vu algorithm by the names of its creators [Con13, Vu13], and is described by the following pseudo-code:

Algorithm 1 PDHG - Primal-Dual Hybrid Gradient

Choose the parameters $\tau, \sigma, \rho > 0$ and the initial estimates $\mathbf{x}_0 \in \mathbb{C}^N$, $\mathbf{u}_0 \in \mathbb{C}^K$. Set $\tilde{\mathbf{x}}_0 = \mathbf{x}_0$ and $\tilde{\mathbf{u}}_0 = \mathbf{u}_0$.

```

1: for  $i \geq 0$  do
2:    $\tilde{\mathbf{x}}_{i+1} = \text{prox}_{\tau g}(\mathbf{x}_i - \tau \nabla f(\mathbf{x}_i) - \tau \mathbf{L}^* \mathbf{u}_i)$ 
3:    $\tilde{\mathbf{u}}_{i+1} = \text{prox}_{\sigma h^*}(\mathbf{u}_i + \sigma \mathbf{L}(2\tilde{\mathbf{x}}_{i+1} - \mathbf{x}_i))$ 
4:    $\mathbf{x}_{i+1} = \rho \tilde{\mathbf{x}}_{i+1} + (1 - \rho) \mathbf{x}_i$ 
5:    $\mathbf{u}_{i+1} = \rho \tilde{\mathbf{u}}_{i+1} + (1 - \rho) \mathbf{u}_i$ 
6: end for

```

At every iteration, it only requires the calculation of the gradient $\nabla f(\cdot)$, the proximity operators $\text{prox}_g(\cdot)$ and $\text{prox}_h(\cdot)$, as well as the linear operator L and its adjoint operator L^* . It proceeds without any nested loops, providing a compact formulation and easy implementation. Note that the function $h(\cdot)$ is actually involved through the proximity operator of its Legendre-Fenchel conjugate $\text{prox}_{h^*}(\cdot)$, which can be easily obtained from $\text{prox}_h(\cdot)$ by using Moreau decomposition. The algorithm relies on two proximal parameters $\tau, \sigma > 0$ that act as respective step-sizes for the update of the primal and dual variable. They must be chosen interdependently within suitable bounds in order to ensure the convergence of the algorithm. Convergence is proven under the following condition for the proximal parameters

$$\tau \left(\frac{\beta}{2} + \sigma \|L^* L\| \right) < 1, \quad (5.44)$$

where β denotes the Lipschitz constant of the gradient $\nabla f(\cdot)$ and $\|\cdot\|$ is defined as the operator norm. The relaxation parameter ρ controls the trade-off between how much we trust the new estimates versus how much we rely on the old information. Convergence is guaranteed [Con13, Con87] only if the relaxation parameter lies in the interval $\rho \in]0, 1]$, thus an over-relaxation of this algorithm is not permitted.

The splitting algorithm proceeds in a forward-backward manner [CP09]. At each iteration, the update of the primal and dual variable involves an explicit forward gradient step and an implicit backward gradient step.

$$\begin{aligned}
\tilde{\mathbf{x}}_{i+1} &= \underbrace{\text{prox}_{\tau g}}_{\text{backward (implicit) step}} \left(\underbrace{\mathbf{x}_i - \tau[\nabla f(\mathbf{x}_i) + \underbrace{\mathbf{L}^* \mathbf{u}_i}_{\text{primal-dual coupling}}]}_{\text{forward (explicit) step}} \right) \\
\tilde{\mathbf{u}}_{i+1} &= \underbrace{\text{prox}_{\sigma h^*}}_{\text{backward (implicit) step}} \left(\underbrace{\mathbf{u}_i + \sigma \mathbf{L}(2\tilde{\mathbf{x}}_{i+1} - \mathbf{x}_i)}_{\text{primal-dual coupling}} \right)
\end{aligned} \tag{5.45}$$

The forward step is an actual gradient step, whereas the backward step is achieved by the use of the proximity operator, which is equivalent to a subgradient step for non-smooth functions. The difference between an explicit and implicit step lies in the evaluation point of the (sub)gradient. While carrying out an explicit step involves the computation of the gradient at the initial starting point \mathbf{x}_i , an implicit step depends on the gradient evaluated at the endpoint \mathbf{x}_{i+1} , which corresponds to the proximal mapping of \mathbf{x}_i . When considered separately, the explicit step forms the standard *gradient descent* method

$$\mathbf{x}_{i+1} = \mathbf{x}_i - \tau \nabla f(\mathbf{x}_i), \tag{5.46}$$

while the backward step yields the *proximal point* algorithm

$$\begin{aligned}
\mathbf{x}_{i+1} &= \text{prox}_{\tau f}(\mathbf{x}_i) = \arg \min_{\mathbf{z}} \left\{ \frac{1}{2} \|\mathbf{z} - \mathbf{x}_i\|_2^2 + \tau f(\mathbf{z}) \right\} \\
\rightarrow \nabla \left[\frac{1}{2} \|\mathbf{x}_{i+1} - \mathbf{x}_i\|_2^2 + \tau f(\mathbf{x}_{i+1}) \right] &= (\mathbf{x}_{i+1} - \mathbf{x}_i) + \tau \nabla f(\mathbf{x}_{i+1}) = 0 \\
\rightarrow \mathbf{x}_{i+1} &= \mathbf{x}_i - \tau \nabla f(\mathbf{x}_{i+1}).
\end{aligned} \tag{5.47}$$

Also note that every iteration cycle involves an alternation between minimizing with respect to the primal variable and maximizing with respect to the dual variable, indicated by a gradient descent and ascent step and the proximity operator of the primal and the convex conjugate function respectively. The coupling term incorporates knowledge about the dual variable into the update of the primal variable and vice versa, hence it provides the link between the primal and dual problems. It contains either the linear operator \mathbf{L} , which maps the primal variable to the dual variable space, or its adjoint \mathbf{L}^* , serving as the backward mapping. Note that the variable update is not symmetrical as the dual variable uses the over-relaxed extrapolation $2\tilde{\mathbf{x}}_{i+1} - \mathbf{x}_i$ for the calculation of its new estimate. The algorithm proceeds until some predefined stopping criterion is satisfied, which usually depends on the duality gap or the distance between consecutive iterates.

5.5.3 Practical application of the algorithm

This section regards the practical application of the algorithm to solve the inverse problem in (3.3) regularized by the various derivative-based methods proposed in this work. At first, we formulate the inverse problem as a minimization task of an objective function that exhibits the generic form of (5.37). This requires an appropriate definition of $f(\cdot)$, $g(\cdot)$ and $h(\cdot)$. Then all the necessary ingredients of the algorithm are specified, which involve the gradient ∇ of $f(\cdot)$, the proximity operators $\text{prox}_{\tau g}(\cdot)$ of $g(\cdot)$ and $\text{prox}_{\sigma h^*}(\cdot)$ of $h^*(\cdot)$, the linear operator \mathbf{L} and its adjoint \mathbf{L}^* .

In this thesis, the differentiable term of the objective function always describe the data fidelity between the observed and estimated data, such that $f(\mathbf{q}) = \|\mathbf{G}\mathbf{q} - \mathbf{p}\|_2^2$. This least-squares type function has a well-known complex gradient expression of $\nabla f(\mathbf{q}) = 2\mathbf{G}^H(\mathbf{G}\mathbf{q} - \mathbf{p})$. The Lipschitz constant of the gradient is given by the squared operator norm of the propagation matrix, hence $\beta = 2\|\mathbf{G}\|^2$.

The composite term of the objective function always involves a linear operator \mathbf{L} , which can take the form of the gradient ∇ in case of first-order TV, the vectorized Laplacian Δ in case of second-order TV2, the Hessian \mathcal{H} in case of the Schatten-norm regularizers or the identity operator \mathbf{I} when considering the classical ℓ_1 - or ℓ_2 -norm approaches. The corresponding adjoint operator \mathbf{L}^* must be chosen accordingly. The differential operators and its adjoints are implemented via forward and backward finite differences, respectively. The composite function $\lambda h(\cdot)$ is then set to the regularizer of choice. All derivative-based regularizers exhibit the general form of a mixed p - ℓ_1 -norm, where the p -norm groups the involved derivatives at each evaluation point and the ℓ_1 -norm acts on the resulting map. Without loss of generality, we can choose $\lambda h(\mathbf{L}\mathbf{q}) = \lambda\|\mathbf{L}\mathbf{q}\|_{p,1}$. To compute the associated proximity operator $\text{prox}_{\sigma\lambda h^*}$, we first need to define the Legendre-Fenchel conjugate $h^*(\cdot)$. According to chapter 5.1.1, any norm can be alternatively viewed as the support function of its dual norm ball. By equation (5.10), the Legendre-Fenchel conjugate of a support function of a convex set yields the indicator function of that set. The proximity operator of an indicator function reduces to a simple projection onto the associated set. This leads to the following general expressions for the proximity operator $\text{prox}_{\sigma\lambda h^*}$ for the respective cases of conventional vector- and Hessian Schatten-norm-based regularizers

$$\begin{aligned} \lambda h(\mathbf{L}\mathbf{q}) = \lambda\|\mathbf{L}\mathbf{q}\|_{p,1} &= \sup_{\mathbf{v} \in \mathcal{B}_{\infty,q \leq \lambda}} \langle \mathbf{v}, \mathbf{L}\mathbf{q} \rangle = S_{\mathcal{B}_{\infty,q \leq \lambda}} \quad \rightarrow \quad \lambda h^*(\mathbf{L}\mathbf{q}) = i_{\mathcal{B}_{\infty,q \leq \lambda}} \\ &\rightarrow \text{prox}_{\sigma\lambda h^*} = \text{prox}_{\sigma i_{\mathcal{B}_{\infty,q \leq \lambda}}} = P_{\mathcal{B}_{\infty,q \leq \lambda}}, \end{aligned} \quad (5.48)$$

$$\begin{aligned} \lambda h(\mathcal{H}\mathbf{q}) = \lambda\|\mathcal{H}\mathbf{q}\|_{\mathcal{S}_p,1} &= \sup_{\mathbf{V} \in \mathcal{B}_{\infty,\mathcal{S}_q \leq \lambda}} \langle \mathbf{V}, \mathcal{H}\mathbf{q} \rangle = S_{\mathcal{B}_{\infty,\mathcal{S}_q \leq \lambda}} \quad \rightarrow \quad \lambda h^*(\mathcal{H}\mathbf{q}) = i_{\mathcal{B}_{\infty,\mathcal{S}_q \leq \lambda}} \\ &\rightarrow \text{prox}_{\sigma\lambda h^*} = \text{prox}_{\sigma i_{\mathcal{B}_{\infty,\mathcal{S}_q \leq \lambda}}} = P_{\mathcal{B}_{\infty,\mathcal{S}_q \leq \lambda}}, \end{aligned} \quad (5.49)$$

where q stands to p as stated in chapter 5.1.1, which also holds for the considered Schatten-norms. The projection onto the vector-norm ball $\mathcal{B}_{\infty,q \leq \lambda} = \{\mathbf{v} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N] : \|\mathbf{v}_n\|_q \leq \lambda, \forall n = 1, \dots, N\}$ can be achieved by the elementwise projection of the vector components onto the \mathcal{B}_q norm ball. The projection onto the matrix-norm ball $\mathcal{B}_{\infty,\mathcal{S}_q \leq \lambda} = \{\mathbf{V} = [\mathbf{V}_1, \mathbf{V}_2, \dots, \mathbf{V}_N] : \|\mathbf{V}_n\|_{\mathcal{S}_q} \leq \lambda, \forall n = 1, \dots, N\}$ can be achieved by separately projecting each submatrix onto the $\mathcal{B}_{\mathcal{S}_q}$ norm ball [LWU13]. Moreover, the orthogonal projection is independent of the proximal parameter σ . Since the indicator function of a convex set always forces a hard constraint on the solution, it is either zero or ∞ regardless of the weighting. Closed-form expressions for the projections onto the considered norm balls are readily available in chapter 5.2.1.

Recall that for the Hessian Schatten-norm regularizers, the projection of a Hessian onto the \mathcal{S}_q -norm ball is realized by the projection of its singular value vector on the associated ℓ_q -norm ball. Since the singular values of a matrix are always non-negative, real numbers, there is no need for an extension of the proximity operator to complex values.

Regarding the first- and second-order Total Variation, as well as the alternative formulation of the \mathcal{HS}_2 regularizer, a projection onto the \mathcal{B}_2 norm ball is required. The involved complex-valued vectors in \mathbb{C}^M can be treated as real-valued vectors in \mathbb{R}^{2M} of twice the size, as there exist an isomorphic relation between the vector spaces [FR06]. This way, the standard proximity operator for real-valued functions can still be evaluated for the actual complex-valued data. A backmapping $\mathbb{R}^{2M} \mapsto \mathbb{C}^M$ is performed afterwards.

When only the derivative-based regularizer shall be applied without any secondary regularization, we need to set the remaining term in the objective function to zero, i.e. $g(\mathbf{q})=0$. The resulting proximity operator then simply returns its argument as $\text{prox}_{\tau g}(\mathbf{q}) = \mathbf{q}$. An additional sparsity constraint on the solution can be imposed by setting $\lambda g(\mathbf{q}) = \lambda \|\mathbf{q}\|_1$. In this case, no remapping of the complex-valued vector is necessary as the definition of the involved proximity operator provided by equation (5.28) also works for complex-valued data.

The implementation of the classical ℓ_1 - and ℓ_2 -norm approaches is best carried out by using $\lambda g(\cdot)$ and discarding the composite term $(h \circ L) = 0$ in the objective function. The algorithm then neglects the information of the dual variable and reduces to a simple forward-backward splitting, given by

Algorithm 2 Forward-Backward Splitting Scheme

Choose the parameters $\tau, \rho > 0$ and the initial estimates $\mathbf{x}_0 \in \mathbb{C}^N$. Set $\tilde{\mathbf{x}}_0 = \mathbf{x}_0$.

- 1: **for** $i \geq 0$ **do**
 - 2: $\tilde{\mathbf{x}}_{i+1} = \text{prox}_{\tau g}(\mathbf{x}_i - \tau \nabla f(\mathbf{x}_i))$
 - 3: $\mathbf{x}_{i+1} = \rho \tilde{\mathbf{x}}_{i+1} + (1 - \rho) \mathbf{x}_i$
 - 4: **end for**
-

To ensure convergence of the PDHG algorithm, the proximal parameters τ and σ must be chosen according to (5.44). For convenience, the parameters are chosen according to their respective maximum value, such that $\tau = \tau_{max} = \frac{2}{\beta}$ and $\sigma = \sigma_{max} = \frac{1}{\|\mathbf{L}^* \mathbf{L}\|} (\frac{1}{\tau} - \frac{\beta}{2})$. Convergence of the forward-backward splitting scheme is guaranteed for $0 < \tau < \frac{\beta}{2}$. The relaxation parameter is chosen in the interval $\rho \in]0, 1]$. Now we are equipped with all the necessary ingredients to use the algorithm for the study of the proposed regularizers.

A concluding remark may be offered concerning the extension of the proximity operator to the complex-valued case. Although the workarounds stated above can be successfully applied for the purpose of this study, a more general extension of the proximity operator to a class of convex functions with complex-valued variables is provided. This shall highlight that their use is not limited to the present cases. For a function $\Phi : \mathbb{C}^M \rightarrow]-\infty, \infty]$, $\mathbf{x} \mapsto \phi^{Re}(\text{Re}\{\mathbf{x}\}) + \phi^{Im}(\text{Im}\{\mathbf{x}\})$, where ϕ^{Re} and ϕ^{Im} are proper, lower semicontinuous and convex functions in \mathbb{R}^M , and $\text{Re}\{x\}$ and $\text{Im}\{x\}$ denote vectors composed of the real and imaginary parts of the elements of $\mathbf{x} \in \mathbb{C}^M$, the proximity operator is given by [CCMP14, Lar16, LCT⁺16]

$$\text{prox}_{\Phi} : \mathbb{C}^M \mapsto \mathbb{C}^M, x \mapsto \text{prox}_{\phi^{Re}}(\text{Re}\{\mathbf{x}\}) + j \text{prox}_{\phi^{Im}}(\text{Im}\{\mathbf{x}\}). \quad (5.50)$$

For more information regarding extensions of the proximity operator to more specific complex-valued cases, the reader is referred to [STS⁺16, GGC16, MAYGB11, Lar16].

Regularizer Characteristics - Simulation Study

In this chapter, a numerical simulation study is conducted in order to assess the characteristics and abilities of the proposed derivative-based regularization methods in realistic scenarios. The main objective is the investigation of the quality of the estimated solution in terms of reconstruction error, as well as their graphical inspection and interpretation. To demonstrate the extent of applicability of the regularizers, four examples with different spatial source distributions are considered, covering the most general cases of acoustic sources. They include the case of a vibrating plate, a baffled circular piston, a single isolated elementary source and a cluster of sparsely distributed monopoles, each representing either a spatially extended or concentrated source distribution. Not only are we keen on characterizing these sources by making use of different regularization terms, but we are also concerned with the retrieval of certain acoustic quantities from their radiated sound fields. Moreover, the effect of measurement noise on the accuracy of the various solutions shall be investigated. The comparative study encompasses all the proposed derivative-based methodologies, including first- and second-order Total Variation together with the Hessian Schatten-Norm regularizers, as well as the state-of-the-art techniques of least-squares estimation and compressive sensing. The research in [FGD18, FD16] served as a useful template on how to properly conduct the study of the various derivative-based regularizers to highlight their special features.

6.1 Simulation setup

The acquisition of the pressure field is simulated by a uniform rectangular array, consisting of $8 \times 8 = 64$ microphones that are distributed equidistantly over an area of size 50×50 cm². The size of the source surface is chosen accordingly, such that the array covers its whole extent. The inter-element spacing of 6.25 cm results in theoretical frequency limit of 2744 Hz, providing the limit above which aliasing effects are introduced. The measurements are taken at a distance of 6 cm from the source surface. The equivalent sources used to model the sound field are disposed on a grid of dimension 41×41 , covering the same area as the source and hence enclosing the radiator. When investigating the ability to retrieve the source strength distribution along the source plane, the equivalent sources are placed immediately on the actual source surface at $z_0 = 0$ m, revealing a retreat distance of zero. In case of reconstructing certain quantities of the sound field above the plane from the estimated source strengths, the equivalent sources are retracted 1 cm in the interior volume of the source, regularly distributed on the xy-plane. The reconstruction then takes place at a distance of 3 cm above the actual source surface. In order to simulate real acquisition conditions, the measurement data is degraded by complex additive white Gaussian noise (AWGN) with a signal-to-noise ratio (SNR) of 30 dB. Based on the resulting dynamic range, an ad-hoc choice of the regularization parameter is made, which then is used for all regularizers with no further intention of individual tuning.

