Vom Fachbereich Mathematik der Technischen Universität Kaiserslautern
zur Verleihung des akademischen Grades
Doktor der Naturwissenschaften (Doctor rerum naturalium, Dr. rer. nat.)
genehmigte

Dissertation

Convex Analysis for Processing Hyperspectral Images and Data from Hadamard Spaces

Martin Julian Montag

Gutachter: Prof. Dr. Gabriele Steidl
Prof. Dr. Jesús Angulo

Datum der Disputation:
20. Februar 2017

D 386

Promotionskommission

Vorsitzender:	Prof. Dr. René Pinnau, TU Kaiserslautern
Erstgutachterin:	Prof. Dr. Gabriele Steidl, TU Kaiserslautern
Weitere Prüferin:	Prof. Dr. Claudia Redenbach, TU Kaiserslautern

Acknowledgment

I would like to thank

my academic supervisor **Gabriele Steidl** for her guidance and constant support during the last three years, for taking interest in hyerspectral imaging and for her detailed feedback to a beginning scientific writer,

Henrike Stephani, my supervisor at the Fraunhofer ITWM, for collaboration and support and for allowing me to work in the fascinating field of hyperspectral imaging and to take part in the Hypermath BMBF meetings,

my **colleges from the Imaging Department at Fraunhofer ITWM** for an atmosphere which is very friendly, helpful and welcoming to new PhD students, in particular **Ronald Rösch**, who led the department when I started, and has certainly contributed to this atmosphere, **Kai Teubner** and **Franz Schreiber** from the imaging lab for quick and practical solutions and advice, **Ali Moghiseh** and **Martin Braun** for programming support in setting up the development environment,

all **fellow PhD students** at the Fraunhofer ITWM for discussions and for enriching my time here, in particular **Prakash Easwaran** and my office mates **Martina Sormani** and **Michael Arnold**,

Markus Rauhut for generously supporting the participation of PhD students in inhouse seminars and scientific conferences as head of the department,

Bettina Heise from the Johannes Kepler Universität Linz, Austria, for inviting me to a research visit and generously offering her time to make my stay a very enriching experience,

Miroslav Bačák for collaboration on the results in Section 6.3.3, for his excellent introductory book on Hadamard spaces and for interesting suggestions on further work,

all colleagues at the Felix Klein Center, in particular our secretary **Kirsten Höffler**, my former colleagues from the imaging group, **Sören Häuser**, **René Ciak**, **Behrang Shafei**, **Stanislav Harizanov**, **Xiaohao Cai**, and current colleagues **Ronny Bergmann**, **Friederike Laus**, **Jan Henrik Fitschen**, **Max Nimmer**, **Johannes Persch**, for the pleasant working atmosphere, many discussions and stimulating exchange,

for partial proofreading: **Johannes Persch**, **Friederike Laus**, **Rasmus Schroeder**, **Ronald Rösch**, **Henrike Stephani**,

my **friends** for support and enriching my time here and all whom I have forgotten to mention here,

last, but definitely not least, my **parents** for their constant support for so many years and for always being there for me – without your encouragement this thesis would not have been possible.

Abstract

This thesis brings together convex analysis and hyperspectral image processing. Convex analysis is the study of convex functions and their properties. Convex functions are important because they admit minimization by efficient algorithms and the solution of many optimization problems can be formulated as minimization of a convex objective function, extending much beyond the classical image restoration problems of denoising, deblurring and inpainting.

At the heart of convex analysis is the duality mapping induced within the class of convex functions by the Fenchel transform. In the last decades efficient optimization algorithms have been developed based on the Fenchel transform and the concept of infimal convolution.

The infimal convolution is of similar importance in convex analysis as the convolution in classical analysis. In particular, the infimal convolution with scaled parabolas gives rise to the one parameter family of Moreau-Yosida envelopes, which approximate a given function from below while preserving its minimum value and minimizers. The closely related proximal mapping replaces the gradient step in a recently developed class of efficient first-order iterative minimization algorithms for non-differentiable functions. For a finite convex function, the proximal mapping coincides with a gradient step of its Moreau-Yosida envelope. Efficient algorithms are needed in hyperspectral image processing, where several hundred intensity values measured in each spatial point give rise to large data volumes.

In the **first part** of this thesis, we are concerned with models and algorithms for hyperspectral unmixing. As part of this thesis a hyperspectral imaging system was taken into operation at the Fraunhofer ITWM Kaiserslautern to evaluate the developed algorithms on real data. Motivated by missing-pixel defects common in current hyperspectral imaging systems, we propose a total variation regularized unmixing model for incomplete and noisy data for the case when pure spectra are given. We minimize the proposed model by a primal-dual algorithm based on the proximum mapping and the Fenchel transform. To solve the unmixing problem when only a library of pure spectra is provided, we study a modification which includes a sparsity regularizer into model.

We end the first part with the convergence analysis for a multiplicative algorithm derived by optimization transfer. The proposed algorithm extends well-known multiplicative update rules for minimizing the Kullback-Leibler divergence, to solve a hyperspectral unmixing model in the case when no prior knowledge of pure spectra is given.

In the **second part** of this thesis, we study the properties of Moreau-Yosida envelopes, first for functions defined on Hadamard manifolds, which are (possibly) infinite-dimensional Riemannian manifolds with negative curvature, and then for functions defined on Hadamard spaces.

In particular we extend to infinite-dimensional Riemannian manifolds an expression for the gradient of the Moreau-Yosida envelope in terms of the proximal mapping. With the help of this expression we show that a sequence of functions converges to a given limit function in the sense of Mosco if the corresponding Moreau-Yosida envelopes converge pointwise at all scales.

Finally we extend this result to the more general setting of Hadamard spaces. As the reverse implication is already known, this unites two definitions of Mosco convergence on Hadamard spaces, which have both been used in the literature, and whose equivalence has not yet been known.

Zusammenfassung

Diese Arbeit vereint konvexe Analysis und hyperspektrale Bildverarbeitung. Konvexe Analysis untersucht die Eigenschaften konvexer Funktionen. Konvexe Funktionen besitzen zentrale Bedeutung in der Optimierung: Unter praktischen Gesichtspunkten lassen sie sich effizient minimieren und beschreiben gleichzeitig eine Vielzahl realer Problemstellungen, weit über die klassischen Bildverarbeitungsaufgaben des Entrauschens, Schärfens und Wiederherstellens hinaus.

In gleichem Maße wie die Regularisierungstheorie partieller Differentialgleichungen aus dem Studium des Spezialfalles elliptischer Gleichungen hervorgegangen ist, so entwickeln sich eine Vielzahl von Minimierungsmethoden für nicht-konvexe Probleme ausgehend von den Methoden der konvexen Analysis.

Im Mittelpunkt der konvexen Analysis steht die Dualität, welche durch die Fenchel-Konjugation auf den konvexen Funktionen induziert wird. Aufbauend auf der Dualitätsabbildung und der infimalen Faltung mit quadratischen Funktionen, ist in den letzten Jahren die Klasse der proximalen Algorithmen entstanden. Proximale Algorithmen erlauben die effektive iterative Minimierung nicht-differenzierbarer Funktionen, die in der Bildverarbeitung durch räumliche TV-Regularisierung entstehen.

Die infimale Faltung mit skalierten quadratischen Funktionen erzeugt die einparametrische Familie der Moreau-Yosida-Regularisierungen, welche eine gegebene Funktion von unten approximieren und dabei Minimum und Minimierer beibehalten. Für reellwertige Funktionen stimmt die Proximumsabbildung mit dem Gradientenschritt auf der Moreau-Yosida-Regularisierung überein.

Effiziente Algorithmen sind in der hyperspektralen Bildverarbeitung unabdingbar, da mehrere hundert spektrale Intensitätswerte in jedem Bildpunkt zu großen Datenmengen führen.

Im **ersten Teil** dieser Arbeit ist es unser Anliegen, effiziente Algorithmen für hyperspektrale Entmischung zu entwickeln. Um diese auf realen Messungen zu testen, wurde im Rahmen dieser Arbeit eine hyperspektrale Nahinfrarotkamera am Fraunhofer ITWM in Betrieb genommen. Motiviert durch fehlende Pixel in gängigen Kamerasensoren präsentieren wir ein Entmischungsmodel für unvollständige und verrauschte Daten, wenn die Reinmaterialspektren bekannt sind. Wir minimieren dieses Modell mit einem primal-dualen Algorithmus, welcher als Teilschritt die Proximumsabbildung verwendet.

Um das Entmischungsproblem zu lösen, wenn die Reinmaterialspektren nicht direkt bekannt, sondern zwischen weiteren Spektren einer Spektrendatenbank versteckt sind, untersuchen wir eine Abwandlung des Modells durch einen additiven Sparsity-Term.

Wir beenden den ersten Teil mit der Konvergenzanalyse für einen alternierenden multiplikativen Algorithmus. Dieser erweitert bekannte multiplikative iterative Verfahren

zur Minimierung der Kullback-Leibler-Distanz, um das Entmischungsproblem ohne Vorabinformationen über die vorkommenden Spektren zu lösen.

Im **zweiten Teil** dieser Arbeit untersuchen wir die Eigenschaften von Moreau-Yosida-Regularisierungen, zuerst für Funktionen, die auf Hadamard-Mannigfaltigkeiten definiert sind, d.h. möglicherweise unendlichdimensionalen riemannschen Mannigfaltigkeiten mit negativer Krümmung, und anschließend für Funktionen, die auf Hadamard-Räumen definiert sind.

Insbesondere erweitern wir den klassischen Ausdruck für den Gradienten der Moreau-Yosida-Regularisierung, in Abhängigkeit von der Proximumsabbildung, auf Hadamard-Mannigfaltigkeiten. Mithilfe dieses Ausdruckes zeigen wir, dass eine Folge von Funktionen genau dann im Sinne von Mosco gegen eine Grenzfunktion konvergiert, wenn die zugehörigen Moreau-Yosida-Regularisierungen aller Skalen punktweise konvergieren.

Dieses Resultat erweitern wir schließlich auf allgemeine Hadamard-Räume. Da die umgekehrte Implikation bereit bekannt ist, vereint dies zwei Definitionen der Mosco-Konvergenz auf Hadamard-Räumen, welche beide in der Literatur verwendet werden, ohne dass ihre Äquivalenz bislang bekannt war.

Contents

List of Publications

During the work on this thesis the following articles and preprints have been published. The results of Chapter 2 have been published in [81]. The results on quotient functionals are not included in this thesis and can be found in the publication [80]. The results from Section 6.3.3 have been submitted for publication [15]. The numbers of the respective articles are the same as in the full bibliography at the end of the thesis.

[15] M. Bačák, M. J. Montag and G. Steidl. Convergence of functions and their Moreau-Yosida envelopes on Hadamard spaces. *Preprint arXiv: 1604.08047*, 2016.

[80] G. Moerkotte, M. J. Montag, A. Repetti and G. Steidl. Proximal operator of quotient functions with application to a feasibility problem in query optimization. *Journal of Computational and Applied Mathematics*, 285:243–255, 2015.

[81] M. J. Montag and H. Stephani. Hyperspectral unmixing from incomplete and noisy data. *Journal of Imaging*, 2(1):7, 2016.

Notation and Symbols

Basic Symbols

$\mathbb{N}$	Set of natural numbers without 0
$\mathbb{Z}$	Set of integers
$\mathbb{R}$	Set of real numbers
$\mathbb{R}_{\geq}$	Set of nonnegative real numbers
$\mathbb{R}_{>}$	Set of strictly positive real numbers
$\mathbb{R}_{>}^{d}$	Set of d-dimensional vectors with strictly positive real entries
$\mathbb{R}_{\geq}^{L,N}$	Set of $L \times N$ matrics whose entries are nonnegative real numbers

Vectors, Matrices and Discrete Structures

$\mathrm{vec}(A)$	Column-wise reshaping the matrix A into a vector
$\mathrm{diag}(v)$	Diagonal matrix with diagonal entries v
$\mathrm{tr}(A)$	Trace of the matrix A
$A^{\top}$	Transpose of the matrix A
I	Identity matrix, dimension is usually clear from context
$\mathbb{1}_d$	One-vector in $\mathbb{R}^d$, i.e. $\mathbb{1}_d := (1, 1, \ldots, 1)^{\top}$
$\mathbb{1}_{L,N}$	Matrix of ones with size $L \times N$
$\triangle_p$	Probability simplex defined by $\triangle_p := \{v \in \mathbb{R}^p \colon \langle v, \mathbb{1}_p \rangle = 1,\ v \geq 0\}$
$(\triangle_p)^N$	Set of $N \times p$ matrices whose column vectors belong to $\triangle_p$
$\|\cdot\|_1$	l_1-norm of a vextor
$\|\cdot\|_2$	l_2-norm of a vextor
$\|\cdot\|_F$	Frobenius norm, $\|A\|_F^2 := \mathrm{tr}(A^{\top}A) = \sum_{i,j} A_{i,j}^2$
$\|\cdot\|_{1,\infty}$	Row-sparsity norm, $\|A\|_{1,\infty} := \sum_i \max_j \lvert A_{i,j}\rvert$
$\langle a, b \rangle$	Inner product/scalar product of two vectors a, b
$\langle A, B \rangle_F$	Frobenius inner product of two matrices, $\langle A, B \rangle_F := \sum_{i,j} A_{i,j} B_{i,j}$
$\log A$	Convex extension of the natural logarithm, defined in Section 4.2
$\mathrm{KL}(a, b)$	Kullback-Leibler divergence or I-divergence between vectors a, b; for matrices A, B the sum is taken over corresponding columns
∇	Discrete gradient operator, defined on page 27
D_x, D_y	Forward difference matrices, defined on page 27
$\|\|\|\cdot\|\|\|_1$	See page 27
TV	Isotropic total variation operator, defined in Section 2.2.1
$A \circ B$	Hadamard product of matrices, defined by $(A \circ B)_{i,j} = A_{i,j} B_{i,j}$
$A \otimes B$	Kronecker product or tensor product, defined in Section 2.2.1

Hyperspectral Unmixing

Y	Data matrix, each column is one measured spectrum
K	Endmember matrix
X	Abundance matrix
$\underline{k}_1, \dots, \underline{k}_p$	Pure spectra, called endmembers, columns of the matrix K
$\underline{y}(j), \underline{x}(j)$	Columns j of the matices Y, X respectively
$x_r(j)$	Abundance, i.e. estimated mixture fraction, of the r-th material in the j-th data spectrum
$\bar{x}_r$	Row of X, containing the abundances of the r-th material
$\mathcal{N}(\mu, \sigma)$	Normal distribution with mean μ and variance σ

Functions and Topology

$\mathrm{B}_r(x)$	Open metric ball of radius r around the point x
$f \circ g$	Concatenation of functions, defined by $\big(f \circ g\big)(x) := f\left(g(x)\right)$
ι_A	Indicator function $\iota_A \colon \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$ of A defined by $\iota_A(x) := \begin{cases} 0 & \text{for } x \in A, \\ +\infty & \text{otherwise} \end{cases}$
χ_A	Characteristic function $\chi_A \colon \mathbb{R}^d \to \mathbb{R}$ of A defined by $\chi_A(x) := \begin{cases} 1 & \text{for } x \in A, \\ 0 & \text{otherwise} \end{cases}$
$\operatorname{dom} f$	Effective domain of f, defined by $\operatorname{dom} f := \{x \colon f(x) < +\infty\}$
$\operatorname{int} S$	Interior of the set S
$\overline{S}$	Metric closure of the set S
$\mathcal{P}_C$	Metric projection to the closed convex set C
$\xrightarrow{w}$	Weak convergence of points in a Banach space or Hadamard space
$\xrightarrow{\Gamma}$	Gamma convergence of functions
$\xrightarrow{M}$	Mosco convergence of functions or sets
$\frac{\mathrm{d}}{\mathrm{d}t} f$	Derivative of the univariate function f with respect to t
$\partial_r f$ or $\frac{\partial}{\partial x_r} f$	Partial derivative with respect to the r-th variable
$\operatorname{Hess}_f(x)$	Hessian matrix of a twice differentiable function f at the point x
$\nabla_K \mathcal{E}(\cdot, \cdot)$	Vector or matrix of partial derivatives corresponding to the designated block of variables
∂f	Multivalued subdifferential of the function f defined by (5.4) in Banach spaces and by (6.3) in Riemannian Manifolds
Prox_f	Proximum mapping defined by (5.1), and for Hadamard spaces in Section 5.3.1
f_λ	Moreau-Yosida envelope of f defined by (5.2) and in Section 5.3.1

Manifolds and Hadamard Spaces

Unless stated otherwise the following symbols are introduced in Section 6.1.2.