6.2 Figures of merit

For better comparison of the regularization strategies, we need to specify some suitable figures of merit for their accuracy and performance. The quality of a regularizer is determined by the amount of deviation of its estimated solution from the ground truth. This deviation may either be assessed locally for each individual equivalent source point or in form of a spatially-averaged single value error. Which reconstruction error is the more appropriate choice depends on the information one is trying to retrieve. For the purpose of an initial assessment of the overall accuracy of the methods over a broad range of frequencies or different experimental setups, the spatially-averaged reconstruction error ϵ can be employed as a meaningful metric. The relative reconstruction error of each individual distribution point is defined as the normalized difference between the true and reconstructed quantity of interest (source strength, pressure, velocity, intensity). For source reconstruction, the element error is thus given by

$$\epsilon_i = \frac{|q_{i,\text{true}} - q_{i,\text{rec}}|}{|q_{i,\text{true}}|}. \quad (6.1)$$

The spatially-averaged or normalized root-mean-square (RMS) error is then obtained by

$$\epsilon = \sqrt{\frac{\sum_{i=1}^N (q_{i,\text{true}} - q_{i,\text{rec}})^2}{\sum_{i=1}^N (q_{i,\text{true}})^2}} = \frac{\|\mathbf{q}_{\text{true}} - \mathbf{q}_{\text{rec}}\|_2}{\|\mathbf{q}_{\text{true}}\|_2}. \quad (6.2)$$

To further classify the mismatch in terms of its spatial distribution, an error location vector \mathbf{e} may be determined that provides additional information about local peculiarities of the error pattern

$$\mathbf{e} = \sqrt{\frac{(|\mathbf{q}_{\text{true}}| - |\mathbf{q}_{\text{rec}}|)^2}{(\max_i |q_{i,\text{true}}|)^2}}. \quad (6.3)$$

The percentage value of the reconstruction error can simply be obtained by multiplying the associated error equation with the marginal total of 100 %.

6.3 Vibrating plate

We first consider the reconstruction of the source distribution of a vibrating rectangular plate based on measurements of the emitted sound field. A baffled, simply supported steel plate of thickness 1 mm is chosen, which vibrates at a resonance frequency of about 130 Hz without explicit excitation of an external force. The necessary material-specific constants include Young's modulus $E = 210 \frac{\text{GN}}{\text{m}^2}$, Poisson's ratio $\nu = 0.3$ and the density $\rho = 7850 \frac{\text{kg}}{\text{m}^3}$. The plate is resonating in a 3-2 mode shape at a normal vibrational velocity according to a selected maximum equivalent source strength of $Q = 10^{-3} \frac{\text{m}^3}{\text{s}}$. The plate is of dimension $50 \times 50 \text{ cm}^2$ and it is discretized on a uniform grid with 41×41 points to establish the ground truth. The resulting normal velocity distribution is displayed in Figure 6.1. The generated acoustic field is then measured and the sound waves are traced back to the source surface, based on the use of the proposed regularization strategies. For the derivative-based regularizers, periodic boundary conditions were chosen, as it seems most appropriate for the modal structure. The parameter D for the selection of the regularization parameter is empirically chosen as $D=45 \text{ dB}$. At first, a visual inspection of the source strengths estimated by the respective methods is conducted using Figure 6.1. The results are compared to the ground truth and among each other, in order to highlight their characteristics and abilities, as well as the nature of the introduced artifacts.




Figure 6.1: Estimated equivalent sources for the plate resonating in a 3-2 mode shape.

regularizer	ℓ_1 -norm	ℓ_2 -norm	∇TV	ΔTV2	\mathcal{HS}_1	\mathcal{HS}_2	\mathcal{HS}_∞
avg. error [%]	429.55	20.31	28.29	14.44	15.75	13.77	15.33

Table 6.1: Spatially-averaged reconstruction error for the plate resonating in a 3-2 mode shape.

At first sight, the spatially-averaged reconstruction error reveals the superiority of the second-order methods over the classical least-squares solution. In particular, the \mathcal{HS}_2 and ΔTV2 regularizers yield the most accurate reconstruction, very closely followed by the \mathcal{HS}_1 and \mathcal{HS}_∞ approaches in the average error shown in Table 6.1.




Figure 6.2: Spatial distribution of the reconstruction error for the plate resonating in a 3-2 mode shape.

The sparse solution obtained by the ℓ_1 -norm consists of a combination of spatially focal source points that are situated around the centre of each modal sub-pattern, see also Figure 6.2. The active clusters of equivalent sources greatly exceed the maximum magnitude of the ground truth, causing the large contiguous regions to be free of any source activity. Although the acquired sound field can be approximated by the one radiated from this estimate, the sparse representation does not conform with the physical structure of the source. It is noteworthy that structural mode

shapes exhibit a sparse spectrum in the wavenumber domain, hence the application of the ℓ_1 -norm on the k-space would provide an accurate reconstruction of the source distribution in space.

The least-squares solution, associated with the ℓ_2 -norm, is capable of capturing the characteristic shape of the mode, although it induces additional spatial variations in the distribution. This is the visual expression of the occurring energy spread of the solution vector among its components. The deviation of the ground truth becomes more apparent towards the boundaries and in the intermediate spaces of the pattern, see Figure 6.2. Since it provides the solution with minimum energy, it slightly underestimates the magnitude of the source.

The gradient-based Total Variation ∇ TV provides an estimate displaying regions of constant magnitude that only recovers the coarse structure of the mode pattern in Figure 6.1. The continuous transitions between the extrema are lost during the minimization process. Instead, each section of the mode is represented by a distinctly separated block of constant value that can be considered individually as flat pistons. The appearance of the staircasing effect can be observed in the immediate neighborhood of the blocks, representing the former continuous edge as a collection of constant segments that erratically decrease in value. It becomes more pronounced towards the horizontal boundaries, leading to a spatial extension of the affected pistons. This is caused by the prescribed boundary conditions chosen for the forward and backward finite difference implementation of the gradient and its adjoint, respectively. The reason for this not to occur at the vertical boundaries is the less steep slope of the vertical profile towards the border, leading to a different evolution of the gradient field. However, when exchanging the order of the mode shape for both directions, the larger staircase artifacts occur at the vertical borders. This method yields a strong overall underestimation of the magnitude, where the block sources on both sides exhibit a different plateau than the pair in the middle.

As expected, the Laplacian-based Total Variation Δ TV2 successfully recovers the continuous variations on the plate surface. Unfortunately, it does not accurately reflect the shape of the individual modes. They appear as diamond-shaped areas or rhombs with sharp edges, tapering in magnitude from the sides towards a gently curved tip at the centre. The solution does not suffer from the staircasing effect as in the case of ∇ TV. Although very subtle, the demanded smoothness of the transitions causes an increased coefficient activity in the intermediate spaces when compared to the ground truth. Moreover, a slight overestimation in magnitude occurs.

The results obtained by the Hessian Schatten-norm-based penalties are similar to the Laplacian-based solution, as these regularizers also act on the second-order differential domain. The primary spatial features of the source are correctly revealed. The estimate exhibits a smooth surface morphology without introducing any severe deformations or other artifacts. There are some clear distinctions to the modal shapes of the Laplacian-based solution. The \mathcal{HS}_2 regularizer can be considered the closest fit to the Laplacian penalty and it differs only by incorporating the cross-derivatives into its definition. This results in additional smoothing along the diagonal directions of the modes, which blurs the sharp ridges in the Laplacian solution. Minimum curvature is achieved along any discrete direction, leading to a slight elevation of the surface intensity at the crossing points between adjacent modes. In addition, the amount of overestimation of the true magnitude decreases. Altogether, the \mathcal{HS}_2 regularizer provides a fairly close approximation of the ground truth.

The \mathcal{HS}_1 regularizer yields rectangular-shaped modes with well defined edges and corners. This is due to the anisotropic smoothing properties of the functional. It exhibits a preference for distinct edges or ridges as opposed to isotropic smoothing or distinct conical dips or peaks. It sharpens edge profiles while smoothing along lines of the edge. As a result, we observe magnitude correlations along the lines forming the rectangles. They seem notably sharper and show decreased blur at the corners. When considering the magnitude map as a 3D surface, the individual modes appear roof- or pyramid-shaped to provide the most pronounced tapering of the tip. Hence of the Hessian-Schatten regularizers it exhibits the highest overestimation in

magnitude, together with a little increase of source activity along the nodal lines of the mode shape.

While the solution associated with the \mathcal{HS}_∞ regularizer is not much different from the \mathcal{HS}_2 solution, a closer look reveals a distinct difference in the way the curvature of the surface is shaped. \mathcal{HS}_∞ makes the oval shape of the modes more rounded. When looking at the elevation in profile, we observe a broader, rounded off peak, giving the least overestimation of the true magnitude in Figure 6.3, which gives a nice overview of some of the shaping effects.

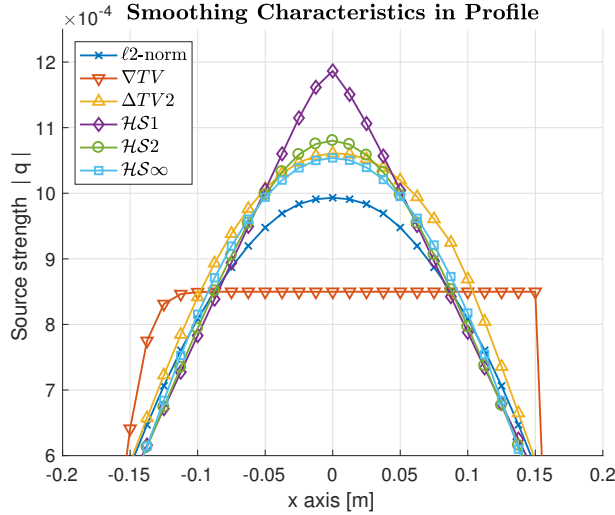


Figure 6.3: Smoothing characteristics of the methods.

Since all regularizers are norm functionals with scalar output, a cross comparison of their output value may lend itself to reveal potential inter-relations. More specifically, the active minimization of a certain norm may yield to an implicit, passive minimization of another norm. To further explore this behavior, all considered norm values of the respective solution estimates are listed in Table 6.2.

		passive						
		ℓ_1 -norm	ℓ_2 -norm	∇ TV	Δ TV2	\mathcal{HS}_1	\mathcal{HS}_2	\mathcal{HS}_∞
active	Ground Truth	$64.68 \cdot 10^{-2}$	$20.00 \cdot 10^{-3}$	$21.52 \cdot 10^{-2}$	$67.33 \cdot 10^{-3}$	$96.44 \cdot 10^{-3}$	$79.15 \cdot 10^{-3}$	$22.98 \cdot 10^{-3}$
	ℓ_1 -norm	$47.18 \cdot 10^{-2}$	$87.93 \cdot 10^{-3}$	1.27	2.50	4.65	4.09	$73.87 \cdot 10^{-2}$
	ℓ_2 -norm	$64.87 \cdot 10^{-2}$	$19.05 \cdot 10^{-3}$	$20.81 \cdot 10^{-2}$	$69.85 \cdot 10^{-3}$	$11.38 \cdot 10^{-2}$	$92.12 \cdot 10^{-3}$	$28.19 \cdot 10^{-3}$
	∇ TV	$66.10 \cdot 10^{-2}$	$20.18 \cdot 10^{-3}$	$17.56 \cdot 10^{-2}$	$14.99 \cdot 10^{-2}$	$20.81 \cdot 10^{-2}$	$18.30 \cdot 10^{-2}$	$29.76 \cdot 10^{-3}$
	Δ TV2	$64.03 \cdot 10^{-2}$	$19.60 \cdot 10^{-3}$	$21.37 \cdot 10^{-2}$	$48.78 \cdot 10^{-3}$	$84.91 \cdot 10^{-3}$	$66.96 \cdot 10^{-3}$	$25.82 \cdot 10^{-3}$
	\mathcal{HS}_1	$66.66 \cdot 10^{-2}$	$19.97 \cdot 10^{-3}$	$21.47 \cdot 10^{-2}$	$60.65 \cdot 10^{-3}$	$73.42 \cdot 10^{-3}$	$66.58 \cdot 10^{-3}$	$8.47 \cdot 10^{-3}$
	\mathcal{HS}_2	$65.76 \cdot 10^{-2}$	$19.79 \cdot 10^{-3}$	$20.63 \cdot 10^{-2}$	$54.99 \cdot 10^{-3}$	$75.44 \cdot 10^{-3}$	$61.72 \cdot 10^{-3}$	$17.88 \cdot 10^{-3}$
	\mathcal{HS}_∞	$65.80 \cdot 10^{-2}$	$20.05 \cdot 10^{-3}$	$21.59 \cdot 10^{-2}$	$59.51 \cdot 10^{-3}$	$89.55 \cdot 10^{-3}$	$69.35 \cdot 10^{-3}$	$29.61 \cdot 10^{-3}$

Table 6.2: Cross comparison of norm values for the plate resonating in a 3-2 mode shape.

The norm values of the ground truth serve as the guiding reference that reveal certain properties of its spatial structure. In order to achieve a reasonable interpretation of the norm values, a quick reminder of their meaning is given. The smallest ℓ_1 -norm is obviously provided by the ℓ_1 -norm solution, indicating a sparse distribution with the least number of active coefficients. Since the other regularizers do not directly promote spatial sparsity, they reveal a much higher ℓ_1 -norm that better fit the reference value of the ground truth, that is far away from sparse. The smallest ℓ_2 -norm indicates the solution with minimum energy, naturally obtained by the ℓ_2 -norm minimization. The sparse solution shows the largest ℓ_2 -norm, as very high energy is concentrated

in the few active source components. The energy level of the derivative-based solutions is quite similar to the ℓ_2 -norm estimate, but with a tendency towards higher values. The smallest ∇TV -norm indicates the least amount of active gradients and consequently a low spatial variation with a majority of constant regions. The ground truth, as well as the second-order estimates, exhibit a modal pattern that varies continuously in magnitude in every direction, preventing the gradient from vanishing and thus yielding a higher value. The ℓ_1 -norm solution entails huge gradients due to isolated coefficient activities with very high energy, leading to an enormous ∇TV -norm. The minimum ΔTV2 -norm is obtained by the solution with the smallest second-order derivatives along the cardinal directions of the coordinate system, whereas the minimum \mathcal{HS}_2 -norm also takes the cross-derivatives into consideration. The minimum \mathcal{HS}_1 -norm points out the solution with the lowest number of active eigenvalues, revealing a preference of smoothing along a certain principal direction, usually along an edge. The minimum \mathcal{HS}_∞ -norm suggest a higher level of smoothing across edges as opposed to smoothing along them, as it gives the solution with the least sum of its larger eigenvalues. In the latter four cases, the second-order methods distinguish themselves clearly from the other regularizers by generally showing quite similar and much smaller values. It is remarkable that, against expectation, the \mathcal{HS}_∞ -norm is not minimized the most by the associated regularizer itself, but rather by every other second-order method. The \mathcal{HS}_∞ regularizer only penalizes the larger eigenvalues while ignoring the smaller ones completely. They may keep larger values in order to satisfy the data fidelity constraint, which prevents a sufficient minimization and results in a solution that may not exhibit the smallest \mathcal{HS}_∞ -norm. By contrast, the \mathcal{HS}_2 and ΔTV2 regularizer directly and indirectly call for eigenvalue pairs with minimum energy as both require a solution with minimum curvature. The implicit minimization of the smaller eigenvalues may permit a higher number of low level eigenvalues, yielding a smaller \mathcal{HS}_∞ -norm. Since the \mathcal{HS}_1 regularizer attempts to directly force a multitude of the larger eigenvalues to be zero, it is only natural that it implicitly minimizes the \mathcal{HS}_∞ -norm to the greatest extent. This also applies for the ∇TV regularizer as it yields a solution with well defined, sharp edges and large regions of constant magnitude, which implies a high number of eigenvalues that tend to zero.

To further emphasize the superior performance of derivative-based optimization, the same study is repeated for another modal pattern. This time, a central driven aluminum plate of thickness 5 mm is considered as radiator, while maintaining the same simulation setup. The necessary material-specific constants include Young's modulus $E = 69 \frac{\text{GN}}{\text{m}^2}$, Poisson's ratio $\nu = 0.334$ and the density $\rho = 2712 \frac{\text{kg}}{\text{m}^3}$. The applied point force of $F = 10$ N oscillates at 800 Hz. The regularization parameter is chosen with $D=40$ dB. Figure 6.4 depicts the estimated equivalent source distributions and Figure 6.5 provides the corresponding spatial reconstruction error. As indicated by Table 6.3, the second-order-differential regularizers outperform the other ones and provide comparably small overall-errors.




Figure 6.4: Estimated equivalent sources for the central driven plate.

regularizer	ℓ_1 -norm	ℓ_2 -norm	∇TV	$\Delta\text{TV}2$	\mathcal{HS}_1	\mathcal{HS}_2	\mathcal{HS}_∞
avg. error [%]	439.37	30.72	36.56	21.38	22.81	20.48	24.03

Table 6.3: Spatially-averaged reconstruction error for the central driven plate.

In this example, the difference in reconstruction error between the classical least-squares estimate and the second-order differential solutions becomes even more significant, with the \mathcal{HS}_2 again providing the most accurate result. As before, ℓ_1 -norm is unpractical for the reconstruction of an extended source as it produces a sparse solution. Although the ∇TV estimate can emulate the general structure of the pattern, it fails at providing a sufficient resolution to capture the details of the source. This time, the minimum energy demand of the ℓ_2 -norm also seems inappropriate for the recovery. Although the reconstructed distribution reasonably reflects the shape of the modal pattern, it exhibits strong warping artifacts and prevailing side lobes, which become most apparent at the boundaries. By contrast, the second-order methods do all restore a well resolved image of the source that closely approximates the ground truth. Besides a slight blur and some warping artifacts towards the borders, supposedly due to the prescribed boundary conditions

along with the corruptive noise, the structural characteristics of the source are nearly perfectly retrieved. The request of a solution with minimum curvature seems to be more appropriate for the reconstruction of mode shapes, representing the general case of a spatially extended source.




Figure 6.5: Spatial distribution of the reconstruction error for the central driven plate.

For the sake of completeness, the different norms of the respective solution estimates are presented in Table 6.4, from which similar conclusions can be drawn.

		passive						
		ℓ_1 -norm	ℓ_2 -norm	∇TV	$\Delta TV2$	\mathcal{HS}_1	\mathcal{HS}_2	\mathcal{HS}_∞
active	Ground Truth	$1.79 \cdot 10^{-3}$	$54.64 \cdot 10^{-6}$	$57.21 \cdot 10^{-5}$	$18.10 \cdot 10^{-5}$	$34.30 \cdot 10^{-5}$	$27.72 \cdot 10^{-5}$	$90.09 \cdot 10^{-6}$
	ℓ_1 -norm	$1.26 \cdot 10^{-3}$	$24.57 \cdot 10^{-5}$	$3.30 \cdot 10^{-3}$	$6.37 \cdot 10^{-3}$	$12.43 \cdot 10^{-3}$	$10.87 \cdot 10^{-3}$	$2.05 \cdot 10^{-3}$
	ℓ_2 -norm	$1.89 \cdot 10^{-3}$	$53.85 \cdot 10^{-6}$	$65.01 \cdot 10^{-5}$	$26.80 \cdot 10^{-5}$	$47.14 \cdot 10^{-5}$	$38.68 \cdot 10^{-5}$	$10.84 \cdot 10^{-5}$
	∇TV	$1.94 \cdot 10^{-3}$	$54.39 \cdot 10^{-6}$	$44.49 \cdot 10^{-5}$	$42.70 \cdot 10^{-5}$	$69.48 \cdot 10^{-5}$	$61.19 \cdot 10^{-5}$	$99.00 \cdot 10^{-5}$
	$\Delta TV2$	$1.82 \cdot 10^{-3}$	$54.31 \cdot 10^{-6}$	$55.98 \cdot 10^{-5}$	$14.03 \cdot 10^{-5}$	$32.18 \cdot 10^{-5}$	$24.98 \cdot 10^{-5}$	$10.27 \cdot 10^{-5}$
	\mathcal{HS}_1	$1.83 \cdot 10^{-3}$	$53.93 \cdot 10^{-6}$	$56.06 \cdot 10^{-5}$	$16.54 \cdot 10^{-5}$	$27.19 \cdot 10^{-5}$	$24.29 \cdot 10^{-5}$	$24.29 \cdot 10^{-5}$
	\mathcal{HS}_2	$1.81 \cdot 10^{-3}$	$53.26 \cdot 10^{-6}$	$53.60 \cdot 10^{-5}$	$15.27 \cdot 10^{-5}$	$27.79 \cdot 10^{-5}$	$22.19 \cdot 10^{-5}$	$73.96 \cdot 10^{-5}$
	\mathcal{HS}_∞	$1.81 \cdot 10^{-3}$	$53.41 \cdot 10^{-6}$	$54.95 \cdot 10^{-5}$	$18.29 \cdot 10^{-5}$	$33.04 \cdot 10^{-5}$	$24.45 \cdot 10^{-5}$	$13.06 \cdot 10^{-5}$

Table 6.4: Cross comparison of norm values for the central driven plate.

6.4 Spatial response to a monopole

This section is dedicated to the quantification of the regularizers ability to recover a single isolated monopole source. In order to provide a good approximation, they are required to produce a sufficiently sparse solution vector. To address this issue, we first explore the spatial response to a point source by means of a point spread function (PSF), which can be considered the impulse response of the environment. For this purpose, a monopole source is centered at $(x_0, y_0, z_0) = (0, 0, 0)$ on the source surface, exciting the surrounding medium with a volume velocity of $Q = 10^{-2} \frac{m^3}{s}$. The spatial response of the different methods is determined along a vertical cut through the source position at height of the source, i.e. for a fixed $x = 0$ along the y -axis interval $y = [-0.25m, 0.25m]$ on the source plane at $z = 0$. The resulting PSF shows the magnitude profile of the solution along the evaluation line, which provides an illustration of the spatial extension of the different estimates. The simulations are conducted for a low (500 Hz), mid (1500 Hz) and high frequency (3500 Hz) excitation.




Figure 6.6: Point spread function (PSF) - Spatial response to a monopole source radiating at 500 Hz (left), 1500 Hz (center) and 3500 Hz (right).

Figure 6.6 depicts the various solutions for the different excitation frequencies, with the most narrow peak indicating magnitude and location of the true source. The ℓ_1 -norm solution provides the most accurate estimation, as it promotes direct sparsity in the spatial domain. It resembles an almost perfect Dirac delta distribution. This behavior remains constant for the selected frequencies. The least-squares solution shows a sinc-like spatial response which is decreasing towards the borders. It exhibits a broad main lobe around the source location where the main energy is concentrated, accompanied by the appearance of side lobes that result from the energy smearing induced by the ℓ_2 -norm. When looking at the reconstructed spatial distribution as a whole, these side lobes manifest as concentric circles surrounding the monopole and expanding its spatial confinement. An increase in frequency results in a constriction of the main lobe, which improves the spatial resolution and localization. On the other hand, the distribution shows an increasing level of side lobes that lead to spurious artifacts, appearing as lateral blurring of the actual point source. This makes the least-squares approach inadequate for the reconstruction of sources vibrating at frequencies above the sampling limit. The gradient-based ∇TV solution shows a wide rectangular window that encloses the true position of the source. The monopole is imperfectly replicated in form of a rigid piston. Due to the inherent introduction of staircasing artifacts, the edges appear slightly bevelled in the interpolated profile. At higher frequencies, an increasing number of distinct magnitude plateaus emerge. The apparent narrowing of the main window, together with the relatively low magnitude level at the sides, lead to an improved accuracy of identification towards high frequencies. The solutions obtained by penalizing second-order derivatives do all show a very similar spatial structure. Just like the least-squares solution, they all exhibit a widened profile due to the occurrence of side lobes. In direct comparison, the side lobes are broader and thus there are fewer less pronounced ones, leading to a smoother

transition. Moreover, with the higher magnitude of the main lobe, a slight improvement in localization is achieved. At higher frequencies, we observe a narrowing of the main lobe which increasingly exceeds the least-squares estimate in magnitude around the source point. At the same time, the level of elevation of the side lobes consistently remains smaller in comparison, leading to a less blurred solution. The differences among the second-order methods are very subtle and mostly concerned with the extent of the gaps between the side lobes.

Fused regularizers

Although the derivative-based regularizers outperform the classical least-squares approach, especially at high frequencies where aliasing errors occur, they do not produce a sparse representation of the source. Hence they fail to provide a meaningful solution for the reconstruction of spatially compact sources. In order to keep their many advantages while improving sparsity, we consider the use of a fused method that combines the promotion of sparse derivatives with additional sparsity in the spatial domain. This can either be achieved by introducing an additional ℓ_1 -norm penalty to the objective function that acts on the solution, or by directly incorporating the spatial information into the derivative-based regularizer itself. For a more detailed explanation of these two variants, see section 4.2.5. To illustrate the effects of the additional spatial sparsity on the reconstruction, the same study is repeated for the fused approaches. Let's first consider the case of the externally applied regularization term. Figure 6.7 illustrates the PSF for the different solutions with the unaltered ℓ_2 -norm solution that is kept as a visual reference.




Figure 6.7: Point spread function (PSF) - Spatial response to a monopole source radiating at 500 Hz (left), 1500 Hz (center) and 3500 Hz (right) for the fused regularizers with external ℓ_1 -norm.

The region of source activity in the reconstructed profile is significantly reduced for all methods, leading to an improved approximation of the point source. The suppression of side lobes and the reduction in the width of the main lobe both yield a considerable enhancement of the spatial resolution. At medium to high frequencies, the impact of the first order side lobes start to increase, leading to a widening of the source profile up to a certain amount. However, they still remain decisively below the level of the main lobe. The gradient-based approach still exhibits its characteristic piecewise constant profile, while the second-order methods do all show smooth variations with nearly perfect lateral symmetry, with only very nuanced differences between them.

Another way to impose an additional sparsity constraint on the solution is to directly embed it into the differential operator of choice. This approach can be very appealing, because the majority of the common algorithms don't permit the use of multiple regularization terms, especially when they are composed of linear operators. Furthermore, tuning of the regularization parameters to find the right trade-off may become more involved with multiple penalties. The different methods based on this approach are again expressed in terms of their associated PSFs in Figure 6.8, to

investigate the effects of the direct implementation on the solution.




Figure 6.8: Point spread function (PSF) - Spatial response to a monopole source radiating at 500 Hz (left), 1500 Hz (center) and 3500 Hz (right) for the fused regularizers (Sparse Variation) with internal ℓ_1 -norm.

We now observe a quite different behavior of certain methods, concerning the spatial variations of the profile. The solution to the first-order Total Variation, which is usually characterized by its piecewise-constant trajectory, now exhibits a piecewise-linear evolution very similar to the one obtained by the second-order approach. The reason for this lies in the way the spatial weights are incorporated in the regularizer. The squared weights are nested within the inner ℓ_2 -norm of equation (4.13), providing a grouping of spatial and differential activity. By this approach, the coefficients and its derivatives are always jointly forced to be zero. Thus any source activity is always linked to a non-zero spatial derivative, which indicates a change in magnitude of the adjacent coefficient. As a consequence, the active regions of constant magnitude are replaced by smooth transitions. This makes the gradient-based approach much more suitable for the reconstruction of acoustic sources, as the smooth variations allow to capture their physical nature more appropriately.

Another noticeable feature is the sudden asymmetry of the magnitude profile. There are two factors that contribute to this effect. First, recall that all spatial derivatives are approximated using a forward finite difference scheme. Hence, the derivative evaluated at a given point is always based on the magnitude of the next coefficient in row. Second, the promotion of joint sparsity of spatial weights and derivatives leads to an asymmetric penalization of transition points between contiguous active and inactive regions when using forward finite differences. A coefficient previous to a transition to an active region cannot be forced to zero since the derivative at this point is most certainly unequal to zero. The relaxed demand of a small spatial weight, together with a minimum slope or curvature may be contradictory and a compromise must be found. This leads to the progressive increase of the profile over several coefficients at the left hand side of the source. On the other hand, a coefficient that lies right next to a transition from an active to an inactive zone has a vanishing derivative because of the successive inactivity of the following coefficients. Therefore, it can be successfully forced to zero which leads to an immediate descent of the magnitude profile. Note that the use of backward finite differences lead to a mirrored profile, whereas central finite differences can be used to prevent this asymmetry completely, as shown in Figure 6.9 for the Laplacian-based solution. When using central finite differences, the Laplacian operator remains self-adjoint, whereas adjoint calculations get more complicated for other differential operators.




Figure 6.9: Effect of different finite difference schemes on the magnitude profile of the solution obtained by the fused (SV) regularizers.

With regard to the magnitude, the fused approaches that directly contain the spatial weights show a slight increase of the main lobe height, along with an improved reduction of the side lobes. In particular, the gradient-based method seems to benefit the most from this type of implementation.

However, these effects do only concern regularizers that involve an ℓ_2 -norm penalty, including ∇TV , ΔTV2 and \mathcal{HS}_2 . By contrast, when the spatial weights are included in the fully separable ℓ_1 -norm of the \mathcal{HS}_1 penalty, the solutions behavior remains unchanged compared to an external application of the ℓ_1 -norm on the solution. In fact, it results in the exact same objective function when the hyperparameter in (4.22) is chosen accordingly. The \mathcal{HS}_∞ regularizer sparsifies the maximum entry of the singular value vector at each source point. When embedding the spatial information directly into the singular value vector, sparsity is either promoted on the largest eigenvalue or on the spatial domain, depending on which provides the larger magnitude, but never on both simultaneously, as

$$\mathcal{R}_{\mathcal{HS}_\infty-SV}(\mathbf{q}) = \sum_{i=1}^N \|\mathcal{H}q(\mathbf{r}_i)\|_{S_\infty-SV} = \sum_{i=1}^N \max \left[\sigma_1(\mathcal{H}q(\mathbf{r}_i)), \sigma_2(\mathcal{H}q(\mathbf{r}_i)), q(\mathbf{r}_i) \right]. \quad (6.4)$$

This implementation yields unsatisfactory results as it does not induce sufficient spatial sparsity nor can it ensure a consistent behavior. However, it is still possible to incorporate the spatial information into the regularizer by simply adding it to the singular value vector.

$$\begin{aligned} \mathcal{R}_{\mathcal{HS}_\infty-SV}(\mathbf{q}) &= \sum_{i=1}^N \|\mathcal{H}q(\mathbf{r}_i)\|_{S_\infty-SV} = \sum_{i=1}^N \max_{k=1,2} \left[\sigma_k(\mathcal{H}q(\mathbf{r}_i)) + |q(\mathbf{r}_i)| \right] \\ &= \sum_{i=1}^N \max_{k=1,2} \left[\sigma_k(\mathcal{H}q(\mathbf{r}_i)) \right] + |q(\mathbf{r}_i)| = \mathcal{R}_{\mathcal{HS}_\infty}(\mathbf{q}) + \mathcal{R}_{\ell_1}(\mathbf{q}) \end{aligned} \quad (6.5)$$

This way, the \mathcal{HS}_∞ -SV penalty generates the same objective function as the \mathcal{HS}_∞ - ℓ_1 regularizer, but without the necessity of a secondary regularization term. This type of regularizer is used for all the previous and upcoming examples. This improved associated solution shows no mentionable differences in the magnitude profile, compared to an external sparsification.

6.5 Rigid piston

At first, we consider the case of a spatially extended source. A vibrating piston is considered that is mounted in a rigid baffle. We assume ideal conditions, where all points at the surface of the piston are vibrating in phase with identical displacements. The piston is located at the centre of the plate, which coincides with the origin $(x, y) = (0, 0)$ of the coordinate system. It is of circular shape and has a diameter of 40 cm. The sound pressure field is measured at a distance of $z = 0.06$ m and the equivalent source grid is placed detached from the actual source surface at a retreat distance of $z_0 = -0.01$ m. After the source strengths are determined, the components of the sound field are reconstructed at a distance of $z_S = 0.03$ m. To evaluate the performance of the different methods in terms of reconstruction accuracy, the retrieval of the sound field is conducted over a frequency range from 100 Hz to 4.5 kHz in steps of 100 Hz. The regularization parameter is chosen with $D = 30$ dB. The reconstructed spatial pattern is then compared with the sound field generated by the actual source in terms of the spatially averaged reconstruction error. This analysis is carried out for the acoustic pressure, normal velocity and normal intensity field. Moreover, a visual examination of the estimated equivalent sources is performed.

According to Table 6.5 the ∇ TV regularizer that favors piecewise-constant patterns outperforms the other regularizers in terms of mean-square error. Figure 6.10 shows the equivalent source distributions of the rigid piston provided by the different methods for an exemplary frequency of 1 kHz.




Figure 6.10: Estimated equivalent sources for the rigid piston vibrating at 1 kHz.

regularizer	ℓ_1 -norm	ℓ_2 -norm	∇ -TV	Δ -TV2	\mathcal{HS}_1	\mathcal{HS}_2	\mathcal{HS}_∞
avg. error [%]	672.82	36.59	22.00	32.87	33.37	32.27	34.02

Table 6.5: Spatially-averaged reconstruction error for the rigid piston vibrating at 1 kHz.

The portrayed solutions highlight the characteristic properties of the associated methods, as already discussed in more detail for the example of the vibrating plate. For the ℓ_1 -norm, the surface of the piston is decomposed into a sparse cluster of point sources with a few outliers, only merely hinting at its original form. The ℓ_2 -norm correctly identifies the spatial structure of the piston, but we obtain a relatively high level of side lobes surrounding the target source. All derivative-based methods introduce sporadic artifacts or ghost sources along the borders, all to differing degrees and at different locations. They are induced by the applied boundary conditions used for the calculation of the finite difference derivatives. The ∇ TV is the method that best preserves the contours of the piston while simultaneously keeping its interior surface at a constant magnitude. The source appears slightly extended due to the bright shading artifacts around the edges, as a result of the inherent staircasing effect. When looking at the linear image distribution, the Laplacian-based approach Δ TV2 again reveals a continuously varying and somewhat rhomb-shaped estimate of the source due to the lack of smoothing along the diagonal direction. This deficiency is compensated by the Hessian Schatten-norm regularizers, whose mutual differences are very subtle.

These equivalent sources are now used to model the sound field at the specified reconstruction distance, respectively. The observed sound field generated by the actual piston at a distance of 3 cm is illustrated in Figure 6.11. Note that the depicted quantities are all related to their respective reference value.




Figure 6.11: True sound field radiated by the rigid piston 3 cm above the source plane.

For reasons of clarity, the spatial distribution of the reconstruction error of the restored acoustic fields is not examined in detail. For each regularizer, the average error, regarding the recovery of the pressure, velocity and intensity field, has been respectively determined as a function of frequency. For all cases, an ensemble average value is given, using four different realizations of additive noise per frequency. The obtained results are depicted in Figure 6.12.




Figure 6.12: Spatially-averaged reconstruction error for the reconstructed sound pressure (left), normal velocity (center) and normal intensity field (right).

The observed error behavior allows to conclude some general statements regarding the reconstruction accuracy of the regularizers. For the considered scenario, the derivative-based schemes outperform the state-of-the-art approaches of ℓ_1 - and ℓ_2 -norm minimization at almost the whole frequency range of interest. The error is considerably lower with high consistency for all three sound field quantities. The sound field recovered by the sparse ℓ_1 -norm estimate shows the largest deviation from the true observations, as the piston is not sparse itself. The higher the frequency gets, the stronger becomes the disintegration of the piston surface into isolated point sources, which leads to a collapse in error towards high frequencies. The ℓ_2 -norm leads to an accurate reconstruction of the sound field at low to mid frequencies, as the estimated equivalent sources yield a fairly good approximation of the piston structure. Concerning frequencies above the spatial sampling limit, the restoration quality radically decreases due to aliasing effects and other magnitude distortions. This makes the least-squares approach unsuitable for wideband reconstruction. By contrast, all derivative-based methods are able to provide a good approximation of the sound field way beyond the sampling limit with a still tolerable amount of error. A visual assessment of the recovered sound fields at high frequencies emphasize their superior abilities as wideband regularizers. They are able to properly estimate the shape and extent of the source, where the classical approaches already fail to provide a valid solution. The source estimate of ∇TV best complies with the non-smooth shape and constant magnitude of the piston, hence it mostly succeeds at best in emulating its radiation pattern. The second-order methods consistently show a very similar behavior. At low to mid frequencies, the error curves are nearly congruent, whereas they slightly drift apart at frequencies above the spatial sampling limit, revealing some minor differences in error. Among the second-order methods, the \mathcal{HS}_1 regularizer consistently yields the lowest average error, closely followed by the \mathcal{HS}_2 solution.

Fused regularizers

As already mentioned, to make the proposed regularizers also usable for the recovery of sparse sources, an additional ℓ_1 -norm acting on the solution is beneficial. It is very likely to encounter sound fields radiated from extended vibrating structures in conjunction with confined local excitations in the immediate neighbourhood. For this reason, we are also interested in discovering the impact of the added sparsity constraint on the reconstruction of a spatially extended source. The same study is repeated for the fused approaches introduced in the previous example. The hyperparameter controlling the amount of spatial sparsity is set to $\mu = 0.5$. Again, the disparities between the explicit and implicit implementation of the ℓ_1 -norms are of additional interest. At first, the estimated source coefficients for a radiation frequency of 1kHz are presented in Figure 6.13.




Figure 6.13: Estimated equivalent sources for the rigid piston vibrating at 1 kHz retrieved with the fused regularizers with external ℓ_1 -norm.

regularizer	$\nabla\text{TV}-\ell_1$	$\Delta\text{TV}2-\ell_1$	$\mathcal{HS}_1-\ell_1$	$\mathcal{HS}_2-\ell_1$	$\mathcal{HS}_\infty-\ell_1$
avg. error [%]	23.59	31.91	32.93	31.70	33.78

Table 6.6: Spatially-averaged reconstruction error for the rigid piston vibrating at 1 kHz retrieved with the fused regularizers with external ℓ_1 -norm.

The most notable feature of the additional sparsity constraint is the significantly reduced degree of coefficient activity in the surrounding of the piston. This allows a more precise distinction of the actual source as it minimizes the level of side lobes and suppresses misleading noise artifacts. The applied sparsification seems to work best at the corners of the aperture. However, there are still contiguous regions of remaining coefficient activity, which are mainly related to the effects of the prescribed image boundary conditions. While the $\nabla\text{TV}-\ell_1$ is still able to recover a source of constant magnitude over the whole extent of the piston, the second-order solutions literally tear a hole into the piston's centre. As second-order methods do all promote piecewise-linear

solutions, they intend to model the flat piston head as a series of ridges and valleys, where the ripple increases with frequency. The additional sparsity constraint tries to minimize the source components wherever possible, which produces a dip in the middle of the extended piston source for the observed frequency. And yet, except the gradient-based approach, all regularizers bring a slight improvement in terms of the spatially averaged reconstruction error. To further evaluate their performance for the retrieval of the acoustic field, the radiation due to the equivalent source estimates is calculated. Figure 6.14 illustrates the resulting spatially averaged error versus frequency for the reconstruction of the pressure, velocity and intensity field.




Figure 6.14: Spatially-averaged reconstruction error for the reconstructed sound pressure (left), normal velocity (center) and normal intensity field (right) retrieved with the fused regularizers with external l_1 -norm.