$(\mathcal{H}, \mathrm{d})$	Finite or infinite-dimensional Hadamard manifold or Hadamard space with distance d
$T_x\mathcal{H}$	Tangent space to $\mathcal{H}$ at the point x
$\langle \cdot, \cdot \rangle_x$	Riemannian metric on $T_x\mathcal{H}$
∇f	Riemannian gradient
$\log_x y$	Riemannian logarithm
$\exp_x v$	Riemannian exponential map
$\gamma_{\widehat{x,y}}$	Geodesic from x to y, defined in Section 5.3.1
$\mathcal{T}_{xy}\xi$	Parallel transport of the tangent vector $\xi \in T_x\mathcal{H}$ to $T_y\mathcal{H}$
$(\mathbb{M}_\kappa, \mathrm{d}_\kappa)$	Model space with constant curvature κ, defined in Section 6.1.4
$\measuredangle ACB$	Euclidean angle; in Hadamard spaces the angle between the geodesics $\gamma_{\widehat{C,A}}$ and $\gamma_{\widehat{C,B}}$ at the point C, defined by (6.8)
$\measuredangle_\kappa \overline{ACB}$	Same as $\measuredangle \overline{ACB}$ for points $\overline{A}, \overline{B}, \overline{C}$ in $\mathbb{M}_\kappa$
$\measuredangle_\kappa(A, C, B)$	Comparison angle in $\mathbb{M}_\kappa$, for points A, B, C in a metric space, defined on page 104

Introduction

This thesis consists of two parts, a first applied part on hyperspectral image processing and a second part with more fundamental results on proximal mappings and Moreau-Yosida envelopes in Hadamard spaces.

The joining element of both parts is the proximal mapping. Proximal mappings and Moreau-Yosida envelopes of convex functions play a central role in convex analysis. In particular, they appeared in various minimization algorithms which have recently found application in image processing and machine learning. For overviews, see for instance [16, 28, 87].

For a convex lower-semicontiuous function f from a Hilbert space H into $\mathbb{R} \cup \{+\infty\}$ the *proximal mapping* $\operatorname{Prox}_{\lambda f} \colon H \to H$ is defined by

$$\operatorname{Prox}_{\lambda f}(x) = \operatorname*{argmin}_{y \in H} \Big\{ f(y) + \tfrac{1}{2\lambda} \operatorname{d}(x,y)^2 \Big\}, \tag{0.1}$$

where $\operatorname{d}(x,y) := \|x - y\|$. For $\lambda > 0$, the *Moreau-Yosida envelope* of f is given by

$$f_\lambda(x) := \min_{y \in H} \Big\{ f(y) + \tfrac{1}{2\lambda} \operatorname{d}(x,y)^2 \Big\}.$$

The one-parameter family of Moreau-Yosida envelopes approximates f from below and satisfies

$$\operatorname{Prox}_{\lambda f}(x) = x - \lambda \nabla f_\lambda(x).$$

In this sense the proximal mapping is the next iterate after a gradient descent step on the regularized version f_λ of f. An efficient algorithm for the minimization of non-differentiable convex functions, which uses proximal mappings as subiterations and alternates primal-dual optimization, is the *primal-dual hybrid gradient algorithm (PDHG)*. We are going to use this algorithm to solve problems from hyperspectral image processing.

In the first part of this thesis we start with an introduction to hyperspectral imaging, an exciting technology which combines the discriminative power of spectroscopy with fast and versatile imaging systems. Hyperspectral imaging is the next step after multispectral imaging, which has up to now been reserved to specialist applications including earth observing satellites, industrial conveyor belt sorting and process control, yet can be expected to become cheaply available to the market in the near future. Multispectral sensors that are small enough for smartphones, are currently being developed at the Leibniz Institute of Photonic Technology in Jena (R. Riesenberg, personal communication).

The abundance of information in a hyperspectral image requires a condensed representation. Hyperspectral unmixing represents a hyperspectral image with L channels by $p \ll L$ abundance channels and p corresponding pure spectra, $\underline{k}_1, \ldots, \underline{k}_p$, called *endmembers*. According to the commonly adopted *linear mixing model* the spectrum at spatial position i is given by

$$\underline{y}(i) = \sum_{r=1}^{p} \underline{k}_r x_r(i),$$

where $\underline{x}(i) \in \mathbb{R}^p$, the vector of abundances at position i, with components $x_r(i), r = 1, \ldots, p$, gives the relative mixture coefficients of the endmembers.

If the endmembers are known, hyperspectral unmixing infers the abundances $\underline{x}(i)$ at all points $i = 1, \ldots, N$ of the image. As real images contain noise and outliers, the reconstruction is improved by spatial regularization, see e.g. [59] and the unmixing model with total variation (TV) regularization in Section 2.

Motivated by missing-pixel defects common in current hyperspectral imaging systems, we propose a TV-regularized unmixing model for incomplete and noisy data for the case when pure spectra are given. We minimize the proposed model by the PDHG algorithm with (0.1) as subiteration.

Library unmixing refers to the case where the endmembers are not known a priori and have to be selected from a library. This problem is more difficult because the spectra in hyperspectral libraries are commonly strongly correlated, and the unmixing can be improved by enforcing sparsity of the abundance vector in each pixel, see e.g. [59], where the l_1-norm $\|\underline{x}(i)\|_1$ has been used. A more sophisticated approach is to enforce a sparsity pattern jointly on all pixels by the $l_{\infty,1}$-norm

$$\|X\|_{\infty,1} := \sum_{r=1}^{p} \max_{i \in \{1,\ldots,N\}} |x_r(i)|.$$

We present a novel unmixing model that combines for the first time spatial TV-regularization with $l_{\infty,1}$-sparsity regularization and we minimize the model using again the PDHG algorithm.

Finally, for *blind unmixing*, where no prior information about the pure spectra is given, we consider the following unmixing model with the Kullback-Leibler divergence in the data term,

$$\operatorname*{argmin}_{K,X} \ \widetilde{\mathrm{KL}}\,(Y, KX) + \frac{\mu}{2}\|K\|_F^2 + \frac{\nu}{2}\|X\|_F^2$$

subject to

$$K \in \mathbb{R}^{L,p}_{\geq}, \ X \in \triangle_p^N, \qquad \triangle_p := \big\{x \in \mathbb{R}^p \colon x \geq 0, \ \sum_{i=1}^{p} x_i = 1\big\} \tag{0.2}$$

with μ, ν positive regularization parameters. Extending the surrogates by D. D. Lee and H. S. Seung [70], we derive a multiplicative algorithm which automatically preserves the positivity constraints. The derivation by optimization transfer implies a decrease of

the objective in each iteration. While the original multiplicative update rules suggested in [70] for the functional $\mathcal{E}(K,X) = \widetilde{\mathrm{KL}}\,(Y, KX)$ with pure nonnegativity constraints on K and X, received much attention, it has often been overlooked, that the convergence of the objective does not imply convergence of the iterates. Using Zangwill's convergence theory [109], as suggested for the particular case of blind deblurring in [69], requires continuity of the update operator on the set of limit points of the sequence of iterates, which to the best of our knowledge has never been identified for the general case.

For the original update rules from [70] a convergence proof has been given in [45], by lifting the involved matrices to stochastic tensors.

Our approach for the update rules derived for the extended objective with $\mu, \nu > 0$ and the constraints (0.2) is independent from theirs. We give here a first rigorous convergence analysis for the extended update rules. In particular, we provide an example which shows that the update mapping is discontinuous on the boundary of the nonnegativity constraint set, we identify a smaller compact set containing all limit points, to which the update mapping extends continously, and conclude that the limit point of every convergent subsequence is a fixed point.

In the second part of this thesis we study Moreau-Yosida envelopes for functions defined on infinite-dimensional Hadamard manifolds and Hadamard spaces. In particular we generalize to this setting a well known result of H. Attouch which states that Mosco convergence of a sequence of proper convex lower semicontinuous functions defined on a Hilbert space is equivalent to pointwise convergence of the associated Moreau-Yosida envelopes. More precisely, while it has already been known that the Mosco convergence of a sequence of convex lower semicontinuous functions on a Hadamard space implies the pointwise convergence of the corresponding Moreau-Yosida envelopes, the converse implication was an open question. We now fill this gap.

To understand the statement, let us briefly recall the result in a Hilbert space H. First we need the definition of Mosco convergence, which goes back to U. Mosco [83]. A sequence $\{f_n\}_n$ of functions $f_n \colon H \to \mathbb{R} \cup \{+\infty\}$ *Mosco-converges* to $f \colon H \to \mathbb{R} \cup \{+\infty\}$, abbreviated $f_n \xrightarrow{M} f$, if, for each $x \in H$, the following two conditions are fulfilled:

i) $f(x) \leq \liminf_{n\to\infty} f_n(x_n)$ whenever $x_n \overset{w}{\to} x$,

ii) there is a sequence $\{y_n\}_n$ such that $y_n \to x$ and $f_n(y_n) \to f(x)$,

where $x_n \overset{w}{\to} x$ stands for weak convergence. A weaker type of convergence particularly suited for minimization problems is given by Γ-convergence, see, e.g. [26, 37], for which we just have to replace the first statement in the above definition by

i) $f(x) \leq \liminf_{n\to\infty} f_n(x_n)$, whenever $x_n \to x$.

The following theorem has been proved by H. Attouch, see [7, 8].

Theorem 1 *Let H be a Hilbert space and let $f_n \colon H \to \mathbb{R} \cup \{+\infty\}$, $n \in \mathbb{N}$, be proper convex lower semicontinuous functions. Then the following statements are equivalent:*

i) $\{f_n\}_n$ *converges to a function* $f\colon X \to \mathbb{R}\cup\{+\infty\}$ *in the sense of Mosco,* $f_n \xrightarrow{M} f$.

ii) The sequence of Moreau-Yosida envelopes $\{f_{n,\lambda}\}_n$ *of* $\{f_n\}_n$ *converges pointwise to the Moreau-Yosida envelope* f_λ *of* f*, for all* $\lambda > 0$.

Our aim in the second part of this thesis is to generalize Theorem 1 to Hadamard spaces. The implication i) ⇒ ii) has already been extended to Hadamard spaces, first for nonnegative functions by K. Kuwae and T. Shioya [66], and then without the assumption of nonnegativity by M. Bačák [13]. We therefore focus on the inverse implication, which was left open [11, Question 5.2.5] and is our main result.

Our presentation is in two steps. First we give a proof in simply-connected Riemannian manifolds of nonpositive curvature. For this we need several intermediate results on infinite-dimensional Riemannian manifolds which can not be found in the literature. Several of these results are interesting in their own right. In particular we prove a subdifferential sum rule for convex functions defined on infinite-dimensional manifolds.

We prove that Hadamard manifolds with a lower bound on the sectional curvature satisfy for every sequence $\{x_n\}_n$ the following relation,

$$\text{(A)} \qquad x_n \xrightarrow{w} x \quad \Rightarrow \quad \log_x x_n \xrightarrow{w} 0.$$

I.e., we show that that weak convergence implies weak convergence in the tangent space, under the assumption that the manifold has bounded nonpositive curvature. We give a new proof in Hadamard spaces of the following result on the parameter dependence of the Moreau-Yosida envelopes, known in Banach spaces,

$$\frac{\mathrm{d}}{\mathrm{d}\lambda}\lambda f_\lambda(x) = f(\operatorname{Prox}_{\lambda f} x),$$

which builds on the definition of weak convergence in Hadamard spaces, along with Γ-convergence and known compactness criteria for weak convergence.

Finally we prove the following representation for the gradient of the Moreau-Yosida envelope, see Theorem 6.22.

Theorem 2 *Let* $\mathcal{H}$ *be a Hadamard manifold and let* $f\colon \mathcal{H} \to \mathbb{R}\cup\{+\infty\}$ *be a proper convex lower semicontinuous function. Then, for all* $\lambda > 0$*, the Moreau-Yosida envelope* f_λ *is differentiable with Riemannian gradient*

$$\nabla f_\lambda(x) = -\tfrac{1}{\lambda}\log_x\left(\operatorname{Prox}_{\lambda f}(x)\right).$$

With the help of these results we build on the idea originally used by Attouch in Banach spaces, to prove the implication ii) ⇒ i) on Hadamard manifolds satisfying (A).

In a second part of the presentation we show that the proof can be generalized further to Hadamard spaces and is simplified considerably by working directly from the curvature bounds and weak convergence. More precisely we prove the following main result.

Theorem 3 *Let* $\mathcal{H}$ *be a Hadamard space,* $\{f_n\}_n$ *a sequence of proper convex lsc functions* $f_n\colon \mathcal{H} \to \mathbb{R}\cup\{+\infty\}$*, and* $f\colon \mathcal{H} \to \mathbb{R}\cup\{+\infty\}$ *a proper convex lsc function. Assume that for each* $\lambda > 0$ *the sequence of Moreau-Yosida envelopes* $\{f_{n,\lambda}\}_n$ *converges pointwise to the Moreau-Yosida envelope* f_λ*. Then* $f_n \xrightarrow{M} f$ *as* $n \to \infty$.

Both Γ- and Mosco convergence have already been used in the framework of Hadamard spaces. In [61], J. Jost introduced Γ-convergence on Hadamard spaces as a tool and defined Mosco convergence by saying that a sequence of convex lsc functions on a Hadamard space Mosco converges if their Moreau-Yosida envelopes converge pointwise. In [66], K. Kuwae and T. Shioya give an in depth study of both Γ- and Mosco convergence in Hadamard spaces. They already gave the standard definition of the Mosco convergence, relying on the notion of weak convergence, and right after their Definition 5.7 in [66] they note "Jost's definition of Mosco convergence... seems unfitting in view of Mosco's original definition." By our main result it follows that both definitions are equivalent. As a corollary we obtain that Mosco convergence of convex closed sets is equivalent to Frolík-Wijsman convergence.

Part I.

Hyperspectral Image Processing

1. Introduction to Hyperspectral Image Processing

Our major source of information from the world around us, is to observe its interaction with electromagnetic radiation. The human eye perceives the world by averaging the intensity of backscattered radiation over three ranges in the electromagnetic spectrum corresponding to red, green and blue. Colour photography has been one of the major breakthroughs in photography, yet neither nature nor technology have stopped at three colour channels. The mantis shrimp has adapted to its colourful environment in the top few meters of coral reefs by developing a retina with four colour receptors, and some species have up to twelve [76].

Similarly, the number of channels recorded by satellites has increased. While the first multispectral scanner launched to space on the Landsat I satellite in 1972 had four bands, the Thematic Mapper had seven spectral bands in 1982 and the Hyperion sensor system on-board an experimental NASA spacecraft in 2000 measured 220 bands [67]. At each point of the scene, the Hyperion sensor measures 220 reflectance values, which form a spectrum that can act as a fingerprint of the materials which reflected the light.

Being on the one hand an extension of colour imaging to more than three colour channels, hyperspectral imaging can be regarded as a marriage of imaging, the acquisition of 2D or 3D grids of measurements, with spectroscopy.

Classical spectroscopy is used to acquire single spectra. With imaging systems developed in the last years, it has become possible to measure images that have a spectrum in each pixel, containing several hundred reflectance values. They are called **hyperspectral images**.

Paralleling the two perspectives, a hyperspectral image can be regarded either as a stack of images, one for each wavelength range, see Figure 1.1 (left) or as an assembly of spectra, one for each point of the measured scene or object, see Figure 1.1 (right).

Hyperspectral images can also be obtained from classical spectrometry techniques such as mass spectroscopy, by measuring pixel by pixel, one at a time.

Common to all kinds of hyperspectral images is a mass of available data that must be processed to extract the information of interest. The first part of this thesis is devoted to one kind of dimension reduction, called hyperspectral unmixing.

Dimension reduction is necessary for several reasons. The first and simplest is that a human observer can only view one, or, with false colour combinations, at most three channels at a time. In data exploration it may be feasible to view a few handful of channels, yet not hundreds. Second, information may be spread over many bands and may become visible only after a basis change in the vector space of spectra. This is the case for near-infrared (NIR) spectroscopy. Third, the dimension reduction may be

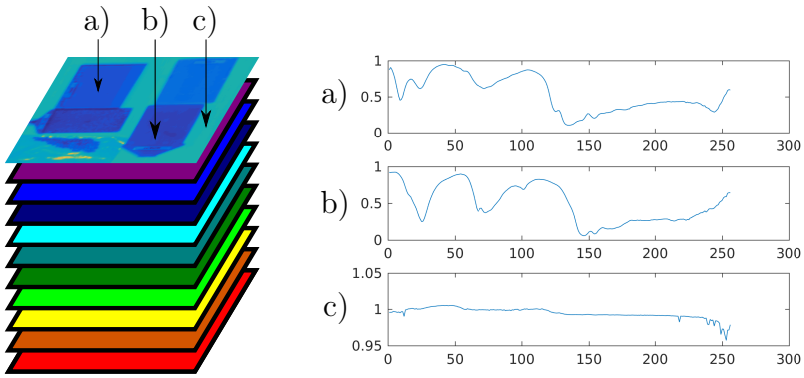

Figure 1.1.: A hyperspectral image can be regarded as a stack of images, one for each wavelength range (left) or as an assembly of spectra, one for each point (right).

required to achieve a good performance of sequential processing steps such as training of classifies. E.g. support vector machines (SVMs) perform better on few channels than on many channels with high redundancy. Finally the noise level is lower after the dimension reduction.