The error trend of the reconstruction remains largely unchanged in the presence of an additional sparsity constraint on the solution. A slight offset of the curves towards a higher error level can be observed but only is of little relevance regarding the outcome. Although the fused regularizers inherit the ability of wideband reconstruction, they get slightly worse at high frequencies. This results in a somewhat sparse representation of the piston, comprised of many single coefficients of high magnitude to fit the data fidelity energy. When comparing the methods to the classical approaches, they still provide either equally accurate or more accurate results across the whole spectrum.

At last, we are eager to assess possible benefits or drawbacks from an implicit implementation of the ℓ_1 -norm penalty within the derivative-based regularizer itself. Figure 6.15 shows the estimated equivalent source distribution of the associated fused regularizers.




Figure 6.15: Estimated equivalent sources for the rigid piston vibrating at 1 kHz retrieved with the fused regularizers (SV) with internal ℓ_1 -norm.

regularizer	$\nabla\text{TV-SV}$	$\Delta\text{TV2-SV}$	$\mathcal{HS}_1\text{-SV}$	$\mathcal{HS}_2\text{-SV}$	$\mathcal{HS}_\infty\text{-SV}$
avg. error [%]	34.90	36.24	32.85	35.38	33.95

Table 6.7: Spatially-averaged reconstruction error for the rigid piston vibrating at 1 kHz retrieved with the fused regularizers (SV) with internal ℓ_1 -norm.

The embedding of the spatial sparsity into the regularizer yields a significant loss of quality in the reconstruction of the source. This mainly affects the regularizers, where the spatial information is grouped along with the spatial derivatives in terms of energy by a ℓ_2 -norm mapping, i.e. the $\nabla\text{TV-SV}$, $\Delta\text{TV2-SV}$ and the $\mathcal{HS}_2\text{-SV}$ penalty. In the case of $\nabla\text{TV-SV}$, coefficient activity is always accompanied by a non vanishing gradient, which yields an enforcement of continuous variations within active regions. The piston head is no longer of constant magnitude but varies smoothly between adjacent equivalent sources. The constant plateaus at the boundaries of the source plane have been replaced by smooth transitions that fade out naturally. Moreover, the former well preserved, sharp edges of the piston are smeared out which effectively eliminates the staircasing artifacts. In fact, the fused first-order $\nabla\text{TV-SV}$ shows a behavior very similar to that of a second-order derivative penalty. The main difference lies in the slightly decentralized appearance of the source distribution due to the forward finite difference implementation of the derivatives, as already discussed in the previous section. When including spatial sparsity in $\Delta\text{TV2-SV}$ or $\mathcal{HS}_2\text{-SV}$, coefficient activity only comes along with non-zero second-order derivatives. As soon as a coefficient is active, it prevents, or at least alleviates the enforcement of a minimal curvature of the magnitude surface at this point. This causes an amplification of the ripple on the piston head, as ridges and valleys that approximate the constant of the source

provide sharper and sparser edges with increased curvature. The ripple forms a curl around the center of the piston, which otherwise would only emerge at higher frequencies due to smaller wavelengths. For the \mathcal{HS}_2 -SV penalty, this effect gets slightly attenuated presumably due to the additional diagonal smoothing. The increased waviness also becomes noticeable at the contours of the piston as well as through the large variations in the boundary pattern. The reconstruction error associated with the ℓ_2 -norm-based methods consequently suffers a large increase in value. By contrast, the \mathcal{HS}_1 -SV and \mathcal{HS}_∞ -SV regularizer produce a result almost identical to the one obtained by their \mathcal{HS} - ℓ_1 -norm equivalent. The mere distinctions result from the different ways the spatial sparsity is incorporated into the algorithm and from the ensemble-averaging. The explicit implementation of the ℓ_1 -norm in form of a secondary regularization term always depends on an additional proximal parameter to ensure convergence. This leads to a slightly different weighting of the norm, which typically can be tuned but must be accepted when working with an ad hoc choice of the regularization parameter. It reflects in a minor change of the average error for the \mathcal{HS}_1 -SV and \mathcal{HS}_∞ -SV penalty, which is no great issue in practice.

To investigate the indirect implementation on a broader perspective, their estimated source distributions are again used to recover the pressure, velocity and intensity field. The resulting spatially averaged error is depicted in Figure 6.16.




Figure 6.16: Spatially-averaged reconstruction error for the reconstructed sound pressure (left), normal velocity (center) and normal intensity field (right) retrieved with the fused regularizers (SV) with internal ℓ_1 -norm.

It becomes apparent that the successful application of ∇TV -SV, $\Delta\text{TV}2$ -SV and \mathcal{HS}_2 -SV is strictly confined to frequencies below the spatial sampling limit. In this part of the spectrum, the reconstruction error remains within a tolerable limit and the methods can still compete with the least-squares solution. Above the sampling limit, the methods fail to produce meaningful results. The physical shape of the source is no longer recognizable as the ripple and disintegration of the source increases with frequency. Contrarily, the \mathcal{HS}_1 -SV and \mathcal{HS}_∞ -SV penalty provide a more stable behavior over the depicted frequency range. They maintain their ability of wideband reconstruction far beyond the spatial sampling limit but ultimately break down at very high frequencies where the spatial sparsity becomes too dominant.

6.6 Point source distribution

This section is dedicated to the evaluation of the methods performance of reconstructing spatially compact sources and their associated acoustic fields. The source under study is a longitudinal quadrupole, composed of an arrangement of three monopole sources, with the middle one oscillating with an inverted phase and twice the amplitude. The quadrupole provides a length of about 25 cm and is placed in the centre of the source plane. The volume velocity of the center monopole is $Q_c = 2 \cdot 10^{-5} \frac{m^3}{s}$, which results in a source strength of $Q_{l,r} = 10^{-5} \frac{m^3}{s}$ for the two adjacent monopoles. To explore the impact of the boundary conditions on the reconstruction, two additional point sources are positioned on the diagonal axis near the opposite corners of the source frame. Their volume velocity is equally chosen to be $Q_d = 1.5 \cdot 10^{-5} \frac{m^3}{s}$ and both are radiating with opposite phase. For a sparse source distribution, the used regularization parameter choice method (4.27) reveals its weaknesses. Although the additive noise still generates the same SNR of 30 dB, the estimated dynamic range necessary for the selection of the regularization parameter is chosen $D = 50$ dB. Otherwise, the parameter yields an overregularized solution for all methods. Thus the impact of a possible mismatch between the true and the estimated SNR on the reconstruction error is the subject of a separate chapter. The hyperparameter that controls the amount of spatial sparsity is set to $\mu = 1$. The analysis is conducted with the same simulation setup as for the piston, but with the equivalent sources directly placed on the source plane at $z_0 = 0$. The following study aims at investigating the regularizers performance of localization and separation of isolated compact sources. The prescribed source distribution is estimated by all methods and the accuracy of the solutions is verified visually and in terms of reconstruction error. In addition, the methods capability of reproducing the sound field of the source is assessed in a forward radiation scenario involving the estimated sources. The considered reconstruction distance remains at $z_S = 0.03$ m. To reveal the methods ability to overcome the aliasing limit, the recovery of the field is explored over a broad frequency range.

Table 6.8 provides the spatially averaged reconstruction error of the methods, according to which the ℓ_1 -norm clearly outperforms the other regularizers. Figure 6.17 illustrates the true sparse distribution of the source and the equivalent sources estimated by the different methods at a radiation frequency of 3 kHz. The solutions are depicted with a 20 dB dynamic range.

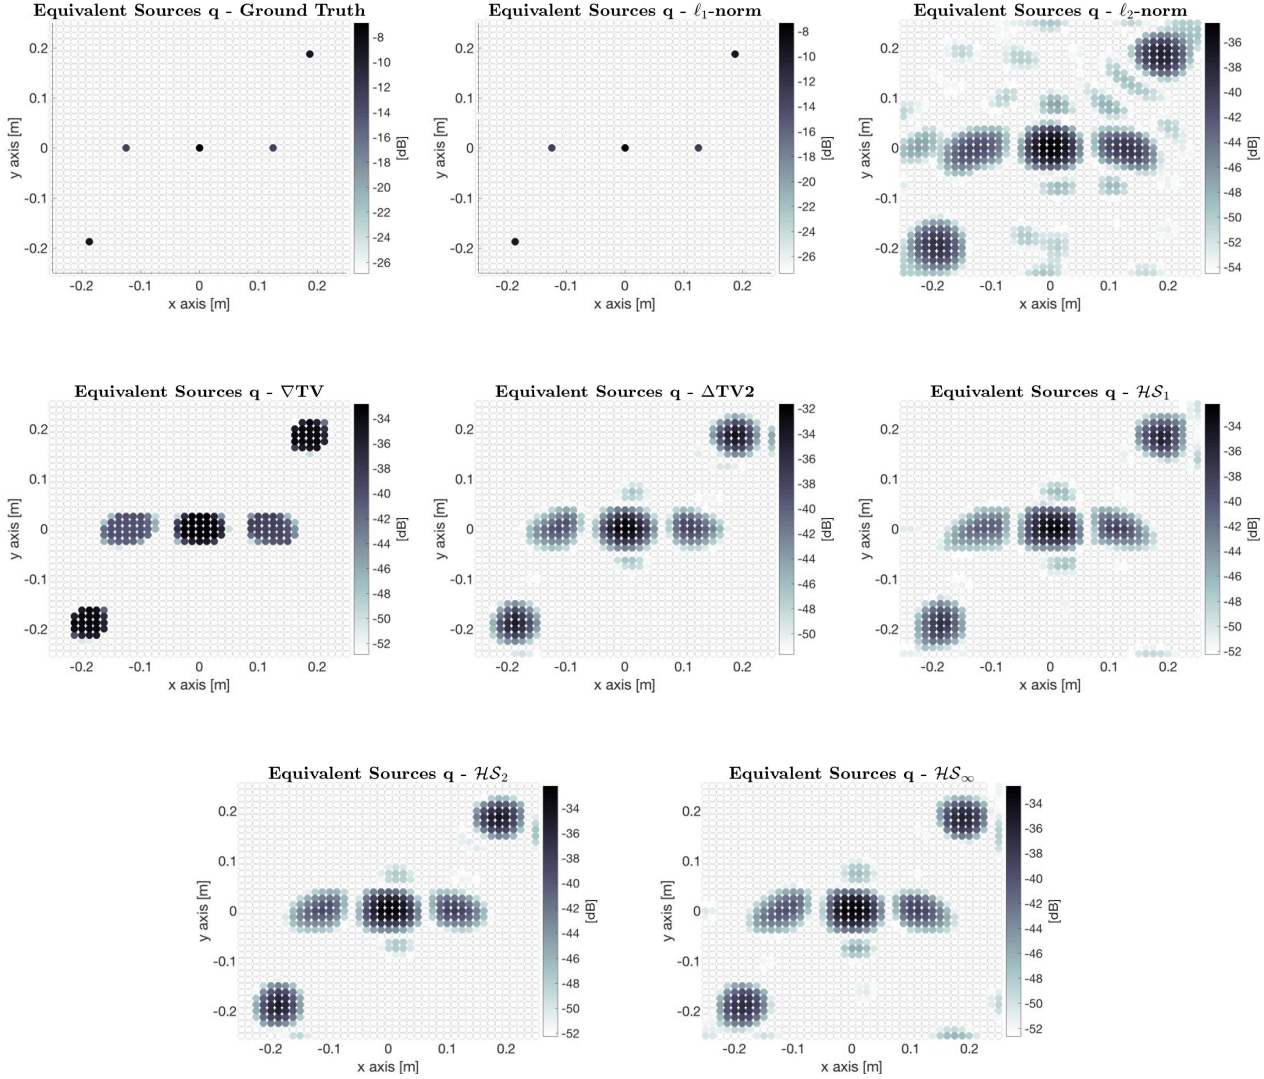


Figure 6.17: Estimated equivalent sources for the quadrupole and monopole sources vibrating at 3 kHz.

regularizer	ℓ_1 -norm	ℓ_2 -norm	∇ -TV	Δ -TV2	\mathcal{HS}_1	\mathcal{HS}_2	\mathcal{HS}_∞
avg. error [%]	5.79	98.03	97.65	96.54	97.37	96.97	97.44

Table 6.8: Spatially-averaged reconstruction error for the quadrupole and monopole sources vibrating at 3 kHz.

As expected, the ℓ_1 -norm minimization yields the most accurate solution as the true source distribution is sparse. It reveals the correct location of the monopoles and is able to closely approximate the true magnitude level of the source, as it concentrates the whole energy in very few coefficients. By contrast, the ℓ_2 -norm tends to evenly spread the energy throughout the solution vector space, which yields a poor reconstruction. Most of the energy is divided into several coefficients surrounding the actual source points, making it difficult to identify their true position. Also the high level of side lobes disturbs the quality of the reconstruction. Since the method provides the solution with minimum energy, a severe difference in magnitude level can be observed, resulting in a very high average error. The ∇ TV approach depicts the point sources as rectangular patches with a slightly increased magnitude compared to the least-squares solution. This yields an improved localization as it narrows down the area of the possible source locations

to a better extent. Again, staircasing artifacts become apparent at the borders of the blocks, leading to a slight spatial extension. The asymmetry in the source distribution contributes to the antiphase radiation of the monopoles in the corners. The second-order approaches all yield very similar estimates, where the differences are just about recognizable with the bare eye. As they enforce smooth variations with minimum curvature, their display of the sources closely resembles that of the ℓ_2 -norm, but with a substantially lower level of side lobes. The resulting energy surplus in the main lobes of the sources leads to a better approximation of their true magnitude, but there still remains a large gap between. Moreover, sharp edged artifacts emerge at the corners where the point sources are located, due to the applied boundary conditions. A final remark must be made regarding the spatially-averaged reconstruction error of the source estimates. Since only the ℓ_1 -norm regularization directly enforces sparsity in the spatial domain, it is the one single method that can successfully recover a truly sparse source as in the present case within a suitable error bound. The other methods all involve multiple active coefficients to model the source, which leads to a spread of the concentrated energy of the point sources. This results in a very large difference in magnitude at the source points which yields an excessive boost of the average relative error. Nevertheless, the methods can still be used for the purpose of source localization, as well as for the recovery of the radiated sound field at a location different from the measurement plane. To explore the methods abilities in this regard, the radiated sound field of the point sources is reconstructed at half the measurement distance by their respective equivalent sources. Figure 6.18 shows the true sound pressure, normal velocity and intensity level across the reconstruction plane at 3 kHz.




Figure 6.18: True sound field radiated by the quadrupole and monopole sources 3 cm above the source plane.

The associated spatially-averaged error is depicted in Figure 6.19 for the frequency range of interest.




Figure 6.19: Spatially-averaged reconstruction error for the reconstructed sound pressure (left), normal velocity (center) and normal intensity field (right).

As we can see, the error for the reconstruction of the radiated sound field is far less severe than on the source plane, as the model is better fitted closer to the measurement plane. At least for frequencies below the sampling limit, the pressure field can be recovered with a still reasonable amount of error. The ℓ_1 -norm optimized solution consistently shows the smallest error across the whole spectrum.

Fused regularizers

To cope with the sparse nature of the source, an additional sparsity constraint on the spatial domain seems beneficial. At first, we examine the use of the secondary sparsifying regularization term. The associated equivalent sources are depicted in Figure 6.20. Table 6.9 provides the average reconstruction error of the methods.

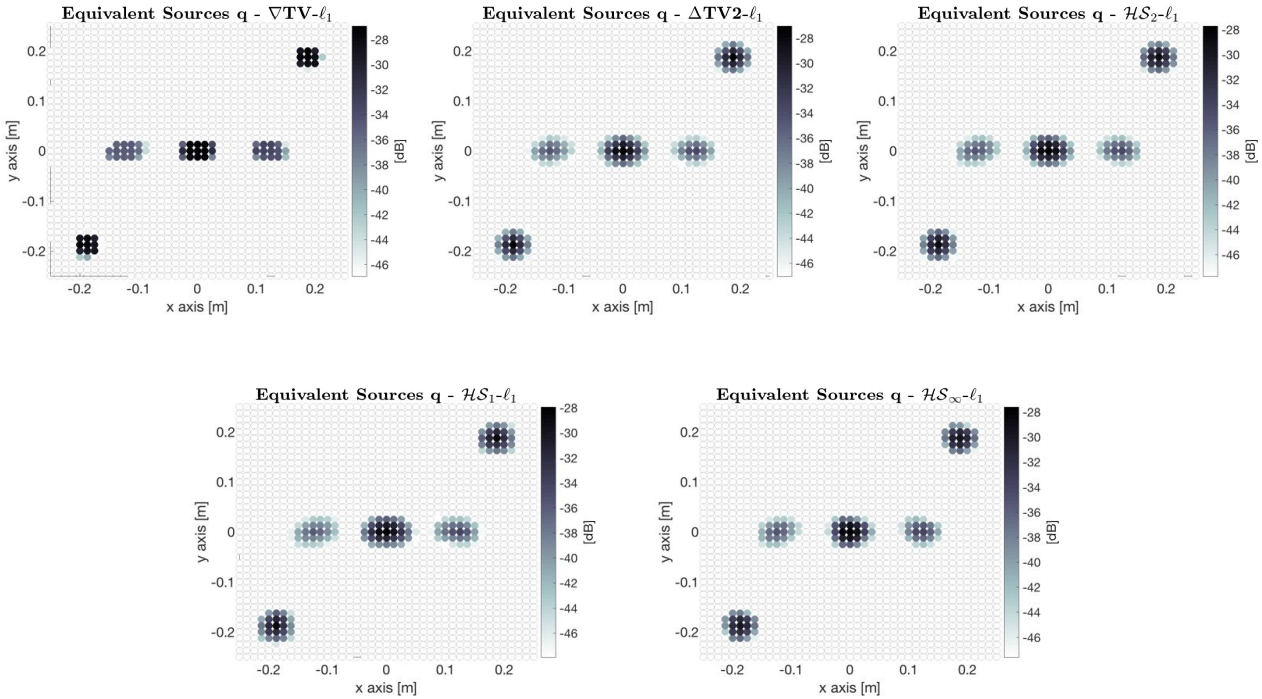


Figure 6.20: Estimated equivalent sources for the quadrupole and monopole sources vibrating at 3 kHz retrieved with the fused regularizers with external ℓ_1 -norm

regularizer	$\nabla TV-\ell_1$	$\Delta TV2-\ell_1$	$\mathcal{HS}_1-\ell_1$	$\mathcal{HS}_2-\ell_1$	$\mathcal{HS}_\infty-\ell_1$
avg. error [%]	93.99	92.42	93.57	93.51	93.87

Table 6.9: Spatially-averaged reconstruction error for the quadrupole and monopole sources vibrating at 3 kHz retrieved with the fused regularizers with external ℓ_1 -norm.

In general, the equivalent sources appear more spatially confined around the location of the point sources, yielding improved localization. Moreover, they experience amplified energy in that region, which leads to a better approximation of the real volume velocities. The additional sparsity constraint strongly attenuates the level of side lobes and other artifacts to way below the depicted dynamic range. Thus, the separation of the quadrupole source components has improved, as well as the suppression of the ghost sources in their vicinity. The patches of the $\nabla TV-\ell_1$ solution uniformly shrank in size but still exhibit the staircasing artifacts along their edges. The second-order methods maintain their smooth magnitude variations, but yet with

a much steeper ascent towards the source points. This is the result of the compromise made between minimum coefficients and minimum curvature. Figure 6.21 reflects this improvements on a broader frequency range.