Depending on the amount of prior knowledge, hyperspectral unmixing can be used in various scenarios, from replacing classical dimension reduction techniques such as PCA or ICA, to continuous relaxations of segmentation with given centres.

Before commenting on the outline and highlighting our contribution, we briefly name three sources of hyperspectral data used in the sequel. With advancing sensor techniques, large parts of the electromagnetic spectrum in Figure 1.2, from X-rays to infrared are now covered. The classic **near-infrared (NIR)** radiation will be described in detail in the next section.

A more recent technology is **terahertz spectroscopy**. For a thorough introduction to terahertz-time domain spectroscopy (THz-TDS) we refer to [100].

MALDI imaging mass spectroscopy In mass spectrometry imaging, the sample is moved in two dimensions and a mass spectrum is recorded in each point by locally vapourizing and ionizing the sample with a laser beam. The term MALDI imaging mass spectroscopy refers to the use of matrix-assisted laser desorption/ionization (MALDI) in this setting. MALDI is a technique for soft ionization (ionization without fragmentation) of biopolymers and other macromolecules, developed by M. Karas and F. Hillenkamp [63]. A pulsed laser is used for vapourization and ionization of the sample. As direct radiation of the analyte would fragment the macromolecules, the analyte is first co-crystalized with a large molar excess of a matrix compound. Laser radiation of the analyte-matrix mixture then vapourizes the matrix, which carries the analyte with it. The matrix further acts as a proton donor and acceptor, ionizing the macromolecules of the analyte without fragmenting them. A brief summary of the early development can be found in [62].

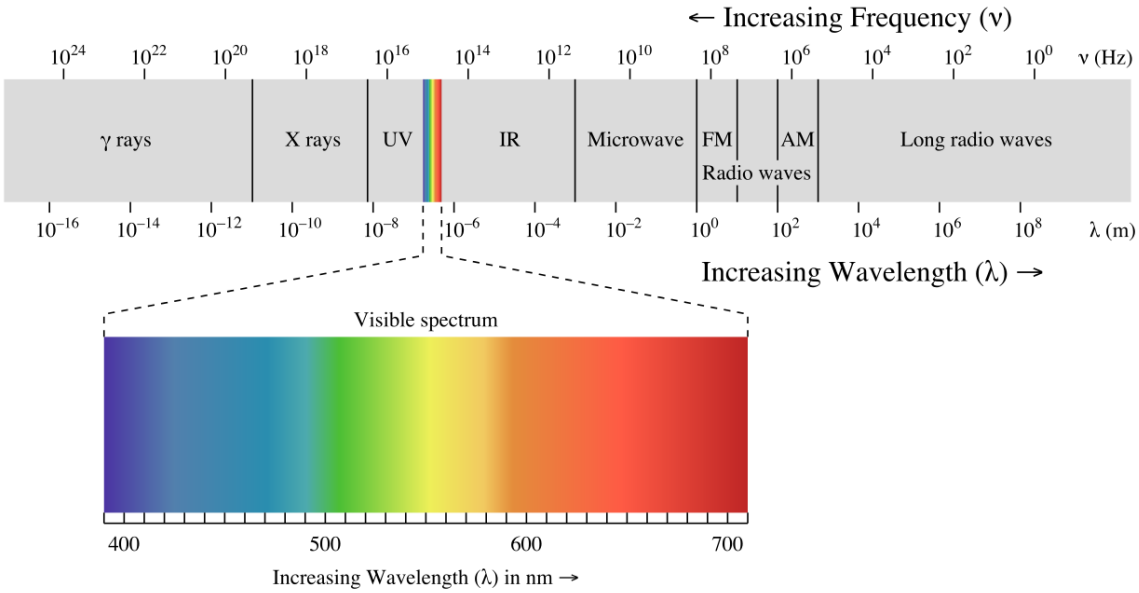

Figure 1.2.: The electromagnetic spectrum. Different types of interaction between the electromagnetic radiation and matter are dominant in different parts of the spectrum. Source of the image: `https://commons.wikimedia.org/wiki/File:EM_spectrum.svg`, called 22 November 2016.

This introductory chapter contains two further sections, before we present our first unmixing model in the next chapter. In Section 1.1 we give an introduction to near-infrared (NIR) imaging, including calibration and applications.

In Section 1.2 we describe dimension reduction techniques used across to different fields and the advantages of hyperspectral unmixing for processing of spectra.

1.1. Near-Infrared Imaging

In the last decades, NIR-imaging systems have become a robust and versatile technique to measure without sample preparation from different observation distances. All near-infrared images in this thesis have been measured with the hyperspectral camera at the Fraunhofer ITWM Kaiserslautern[1], which measures 256 channels in the wavelength range 1073-2300nm. It operates as a line camera with a line width of 320 pixels, as described in the next section. The camera is shown in Figure 1.3 (left).

[1]cofinanced by the European Regional Development Fund (ERDF)

1.1.1. Line Camera

A standard approach in hyperspectral image acquisition is to measure lines of the image one by one simultaneously at all wavelengths. In earth observation, this is referred to as **push-broom scanning**. The method is also used in industrial applications, such as plastics sorting for recycling purposes.

At any time, one line of the image is measured at all wavelengths simultaneously. The incoming light from the current line is diffracted by an optical grid onto the area sensor of the camera (see Figure 1.3 (right)) and recorded as one "sensor frame". The full 2D object is measured by moving the object relative to the camera. While the lines add up to the full 2D object, the sensor frames with spectral and along-the-line directions are stacked along the second image direction to form the full 3D hypercube.

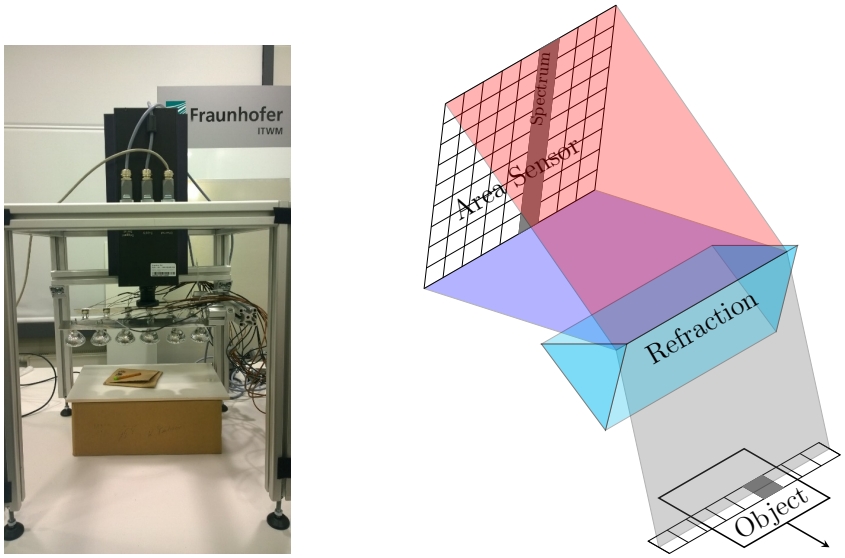

Figure 1.3.: Hyperspectral line camera (**left**) and principle of measuring a line simultaneously at all wavelengths **(right)**.

1.1.2. Calibration

We now describe the calibration which is required to obtain the hypercube from the raw sensor counts. The snapshot of the area sensor taken for the line $y = y_j$ of the object yields a section $Y^{\text{cube}}_{:,y_j,:}$ of the hypercube $Y^{\text{cube}} \in \mathbb{R}^{m,n,L}$ as depicted in Figure 2.1 (right), where the z-axis is the spectral direction. As the sensor frames are measured one after another in the same way, we just consider the first sensor frame $r_{il} := Y^{\text{cube}}_{i1l}$, for $i = 1, \ldots, m$, $l = 1, \ldots, L$. Then $r_i = (r_{i1}, \ldots, r_{iL})^\top$ is the spectrum measured for the light entering the camera at the ith position of the line slit.

The response z_{il} of the (i,l)th sensor pixel increases approximately linearly with the amount of light l_{il} in the lth wavelength range (centered at ν_l) which enters the ith position of the line slit. The slope corresponds to the **sensitivity** φ_{il} of the sensor pixel. Denote by d_{il} the dark counts without incoming light, due to thermal excitations. The

relationship between the sensor counts and the incoming light is given by

$$z_{il} = d_{il} + \varphi_{il} l_{il} = d_{il} + \varphi_{il} \lambda_{il} r_{il}, \tag{1.1}$$

using for the second equality that the incoming light intensity l_{il} is the product of the illumination λ_{il} and reflectance r_{il} in the lth wavelength range at the corresponding point of the object.

If a reference material with reflectances r^w is measured under the same illumination λ_{il}, we can solve the equation $z^w = d + \varphi \circ \lambda \circ r^w$ for $\varphi \circ \lambda$ and obtain from (1.1) that

$$r = \frac{1}{\varphi \circ \lambda} \circ (z - d) = r^w \circ \frac{z - d}{z^w - d}, \tag{1.2}$$

where the **Hadamard matrix product** is defined by $(A \circ B)_{ij} = A_{ij} B_{ij}$. The sensitivities φ_{il} vary slightly from sensor pixel to sensor pixel due the manufacturing process, yet depend strongly on l (i.e. the wavelength), hence the same holds for the quantities $\frac{r^w_{il}}{z^w_{il} - d_{il}} = \frac{1}{\varphi_{il} \circ \lambda_{il}}$. Figure 1.4 shows the vectors $r_i = (r_{il}, \ldots, r_{il})^\top$ for three columns $x = x_i, i = 1, 2, 3$ of an InGaAs-detector grid.

The sensitivity of the InGaAs-sensors declines to zero for wavelengths close to the upper and lower end of the sensitivity interval $[\nu_{min}, \nu_{max}]$. Therefore the multiplication by $\frac{1}{\lambda \circ \varphi}$ in (1.2) increases the noise in the measurement for wavelength close to ν_{min}, ν_{max}.

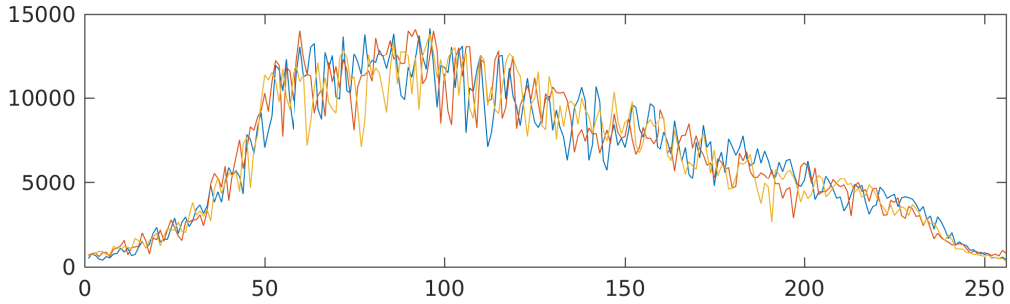

Figure 1.4.: The raw-counts of three columns of the InGaAs-detector grid corresponding to the spectra measured at three points of a white calibration tile. The wavelength increases linearly with the channel index on the x-axis, in the sensitivity range $[\nu_{min}, \nu_{max}]$ of the InGaAs-detector pixels. Towards both ends, ν_{min}, ν_{max}, the sensitivity decays to zero. After calibration these spectra are approximately constant, with increased noise towards the smallest and largest channel index.

Modelling the sensors by the usual model of Poisson random variables, which for usual counts can be replaced by its Gaussian approximation, equation (1.1) becomes

$$z = d + n + \varphi \circ \lambda \circ r$$

where the n_{il}, $i = 1, \ldots, m$, $l = 1, \ldots, L$ are i.i.d. random variables $n_{il} \sim \mathcal{N}(0, \sigma)$ for some $\sigma > 0$. Figure 1.4 suggests the approximation $\frac{1}{\varphi_{il}\,\lambda_{il}} \approx \frac{1}{\overline{\lambda}_l}$ for $\overline{\lambda} \in \mathbb{R}^L$ defined as the spatial average of $\varphi \circ \lambda$ over i. This yields

$$r_{il} = \frac{r_{il}}{z_{il}^w - d_{il}}(z_{il} - d_{il}) + \overline{\lambda}_l n_{il},$$

i.e. r_{il} follows a Gaussian distribution with wavelength dependent variance $\sigma\overline{\lambda}_l$.

1.2. Dimension Reduction

In this section we introduce standard dimension reduction techniques and explain the necessity of a new technique such as nonnegative matrix factorization. We comment on the advantages of the variational approach to nonnnegative matrix factorization.

1.2.1. Wavelets

A wavelet transformation represents a signal in a diffrent basis or, when redundancy is permitted, in a socalled frame. The simplest case of orthogonal wavelets shares with the Fourier transformation the property that the basis transformation is orthogonal, with the difference that wavelets are "localized", i.e. for lower frequencies there there is an increasing number of wavelets which pick up features at different positions. This makes wavelets suitable for spectral analysis, where, in addition to the spectral width and amplitude of a peak in the spectrum, also the spectral position of the peak is relevant, see e.g. the analysis of terahertz spectra in [100].

In wavelet shrinkage the wavelet frequencies which contain most of the noise, are discarded and the spectra are represented in the smaller number of remaining wavelet coefficients. This can lead to a considerable dimension reduction, and in addition features may be more separated into coefficient channels, which can improve segmentation [100].

Wavelets have been used successfully for the tasks of hyperspectral inpainting [40] and estimating the dimension of the hyperspectral subspace [91]. However, wavelets have the disadvantage of a nonadaptive basis and especially if spectral features of both coarse and small scales are relevant, the number of wavelet channels required for a good representation may still be too large.

1.2.2. PCA

Principal Component Analysis (PCA) refers to representing a set of data vectors $\{y_i\}_{i=1}^N \subseteq \mathbb{R}^L$, assembled in a data matrix $Y = (y_1, \ldots, y_N)$, in an orthogonal coordinate system, in which the components of the data vectors are uncorrelated and ordered according to decreasing variance. The axes are called **principle components** and are computed via the **singular value decomposition (SVD)** of the centered data matrix

$$Y_c := Y - (\overline{y}, \ldots, \overline{y}), \tag{1.3}$$

where the $\overline{y}$ denotes the data mean $\overline{y} := \frac{1}{N}\sum_{i=1}^{N} \underline{y}(i)$. We recall the following basic fact, see e.g. [54, Satz 12.1], where the theorem is stated with $\mathbb{R}$ replaced by either $\mathbb{R}$ or $\mathbb{C}$ and the Hermitian conjugate Y^* replacing the transpose $Y^\top$.

Theorem 1.1 *Every matrix $Y \in \mathbb{R}^{L,N}$ of rank r has an SVD, i.e. a system*

$$\big\{\sigma_i, y_j, v_k \colon i = 1, \dots, r, \quad j = 1, \dots, L, \quad k = 1, \dots, N\big\}$$

with $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r > 0$ and orthonormal bases $\{u_j\}_{j=1}^{L}$ and $\{v_k\}_{k=1}^{N}$ of $\mathbb{R}^L$ respectively $\mathbb{R}^N$, such that

$$\begin{aligned} Y v_i &= \sigma u_i, \quad Y^\top u_i = \sigma_i v_i, \quad i = 1, \dots, r, \\ Y v_k &= 0 \qquad Y^\top u_j = 0, \qquad j, k > r. \end{aligned}$$

*The σ_i are called **singular values** of Y and their squares σ_i^2 are, with their respective multiplicities, exactly the nonzero eigenvalues of $Y^\top Y$.*

This can be written concisely in matrix notation. To this end we introduce the matrices

$$U = (u_1, \dots, u_L) \in \mathbb{R}^{L,L} \quad \text{and} \quad V = (v_1, \dots, v_N) \in \mathbb{R}^{N,N},$$

which satisfy

$$U^\top U = \mathrm{I} \in \mathbb{R}^{L,L} \quad \text{and} \quad V^\top V = \mathrm{I} \in \mathbb{R}^{N,N},$$

and

$$\Sigma = \begin{pmatrix} \sigma_1 & & & \mathbf{0} \\ & \ddots & & \vdots \\ & & \sigma_r & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \end{pmatrix} \in \mathbb{R}^{L,N}.$$

Thus the SVD reads

$$\begin{aligned} Y &= U\Sigma V^\top, \qquad Y^\top = V\Sigma^\top U^\top \qquad \text{or} \\ Y &= \sum_{i=1}^{r} \sigma_i u_i v_i^\top, \quad Y^\top = \sum_{i=1}^{r} \sigma_i v_i u_i^\top. \end{aligned}$$

The variance of the centred data along the axis u_j is

$$\|Y_c^\top u_j\|_2^2 = \|\sum_{i=1}^{r} \sigma_i v_i u_i^\top u_j\|_2^2 = \sigma_j^2.$$

When using PCA for dimension reduction, the centred data is projected onto subspaces spanned by principle components. For $\mathcal{I} \subset \{1, \dots, L\}$, the projection $\mathcal{P}_\mathcal{I}$ of a vector $m = \sum_{i=1}^{L} m_i u_i$ onto the subspace spanned by $\{u_i\}_{i\in\mathcal{I}}$, is given by $\sum_{i\in\mathcal{I}} m_i u_i$. To keep only the $p < L$ components with respect to which the data has the largest variations $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_p$, define the projected data $Y_P \approx Y$ by

$$Y_P = (\overline{y}, \dots, \overline{y}) + \mathcal{P}_\mathcal{I}\, Y_c, \qquad \mathcal{I} = \{1, \dots, p\}.$$

Due to the criterion of maximum variance in the choice of principal components, PCA can be strongly affected by noise. If some feature channels have a large variance due to noise, principal directions may become aligned with them. Let us illustrate this with a simple example.

Example 1.2 *We simulate qualitatively the situation for the hyperspectral NIR camera described in Section 1.1.1 with increased noise level towards the ends of the spectral range.*

Two groups containing 23 *spectra with* 120 *values are generated by adding weighted random noise to the mean spectra g, h depicted in Figure 1.5 (a). One sample spectrum from each group is shown in Figure 1.5 (b).*

More precisely, the noise added to the ith component is sampled uniformly at random from the interval $[-w_i, w_i]$, independently from the other components, with weight vector $w \in \mathbb{R}^{120}$ defined by the matlab command `w =[ones(1,19), linspace(1,1/10,22), 1/10*ones(1,38), linspace(1/10,1,22), ones(1,19)]`.

All generated spectra $\{y_1, \ldots, y_{23}\} \cup \{y_{24}, \ldots, y_{46}\}$ are visualized by plotting the matrix $(y_1, \ldots, y_{46})$ as an image in Figure 1.5 (c). The spectra from the second group correspond to the vertical lines of the right half of the plot, where the light bar in the middle reflects the bump present in each of the spectra in the second group.

The PCA coefficients corresponding to the first two principal directions v_1, v_2 are plotted for all 46 *spectra in Figure 1.6. There is no obvious difference between spectra* 1–23 *and spectra* 24–46*, i.e. the first two PCA coefficients do not seem to separate the classes. This is confirmed when looking at the projection to the subspace spanned by v_1, v_2, see Figure 1.6 (b). In the projection the distinguishing bump has vanished.*

In the above example, where the noise variance was channelwise detector noise, independent between the euclidean components of the spectrum, the noise variance in each channel may be known from calibration or previous experiments with the measurement system. In case of the NIR hyperspectral camera described in Section 1.1.2, the variance of channel i in the calibrated image is proportional to the reciprocal sensitivity of the sensor pixel for the associated wavelengths and —up to small variations between the sensor pixels— it is roughly a function of the wavelength/channel. Whenever the noise is independent between the channels and its strength per channel is known for the given measurement system from calibration measurements, in this case we can still use a channelwise weighted PCA, see [89] or the much more general [41], where not only channelwise, but also pointwise weights are considered.

In other cases, where the noise does not follow a statistics with principal directions aligned with the channels, this is no longer possible. If the noise stems from variance of the data within classes, the principal directions of variance of the noise will typically not be aligned with channels.

Consider the example where we seek to detect the weed percentage between agricultural plants from an airborne image. One source of data variance will stem from the amount of leaf cover within the pixel, i.e. the percentage of the pixel where the plant covers the soil. The joined percentage of plant cover within the ground area corresponding to one pixel is likely to dominate the spectral difference between the the agricultural

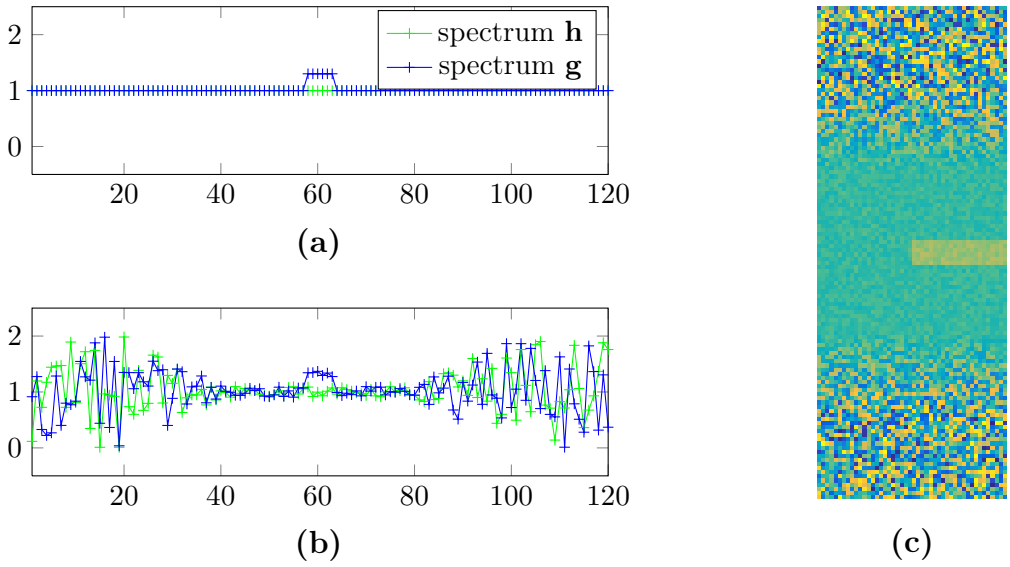

Figure 1.5.: Two classes of spectra simulated by adding channelwise weighted noise to the two respective class means: **(a)** The two mean spectra g, h; **(b)** one sample spectrum from each class and **(c)** all 46 generated spectra visualized by plotting the matrix $(y_1, \ldots, y_{46})$.

plant and the weed. In this example one would first need to find the spectral directions of the endmembers corresponding to both plant species and then use the contribution of each of the two plant species as the discriminating factor.

This example has been given to stress that other knowledge needs to be incorporated. Here spatial regularization may be used to distinguish between the fast varying percentages and the endmember directions, which vary slowly and roughly span the same spectral subspace in different parts of the image. One advantage of hyperspectral unmixing by the variational approach to be considered is its greater flexibility compared to an algebraic definition of PCA. Spatial regularizers can be added as summands to the objective functional and algorithms to minimize many of the resulting functional are now available, see e.g. Chapter 2, where after finding the endmembers, we use spatial regularization in the second step to find the abundances by solving a minimization problem.

1.2.3. ICA

Independent component analysis (ICA) is a method to analyze multivariate statistical data which seeks underlying factors, called components, that are statistically independent and nongaussian. In statistical terms, the columns $\{y(j)\}_{j=1}^{N} \subseteq \mathbb{R}^L$ of the data matrix are regarded as observations of a random vector $y \in \mathbb{R}^L$. Each component y_i corresponds to an individual measurement taken in every observation. The components of y are assumed to be linear combinations of the components $\{s_i\}_{i=1}^{p}$ of an unknown

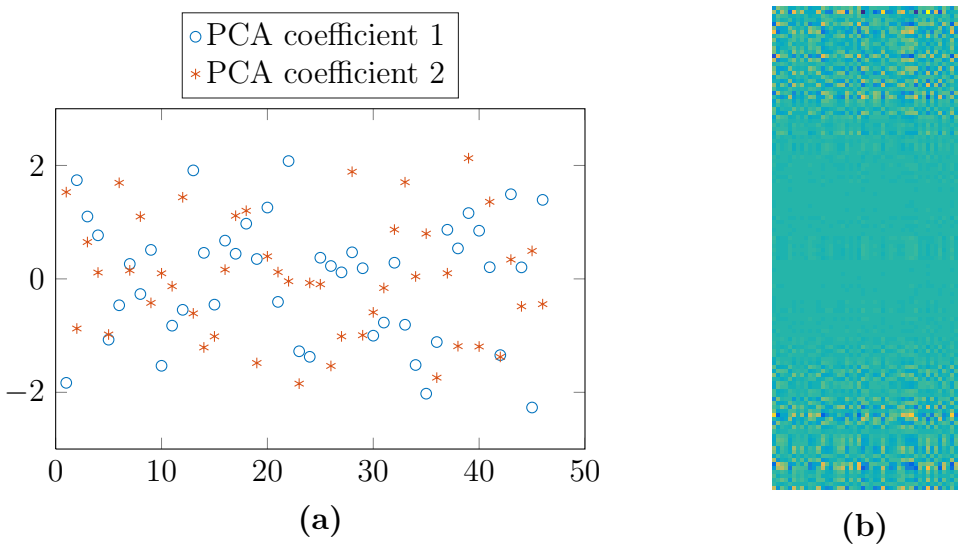

Figure 1.6.: Result of dimension reduction by PCA for $p = 2$. The discriminating information is lost. Shown are the PCA coefficients corresponding to the first two principal directions v_1, v_2 for all 46 spectra (left), and the projections $\mathcal{P}_{\{1,2\}}\, y_i$ of all spectra to the subspace spanned by the first two principal directions, visualized by plotting the matrix $\left(\mathcal{P}_{\{1,2\}}\, y_1, \ldots, \mathcal{P}_{\{1,2\}}\, y_{46}\right)$.

random vector s, which are called **sources**.

In a common illustration the components $\{s_i\}_{i=1}^p$ are different sources of noise produced independently at a cocktail party. Assume that L microphones are positioned at the party (corresponding to the measurements y_i, $i = 1, \ldots, L$). It is natural to assume that each microphone records a weighted sum of the sources, where a weight is stronger if the source is closer to the microphone. Given the recordings of the microphones, can one recover the mixing matrix containing the weights along with the sources?

Assuming that the sources are Gaussian, i.e. normally distributed, it is easy to find independent components, because for Gaussian data, independence is the same as having a diagonal covariance matrix $\Sigma = \operatorname{diag}(\sigma_1, \ldots, \sigma_L)$, and in this case the joint density, here after subtracing the mean, factors

$$f(y_1, \ldots, y_L) \sim e^{-\frac{1}{2} y^\top \Sigma^{-1} y} = \Pi_{i=1}^L e^{-\frac{y_i^2}{2\sigma_i}}.$$

Further, if the mixing matrix and the axes are scaled such that all independent components have unit variance $\sigma_1 = \cdots = \sigma_L$, i.e. if the data has been **sphered**, then the distribution is rotationally symmetric and the independent components are not uniquely defined. Assuming $L = p$ and excluding Gaussian sources essentially makes the problem well-defined [57, Section 7.5]. More precisely, ICA builds on the following assumptions [64]:

(i) For $1 \leq i \neq j \leq p$, s_i and s_j are statistically independent.

(ii) No s_i is normally distributed.

(iii) For $1 \leq i \leq p$, the variance $\sigma_i^2 := E[(s_i - E(s_i))^2]$ of s_i is positive.

Then, given observations $\{y(j)\}_{j=1}^N \subseteq \mathbb{R}^L$, ICA assumes that

$$y(j) = As(j), \qquad j = 1, \ldots, N$$

and finds the **mixing matrix** $A \in \mathbb{R}^{L,p}$ and the $\{s(j)\}_{j=1,\ldots,N} \subseteq \mathbb{R}^p$, where $L = p$, such that the components $s_i(\cdot)$ are maximally independent. From now on we assume $L = p$.

A common approach is based on an idea encountered in the central limit theorem: the sum of two random variables which are not gaussian is more gaussian, i.e., its distribution is "closer to" a normal distribution than both individual distributions. Reformulating ICA as finding the inverse W of the mixing matrix A, a random vector x estimating the sources s is defined as $x := Wy$. Each independent component s_i is thus estimated by a linear combination $x_i := w^\top y$, where $w \in \mathbb{R}^L$, the ith row of W, should be found so that the distribution of x_i is as far from being gaussian as possible. Indeed, using that $y = As$,

$$x_i = w^\top As = z^\top s, \qquad z := A^\top w$$

we see that this will happen if the weight vector is trivial, $z = e_k$, for some $k \in \{1, \ldots, L\}$, and e_k denoting the kth unit vector with entry 1 in the kth position and 0 entries otherwise.

The vector $w_j := (A^\top)^{-1} e_j = W^\top e_j$ corresponding to $z = e_j$, for each $j \in \{1, \ldots, L\}$, should thus be a local maximum of a suitable nongaussianity measure H evaluated at $x = w^\top y$. The independent components are determined as L local maxima of $w \mapsto H(w^\top y)$. Such measures of nongaussianity are the kurtosis and negentropy, see [57] for details.

We note that the negentropy of a random variable y is always nonnegative and zero if and only if y has a Gaussian distribution. Being difficult to optimize, approximations are commonly used [58].

Typical applications of ICA described in [57] are brain imaging, where sources emitting signals in the brain are mixed up in the sensors outside the head, and econometrics, where parallel time series should be decomposed in underlying causes.

1.2.4. Hyperspectral Unmixing and the Variational Perspective

Hyperspectral unmixing assumes the existence of pure spectral signatures, called endmembers and seeks to represent each observed spectrum as a weighted linear sum of these endmembers. The weights are called abundances and are assumed to be nonnegative in all unmixing models, which increases interpretability of the representation.

Denote by $\underline{k}_1, \underline{k}_2, \ldots, \underline{k}_p \in \mathbb{R}^L_{\geq}$ the p endmember spectra, $p \ll L$, each measured at L wavelengths, and by $\underline{x}(j) = (x_1(j), x_2(j), \ldots, x_p(j))^\top \in \mathbb{R}^p$ the associated vector of abundances for the jth observed spectrum $\underline{y}(j) \in \mathbb{R}^L_{\geq}$. Defining the matrices

$$\begin{aligned} Y &:= \big(\underline{y}(1)\ \underline{y}(2)\ \cdots\ \underline{y}(N)\big) \in \mathbb{R}^{L,N} && \text{(data matrix)}, \\ X &:= \big(\underline{x}(1)\ \underline{x}(2)\ \cdots\ \underline{x}(N)\big) \in \mathbb{R}^{p,N} && \text{(abundance matrix)}, \\ K &:= (\underline{k}_1\ \underline{k}_2\ \ldots\ \underline{k}_p) \in \mathbb{R}^{L,p} && \text{(endmember matrix)}, \end{aligned}$$

hyperspectral unmixing seeks an approximate factorization

$$Y \approx KX \qquad \text{s.t.} \quad K \in \mathbb{R}^{L,p}_{\geq},\ X \in \mathbb{R}^{p,N}_{\geq}. \tag{1.4}$$

If only Y and the desired approximation dimension $p \in \mathbb{N}$ are given, the problem is then referred to as **unsupervised unmixing** and is closely related to the problem of **nonnegative matrix factorization**, see Chapter 4. Unsupervised hyperspectral unmixing differs from PCA and ICA in two respects: First, the columns of K are not assumed to be orthogonal as in PCA and the rows of X are not assumed to be statistically independent as in ICA. Second, the nonnegativity assumption on X and hence on the representation of each column of Y as

$$\underline{y}(j) = \sum_{r=1}^{p} x_r(j)\underline{k}_r$$

allows for the interpretation that $\underline{y}(j)$ is similar to $\underline{k}_r$ if the rth entry of the vector $\underline{x}(j)$ is much larger than the other entries of $\underline{x}(j)$.

There are numerous approaches to hyperspectral unmixing, for an overview we refer to the survey by Bioucas-Dias *et al.* [23]. The **variational approach** adopted in this thesis seeks a solution (K, X) of (1.4) by minimizing an objective function $\mathcal{E}$ which typically is the sum

$$\mathcal{E}(K, X) := \mathcal{D}(K, X; Y) + \mathcal{R}(K, X), \tag{1.5}$$

of a **data fidelity term** $\mathcal{D}(K, X; Y)$ that becomes small if the product KX is a good approximation to Y and a **regularization term** $\mathcal{R}(K, X)$ that becomes small if K and X agree with prior knowledge such as spatial or spectral smoothness or sparsity. Hard constraints are incorporated by setting $\mathcal{R}(K, X) = +\infty$ for pairs (K, X) which violate the constraints. In hyperspectral unmixing, $\mathcal{R}(K, X)$ always includes nonnegativity constraints on K and X.

The variational approach has several advantages. One advantage is its flexibility. A given model can be modified by adding further terms to the objective function to account

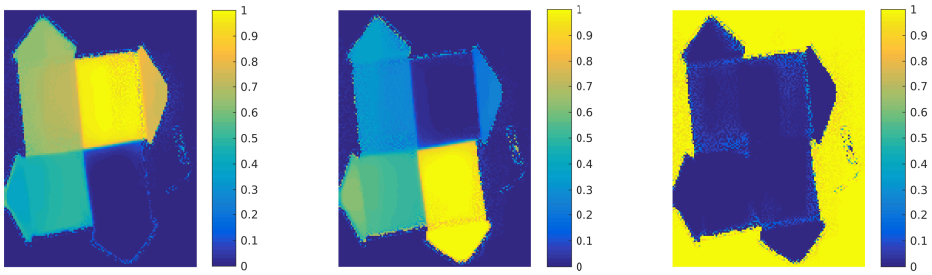

Figure 1.7.: Hyperspectral unmixing for PS-PP mixtures. Knowing the pure spectra of PS, PP and the background, the abundances obtained agree well with the true mixture percentages of both kinds of plastics, also in two tiles tiles on the left that are made of plastic mixtures.

for additional prior knowledge, see e.g. the TV-regularization term on the abundances added in Chapter 2. The data fidelity term can be chosen as the log-likelihood function of the noise, obtaining the squared Frobenius norm $\frac{1}{2}\|Y - KX\|_F^2$ for Gaussian noise, and the Kullback-Leibler divergence $\mathrm{KL}(Y, KX)$ for Poisson noise [96] and multiplicative Gamma noise [99]. We use variational models which cope naturally with some outliers, and an even greater robustness can by achieved with an l_1-data term, compare the recently suggested robust l_1-formulations of PCA and ICA [64].