Figure 6.21: Spatially-averaged reconstruction error for the reconstructed sound pressure (left), normal velocity (center) and normal intensity field (right) retrieved with the fused regularizers with external ℓ_1 -norm.

A significant reduction of the reconstruction error is observable, especially at frequencies above the spatial sampling limit.

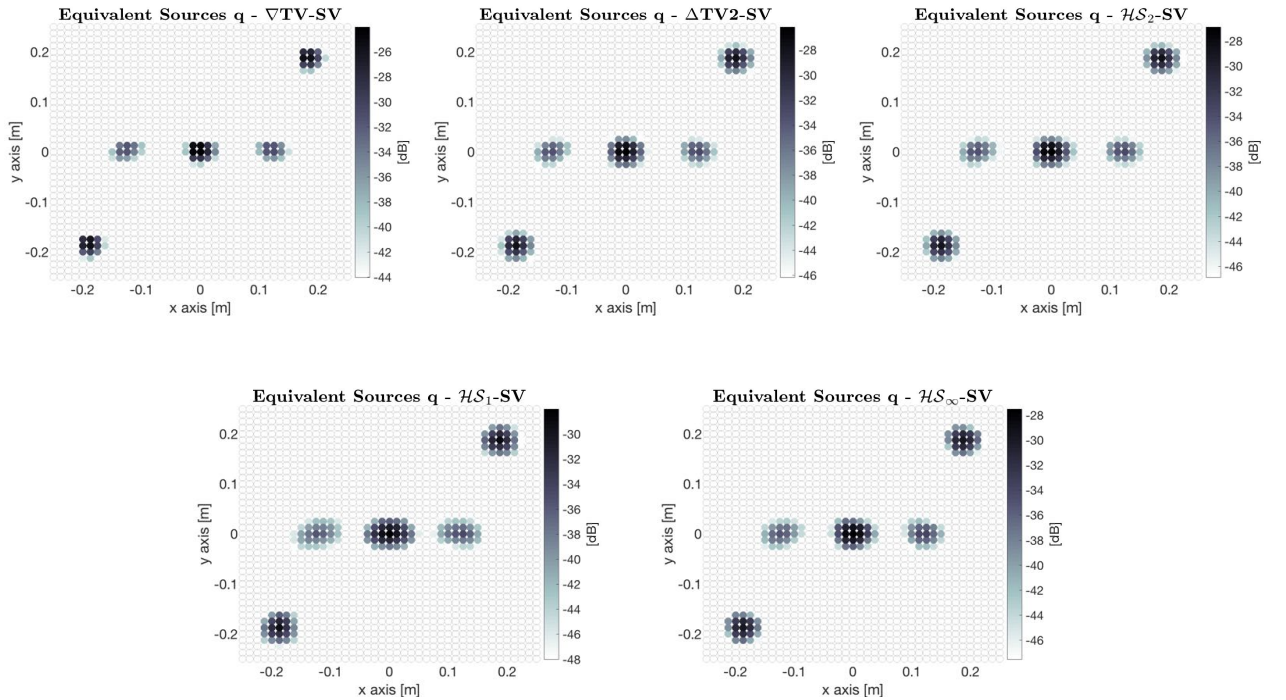


Figure 6.22: Estimated equivalent sources for the quadrupole and monopole sources vibrating at 3 kHz retrieved with the fused regularizers (SV) with internal ℓ_1 -norm.

regularizer	$\nabla\text{TV-SV}$	$\Delta\text{TV2-SV}$	$\mathcal{H}\mathcal{S}_1\text{-SV}$	$\mathcal{H}\mathcal{S}_2\text{-SV}$	$\mathcal{H}\mathcal{S}_\infty\text{-SV}$
avg. error [%]	90.53	92.28	93.61	92.98	93.86

Table 6.10: Spatially-averaged reconstruction error for the quadrupole and monopole sources vibrating at 3 kHz retrieved with the fused regularizers (SV) with internal ℓ_1 -norm.

The embedding of the spatial information into the ℓ_2 -norm-based regularizers introduces the asymmetry discrepancies, already encountered in the example of the piston. However, in case of a sparse source distribution, these effects don't seem as prevalent as they do not affect the reconstruction quality that much. The sparse ∇TV -SV discards its piecewise-constant appearance in favor of a smooth, continuous intensity surface, indicated by the larger variety of magnitude levels at the source points. The remaining regularizers do not display any significant alterations in the spatial pattern compared to the solutions depicted in Figure 6.20 due to the different kind of implementation. This is also reflected in the nearly unchanged averaged error of the source estimate in Table 6.10. Regarding the reconstruction of the sound field, a slight reduction of the average error is achieved with the ℓ_2 -norm-based regularizers, as shown in Figure 6.23.




Figure 6.23: Spatially-averaged reconstruction error for the reconstructed sound pressure (left), normal velocity (center) and normal intensity field (right) retrieved with the fused regularizers (SV) with internal ℓ_1 -norm.

6.7 Sensitivity to noise and parameter choice

A decisive attribute that determines the quality of a regularizer is its robustness to measurement noise. A certain degree of noise is inevitable in any physical measurement setup. When these measurements are used to estimate the unknown equivalent source parameters, the immanent noise becomes the dictating factor for the validity of the reconstruction. For high levels of additive noise, the measurements provide increasingly misleading information that deteriorates the estimation accuracy. Therefore a robust regularization scheme must be applied in order to obtain a meaningful solution. To evaluate the sensitivity of the proposed regularization schemes, they are used to reconstruct the source distribution and sound field quantities under different SNR conditions. Hereby, the regularization parameter D in equation (4.27) is chosen according to the resulting SNR. The source considered is the rigid piston of the preceding example, oscillating at $f = 500$ Hz, as the methods are typically more sensitive to noise at such low frequencies. This time, the equivalent sources are placed right at the source surface at $z_0 = 0$ without using a retreat distance. Figure 6.24 depicts the spatially averaged reconstruction error of the methods for the different noise levels, ensemble-averaged over four different realizations of noise. This time, the ℓ_1 -norm approach is neglected for the non-sparse source distribution as it yields unusable results.

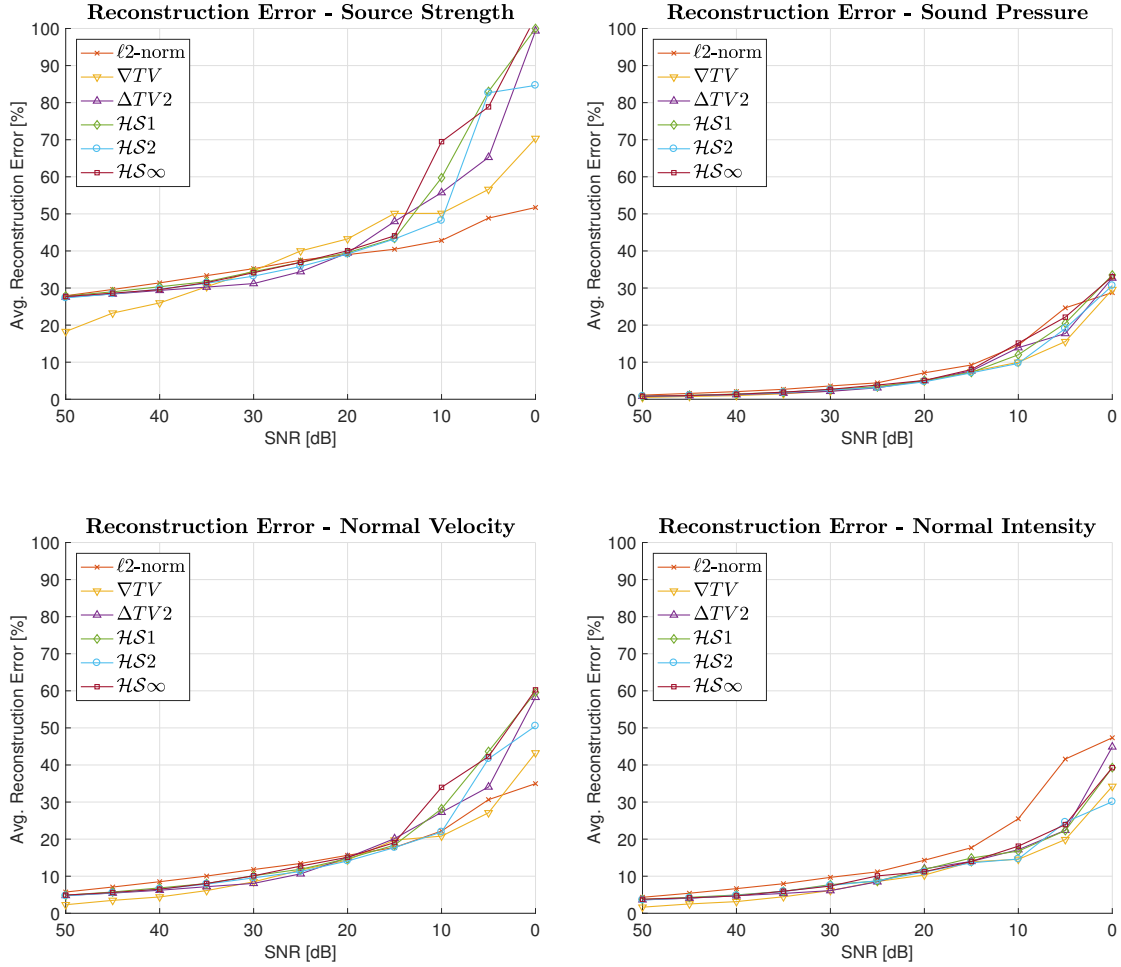


Figure 6.24: Spatially averaged reconstruction error of the proposed regularizers for different levels of SNR.

When retrieving the equivalent source strengths, the ℓ_2 -norm solution is found to be the most robust to noise, as the decreasing SNR yields the least rise in error. However, for low to mid levels of measurement noise, it shows a higher reconstruction error than the derivative-based methods. The most accurate retrieval is obtained by the ∇TV as the solution best reflects the homogenous shape and contour of the piston, provided that the SNR is sufficiently high. The more noise is present, the less accurate the solution will be as it increases staircasing at the edges and introduces occasional plateaus of constant magnitude in the source's surrounding. For low to mid noise exposure, the second-order methods consistently provide lower reconstruction error than the least-squares approach, but they seem to be more susceptible to high levels of noise as the error rapidly increases. Considering the reconstruction of the sound field, the observed error behavior is different. Besides a general decrease of the reconstruction error due to the smaller model mismatch closer to the measurement plane, ℓ_2 -norm minimization almost consistently yields the poorest results, except for the retrieval of the normal velocity from very noisy measurements. The second-order methods all show a very similar error trend, indicating robust reconstruction of the sound field up to very high noise levels. The most accurate solution is obtained by the ∇TV approach, showing effective suppression of noise and artifacts.

Another crucial aspect for an accurate retrieval is to define a suitable amount of regularization. The right balance between trusting the observed data and imposing a certain structure on the solution assumed to fit the true source distribution must be found. Since this is mainly dependent on the contamination of noise in the measurements, it is no trivial task as the noise level is usually unknown a priori and needs to be estimated. The estimated dynamic range is required for a proper selection of the regularization parameter D in (4.27). To show the impact of a wrong parameter choice, the reconstruction scenario is repeated, but this time with a fixed SNR of 30 dB for multiple estimations of the available dynamic range. Figure 6.25 shows the error path for the different methods.




Figure 6.25: Spatially averaged reconstruction error of the proposed regularizers for different estimations of the available dynamic range D for the regularization parameter selection according to Eq. (4.27) for a given SNR=30dB.

The error path reveals the weakness of the proposed parameter choice method, as the optimal regularization parameter for the least reconstruction error is found at an applied dynamic range level D different from the true SNR. Moreover, the location of the optimum value depends on distribution and frequency of the source and the regularization method itself, hence making it impossible to derive general statements about the correct choice of D . However, the parameter choice method empirically yields a slightly over-regularized solution, thus the applied dynamic range D should be chosen higher than the true SNR. Except of the ∇TV solution, where the best result is achieved with $D = 45$ dB, all methods show a minimum reconstruction error at around $D = 40$ dB. A bigger choice of D results in an increasing under-regularization of the problem, which yields a significant increase in error due to the gradually recurring ill-posedness of the

problem. Compared to the ℓ_2 -norm and ∇TV approach, the impact of the second-order methods on the solution seems to decrease faster with an increase of D . Contrarily, they yield far lower reconstruction errors for smaller values of D , in case of over-regularization, where the ℓ_2 -norm and ∇TV start to break down. With decreasing D , hence increasing regularization parameter λ , ℓ_2 -norm minimization provides increasingly oversmoothed results as it steadily shrinks the solution coefficients towards zero. Due to the resulting low amplitudes of the equivalent sources, they can only poorly radiate sound anymore, thus the over-regularized least squares fails at sound field reconstruction. The ∇TV approach gradually groups the solution components, up to a point where the equivalent source distribution is melted to one single group of constant magnitude. For a further decrease of D , the coefficients magnitude and therefore the reconstruction error remain unchanged. In this example, the second-order methods seem more robust to a potential over-regularization and yield lower or equal reconstruction error compared to the classical least squares approach around the optimum parameter choice. The first-order TV may provide the best reconstruction at its optimum value, but is very sensitive to a wrong selection of the regularization parameter.

An additional sparsity constraint on the solution will alter the methods robustness to noise and parameter mismatch. To explore the impact of an additional ℓ_1 -norm on the error behavior, the study is repeated for the sparsifying fused regularizers. The hyperparameter controlling the amount of spatial sparsity is set to $\mu = 0.5$. Figure 6.26 illustrates the average reconstruction error of the fused approaches for different SNR.

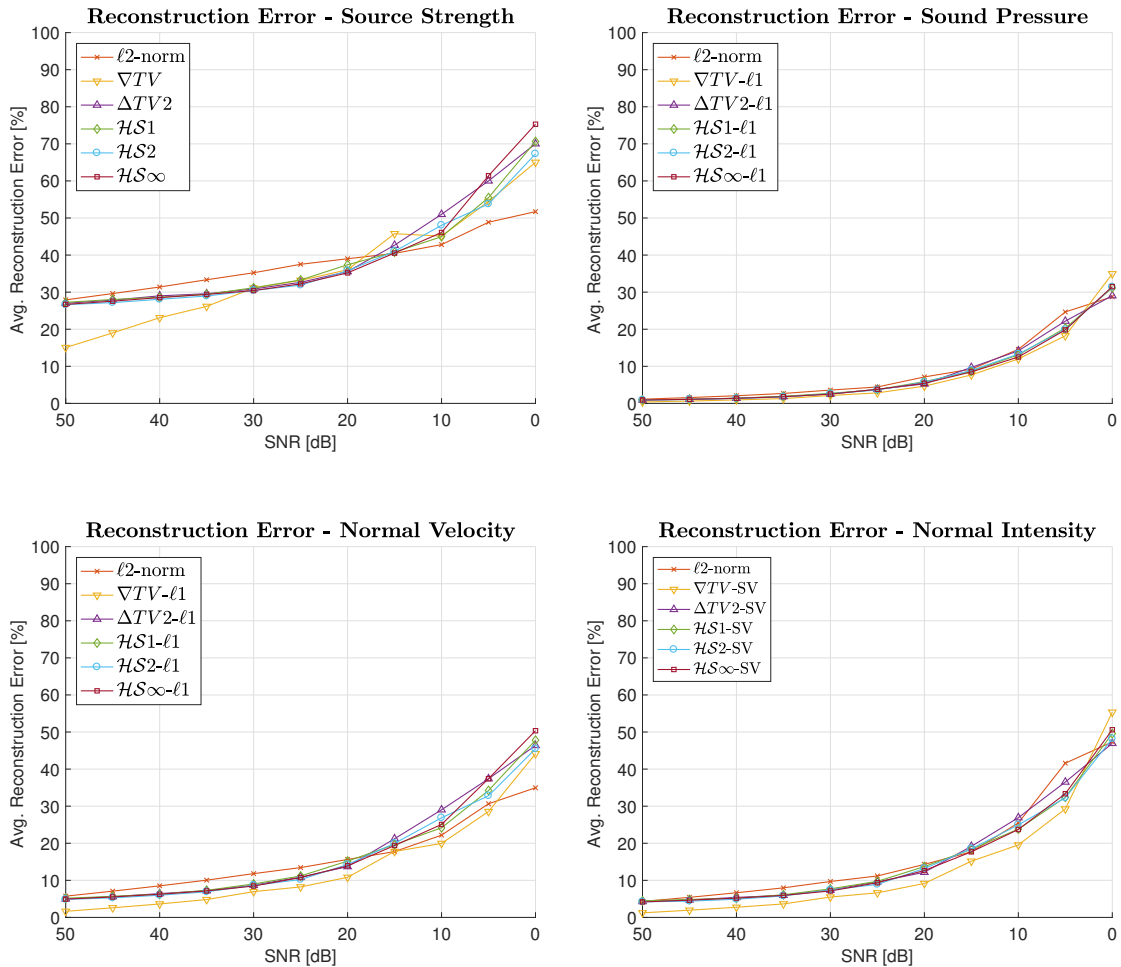


Figure 6.26: Spatially averaged reconstruction error of the proposed regularizers with additional external ℓ_1 -norm penalty for different levels of SNR.

It becomes apparent that the additional requirement of spatial sparsity on the solution provides effective suppression of noise and related artifacts, making the methods more robust towards outliers and spurious ghost sources. Compared to the original methods, lower reconstruction errors are achieved consistently for all SNR levels and the great increase in error at very low SNR is attenuated. Only the retrieval of the normal intensity suffers from the spatial sparsity constraint at higher noise levels as the error slightly increases.

The error path of the methods in Figure 6.27 provides further insights in their solutions behavior for different choices of the regularization parameter.