A last advantage is that the variational approach separates model and algorithm. This is particularly beneficial for strictly convex models, for which the minimizer is unique. The model can first be tested by finding the minimizer with any (simple) optimization algorithm, and if the model yields good results, a faster (more sophisticated) optimization algorithm can be developed.

Recent algorithms for the efficient minimization of convex models include the primal-dual hybrid gradient algorithm (PDHG) [28, 29, 90], the alternating directions method of multipliers (ADMM) [25], with references on recent extensions in [106]. For the Frobenius norm in the data term and sparsity terms on K and X the alternating proximal gradient scheme in [24] has been shown to converge to a stationary point even for non-convex objective functions.

2. Hyperspectral Unmixing from Incomplete and Noisy Data

Outline

In hyperspectral images, once the pure spectra of the materials are known, hyperspectral unmixing seeks to find their relative abundances throughout the scene. We present a novel variational model for hyperspectral unmixing from incomplete noisy data which combines a spatial regularity prior with the knowledge of the pure spectra. The material abundances are found by minimizing the resulting convex functional with a primal dual algorithm. This extends least squares unmixing to the case of incomplete data, by using total variation regularization and masking of unknown data. Numerical tests with artificial and real-world data demonstrate that our method successfully recovers the true mixture coefficients from heavily corrupted data. The results of this chapter have been published in [81].

2.1. Introduction

2.1.1. Hyperspectral Unmixing

Hyperspectral imaging is continuously being suggested for new industrial applications. While it has been used for decades in remote sensing (see, e.g., [95]), it is now being applied in various industrial sorting applications, such as recycling of polymers, paper and quality control of crops and fruit; see the references in [85]. It has also been suggested

for controlling glue coverage in manufacturing of wood strand boards [6] or even gas plume detection [103].

In hyperspectral imaging, the electromagnetic reflectance of an object or scene is measured in several hundred wavelength ranges at once, so that the spectrum measured at one spatial pixel acts as a fingerprint of the mixture of materials at the corresponding point in the scene. A hyperspectral image of size $m \times n$ with L wavelength ranges can be represented by a hypercube $Y^{\text{cube}} \in \mathbb{R}^{m,n,L}$ or, alternatively, by column-wise spatial reshaping, as two-dimensional array $Y \in \mathbb{R}^{L,N}$, where $N = mn$. In the following, we focus on the latter representation.

Mixing of materials can be macroscopic, when, due to low resolution, several materials are measured at one spatial pixel, or microscopic, when the surface layer is an intimate mixture of the pure materials at different concentrations. Usually, hyperspectral unmixing is based on the **Linear Mixing Model**. The model comes from remote sensing, where the ground area corresponding to one pixel is large, e.g. 30×30 meters, and is often made up of several kinds of ground coverage such as grass, soil or asphalt. The intensities of light reflected by the partial areas over the measured electromagnetic range are collected in the sensor. If the partial areas of different coverage are disjoint, it is reasonable to assume that the spectral distribution of light collected in the sensor is the sum of the spectral distributions reflected by the individual types of coverage, multiplied by their respective partial areas. For the limitations of the model and algorithmic approaches we refer to the survey by Bioucas-Dias et al., [23].

Hyperspectral unmixing assumes the existence of pure spectral signatures, called **endmembers**, and seeks to infer the fractions of these endmembers, called **abundances**, in each pixel of the scene. As the number of materials, here denoted p, is usually much smaller than the number of wavelengths, hyperspectral unmixing is also an important data reduction step. Denote by $\underline{k}_1, \underline{k}_2, \ldots, \underline{k}_p \in \mathbb{R}^L$ the endmembers, corresponding to pure (average) reflectance spectra of the individual coverage types and measured at L wavelengths. Let the respective relative areas which they occupy in the ground region corresponding to the jth camera pixel, the abundances, be denoted by $x_1(j)$, $x_2(j), \ldots, x_p(j) \in [0, 1]$. The linear mixing model says that the spectrum $\underline{y}(j) \in \mathbb{R}^L$ observed in the jth position is the convex weighted sum of $\underline{k}_1, \underline{k}_2, \ldots, \underline{k}_p$, weighted by the abundances $x_1(j), x_2(j), \ldots, x_p(j)$, plus some noise vector $\underline{w}(j) \in \mathbb{R}^L$:

$$\underline{y}(j) = \sum_{r=1}^{p} x_r(j)\underline{k}_r + \underline{w}(j), \tag{2.1}$$

where

$$x_r(j) \geq 0, \quad r = 1, \ldots, p, \quad \text{and} \quad \sum_{r=1}^{p} x_r(j) = 1.$$

The two conditions in this constraint are referred to as **abundance non-negativity constraint** (ANC) and **abundance sum-to-one constraint** (ASC) in the literature; see, e.g. [23]. Setting $\underline{x}(j) := (x_1(j), x_2(j), \ldots, x_p(j))^\top \in \mathbb{R}^p$, both conditions can jointly be written as **simplex constraint**

$$\underline{x}(j) \in \triangle_p, \qquad \triangle_p := \{v \in \mathbb{R}^p \colon \langle v, \mathbb{1} \rangle = 1,\, v \geq 0\}. \tag{2.2}$$

Using the vectors in Equation (2.1) for $j = 1, \dots, N$ as columns of corresponding matrices

$$\begin{aligned} Y &:= \big(\underline{y}(1)\ \underline{y}(2)\ \cdots\ \underline{y}(N)\big) \in \mathbb{R}^{L,N} && \text{(data matrix)},\\ X &:= \big(\underline{x}(1)\ \underline{x}(2)\ \cdots\ \underline{x}(N)\big) \in \mathbb{R}^{p,N} && \text{(abundance matrix)},\\ K &:= (\underline{k}_1\ \underline{k}_2\ \dots\ \underline{k}_p) \in \mathbb{R}^{L,p} && \text{(endmember matrix)},\\ W &:= \big(\underline{w}(1)\ \underline{w}(2)\ \cdots\ \underline{w}(N)\big) \in \mathbb{R}^{L,N} && \text{(noise matrix)} \end{aligned}$$

the Linear Mixing Model becomes

$$Y = KX + W, \qquad X \in (\triangle_p)^N. \tag{2.3}$$

2.1.2. Contribution

Now we make the additional assumption that only part of the entries of the data matrix are known; this can be specified by a mask $M \in \{0,1\}^{L\times N}$ of known data, where

$$M_{ij} = \begin{cases} 1 & \text{if } Y_{i,j} \text{ is known},\\ 0 & \text{otherwise.} \end{cases}$$

Then, restricting to the known entries, the Linear Mixing Model (2.3) with missing pixels reads

$$M \circ Y = M \circ (KX + W), \qquad X \in (\triangle_p)^N, \tag{2.4}$$

where the **Hadamard matrix product** is defined by $(A \circ B)_{i,j} = A_{i,j}B_{i,j}$. Hyperspectral unmixing refers to the estimation of the abundance matrix X from the given data Y; the endmember matrix K is either known or has to be found, possibly from a larger dictionary. We assume that K is known. This is a realistic assumption, because in many applications, such as sorting of materials, manufacturing control or remote sensing, reference spectra are available (see, for example, [72] or glue coverage control in [6]).

In this chapter, published in [81], we propose a novel joint unmixing and inpainting approach. Unmixing directly the incomplete hypercube, our approach replaces the two-step procedure of inpainting followed by unmixing, where the inpainting step typically introduces artefacts. These artefacts are avoided in the joint approach, because the knowledge of the signal subspace is used from the beginning, and we successfully recover the abundances knowing only a small part of the data cube.

Our presentation is organized as follows: In Section 2.2, we introduce our novel variational unmixing model based on the Linear Mixing Model with missing pixels and encourage spatial regularity by an edge-preserving discrete total variation regularizer. The minimizer of the corresponding functional is computed by primal-dual optimization methods in Section 2.3. Both on real and simulated data, our approach yields very good results, as can be seen in Section 2.4.

2.1.3. Related Work

The incompleteness of hyperspectral data cubes has different reasons and comes in different forms. In the **first "traditional inpainting" case**, the full spectrum has been measured correctly at most object locations, while at the other locations, nothing is known. A recent inpainting approach in this setting is [77, 78], where missing lines from airborne push-broom scanning are inpainted by anisotropic diffusion. As seems inherent to the PDE approach, the information from other channels enters the restoration of some channel —apart from a normalization step— only in the strength coefficient for the diffusion. Another approach, suggested for example in the survey [23] by Bioucas-Dias et al., is to perform spectral unmixing first and then perform inpainting on the lower dimensional signal subspace. This reduces the computational cost, and the noise level is lower after the dimension reduction. Chen's survey [31] combines PCA with diffusion or TV inpainting methods, successfully inpainting small missing regions.

In the **second case**, there are few object locations where all of the information is lost, but for many object pixels, only part of the spectral information is available. A typical application are hypercubes acquired by so-called push-broom scanning or line cameras with missing information due to faulty sensor pixels in the area sensor. One can attempt to use the methods from the traditional inpainting case. Furthermore, channel-wise inpainting works well if only a few neighbouring lines of pixels are missing. If, however, missing regions become large, its performance degrades. There are inherent problems to channel-wise methods: The affected region may contain spectra which are not in the boundary of the missing region; it may contain intensities smaller than the minimum or larger than the maximum occurring on the boundary of the missing region in that channel, or the missing region may have contour lines not meeting the boundary. If we are in one of these settings, channel-wise diffusion-based inpainting cannot be successful.

In the context of accelerated hyperspectral image acquisition, the authors of [40] try to recover the full cube from a partially-measured cube with a DCT sparsity prior on the spectra. As spatial regularization, they impose spatial sparsity w.r.t. a shift-invariant redundant wavelet basis.

So far, all approaches concentrate only on inpainting of corrupted hyperspectral images. To obtain the abundances, a second unmixing step would be required. An attempt for a joint unmixing and inpainting using beta priors for the coefficient vectors is [104]. They start from a larger endmember dictionary found by other methods (using Hysime [22]) and encourage spatial redundancy by learning a basis of atom cubelets for the abundance cube.

2.2. Mathematical Model

In this section, we derive our unmixing model from incomplete data, combining spatial regularization with the masked unmixing data term. The model allows for the inpainting of missing data during the unmixing.

2.2.1. Unmixing Model with Spatial Regularization and Proposed Model

We start from Model (2.4) and add spatial regularization. If the noise W is independently Gaussian, the maximization of likelihood in Equation (2.4) leads us to seek X as

$$\underset{X\in(\triangle_p)^N}{\operatorname{argmin}} \|M \circ (Y - KX)\|_F^2, \tag{2.5}$$

where for a matrix Y, the square of its Frobenius norm is $\|Y\|_F^2 := \operatorname{tr}(YY^\top)$.

In real-world images, neighbouring pixels are likely to contain the same mixture of materials, i.e., the same abundances. Therefore, the recovery of abundances can be improved by adding a spatial regularity prior, penalizing too much variation of the abundances. This is of particular importance for noisy data. Hence, we now introduce a discrete total variation (TV) functional, see [94], which has become one of the most widely-used regularizers in image restoration models because it preserves sharp edges. We define the discrete gradient operator, as in [97], by

$$\nabla := \begin{pmatrix} D_x \\ D_y \end{pmatrix} \quad \text{with} \quad D_x := D_n \otimes I_m \quad \text{and} \quad D_y := I_n \otimes D_m.$$

Here, I_n denotes the $n \times n$ identity matrix; the **Kronecker product** of two matrices $A \in \mathbb{R}^{\alpha_1\times\alpha_2}$ and $B \in \mathbb{R}^{\beta_1\times\beta_2}$ is given by the matrix

$$A \otimes B := \begin{pmatrix} a_{11}B & \cdots & a_{1\alpha_2}B \\ \vdots & \ddots & \vdots \\ a_{\alpha_1 1}B & \cdots & a_{\alpha_1\alpha_2}B \end{pmatrix} \in \mathbb{R}^{\alpha_1\beta_1\times\alpha_2\beta_2};$$

and

$$D_n := \begin{pmatrix} -1 & 1 & & & 0 \\ & -1 & 1 & & \\ & & \ddots & \ddots & \\ & & & -1 & 1 \\ 0 & & & & 0 \end{pmatrix} \in \mathbb{R}^{n\times n}$$

denotes the forward difference matrix. Note that the zero row corresponds to mirrored boundary conditions. For an image $F \in \mathbb{R}^{m\times n}$, column-wise reshaped into a vector $f \in \mathbb{R}^N$, let

$$\mathrm{TV}(F) := \mathrm{TV}(f) := \| \, |\nabla f| \, \|_1 := \sum_{i=1}^{N} \sqrt{(D_x f)_i^2 + (D_y f)_i^2}.$$

Throughout this chapter, $|\cdot|$ operates on a matrix by replacing, at each pixel, the vector of partial derivatives by its euclidean norm.

Recall that the reshaped abundance image of the rth material, is contained in the rth row $\bar{x}_r$ of X, i.e.,

$$X = \begin{pmatrix} x_1(1) & \dots & x_1(N) \\ \vdots & \ddots & \vdots \\ x_p(1) & \dots & x_p(N) \end{pmatrix} = \begin{pmatrix} \bar{x}_1 \\ \vdots \\ \bar{x}_p \end{pmatrix} = \begin{pmatrix} \underline{x}(1) & \dots & \underline{x}(N) \end{pmatrix}.$$

Summing over all materials, we add the regularization term $\sum_{r=1}^{p} \mathrm{TV}(\bar{x}_r^\top)$ to Equation (2.5), and arrive at the following *Unmixing Model for Incomplete Data:*

$$\operatorname*{argmin}_X \frac{1}{2}\|M \circ (Y - KX)\|_F^2 + \frac{\nu}{2}\|X\|_F^2 + \lambda \sum_{r=1}^{p} \mathrm{TV}(\bar{x}_r^\top) \quad \text{subject to} \quad X \in (\triangle_p)^N, \tag{2.6}$$

for regularization parameters $\nu \geq 0$ and $\lambda > 0$. This corresponds to the isotropic spatial TV-regularization with no coupling between the channels. As already mentioned, the simplex constraint is also known as ANC + ASC.

Remark 2.1 *A related functional has been investigated in [59]. Their model reads, for $\kappa \geq 0$ and $\lambda \geq 0$,*

$$\operatorname*{argmin}_X \frac{1}{2}\|Y - KX\|_F^2 + \kappa\|X\|_{1,1} + \lambda \sum_{r=1}^{p} \mathrm{TV}_{\mathrm{aniso}}(\bar{x}_r^\top) \quad \text{subject to} \quad X \geq 0,$$

where $\|X\|_{1,1} := \sum_{j=1}^{N} \|\underline{x}(j)\|_1$ and

$$\mathrm{TV}_{\mathrm{aniso}}(f) := \|D_x f\|_1 + \|D_y f\|_1 = \sum_{i=1}^{N} \Big(|(D_x f)_i| + |(D_y f)_i| \Big).$$

In contrast to our functional, this functional contains no mask, uses anisotropic TV*, and the additional term $\|X\|_{1,1}$ encourages sparsity. In [59], sparsity is important, because the authors use a larger K with columns corresponding to spectra from a library. In the presence of the ASC, as in our model, this term is constant,*

$$\|X\|_{1,1} = N \qquad \text{for all } X \in (\triangle_p)^N.$$

We have also performed experiments with anisotropic TV*, i.e., replacing* TV *in equation* (2.6) *with* $\mathrm{TV}_{\mathrm{aniso}}$*, and we comment on this briefly in Section 2.4.2. If not mentioned otherwise, all of our results are obtained for isotropic* TV*.*

Lemma 2.2 *The proposed model* (2.6) *is a convex function of X. If $\nu > 0$, then it is strictly convex in X, and therefore has a unique minimizer in $X \in (\triangle_p)^N$.*

Proof. The first term, $\sum_{ij} (m_{ij})^2 \big(y_i(j) - \sum_{r=1}^{p} K_{ir} x_r(j)\big)^2$ is convex, and for each $r \in \{1, \ldots, p\}$, $\mathrm{TV}(\bar{x}_r^\top)$ is a convex function of $\bar{x}_r^\top$. Therefore, the functional without the term $\frac{\nu}{2}\|X\|_F^2$ is convex. The term $\frac{1}{2}\|X\|_F^2$ is strictly convex; thus, for $\nu > 0$, the functional is strictly convex.