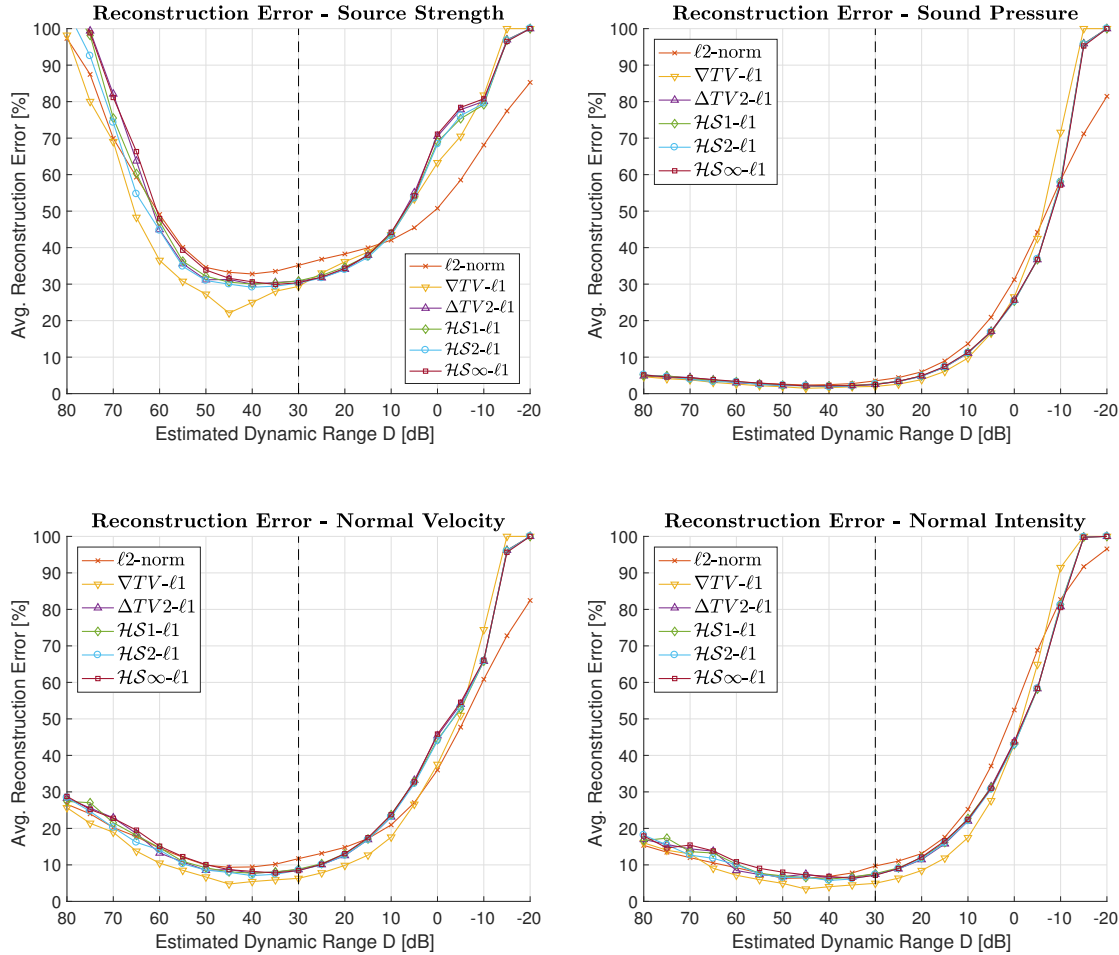


Figure 6.27: Spatially averaged reconstruction error of the proposed regularizers with additional external ℓ_1 -norm penalty for different estimations of the available dynamic range D for the regularization parameter selection according to Eq. (4.27) for a given $\text{SNR}=30\text{dB}$.

For large estimates of the available dynamic range D and thus small values of the regularization parameter λ , the error of the fused methods behaves similar to the original methods, although they seem to be less affected by the measurement noise as the error is slightly smaller and rises less rapidly. Moreover, the minimum error almost remains the same within a broad range around the optimum value, lowering the demand for perfect parameter estimation. For small estimates of D , hence large values of λ , the error rapidly increases as the sparsifying effect of the ℓ_1 -norm penalty gets prevalent in the reconstruction. The piston is gradually represented by fewer active equivalent sources, up until all source amplitudes are forced to zero, leading to enormous reconstruction errors. In general, the additional sparsity constraint enhances the methods robustness to noise and yields more accurate reconstruction around the optimum parameter choice, but makes them very sensitive to a possible over-regularization in return.

When considering the Sparse Variation (SV) regularizers, where the ℓ_1 -norm penalty is incorporated within the regularizer, the error behavior stays roughly the same in case of the \mathcal{HS}_1 -SV and \mathcal{HS}_∞ -SV regularizer. Nonetheless, this does not hold for the ℓ_2 -norm-based regularizers ∇TV -SV, $\Delta TV2$ -SV and \mathcal{HS}_2 -SV, therefore their behavior in the presence of noise deserves their own investigation. Figure 6.28 depicts the reconstruction error of the Sparse Variation methods for different noise levels.




Figure 6.28: Spatially averaged reconstruction error of the proposed regularizers with additional internal ℓ_1 -norm penalty (SV) for different levels of SNR.

For this example, including the spatial information within the inner ℓ_2 -norm of the regularizers, and thus implicitly enforcing spatial sparsity, provides much worse results than an externally applied sparsifying penalty. The error trend of the second-order methods is comparable to the original methods, whereas the first-order ∇TV -SV approach suffers a large increase in error at lower noise levels. With the extension to Sparse Variation, the original method loses its characteristic piecewise-constant magnitude profile in turn for a smoothly varying one and hence is no more able to preserve the sharp edges of the piston, leading to a higher reconstruction error.

For completeness, the influence of a wrong estimation of the available dynamic range for regularization parameter selection on the accuracy of the SV solutions is depicted in Figure 6.29, from which similar conclusions can be drawn. In general, these methods have proven unsuitable for the characterization of extended sources of spatially constant magnitude.

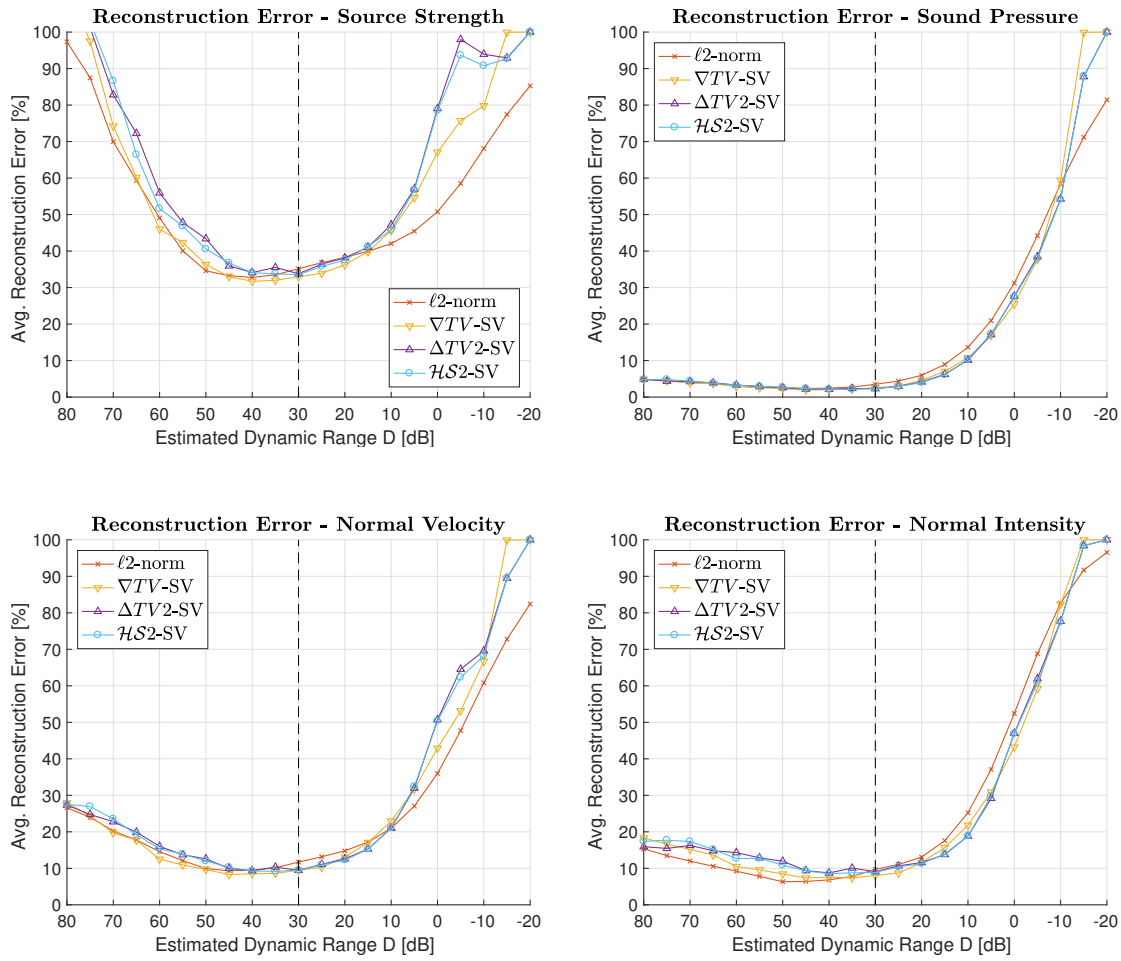


Figure 6.29: Spatially averaged reconstruction error of the proposed regularizers with additional internal ℓ_1 -norm penalty (SV) for different estimations of the available dynamic range D for the regularization parameter selection according to Eq. (4.27) for a given $SNR=30$ dB.

Regularizer Characteristics - Experimental Study

In order to evaluate the proposed regularization methods under real-world conditions, an experimental study is carried out based on already available measurement data, collected in the course of the master's thesis [Mit18] conducted at the IEM (Institute of Electronic Music and Acoustics) at the KUG (University of Music and Performing Arts, Graz). The device under study is a custom made box, equipped with multiple loudspeaker drivers and vibration exciters, which can be driven individually or in groups. With this source, sound fields of different complexity can be generated by either compact or extended source distributions, making it a suitable choice for validation of the methods under laboratory conditions.

7.1 Experimental setup




Figure 7.1: The sound source under study [Mit18].

The sound source is a rectangular box of geometry $70\text{cm} \times 60\text{cm} \times 48\text{cm}$ with loudspeaker drivers and exciters mounted on each side surface. For the purpose of this study, only two sides are considered, one equipped exclusively with drivers to simulate compact source radiation, and the other endowed solely with exciters to generate modal vibrations. The exact positions of the drivers and exciters are depicted in Figure 7.2.




Figure 7.2: Geometric properties of the side surfaces and center positions of the drivers and exciters [Mit18].

The 2 large drivers [Visa] are of diameter $\varnothing = 17$ cm, while the 4 small drivers [Visb] are of diameter $\varnothing = 8$ cm. The 2 exciters used are [Day]. Consecutive nearfield measurements of the sound field were taken with a planar microphone array patch, which was moved manually between measurements to cover the whole surface of the source. The array patch consists of 12 sound intensity probes (PU Mini from Microflown Technologies [Mic]) arranged on an equidistant 3×4 grid with an inter-microphone spacing of 7.5 cm, which was positioned $z_M = 2.5$ cm above the source. The microphone spacing results in a spatial sampling limit of 2286 Hz. The measurements were performed in an anechoic room under free-field conditions. Due to the geometrical mismatch between source and array, some measurement points lie outside the source domain, which may introduce artifacts in the reconstruction process at the affected boundaries of the source. The compound measurement setup at each side surface is comprised of $12 \times 9 = 108$ spatially discrete measurement points, where sound pressure and velocity were captured simultaneously. The microphone positions are illustrated in Figure 7.3. The impulse responses between each sound source component and microphone were determined with a swept sine technique [Far00]. For source reconstruction, only the impulse responses for the sound pressure are considered.




Figure 7.3: Geometry and manual placement of microphone array patch (left) and measurement points above the surface (right) [Mit18].

The equivalent source model consists of 40×50 monopole sources, which are equally distributed over the size of the entire measurement surface of $82.5\text{cm} \times 60\text{cm}$, thus partly exceeding the actual source region. The equivalent sources are placed on the source surface with a zero retreat distance $z_0 = 0$. An appropriate estimation of their complex amplitudes allows to reconstruct the velocity distribution on the source surface, which is goal of the following section.

7.2 Resonating plate

The first target is the reconstruction of the structural vibrations on the radiating surface from the available impulse responses between source and microphone positions. Since the sound field model operates in the frequency domain, the Fourier transform of all impulse responses is required to be able to access the measured sound pressures at particular frequencies. For a preceding analysis of the surface vibrations, the averaged frequency responses of all microphone positions for the respective surface excitation by exciter Nr. 14 and Nr. 15, shown in Figure 7.4, is investigated to localize possible resonance frequencies of the plate.




Figure 7.4: Frequency response between vibrating plate and measurement plane spatially averaged over all microphone positions.

For comparison of the several regularization methods, the retrieval of the normal velocity of an exemplary vibrational mode shape at 75 Hz is conducted. For a first assessment of the mode shape, a simulation of a baffled, simply supported plywood plate of same dimensions is conducted. Due to the constructional characteristics of the box, the freely vibrating portion of the plate reduces to $59\text{cm} \times 49\text{cm}$. When using these smaller dimensions for the simulated baffled plate, the resonant behavior of the screw-mounted plate used for the measurements can be approximated. With a propagation speed on the plate of $c_L = 3000 \frac{\text{m}}{\text{s}}$ and a thickness of $h = 6 \text{ mm}$, the first resonance frequency is obtained, according to [Mös09], at around

$$f_{res} = \frac{\pi}{2} \left[\left(\frac{n_x}{l_x} \right)^2 + \left(\frac{n_y}{l_y} \right)^2 \right] \frac{c_L h}{\sqrt{12}} \quad (7.1)$$

$$= \frac{\pi}{2} \left[\left(\frac{1}{59\text{cm}} \right)^2 + \left(\frac{1}{49\text{cm}} \right)^2 \right] \frac{3000 \frac{\text{m}}{\text{s}} \cdot 6\text{mm}}{\sqrt{12}} \approx 57\text{Hz},$$

which is fairly close to the 75 Hz mode on the real plate considered for investigation. The mounted exciters on the plate additionally affect the position of the resonance frequency, which is not considered for the following simulation. The resulting mode shape for the above settings is depicted in Figure 7.5.

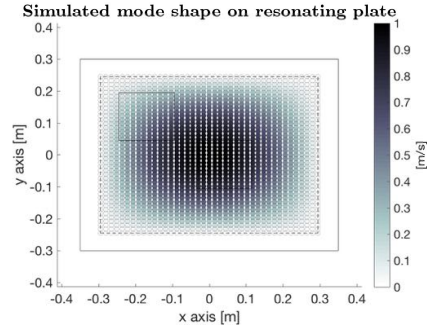


Figure 7.5: Simulation of the vibrational mode shape on a baffled, simply supported plate (dashed area) at a resonance frequency of $f_{res} = 57$ Hz.

From this point of reference, the various regularizers are now used to reconstruct the mode shape normal velocity on the plate from real pressure measurements. Since for real-world data, the SNR due to ambient noise and selfnoise of the microphones is not known in advance, it needs to be estimated for the selection of the regularization parameter. The optimal choice for the applied dynamic range was empirically determined as 25 dB and the regularization parameter was set accordingly for all methods. Figure 7.6 depicts the normal velocity on the source surface, calculated from the different source strength estimates, for the excitation of the mode shape at 75 Hz by exciter Nr. 14.

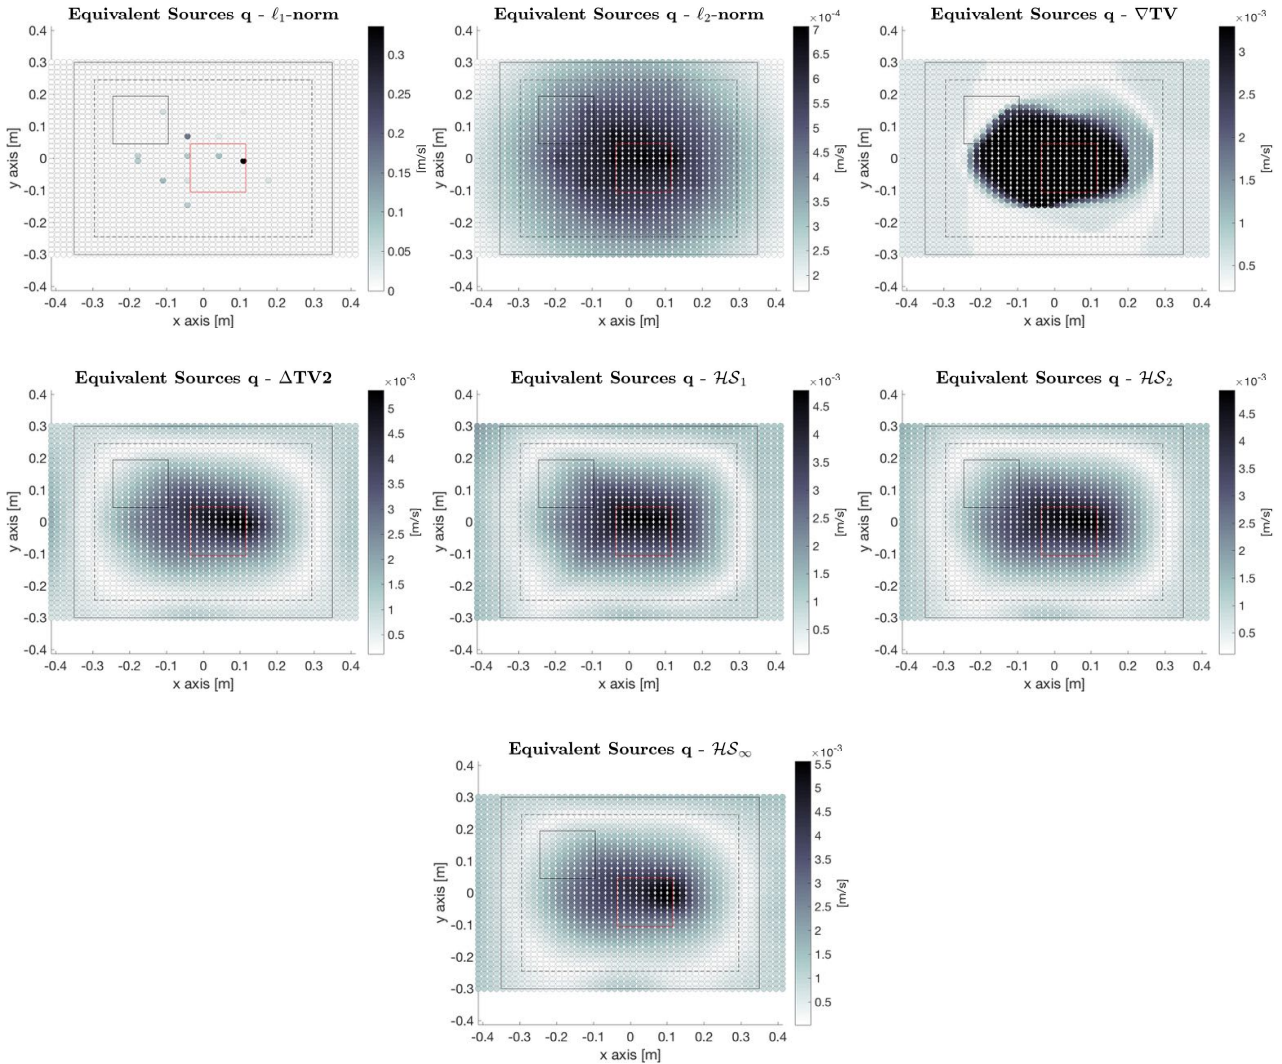


Figure 7.6: Estimated equivalent sources for the plate resonating at 75 Hz.

Since the true normal velocity distribution on the surface is unknown, it is not possible to reliably rank the accuracy of the methods, but only to make general statements regarding structural features of their estimated solutions.

As expected, the ℓ_1 -norm regularizer provides a sparse solution that cannot emulate the continuous variations of the mode shape. The peak values of the source distribution hint at the location of the active exciter, which nearly coincides with the vibrational antinode of the mode shape.

The ℓ_2 -norm regularizer yields a smooth source distribution associated with the smallest normal velocity on the plate. Coefficient activity spreads beyond the freely vibrating region of the plate, marked by the dashed rectangle, where the true normal velocity is zero. The velocity distribution gradually fades out towards the boundaries of the equivalent source model, largely overestimating the spatial extent of the source.

The ∇ TV solution displays a somewhat elliptic area of constant velocity in the center of the plate, merely hinting at the general structure of the modal pattern. The method is not able to correctly retrieve the smooth variations on the surface. The regions near the active exciter suffer from strong staircasing artifacts that unilaterally enlarge the spatial extent of the source. In the source-free region outside the box, the method yields a constant velocity elevation, likely due to the prescribed boundary conditions and measurement noise.

The second-order methods produce very similar results with only subtle differences that manifest themselves mainly in terms of degree of curvature towards the modal antinode. While the \mathcal{HS}_1 and \mathcal{HS}_2 regularizers show a more balanced velocity distribution of the pattern with lower maxima, the Δ TV2 and \mathcal{HS}_∞ exhibit higher values due to a local velocity exaltation close to the active exciter's location. All second-order solutions come close to zero along the dashed rectangle, indicating the effective vibration area of the plate, in accordance to the physically vanishing velocity of the true distribution towards the sides of the box. Unfortunately, source activity does not vanish completely, but instead rises again towards the boundaries, as the methods favor smooth variations over constant source strength. Different choices of boundary conditions did not noticeably reduce the boundary artifacts.

Fused regularizers

A potential improvement may be achieved by additionally promoting spatial sparsity in the solution with the fused regularizers. Figure 7.7 shows the estimated normal velocity distributions based on the different fused approaches with $\mu = 0.2$ for the hyperparameter in equation (4.20).




Figure 7.7: Estimated equivalent sources for the plate resonating at 75 Hz retrieved with the fused regularizers.

The additional sparsity constraint significantly reduces the boundary artifacts, but concurrently decreases the spatial extent of the mode shape. With the Sparse Variation (SV) approach, even the gradient-based ∇TV provides a smoothly varying velocity distribution while providing an entirely source-free background. The regularization parameters must be chosen with caution to not over-regularize the solution, which otherwise leads to an increasingly shrinking deformation of the modal pattern.

7.3 Loudspeaker driver layout

At last, the proposed methods are applied to reconstruct the source strength distribution of a selected driver layout on the box to examine their localization accuracy under real-world conditions. The active driver layout involves one large driver Nr. 3 and the three small drivers Nr. 10, 12 and 13. The center driver Nr. 11 of the loudspeaker triple in the corner is intentionally omitted to test the methods ability of source separation. The spatially averaged frequency responses of the active drivers, depicted in Figure 7.8, are investigated in order to predict the true level ratio between their respective pressure fields, which helps recognizing incorrect source strength estimation.

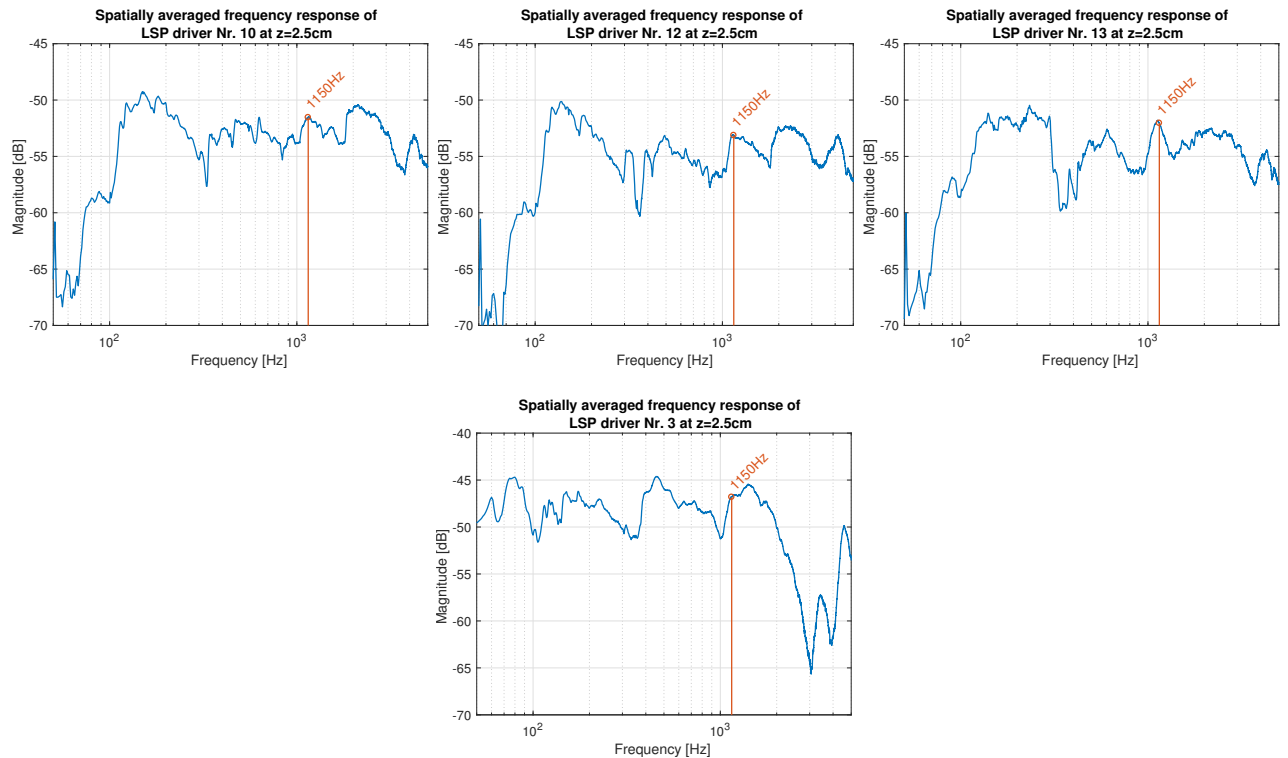


Figure 7.8: Frequency response between loudspeaker drivers Nr. 10 (upper-left), 12 (upper-center), 13 (upper-right), 3 (bottom) and measurement plane, spatially averaged over all microphone positions.

The source retrieval is conducted for an exemplary radiation frequency of 1150 Hz, where the results obtained with the various regularization schemes are presented in Figure 7.9.

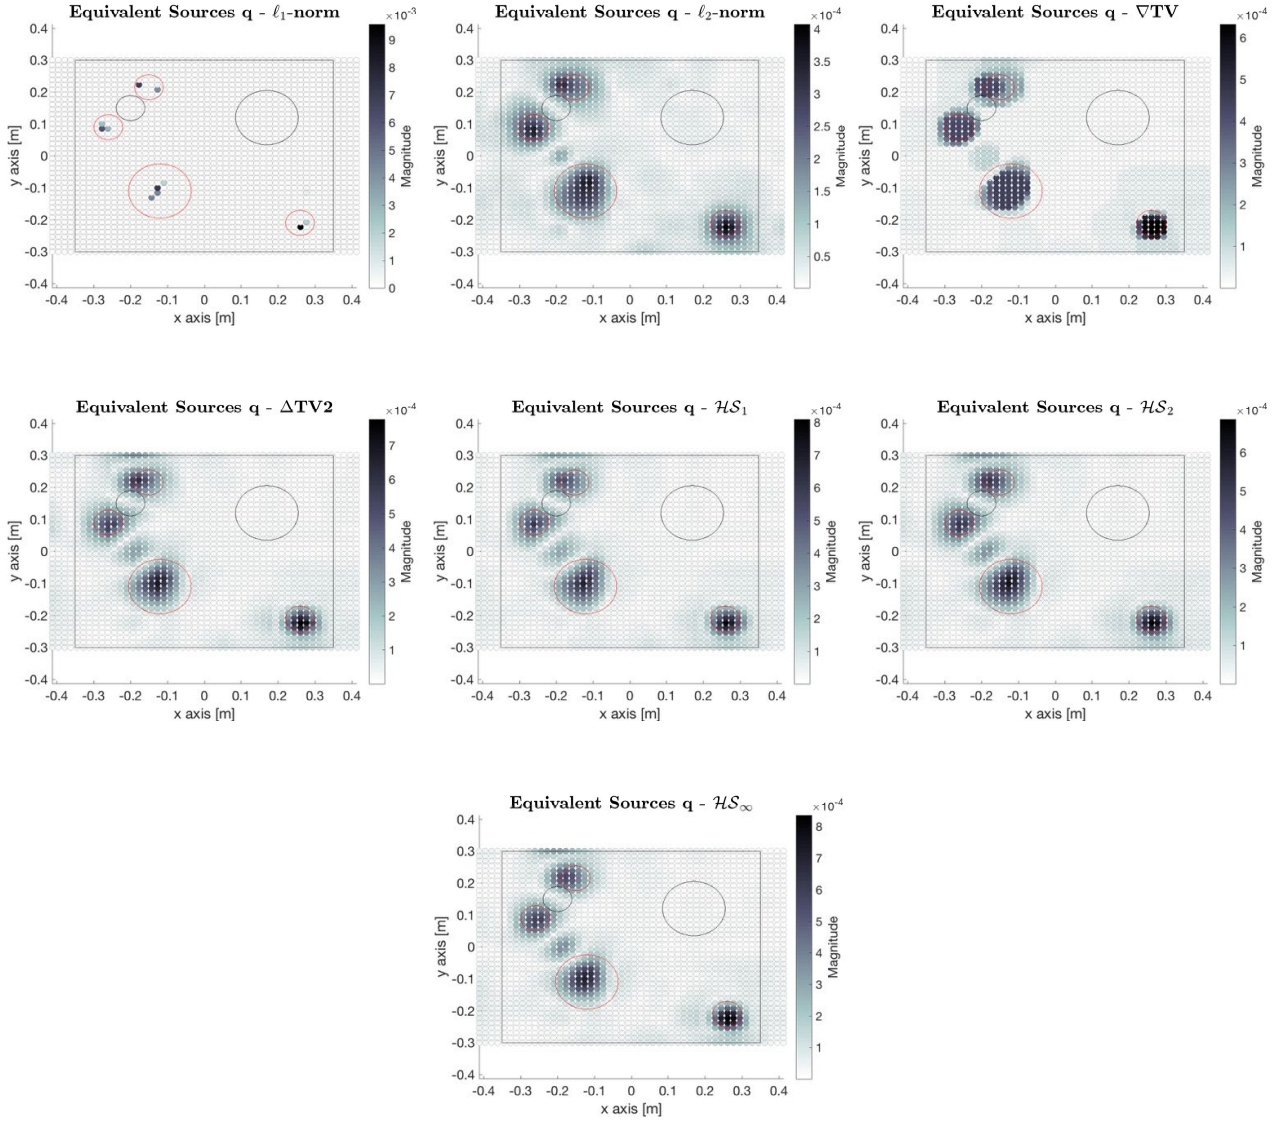


Figure 7.9: Estimated equivalent sources for the driver layout radiating at 1150 Hz.

The ℓ_1 -norm approach yields a sparse representation of the source, consisting of only a few non-zero equivalent sources, positioned at the estimated sound's origin at the center of the individual drivers, indicated by the red circles. The drivers are localized fairly precisely, although the equivalent source activities occur slightly off center. This offset may be caused by small geometric imperfections between measurement and model setup. The manual readjustment of the microphone array patch affects the positioning accuracy, which results in a small mismatch between the true measurement positions and the assumed positions on the equidistant grid. As the exact measurement positions were not accessible, the offset, affecting all methods, has to be tolerated. Moreover, changing the phase relation between the drivers, which often causes slight displacements in the reconstructed source distribution, yields no improvement and can be ruled out as a possible cause. Apart from the offset, the ℓ_1 -norm provides good localization and suppression of noisy artifacts but fails at retrieving the full spatial extent of the drivers.

The ℓ_2 -norm solution exhibits the least source strengths and yet the most activity among the source components. The size of the drivers is uniformly overestimated, thus the method fails to preserve the sharp edges around the bead of the speakers. The appearance of the characteristic sidelobes and strong spatial artifacts further deteriorate the reconstruction quality. Contrarily,

the method performs surprisingly well at separating the individual drivers. Clear separation of the small drivers in the upper corner is achieved, whereas the formation of a ghost source degrades separation of the large and small driver.

The ∇ TV results in its characteristic piecewise constant source model, where the drivers appear as blocky patches of constant source strengths, suffering from staircasing artifacts around the edges. Due to its known noise suppressing properties, the solution exhibits large regions of no source activity, although some low magnitude plateaus emerge in the surrounding of the drivers location. These spatially extended artifacts particularly arise along the boundaries of the equivalent source model as they are encouraged by the prescribed Neumann boundary conditions. Depending on the underlying source distribution, different types of boundary conditions provide the best result. Despite those artifacts, the method provides a fairly accurate localization of the drivers but it cannot reproduce the physical structure of the sound source.

The second order solutions all behave very similar with only nuanced differences between the equivalent source estimates. Compared to least squares, the solutions exhibit a significantly reduced level of sidelobes and erroneous coefficient activity within the source-free region. Sporadic artifacts mainly occur at the boundaries near the sources, attributed to the boundary conditions. The estimated source strengths are generally higher and more concentrated at the actual driver locations, thus improving localization. As the methods promote solutions with minimum curvature, the ghost source in between the drivers is strongly attenuated and prolonged. Source separation performance is still satisfactory as the individual drivers do not fuse into a single source but appear clearly distinguished. The differences between the methods primarily regard the local smoothing of the source variances. All second-order methods outperform the classical ℓ_2 -norm approach in terms of localization and spatial confinement of the drivers and they provide increased suppression of background artifacts.

Fused regularizers

To further improve the reconstruction accuracy, their sparsifying variants are considered next for the source retrieval task, for which the results are shown in 7.10. The hyperparameter controlling the spatial sparsity is chosen as $\mu = 1$.

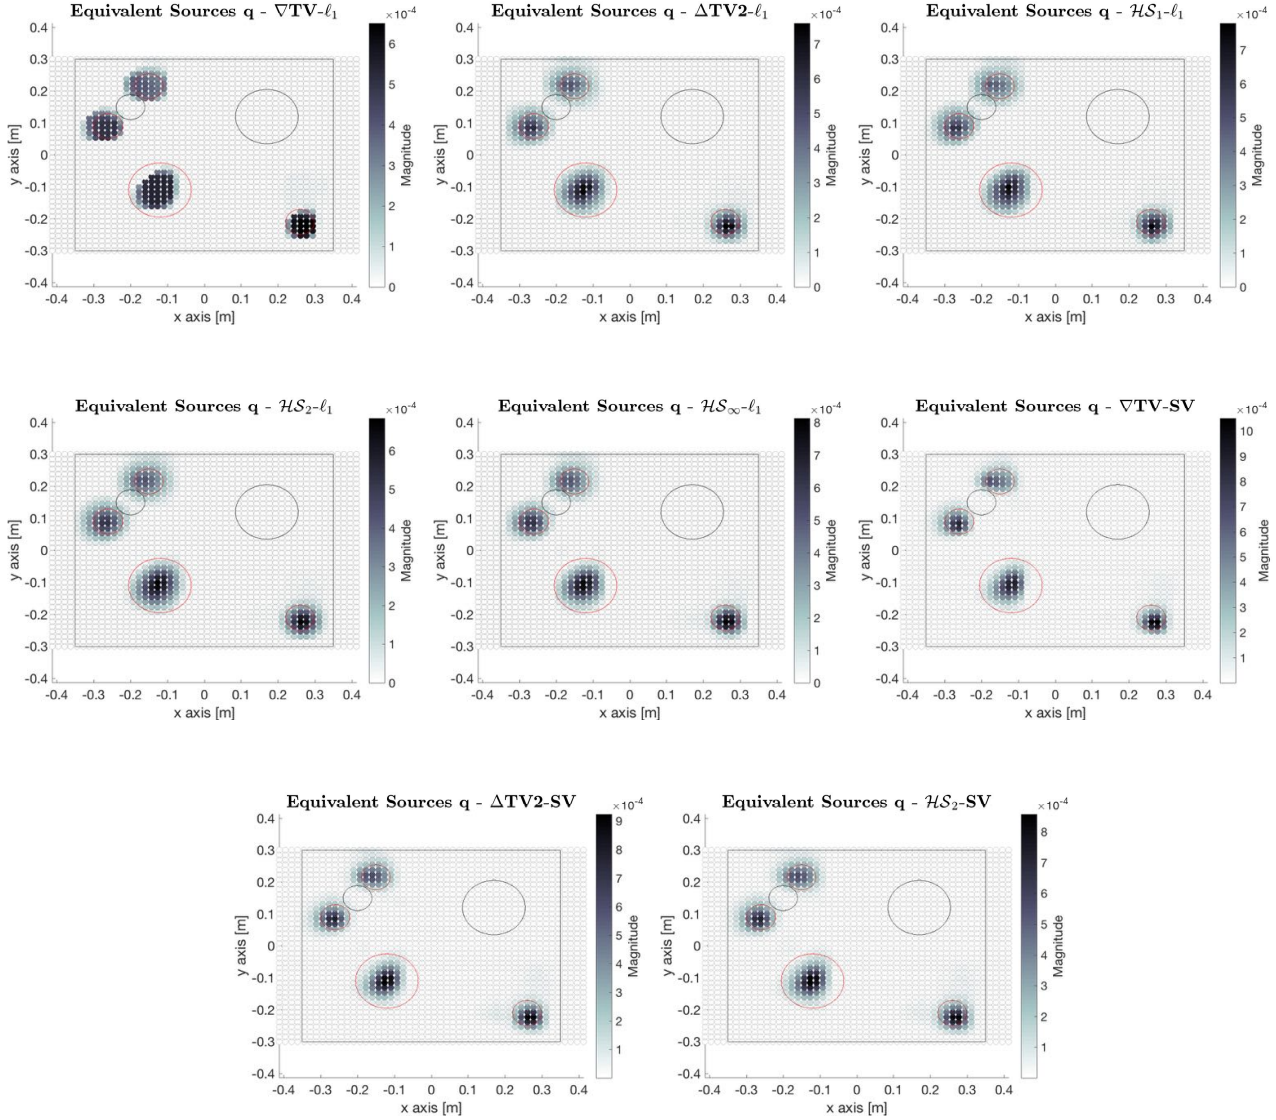


Figure 7.10: Estimated equivalent sources for the driver layout radiating at 1150 Hz retrieved with the fused regularizers.

Similar to the ℓ_1 -norm, the methods favor sparsely distributed sources, which, in contrast, are locally more extended and thus better conform to the actual dimensions of the loudspeakers. The additional promotion of spatial sparsity restricts the equivalent source activity to small areas, partly enclosing the actual driver locations while largely suppressing erroneous artifacts, retaining the source-free background. Moreover, the ghost source in between the drivers completely disappears for all methods and the spurious source activities at the borders, induced by the boundary conditions, can be effectively avoided or at least strongly attenuated. Within the regions of source activity, the characteristic smoothing properties of the original methods largely remain the same, but with a steeper curvature of the edges due to the enforcement of small to zero values in their neighborhood. The maximum source strength slightly decreases for the selected amount of regularization, thus proper adjustment of the regularization parameters

is crucial to not over-regularize the solution. Contrarily, the Sparse Variation (SV) approaches yield amplified source strength estimates, which are spatially more confined, improving localization. On the downside, the asymmetric appearance due to the use of forward finite differences is clearly noticeable, degrading the reconstruction accuracy. Fortunately, this effect can be avoided by considering central finite differences for implementation. Altogether, the fused methods, combining sparse penalization of differential and spatial information, provide the best localization performance (neglecting the error due to geometric imperfections between measurement and model) and they are superior in suppressing sidelobes and artifacts due to noise or boundary conditions. Therefore, they become the first choice when the task is to localize or retrieve the strength distribution of such isolated compact and extended sources.

Summary and Conclusion

This work showed regularization methods for nearfield acoustic holography based on sparse spatial derivatives. The presented derivative-based regularization methods could be shown to largely outperform any of the other typical strategies in the retrieval of structural information of sound sources from a limited number of measurements.