The set $\triangle_p$ and, hence, also the set $(\triangle_p)^N$, which incorporates the ANC and ASC constraints is closed, bounded and convex. The existence of a minimizer is ensured because a continuous function on a compact set attains its minimum. The minimizer is unique for $\nu > 0$, because a strictly convex function has at most one minimizer on a convex set. □

2.2.2. Model Reformulation

In this section, we rewrite the proposed model (2.6) in a sound matrix-vector form to better apply algorithms from convex analysis. Reshaping the matrices in Equation (2.4) appropriately into vectors, i.e., the abundance matrix $X = \begin{pmatrix}\underline{x}(1) & \dots & \underline{x}(N)\end{pmatrix}$, the data matrix $Y = \begin{pmatrix}\underline{y}(1) & \dots & \underline{y}(N)\end{pmatrix} \in \mathbb{R}^{L\times N}$ and the mask $M \in \{0,\,1\}^{L\times N}$ into

$$x := \operatorname{vec} X := \begin{pmatrix}\underline{x}(1)\\ \vdots \\ \underline{x}(N)\end{pmatrix} \in \mathbb{R}^{Np}, \quad y := \begin{pmatrix}\underline{y}(1)\\ \vdots \\ \underline{y}(N)\end{pmatrix} \in \mathbb{R}^{NL}, \quad m := \operatorname{vec} M \in \{0,\,1\}^{NL},$$

we can write the model with related Kronecker products. We make repeated use of the relation

$$\operatorname{vec}(ACB^\top) = (B \otimes A)\ \operatorname{vec}(C). \tag{2.7}$$

First note that

$$\|Y - KX\|_F^2 = \sum_{j=1}^{N} \|\underline{y}(j) - K\underline{x}(j)\|_2^2 = \|y - (I_N \otimes K)x\|_2^2.$$

Next, from $\operatorname{vec} X\nabla^\top = (\nabla \otimes I_p)x$, we have

$$\sum_{k=1}^{p} \mathrm{TV}(\bar{x}_r^\top) = \| \, |\nabla X^\top| \, \|_1 = \| \, |(\nabla \otimes I_p)x| \, \|_1.$$

Inserting these equations in (2.6), and writing the constraint with the help of an indicator function

$$\iota_s(x) := \begin{cases} 0 & \text{if } x \in S, \\ +\infty & \text{otherwise,}\end{cases}$$

we obtain the vectorized **Unmixing Model for Incomplete Data:**

$$\operatorname*{argmin}_x \frac{1}{2}\|m \circ (y - (I_N \otimes K)x)\|_2^2 + \frac{\nu}{2}\|x\|_2^2 + \lambda\| \, |(\nabla \otimes I_p)x| \, \|_1 + \iota_{(\triangle_p)^N}(x). \tag{2.6'}$$

This can equivalently be written as

$$\operatorname*{argmin}_{x,u,v} \frac{1}{2}\|m \circ (y - u)\|_2^2 + \frac{\nu}{2}\|x\|_2^2 + \lambda\| \, |v| \, \|_1 + \iota_{(\triangle_p)^N}(x) \quad s.t. \quad \begin{pmatrix} I_N \otimes K \\ \nabla \otimes I_p \end{pmatrix} x = \begin{pmatrix} u \\ v \end{pmatrix}. \tag{2.6''}$$

2.3. Algorithm

In this section we describe a fast first order algorithm to minimize (2.6"). We briefly recall the Primal-Dual Hybrid Gradient Algorithm (PDHGMp). For $f_1 \in \Gamma_0(\mathbb{R}^N)$, $f_2 \in \Gamma_0(\mathbb{R}^M)$ and $C \in \mathbb{R}^{M\times N}$ the solution of

$$\operatorname*{argmin}_{x,y} f_1(x) + f_2(y) \quad \text{subject to} \quad Cx = y \tag{2.8}$$

can be computed by the PDHGMp Algorithm 1, assuming that a saddle point of the Lagrangian $L(x, y, p) := f_1(x) + f_2(y) + \langle p, Cx - y\rangle$ exists see [29, 90].

Algorithm 1 PDHGMp for problems in standard form (2.8)

Initialization: $\mu > 0$. $\sigma > 0$ with $\mu\sigma < 1/\|C\|^2$, $\theta \in (0, 1]$
$x^{(0)}$, $p^{(0)} = \bar{p}^{(0)}$.
Iterations:
For $r = 0, 1, \ldots$

1. $x^{(r+1)} = \operatorname{argmin}_{x\in\mathbb{R}^N} \; f_1(x) + \frac{1}{2\mu}\|x - \big(x^{(r)} - \mu\sigma C^\top \bar{p}^{(r)}\big)\|^2$
2. $y^{(r+1)} = \operatorname{argmin}_{y\in\mathbb{R}^M} f_2(y) + \frac{\sigma}{2}\|y - (p^{(r)} + Cx^{(r+1)})\|^2$
3. $p^{(r+1)} = p^{(r)} + Cx^{(r+1)} - y^{(r+1)}$
4. $\bar{p}^{(r+1)} = p^{(r+1)} + \theta(p^{(r+1)} - p^{(r)})$

The sequence $(x^{(r)}, p^{(r)})_{r\in\mathbb{N}}$ generated by Algorithm 1 then converges to the saddle point of the Lagrangian [28, Thm. 6.4].

Let us now come back to Model (2.6). We use the Primal-Dual Hybrid Gradient Algorithm (PDHGMp) to solve the equivalent formulation (2.6") by finding a saddle point of the associated Lagrangian

$$L(x, u, v;\; b_u, b_v) \;=\; \frac{1}{2}\|m \circ (y - u)\|_2^2 + \frac{\nu}{2}\|x\|_2^2 + \lambda\| \, |v| \, \|_1 + \iota_{(\triangle_p)^N}(x) + \left\langle \begin{pmatrix} b_u \\ b_v \end{pmatrix} \;,\; \begin{pmatrix} I_N \otimes K \\ \nabla \otimes I_p \end{pmatrix} x - \begin{pmatrix} u \\ v \end{pmatrix} \right\rangle.$$

This leads to Algorithm 2, which is guaranteed to converge to a saddle point of the Lagrangian by [28, Theorem 6.4].

All steps can be computed very efficiently. In Step 1, we compute an orthogonal projection of $\widetilde{x}^{(r)}/(1+\nu\tau)$ to $(\triangle_p)^N$. In Step 2, for $v^{(r+1)}$, we compute a coupled-shrinkage of $(\nabla \otimes I_p)x^{(r+1)} + b_v^{(r)}$, with coupling within each pair of entries corresponding to the forward differences in both spatial directions at one pixel. The update $u^{(r+1)} = \operatorname{vec} U$ is given by

$$\begin{aligned} U[M{=}{=}0] &= (KX + B_u) \\ U[M{=}{=}1] &= \frac{1}{1+\sigma} \cdot Y + \frac{\sigma}{1+\sigma}\big(KX + B_u\big), \end{aligned}$$

where from the expressions on the right hand side only the elements corresponding to the index set on the left are accessed.

The "Kronecker product–vector" multiplications in Algorithm 2 are implemented in "matrix–matrix" form, using relation (2.7); e.g., for $(\nabla \otimes I_p)x$, we compute $X\nabla^\top$.

The number of elementary operations required for one iteration of the algorithm is linear in the number of entries of the hypercube, for constant p. In more detail, each

iteration requires: N vector projections to $\triangle_p$, which are computable with $\mathcal{O}(p^2)$ operations each; the coupled shrinkage of Np pairs of entries, computable in constant time; further, $(2p+1)NL$ multiplications and the same order of additions.

Remark 2.3 *(Parameters and Initialization) As mentioned above, the term* $\frac{\nu}{2}\|X\|_F^2$ *ensures strict convexity of the model. We set* $\nu = 10^{-3}$ *in all experiments below. In practice, setting* $\nu = 0$ *works equally well, e.g., setting* $\nu = 0$ *changes the recovery percentages in Section 2.4.2, Table 2.1, by less than* 0.1%*, where more than* 0.3% *of the pixels are known.*

The parameters τ *and* σ *influence only the convergence speed of the algorithm. Here, we fix them both to* $\left[(\max_j \sum_i K_{ij} + 4)(\max_i \sum_j K_{ij})\right]^{-1/2}$*, which ensures the required bound on the product* $\tau\sigma$*; they could also be chosen adaptively [50].*

The remaining regularization parameter λ *is chosen heuristically taking the strength of the noise and the size of the image features into account.*

As initialization for X*, we start with a uniformly random matrix and then normalize, so that the columns sum to one. Finally, we always use zero initialization for* $u^{(0)}, v^{(0)}, b_u^{(0)}, b_v^{(0)}$*.*

Algorithm 2 Primal-dual hybrid gradient algorithm (PDHGMp) for Equation (2.6").

Initialization: $\tau, \sigma > 0$ with $\tau\sigma \le 1/\left\|\binom{I_N \otimes K}{\nabla \otimes I_p}\right\|_2^2$,

$x^{(0)} \in \mathbb{R}^{Np},\ u^{(0)} \in \mathbb{R}^{NL},\ v^{(0)} \in \mathbb{R}^{2Np}, \quad b_u^{(0)} = \bar{b}_u^{(0)} \in \mathbb{R}^{NL}, \quad b_v^{(0)} = \bar{b}_v^{(0)} \in \mathbb{R}^{2Np}.$

Iterations:

For $r = 0, 1, \ldots$

1. $$\widetilde{x}^{(r)} := x^{(r)} - \tau\sigma\left(I_N \otimes K^\top \,\middle|\, \nabla^\top \otimes I_p\right)\begin{pmatrix}\bar{b}_u^{(r)}\\ \bar{b}_v^{(r)}\end{pmatrix}$$
$$x^{(r+1)} := \operatorname*{argmin}_x \left\{ \iota_{(\triangle_p)^N}(x) + \frac{\nu}{2}\|x\|_2^2 + \frac{1}{2\tau}\left\|x - \widetilde{x}^{(r)}\right\|_2^2 \right\}$$

2. $$u^{(r+1)} := \operatorname*{argmin}_u \left\{ \frac{1}{2}\|m \circ (y-u)\|_2^2 + \frac{\sigma}{2}\|(I_N \otimes K)x^{(r+1)} - u + b_u^{(r)}\|_2^2 \right\}$$
$$v^{(r+1)} := \operatorname*{argmin}_v \left\{ \lambda\| \,|v|\, \|_1 + \frac{\sigma}{2}\|(\nabla \otimes I_p)x^{(r+1)} - v + b_v^{(r)}\|_2^2 \right\}$$

3. $$b_u^{(r+1)} := b_u^{(r)} + (I_N \otimes K)x^{(r+1)} - u^{(r+1)}$$
$$b_v^{(r+1)} := b_v^{(r)} + (\nabla \otimes I_p)x^{(r+1)} - v^{(r+1)}$$

4. $$\bar{b}_u^{(r+1)} := 2b_u^{(r+1)} - b_u^{(r)}$$
$$\bar{b}_v^{(r+1)} := 2b_v^{(r+1)} - b_v^{(r)}$$

2.4. Numerical Results

We give results on real data acquired with a hyperspectral line camera in Section 2.4.1 and on simulated data in Sections 2.4.2–2.4.4.

Recall the hyperspectral image acquisition by a hyperspectral line camera from Section 1.1.1. The snapshot of the area sensor taken for the line $y = y_j$ of the object yields a section $Y^{\text{cube}}_{:,y_j,:}$ of the hypercube $Y^{\text{cube}} \in \mathbb{R}^{m,n,L}$ as depicted in Figure 2.1 (right), where the z-axis is the spectral direction. We call such a section a sensor frame.

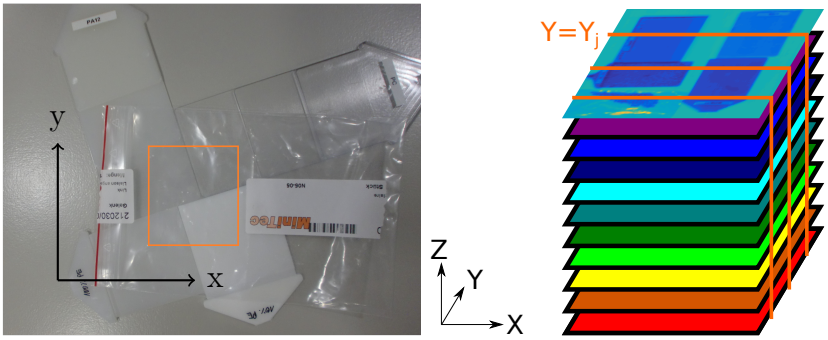

Figure 2.1.: Measured object region (**left**) and sensor frames $y = y_j$ of the hypercube measured each in one sensor snapshot (**right**); they correspond to one line of the object.

The manufacturing process of the sensor commonly produces some defect pixels. For push-broom scanning, each sensor frame $Y^{\text{cube}}_{:,y_j,:}$ has the same pattern of missing pixels. One missing pixel thus creates a missing line, and a cluster of missing pixels creates a cylinder of missing entries in the hypercube.

2.4.1. Numerical Results for Real Data

In this section, we present results for a hyperspectral image measured at the Fraunhofer Institute for Industrial Mathematics ITWM with the line camera shown in Figure 1.3, comparing the restoration obtained as a by-product of our unmixing to a traditional inpainting.

We have measured the marked region of the polymer samples[1] in Figure 2.1 (left), in the wavelength range 1073–2300 nm, at a spectral resolution of $L = 256$ channels. The four regions contain plastic mixtures with different spectral signatures, and the assembly has been covered with a plastic foil.

Figure 2.2c shows the mask of working pixels with a few circular regions of corruption artificially added, marked in violet as "simulated defects". The small circle marked in red is a real defect, and a camera with such a defect is sold at a considerably reduced price.

[1] kindly provided by the Süddeutsche Kunststoff-Zentrum (SKZ)

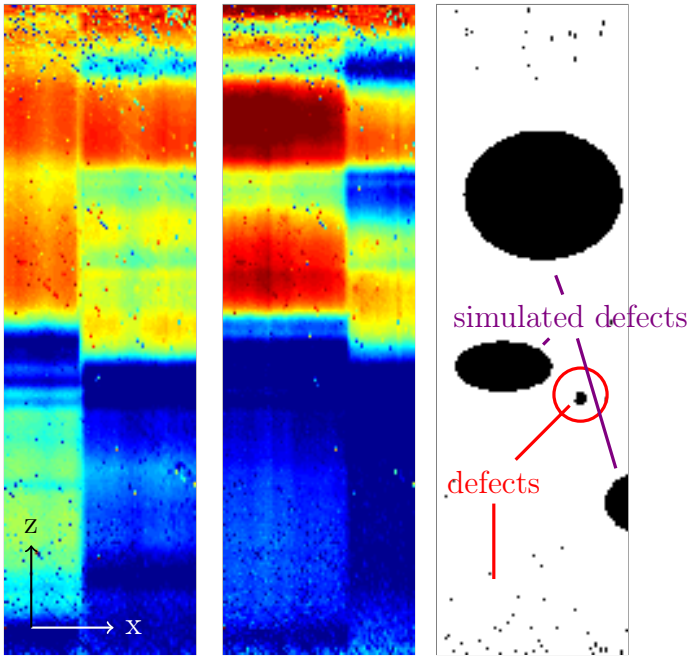

(a) $y = 30$ (b) $y = 60$ (c) mask

Figure 2.2.: (**a**,**b**) Two sensor frames $y = y_0$ with the spectral direction along the z-axis and (**c**) the mask of working sensor pixels.

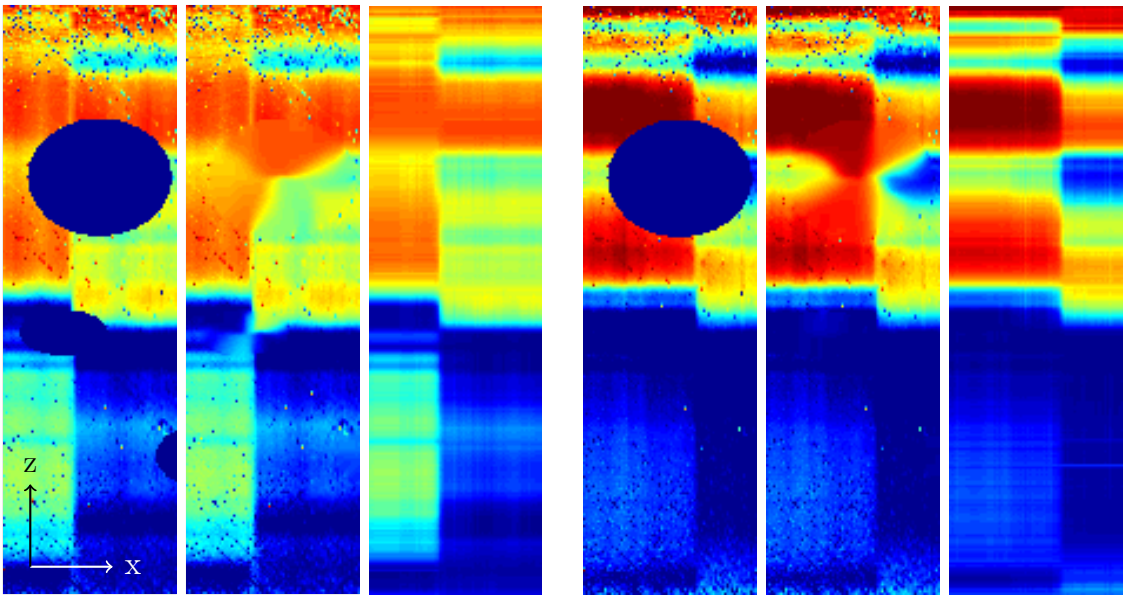

(**a**) input (**b**) NS (**c**) ours (**d**) input (**e**) NS (**f**) ours

Figure 2.3.: Sensor frames of masked noisy original input and after inpainting by Navier–Stokes (NS) and our Model (2.6), respectively; (**a**–**c**): sensor frame $y = 30$, and (**d**–**f**): $y = 60$.