The acoustic background to wave propagation from a distributed source is described in chapter 2. The inverse formulation of the problem is then described by the equivalent source method (ESM) in chapter 3, which is used to calculate the unknown equivalent source strengths from the set of measured pressures. The inverse problem is typically underdetermined, since the number of equivalent sources usually exceeds the number of measurement points. Yet a unique solution can still be obtained in a least-squares manner, which satisfies an additional optimality criterion of minimum energy. However, this solution is not stable with respect to measurement errors, i.e. noise and a lack of evanescent waves. All this contributes to the inherent ill-posedness of the inverse problem, which is analyzed in more detail in this chapter.

Regularization is used to provide numerical stability. It typically requires a reformulation of the ill-posed problem into a convex optimization task, with the goal of finding a solution that minimizes the mismatch between model and measurements while simultaneously minimizing some norm functional, called the regularizer. The selection of an appropriate regularizer is key to obtain a valid solution. The most common regularizer is the weighted squared ℓ_2 -norm, which promotes smooth, minimum-energy solutions. This is suitable for the reconstruction of smoothly varying sources of a certain spatial extent, but it fails at retrieving compact and isolated sources, as it smears the solution over a large number of coefficients. For this purpose, ℓ_1 -norm minimization works best, which promotes sparsity in the solution, meaning that it mostly consists of zeros with only few non-zero coefficients.

To promote more involved solution shapes, the penalization of the spatial derivatives is a promising choice as they characterize evolution and curvature of the solution. In particular, the enforcement of sparse first- or second-order derivatives, resulting in piecewise-constant or piecewise-linear solutions, can effectively model common source structures. This is validated in chapter 4 by one-dimensional simulation results. An intuitive introduction to the use of first- and second-order derivatives for the shaping of two-dimensional source distributions is presented to provide a clear understanding of the derivative-based regularizers introduced thereafter. The most popular derivative-based regularizers, i.e. the first- and second-order Total Variation seminorms, are presented and their characteristics are discussed. They restrict the local changes or curvature in the solution by directly penalizing the univariate directional derivatives. Another, more recent approach of penalizing second-order derivatives is built on the Schatten-norms of the Hessian matrix, which also brings the mixed partial derivatives into play. These Schatten-norm regularizers aim at sparsifying measures based on the local singular values of the Hessian. This indirect way of penalizing derivatives allows different smoothing preferences that are potentially beneficial to the reconstruction accuracy.

Choosing a reasonable amount of regularization is not trivial as the optimal parameter changes with the noise level and measurement setup. The regularization parameter controls the weight between the model mismatch and the regularizer. In this regard, a parameter selection method

based on the estimated SNR is examined, which provides a good ad-hoc choice that can be tuned for further improvement. Moreover, the effect of different parameter choices and regularizers on the solution and its derivatives is investigated, which highlights the importance of a reasonable choice. For the discrete implementation of the derivatives, they are approximated by finite differences. Three types of first- and second-order finite differences and their associated approximation errors are discussed. Furthermore, the effect of different boundary conditions (i.e. zero, periodic and mirrored) is examined, which are necessary to enable the computation of derivatives at the limits of the finite solution domain.

Chapter 5 describes convex optimization problems associated with regularized source reconstruction, and presents the proximal splitting algorithm used to solve the problem. It is decomposed into simpler subproblems that only exploit first-order information of the objective function. The algorithm is primal-dual, hence it utilizes an representation of the original problem in another domain, which is solved simultaneously and provides lower bounds to the original problem. The repeated mapping of the solution between both domains uses differential operators and their adjoints. The minimization of Schatten-norm or sparse regularizers is achieved by evaluation of their proximity operator, which reduces to a simple vector projection onto vector-norm balls. Simple closed-form expressions for all norm ball projections encountered in this work are presented, which allow an efficient implementation.

Chapter 6 considers the most generic types of acoustic sources to test the proposed regularizers by numerical simulations. In the case of structural vibrations, the second-order derivative-based methods consistently outperform the traditional methods. Their demand of a solution with minimum curvature, rather than minimum energy, seems more suitable for the reconstruction of continuously varying source distributions. Only the first-order Total Variation fails to capture the details of the source, as it restricts the variations in the modal pattern, which lowers the resolution. However, this property can be exploited to accurately retrieve sharp-edged, homogeneous source distributions. When using the source estimates to reconstruct the sound field near extended sources, the derivative-based approaches either achieve equal or lower reconstruction errors over the whole frequency range of interest. Moreover, they overcome the spatial sampling limitations of conventional approaches that suffer from aliasing errors, enabling wideband reconstruction within tolerable errors. When considering the reconstruction of sparsely distributed point sources and their sound fields, the derivative-based methods fail to provide satisfactory results, as they always involve multiple active coefficients to model the source.

To improve the reconstruction of spatially compact sources, a fused approach works surprisingly well. It extends the derivative-based regularizers by an additional sparsity constraint in the spatial domain. The simultaneous enforcement of spatial and differential sparsity yields block-sparse solutions, where the few non-zero coefficients appear in groups rather than randomly distributed throughout the solution vector. It suppresses the side lobes and reduces the width of the main lobes, which also increase in magnitude. The enhanced spatial resolution yields a more precise localization of compact sources, to which it is not limited. It can be used for spatially extended sources as well, where it effectively suppresses erroneous artifacts in source-free regions. The additional spatial sparsity constraint also enhances the methods' robustness to noise. The fused methods are superior in almost all simulated cases.

The last chapter 7 regards the experimental validation of the proposed methods to examine their performance under real-world conditions. The source reconstruction concerns a structural mode and a dispersed loudspeaker driver layout, respectively representing a spatially extended and locally concentrated source distribution. In both cases, the fused derivative-based regularizers clearly outperform the traditional methods in terms of localization, source-separation and the suppression of noise, side lobes and other artifacts. Their many advantages and superior performance make them appealing for regularization of inverse problems in practical applications of nearfield acoustic holography.

Bibliography

- [BB98] M. Bertero and P. Boccacci, *Introduction to Inverse Problems in Imaging*. IOP Publishing, 1998.
- [BCJH11] M. Bai, C.-C. Chen, and L. Jia Hong, “On optimal retreat distance of equivalent source method-based near-field acoustical holography,” *The Journal of the Acoustical Society of America*, vol. 129, pp. 1407–16, 03 2011.
- [Bec17] A. Beck, *First-Order Methods in Optimization*. SIAM - Society for Industrial and Applied Mathematics, 2017.
- [BFK91] J. B. Fahnlne and G. Koopmann, “A numerical solution for the general radiation problem based on the combined methods of superposition and singular-value decomposition,” *Journal of The Acoustical Society of America - J ACOUST SOC AMER*, vol. 90, pp. 2808–2819, 11 1991.
- [BIB13] M. R. Bai, J.-G. Ih, and J. Benesty, *Acoustic Array Systems*. John Wiley & Sons, 2013.
- [BKP10] K. Bredies, K. Kunisch, and T. Pock, “Total generalized variation,” *SIAM J. Img. Sci.*, vol. 3, no. 3, pp. 492–526, Sep. 2010. [Online]. Available: <http://dx.doi.org/10.1137/090769521>
- [BV09] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2009.
- [Cap04] S. S. Capizzano, “A note on antireflective boundary conditions and fast deblurring models,” *SIAM J. Scientific Computing*, vol. 25, pp. 1307–1325, 2004.
- [CCMP14] L. Chaari, P. Ciuciu, S. Meriaux, and J.-C. Pesquet, “Spatio-temporal wavelet regularization for parallel mri reconstruction: application to functional mri,” *Magma (New York, N.Y.)*, vol. 27, 03 2014.
- [CCN⁺10] A. Chambolle, V. Caselles, M. Novaga, D. Cremers, and T. Pock, “An introduction to total variation for image analysis,” *Theoretical Foundations and Numerical Methods for Sparse Recovery, De Gruyter, Radon Series Comp. Appl. Math.*, vol. 9, 01 2010.
- [Con13] L. Condat, “A Primal-Dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms,” *Journal of Optimization Theory and Applications*, vol. 158, 08 2013.
- [Con87] L. Condat, “Conjugate duality and optimization,” *IEEE Signal Processing Letters*, vol. 21, no. 8, pp. 985–989, Aug 2087.
- [CP09] P. L. Combettes and J.-C. Pesquet, “Proximal Splitting Methods in Signal Processing,” *arXiv e-prints*, p. arXiv:0912.3522, Dec 2009.
- [Day] Dayton Audio, “TT25-8.” [Online]. Available: <http://www.daytonaudio.com/index.php/tt25-8-puck-tactile-transducer-mini-bass-shaker.html>
- [DMWL85] J. D. Maynard, E. Williams, and Y. Lee, “Nearfield acoustic holography: I. theory of generalized holography and the development of nah,” *Journal of the Acoustical Society of America*, vol. 78, 11 1985.

- [EAS16] V. V. Estrela, H. Aguiar Magalhaes, and O. Saotome, “Total Variation Applications in Computer Vision,” *arXiv e-prints*, p. arXiv:1603.09599, Mar 2016.
- [EDTV15] M. Eickenberg, E. Dohmatob, B. Thirion, and G. Varoquaux, “Total variation meets sparsity: statistical learning with segmenting penalties,” vol. 9349, 07 2015.
- [Ela10] M. Elad, *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*, 1st ed. Springer, 2010.
- [Fah05] F. Fahy, *Foundations of Engineering Acoustics*. Elsevier Academic Press, 2005.
- [Far00] A. Farina, “Simultaneous measurement of impulse response and distortion with a swept-sine technique,” *108th AES Convention, Paris, France*, 2000.
- [FD16] E. Fernandez Grande and L. Daudet, “Near-field acoustic imaging based on laplacian sparsity,” in *Proceedings of the 22nd International Congress on Acoustics*, 2016.
- [FGD18] E. Fernandez-Grande and L. Daudet, “Compressive acoustic holography with block-sparse regularization,” *The Journal of the Acoustical Society of America*, vol. 143, pp. 3737–3746, 06 2018.
- [FGXG17] E. Fernandez-Grande, A. Xenaki, and P. Gerstoft, “A sparse equivalent source method for near-field acoustic holography,” *The Journal of the Acoustical Society of America*, vol. 141, pp. 532–542, 01 2017.
- [FR06] M. Fornasier and H. Rauhut, “Recovery algorithms for vector-valued data with joint sparsity constraints,” *SIAM Journal on Numerical Analysis*, vol. 46, 09 2006.
- [GB14] M. Grant and S. Boyd, “CVX: Matlab software for disciplined convex programming, version 2.1,” <http://cvxr.com/cvx>, Mar. 2014.
- [GGC16] H. Guven, A. Gungor, and M. Cetin, “An augmented lagrangian method for complex-valued compressed sar imaging,” *IEEE Transactions on Computational Imaging*, vol. 2, pp. 1–1, 09 2016.
- [GHO99] G. H. Golub, P. C. Hansen, and D. P. O’Leary, “Tikhonov regularization and total least squares,” *SIAM J. Matrix Anal. Appl.*, vol. 21, no. 1, pp. 185–194, Oct. 1999. [Online]. Available: <http://dx.doi.org/10.1137/S0895479897326432>
- [Gom08] J. Gomes, “A study on regularization parameter choice in near-field acoustical holography,” *The Journal of the Acoustical Society of America*, vol. 123, p. 3385, 06 2008.
- [GXM15] P. Gerstoft, A. Xenaki, and C. F. Mecklenbrauker, “Multiple and single snapshot compressive beamforming,” *Acoustical Society of America Journal*, vol. 138, pp. 1736–1736, Sep 2015.
- [Hal16] J. Hald, “Fast wideband acoustical holography,” *The Journal of the Acoustical Society of America*, pp. 1508–1517, 04 2006.
- [Hal09] —, “Basic theory and properties of statistically optimized near-field acoustical holography,” *The Journal of the Acoustical Society of America*, vol. 125, pp. 2105–20, 05 2009.
- [Hal18] —, “A comparison of iterative sparse equivalent source methods for near-field acoustical holography,” *The Journal of the Acoustical Society of America*, vol. 143, pp. 3758–3769, 06 2018.

- [HJ12] Y. Hu and M. Jacob, “Higher degree total variation (hdtv) regularization for image recovery,” *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, vol. 21, pp. 2559–71, 01 2012.
- [HORJ14] Y. Hu, G. Ongie, S. Ramani, and M. Jacob, “Generalized higher degree total variation (hdtv) regularization,” *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, vol. 23, 04 2014.
- [IROF92] L. I. Rudin, S. Osher, and E. Fatemi, “Nonlinear total variation based noise removal algorithms,” *Physica D: Nonlinear Phenomena*, vol. 60, pp. 259–268, 11 1992.
- [Jal14] K. Jalalzai, “Some remarks on the staircasing phenomenon in total variation-based image denoising,” *Journal of Mathematical Imaging and Vision*, vol. 54, 02 2014.
- [KI96] B.-K. Kim and J.-G. Ih, “On the reconstruction of the vibro-acoustic field over the surface enclosing an interior space using the boundary element method,” *Acoustical Society of America Journal*, vol. 100, pp. 3003–3016, 11 1996.
- [KP14] N. Komodakis and J.-C. Pesquet, “Playing with Duality: An Overview of Recent Primal-Dual Approaches for Solving Large-Scale Optimization Problems,” *arXiv e-prints*, p. arXiv:1406.5429, Jun 2014.
- [Lar16] A. Laruelo, “Integration of magnetic resonance spectroscopic imaging into the radiotherapy treatment planning,” Ph.D. dissertation, Université Paul Sabatier - Toulouse III, 05 2016.
- [LCT⁺16] A. Laruelo, L. Chaari, J.-Y. Tourneret, H. Batatia, S. Ken, B. Rowland, R. Ferland, and A. Laprie, “Spatio-spectral regularization to improve magnetic resonance spectroscopic imaging quantification,” *NMR in biomedicine*, vol. 29, 05 2016.
- [Lee16] S. Lee, “Review: The use of equivalent source method in computational acoustics,” *Journal of Computational Acoustics*, p. 1630001, 12 2016.
- [LeV07] R. J. LeVeque, *Finite Difference Methods for Ordinary and Partial Differential Equations*. SIAM - Society for Industrial and Applied Mathematics, 2007.
- [LU13] S. Lefkimmiatis and M. Unser, “Poisson image reconstruction with hessian schatten-norm regularization,” *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, vol. 22, 07 2013.
- [LWU13] S. Lefkimmiatis, J. Ward, and M. Unser, “Hessian schatten-norm regularization for linear inverse problems,” *Image Processing, IEEE Transactions on*, vol. 22, pp. 1873–1888, 05 2013.
- [MAYGB11] A. Maleki, L. Anitori, Z. Yang, and R. G. Baraniuk, “Asymptotic analysis of complex lasso via complex approximate message passing (camp),” *IEEE Transactions on Information Theory*, vol. 59, 08 2011.
- [Mic] Microflown Technologies, “PU MINI.” [Online]. Available: <http://www.microflown.com/products/standard-probes/pu-mini.html>
- [Mit18] R. Mittmannsgruber, “Leistungs-basierte analyse von luftschallübertragungsanteilen,” Master’s thesis, University of Music and Performing Arts Graz - Institute of Electronic Music and Acoustics, 2018.
- [Mor65] J. J. Moreau, “Proximité et dualité dans un espace hilbertien,” *Bulletin de la Société Mathématique de France*, vol. 93, pp. 273–299, 1965.

- [Mös09] M. Möser, *Technische Akustik*. Springer, 2009, vol. 8.
- [NK01] P. Nelson and Y. Kahana, “Spherical harmonics, singular-value decomposition and the head-related transfer function,” *Journal of Sound and Vibration*, vol. 239, pp. 607–637, 01 2001.
- [NPCK17] H.-N. Nguyen, V. Paveau, C. Cauchois, and C. Kervrann, “Generalized Sparse Variation Regularization for Large Fluorescence Image Deconvolution,” Oct. 2017, working paper or preprint. [Online]. Available: <https://hal.inria.fr/hal-01609810>
- [Pan14] N. Panigrahi, *Computing in geographic information systems*. CRC Press, 2014.
- [Pap15] K. Papafitsoros, “Novel higher order regularisation methods for image reconstruction,” Ph.D. dissertation, University of Cambridge, 2015.
- [PB14] N. Parikh and S. Boyd, “Proximal algorithms,” *Found. Trends Optim.*, vol. 1, no. 3, pp. 127–239, Jan. 2014. [Online]. Available: <http://dx.doi.org/10.1561/2400000003>
- [Per13] A. Pereira, “Acoustic imaging in enclosed spaces,” PhD Thesis, 2013.
- [PG18] F. Pedregosa and G. Gidel, “Adaptive Three Operator Splitting,” *arXiv e-prints*, p. arXiv:1804.02339, Apr 2018.
- [PSS13] K. Papafitsoros, C. Schoenlieb, and B. Sengul, “Combined first and second order total variation inpainting using split bregman,” *IPOLE Journal*, vol. 3, pp. 112–136, 2013.
- [Roc74] R. T. Rockafellar, *Conjugate Duality and Optimization*. Society for Industrial and Applied Mathematics, 1974.
- [Ros07] T. D. Rossing, *Handbook of Acoustics*. Springer, 2007.
- [RP56] D. H. Robey and D. H. Potts, “Derivation of the free-space scalar green’s function,” *Journal of The Acoustical Society of America - J ACOUST SOC AMER*, vol. 28, p. 723, 07 1956.
- [Sar05] A. Sarkissian, “Method of superposition applied to patch near-field acoustic holography,” *Journal of The Acoustical Society of America - J ACOUST SOC AMER*, vol. 118, pp. 671–678, 08 2005.
- [Saw11] A. Sawatzky, “(nonlocal) total variation in medical imaging,” Ph.D. dissertation, Westfälische Wilhelms-Universität Münster, 07 2011.
- [SHBRCH03] A. Schuhmacher, J. Hald, K. Bo Rasmussen, and P. Christian Hansen, “Sound source reconstruction using inverse boundary element calculations,” *The Journal of the Acoustical Society of America*, vol. 113, pp. 114–27, 02 2003.
- [STS⁺16] F. Soulez, É. Thiébaud, A. Schutz, A. Ferrari, F. Courbin, and M. Unser, “Proximity operators for phase retrieval,” *Applied Optics*, vol. 55, pp. 7412–7421, 09 2016.
- [Visa] Visaton, “FRS 8.” [Online]. Available: <http://www.visaton.de/de/produkte/chassiszubehoer/breitband-systeme/frs-8-8-ohm>
- [Visb] —, “WS 17 E.” [Online]. Available: <http://www.visaton.de/de/produkte/chassiszubehoer/tieftoener/ws-17-e-4-ohm>
- [Vu13] B. Vu, “A splitting algorithm for dual monotone inclusions involving cocoercive operators,” *Advances in Computational Mathematics*, vol. 38, 04 2013.

- [VW04] N. Valdivia and E. Williams, “Implicit methods of solution to integral formulations in boundary element method based nearfield acoustic holography,” *Acoustical Society of America Journal*, vol. 116, pp. 1559–1572, 09 2004.
- [Wil99] E. G. Williams, *Fourier Acoustics*. Academic Press, 1999.
- [WW97] Z. Wang and S. Wu, “Helmholtz equation-least-squares method for reconstructing acoustic pressure fields,” *Journal of The Acoustical Society of America - J ACOUST SOC AMER*, vol. 102, 10 1997.
- [XFGG16] A. Xenaki, E. Fernandez-Grande, and P. Gerstoft, “Block-sparse beamforming for spatially extended sources in a bayesian formulation,” *The Journal of the Acoustical Society of America*, vol. 140, pp. 1828–1838, 09 2016.