Figure 2.2a,b shows two sensor frames $y = 30, 60$ of the measured hypercube. Being snapshots of the area sensor, they have the same pattern of missing pixels. Figure 2.3a,d shows the same sensor frames with the artificially added circular defects.

Slices $Y^{\text{cube}}_{:,:,z}$, $z = 1, 2, \ldots, L$ of the hypercube corresponding to a particular wavelength are called channels. Figure 2.4a shows Channels $10, 70, 90$. Here, Channel 10 is noisy and affected by individual broken sensor pixels, each creating one missing line. Figure 2.4b shows the same channels, after the artificial defects have been added, i.e., masked, and while Channel 10 stays unchanged, we see that not much remains of Channels 70 and 90.

For the unmixing, we give to Algorithm 2 the hypercube together with the mask and the four pure spectra present in the scene, i.e., $p = 4$, as columns of K. These have been obtained from averaging spectra over manually selected rectangles and are shown in Figure 2.5a.

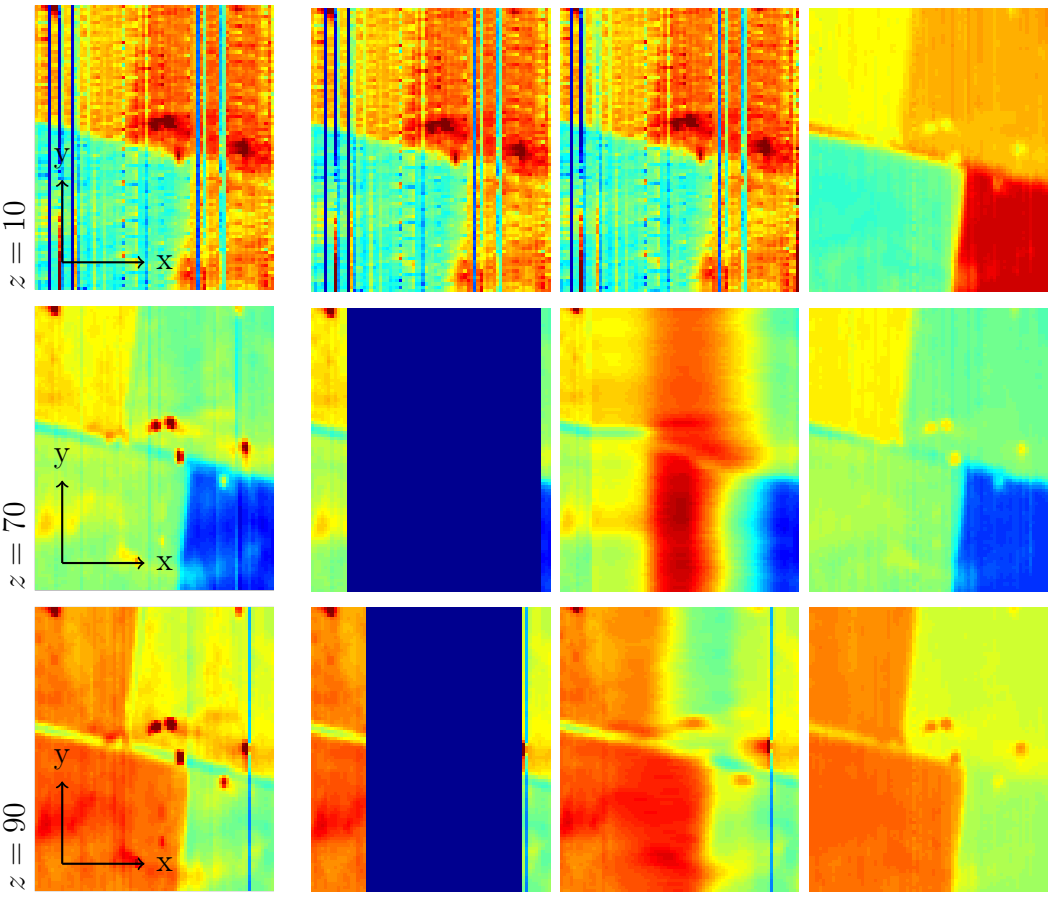

(**a**) noisy original (**b**) input (**c**) NS (**d**) ours

Figure 2.4.: Channels $10, 70, 90$ of: (**a**) the noisy original hypercube; (**b**) the masked original known to the algorithm; (**c**) the restoration by Navier–Stokes; and (**d**) the restoration by our method.

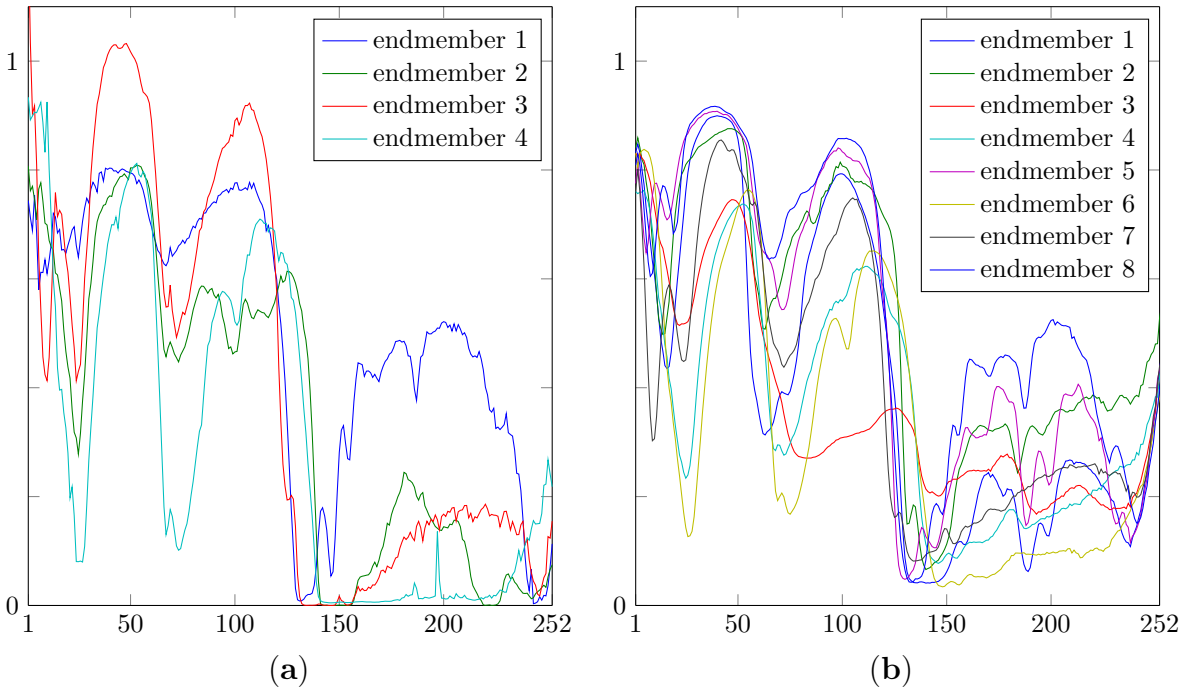

Figure 2.5.: (**a**) The four endmember spectra ($p = 4$) corresponding to each of the plastics blends in Figure 2.1 (left); (**b**) the eight endmember spectra ($p = 8$) used in Section 2.4.4.

For the comparison, we take the following approach: Starting from the unmixing coefficients X obtained by minimizing Equation (2.6) with Algorithm 2, we form the product KX, which is an approximating restoration of the data matrix Y. After reshaping, we compare this restoration to an inpainting of the hypercube Y^{cube} by [20].

We have chosen the Navier–Stokes-based inpainting [20] as representative of neighbourhood-based inpainting methods. The Navier-Stokes inpainting is performed in each x-z-plane of the hypercube and estimates missing data looking at surrounding pixels in that plane. In contrast to the proposed method, such neighbourhood-based inpainting of the data cube lacks the capacity of utilizing the provided information about pure spectra. Runtime is several minutes.

For Algorithm 2, we used the parameters from Remark 2.3 and $\lambda = 0.01$. The relative primal step $\|x^{(r+1)} - x^{(r)}\|_2 / \|x^{(r)}\|_2$ fell below 10^{-3} after 111 iterations, which took 7 s on an Intel Core i7 with 2.93 GHz. The graphics shown are after 1000 iterations.

In Figures 2.3 and 2.4, we compare the performance of Algorithm 2 to a Navier–Stokes inpainting of each x-z-plane of the hypercube.

For larger clusters of defect pixels, the inpainting of the hypercube by Navier–Stokes is not satisfactory: looking at Figure 2.3, the large masked region could still be guessed from the images inpainted with Navier–Stokes. On the other hand, the inpainted sensor images obtained from our method in Figure 2.3c,f agree well with the measured sensor frames in Figure 2.2a,b, also removing some noise, which was not contained in the mask.

In Figure 2.4, we see that the broader missing stripes in Channels 70 and 90 cannot be

restored by Navier–Stokes inpainting and, hence, would introduce errors to any following unmixing step, whereas our joint unmixing remains unaffected and gives a satisfactory (denoised) restoration of the original, because the information from all intact channels is being used.

2.4.2. Numerical Results for Artificial Data (Pure Regions)

We have seen that our unmixing model performs well even if large parts of the hypercube are missing. To analyze which percentage of sensor pixels can be missing, i.e., to quantify the unmixing results in some way, in this section, we use an artificial input image with a small level of noise added.

The artificial image is constructed as follows. The jth of the four regions of the piecewise constant 240×148 ground truth image in Figure 2.6 (left), is filled with copies of the jth endmember spectrum. The four endmember spectra, which form the $p = 4$ columns of the matrix K, are shown in Figure 2.6 (right). The spectra come from our measurements of common plastic blends. The resulting hypercube Y^{pure} belongs to $\mathbb{R}^{240\times148\times256}$. We add independent zero-mean Gaussian noise to each entry $Y^{\text{pure}}_{i,j,r}$, obtaining the hypercube Y^{low}. Here, the standard deviation was taken slightly larger than 1% of the maximal entry of Y^{pure}.

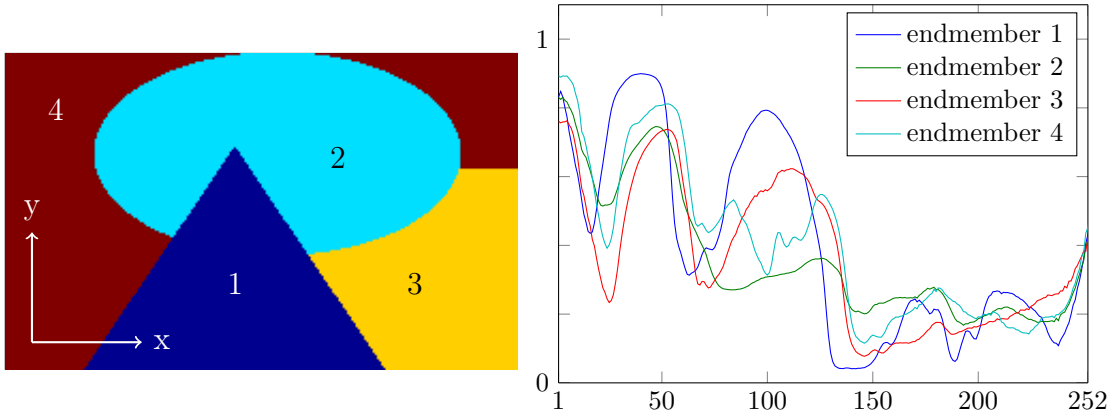

Figure 2.6.: The artificial image is constructed by filling region j of the image on the left with the jth endmember spectrum on the right ($p = 4$).

For several percentages from 100% down to 0.1%, we randomly generate sensor masks of size 240×256 with approximately this percentage of working pixels, by flipping a biased coin for each pixel. Figure 2.7a shows the upper left quarter of the sensor mask for 3% working pixels. Note that the percentage of known sensor pixels and the percentage of entries of Y^{cube} which are known to the unmixing algorithm, are the same, because each masked sensor pixel corresponds to one line in the hypercube along the y-axis.

For each percentage and corresponding mask applied to Y^{low}, we find the minimizer X of equation (2.6) by Algorithm 2 with $\lambda = 0.1$ and other parameters as in Remark 2.3; running 300 iterations took an average of 110 s per image. Then, we reshape X into the cube X^{low} of unmixing coefficients. To quantify the quality of X^{low}, we assign to each

image pixel (i, j) the endmember corresponding to the largest of the four unmixing coefficients $X^{\text{low}}_{ij1}, \ldots, X^{\text{low}}_{ij4}$ at that image pixel. Some of the resulting label images are shown in Figure 2.7b. Table 2.1 lists the percentages of correctly assigned pixels depending on the percentage of known sensor pixels. For more than 10% working sensor pixels, the algorithm found the largest material abundance at the correct material for 100% of the image pixels.

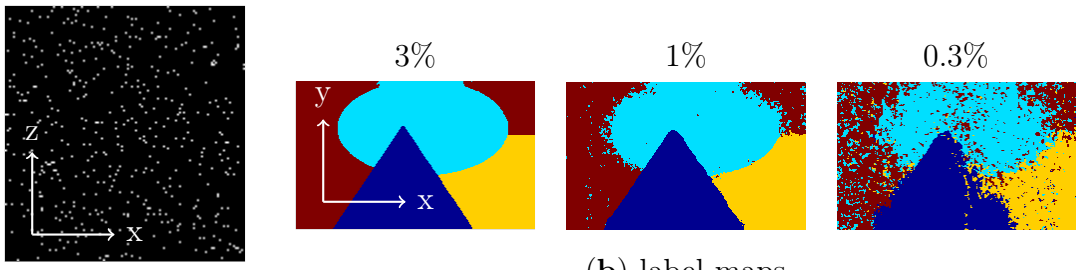

(**a**) sensor 3%

(**b**) label maps

Figure 2.7.: (**a**) Upper left corner of the 240 × 256 mask for 3% working sensor pixels; (**b**) label maps obtained from the unmixing by assigning to pixel (i, j) the index r of the largest coefficient $X^{\text{low}}_{ijr} \in \{X^{\text{low}}_{ij1}, \ldots, X^{\text{low}}_{ij4}\}$ at that location.

Furthermore, in Table 2.1, we give the results from minimizing the model obtained by replacing isotropic TV in model (2.6) with anisotropic TV, as introduced in Remark 2.1. The results are slightly worse, though not much.

sensor pixels known, in %	30	10	3	1	0.3	0.1
correctly assigned image pixels using TV, in %	100	100	99.5	96.3	83.9	54.1
correctly assigned image pixels using TV_{aniso}, in %	100	100	99.4	95.5	82.8	52.7

Table 2.1.: Percentage of image pixels, for which the largest of the material abundances found by the unmixing, correctly identifies the material at that pixel.

2.4.3. Numerical Results for Artificial Data (Mixed Regions)

Next, we construct a test image comprising patches of linear mixtures of the endmembers, to meet the real-world scenario of mixed pixels. We take again the four endmembers shown in Figure 2.6, hence again $p = 4$. As shown in the false colour visualization of the ground truth in Figure 2.8 (left), the corners are filled with pure spectra, and along the sides of the image, we form linear mixtures. Then, a larger amount of noise, with a standard deviation of about 10% of the maximum, is added to the hypercube, and only about 10% of the pixels are retained. As in the previous section, we do this by first sampling a sensor mask under the assumption that each sensor pixel works with a probability of 10% and then simulating the line camera measurement. In Figure 2.8, we

plot Channel 40 of the hypercube, in the middle before masking, and on the right, as known to the algorithm.

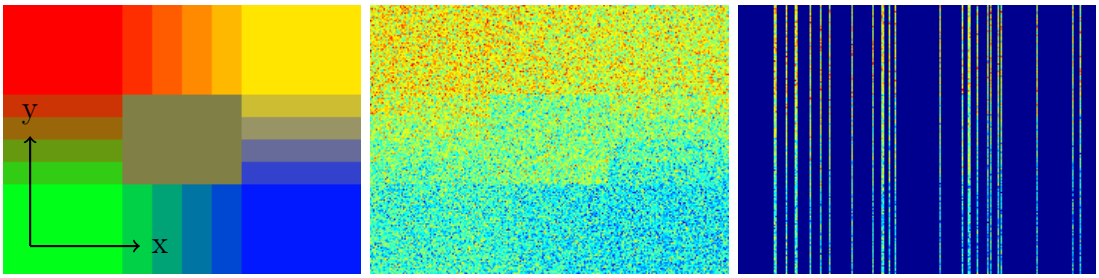

Figure 2.8.: Ground truth in false colours (left); Channel 40 of the noisy hypercube (middle); Channel 40 noisy and masked (right); here, 10% are known.

Figure 2.9 shows the obtained abundances, for parameters as in Remark 2.3 and $\lambda = 0.3$ after 918 iterations, which took 447 s on an Intel Core i7 with 2.93 GHz. The result is visually identical to the ground truth.

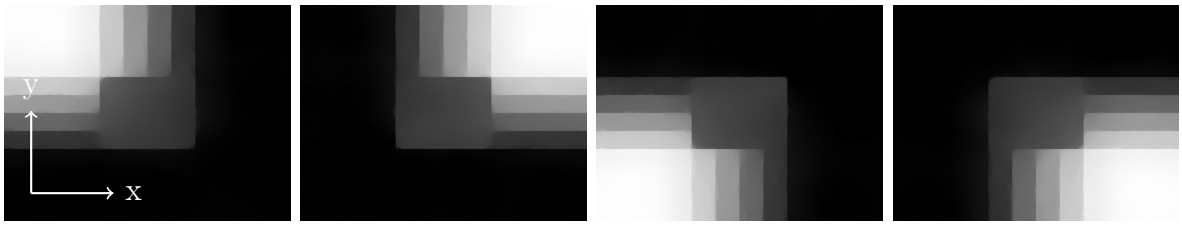

Figure 2.9.: Unmixing result, knowing 10% of the noisy hypercube. The images show the estimated abundances $X^{\text{cube}}_{ij1}, \ldots, X^{\text{cube}}_{ij4}$ corresponding to the four pure spectra.

2.4.4. Numerical Results for Non-Occurring Endmembers

In applications, some of the expected materials might not be present in the image. We proceed as in the previous Section 2.4.3, except that the endmember matrix K now contains the eight spectra shown in Figure 2.5b ($p = 8$), of which four are chosen and mixed using the same abundances as in Section 2.4.3. The abundances obtained from unmixing 10% of the hypercube are shown in Figure 2.10.

Again, the true abundances are perfectly recovered.

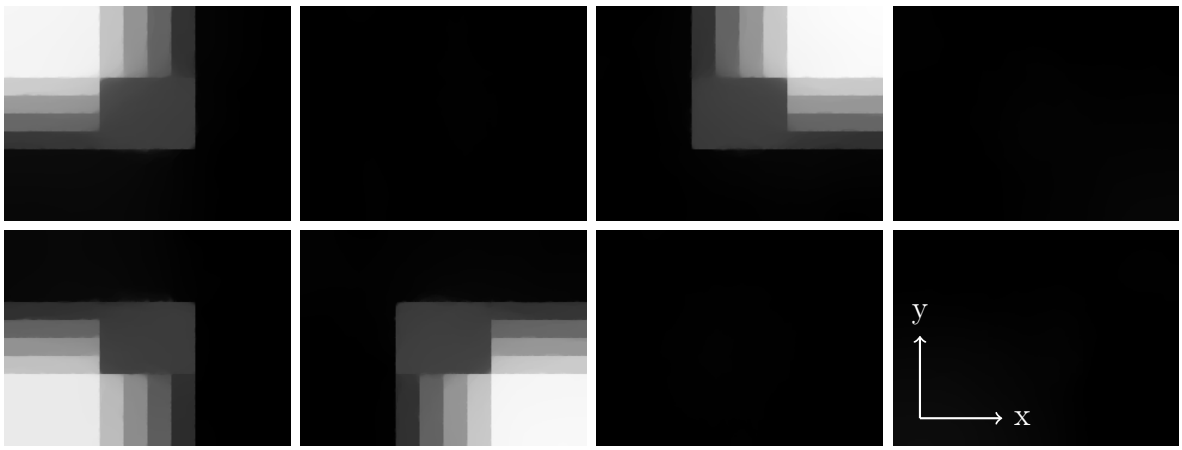

Figure 2.10.: Unmixing result with eight endmembers of which four occur in the image; with 10% of the noisy hypercube known. The images show the estimated abundances $X_{ij1}^{\text{cube}}, \ldots, X_{ij8}^{\text{cube}}$ corresponding to the eight pure spectra.

2.5. Conclusions

In this chapter, we have introduced a novel model for hyperspectral unmixing from incomplete and noisy data, provided that an estimate of the signal subspace is available. The model applies if an arbitrary part of the hypercube is known, e.g., a random selection of entries.

The model allows unmixing in one step from incomplete data cubes, having larger regions of missing data than could be restored by preprocessing methods unaware of the signal subspace.

We demonstrated results from line cameras, for which the mask of known entries is structured: lines along the second spatial direction of the hypercube are either fully known or fully unknown, which leads to a more difficult inpainting problem compared to, say, a random distribution of known entries in the hypercube. For large missing regions, where traditional Navier–Stokes inpainting failed, the data term approximation, obtained by our unmixing, provided a good inpainting.

We simulated artificial data with this special structure of the mask. For an image composed of pure materials, knowing only 3% of the sensor pixels in the simulated measurement, the rounded unmixing is a 99.5% correct assignment of pixels to materials. For a mixed and noisy image, knowing 10% of the sensor pixels, we obtain a visually perfect recovery of the abundances.

Our results show on real data that unmixing results can be perfect in spite of missing regions in the hypercube, which are orders of magnitude larger than for current line cameras. This shows the potential for further applications, such as those arising from novel image acquisition techniques. That the variational model is simple and can be easily extended is certainly a further benefit.

3. Sparse Unmixing

Outline

In real world applications, we often do not know exactly, which materials occur in an image. Instead, we often need to choose suitable endmembers from a spectral library. In this section we are going to extend the model from the previous chapter, (2.6), to this setting.

3.1. Introduction and Proposed Model

Our new problem description reads: Given the spectral library $K_{\text{lib}} = K$, and the reshaped hypercube Y, select a subset K' of the columns of K_{lib}[1] and spatially regular abundance matrix X', such that that $Y \approx K_{\text{lib}}X \approx K'X'$, as seen through the mask M. Here X' can be thought of as being formed by those rows of the full abundance matrix X which correspond to the selected columns of K_{lib}.

Let us repeat: selecting a subset of the columns of K_{lib}, which should occur in the mixture, is equivalent to requiring that only the corresponding selection X' of rows of X contain non-zero entries. A regularization term that induces this **row-sparsity** on X is given by

$$\|X\|_{1,\infty} := \sum_{i=1}^{p} \max_j x_{i,j}.$$

Adding this term to (2.6), we obtain our **Row-Sparse Unmixing Model from In-**

[1] In this paragraph we write K_{lib} instead of K, as a reminder that K_{lib} contains more endmember spectra than occur in the image.

complete Data:

$$\underset{X}{\operatorname{argmin}} \left(\frac{1}{2}\|M \circ (Y - KX)\|_F^2 + \frac{\nu}{2}\|X\|_F^2 + \lambda \sum_{i=1}^{p} \mathrm{TV}(\bar{x}_i^\top) + \zeta \sum_{i=1}^{p} \max_j x_{i,j} \right) \quad \text{subject to} \quad X \in (\triangle_p)^N. \tag{3.1}$$

3.2. Related Work and Exact Relaxation

3.2.1. Exact Relaxation Property

In the setting of [43], the $l_{\infty,1}$-norm is an exact relaxation of the (row-0)-norm, which counts the number of non-zero rows of a matrix. We recall the result as a motivation in our setting and comment on the practical limitation. The following has been proved in [43, Lemma III.1].

Lemma 3.1 *If we remove repeated columns of X and have l_2-normalized data, the set of minimizers of*

$$\min_{T \geq 0} \|T\|_{\mathrm{row}-0} \quad \textit{such that } XT = X \tag{3.2}$$

and

$$\min_{T \geq 0} \|T\|_{1,\infty} \quad \textit{such that } XT = X \tag{3.3}$$

are the same.

Here we still assume the simplex constraint, which incorporates both the ASC and ANC, and says that the data spectra are convex combinations of the endmembers. We note that the above result still holds with l_1-normalization in place of l_2-normalization and extends to the case $Y = KX$, for a library K, as long as those columns of K which are endmembers, are also columns of Y, and both the columns of K and of Y are normalized.

Unfortunately the ASC part of the simplex constraint relies on library and image being calibrated to the same white reference. This calibration is destroyed if the library is normalized afterwards. We are therefore not going to normalize the library or data and the above can only serve as a motivation.

3.3. Reformulation of the Model

Before we apply the PDHGMp-Algorithm let us again rewrite the model in a sound matrix-vector form. We proceed as in Section 2.2.2 and comment only on the differences.

Reshaping again the matrices in (3.1) into vectors $x = \operatorname{vec} X \in \mathbb{R}^{pN}$, $y = \operatorname{vec} Y \in \mathbb{R}^{LN}$, $m = \operatorname{vec} M \in \{0,1\}^{LN}$ and incorporating the simplex constraint via the indicator function $\iota_{(\triangle_p)^N}(x)$ we see that (3.1) is equivalent to the **Vectorized Row-Sparse Unmixing Model for Incomplete Data**:

$$\operatorname*{argmin}_x \frac{1}{2}\|m \circ (y - (I_N \otimes K)x)\|_2^2 + \frac{\nu}{2}\|x\|_2^2 + \lambda\| \, |(\nabla \otimes I_p)x| \, \|_1 + \zeta \sum_i \max_j x_{i,j} + \iota_{(\triangle_p)^N}(x). \tag{3.1'}$$

or

$$\begin{aligned} \operatorname*{argmin}_{x,u,v} \frac{1}{2}\|m \circ (y - u)\|_2^2 + \frac{\nu}{2}\|x\|_2^2 + \lambda\| \, |v| \, \|_1 + \iota_{(\triangle_p)^N}(x) + \zeta \sum_i \max_j x_{i,j} \\ s.t. \quad \begin{pmatrix} I_N \otimes K \\ \nabla \otimes I_p \end{pmatrix} x = \begin{pmatrix} u \\ v \end{pmatrix}. \end{aligned} \tag{3.1''}$$

Note that the term containing x no longer decouples because $\iota_{(\triangle_p)^N}(x)$ is a sum of functions of columns of X and $\|X\|_{1,\infty} = \sum_i \max_j x_{i,j}$ is a sum of maxima along the rows of X. We therefore split again, inserting a copy z of x into the second term:

$$\begin{aligned} \operatorname*{argmin}_{x,u,v,z} \frac{1}{2}\|m \circ (y - u)\|_2^2 + \lambda\| \, |v| \, \|_1 + \big(\zeta\|z\|_{1,\infty} + \iota_{\geq 0}(z)\big) + \frac{\nu}{2}\|x\|_2^2 + \iota_{(\triangle_p)^N}(x) \\ \text{s.t.} \quad \begin{pmatrix} I_N \otimes K \\ \nabla \otimes I_p \\ I_{Np} \end{pmatrix} x = \begin{pmatrix} u \\ v \\ z \end{pmatrix}. \end{aligned} \tag{3.4}$$

3.3.1. Computation of the Row-Sparsity Proximation

In this subsection we explain how the proximum of the $l_{\infty,1}$-norm is computed via convex duality. This will give us the remaining ingredient for Algorithm 3.4 below.

We are going to rewrite the proximum in terms of a projection and to this end recall a standard result by Moreau, see [82]. A proof can be found for example in [16, Theorem 14.3].

Theorem 3.2 (Moreau decomposition) *Let $f\colon \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$ be a proper convex lsc function. The following decomposition holds:*

- $\operatorname{Prox}_f(x) + \operatorname{Prox}_{f^*}(x) = x$,
- $(f)_1(x) + (f^*)_1(x) = \frac{1}{2}\|x\|_2^2$.

Herein $(f)_1$ denotes the Moreau-Yosida regularization of f with parameter one.

In order to compute the update

$$z^{(r+1)} := \operatorname*{argmin}_z \left\{ \zeta\|z\|_{1,\infty} + \iota_{\geq 0}(z) + \frac{\sigma}{2}\left\|x^{(r+1)} - z + b_z^{(r)}\right\|_2^2 \right\},$$

it is convenient to denote by $Z, B_z^{(r)} \in \mathbb{R}^{p,N}$ the matrices before reshaping to $z = \operatorname{vec} Z$, $b_z^{(r)} = \operatorname{vec} B_z^{(r)}$. We observe that the minimization decouples into separate problems for

each row $\bar{z}_i$ of Z, $i = 1, \dots, p$, and use that the Legendre transformation (Fenchel dual) of $h_\zeta(z) := \iota_{\geq 0}(z) + \zeta\|z\|_\infty$ is the indicator function of the set

$$C_\zeta := \{z \in \mathbb{R}^N \colon \|\max(z, 0)\|_1 \leq \zeta\}.$$

By Theorem 3.2 we obtain, understanding $\operatorname{Prox}_{h_{\frac{\zeta}{\sigma}}}$ row-wise, that

$$\begin{aligned} Z^{(r+1)} &= \operatorname{Prox}_{h_{\frac{\zeta}{\sigma}}}\big(X^{(r+1)} + B_z^{(r)}\big) \\ &= \big(X^{(r+1)} + B_z^{(r)}\big) - \operatorname{Prox}_{(h_{\frac{\zeta}{\sigma}})^*}\big(X^{(r+1)} + B_z^{(r)}\big) \\ &= \big(X^{(r+1)} + B_z^{(r)}\big) - \mathcal{P}_{(C_{\frac{\zeta}{\sigma}})^p}\big(X^{(r+1)} + B_z^{(r)}\big), \end{aligned}$$

where $\mathcal{P}_{(C_{\frac{\zeta}{\sigma}})^p}$ orthogonally projects each row of $X^{(r+1)} + B_z^{(r)}$ onto $C_{\frac{\zeta}{\sigma}}$.

If we know that projections commute with scaling, then

$$Z^{(r+1)} = \big(X^{(r+1)} + B_z^{(r)}\big) - \frac{\zeta}{\sigma}\,\mathcal{P}_{(C_1)^p}\Big(\frac{\sigma(X^{(r+1)} + B_z^{(r)})}{\zeta}\Big). \tag{3.5}$$

Let us now prove this simple fact.

Lemma 3.3 *Let $\mathcal{C} \subset \mathbb{R}^n$ be a nonempty closed set, $\lambda > 0$, and $\mathcal{C}' := l(\mathcal{C})$ for $l \colon x \mapsto \lambda x$ its rescaling. Then the orthogonal projection $\mathcal{P}_{\mathcal{C}'}$ to $\mathcal{C}'$ is given by*

$$\mathcal{P}_{\mathcal{C}'}(x) = \lambda\,\mathcal{P}_{\mathcal{C}}\Big(\frac{x}{\lambda}\Big) \qquad \text{for all } x \in \mathbb{R}^n.$$

Lemma 3.3 is a special case of the following result.

Lemma 3.4 *Let $\mathcal{C} \subset \mathbb{R}^n$ be a nonempty closed set and $\mathcal{C}' := l(\mathcal{C})$ for a linear angle preserving bijection $l \colon \mathbb{R}^n \to \mathbb{R}^n$. Then the orthogonal projection $\mathcal{P}_{\mathcal{C}'}$ to $\mathcal{C}'$ is given by*

$$\mathcal{P}_{\mathcal{C}'}(x) = l\left(\mathcal{P}_{\mathcal{C}}\big(l^{-1}(x)\big)\right) \qquad \text{for all } x \in \mathbb{R}^n. \tag{3.6}$$

Proof. The claim (3.6) is equivalent to

$$\begin{aligned} & \big\langle c' - l\left(\mathcal{P}_C(l^{-1}(x')\right)\ ,\ x' - l\left(\mathcal{P}_C(l^{-1}(x')\right)\big\rangle \leq 0 && \forall x' \in \mathbb{R}^n \quad \forall c' \in \mathcal{C}' \\ \Leftrightarrow \quad & \langle l(c) - l\left(\mathcal{P}_C\, x\right)\ ,\ l(x) - l\left(\mathcal{P}_C(x)\right)\rangle \leq 0 && \forall x \in \mathbb{R}^n \quad \forall c \in \mathcal{C} \\ \Leftrightarrow \quad & \langle c - \mathcal{P}_C(x)\ ,\ x - \mathcal{P}_C(x)\rangle \leq 0 && \forall x \in \mathbb{R}^n \quad \forall c \in \mathcal{C}, \end{aligned}$$

where the first equivalence holds because l is surjective, and the second equivalence holds because l is linear and angle preserving. □

Finally, for each row $\tilde{z}$ of $X^{(r+1)} + B_z^{(r)}$ in (3.5), the projection of $\tilde{z}$ to C_1 can be computed by projecting the subvector containing the positive components of $\tilde{z}$ onto the intersection of the unit l_1-ball with the positive orthant. The projection of $\tilde{z}$ to C_1 agrees with $\tilde{z}$ in the nonpositive components of $\tilde{z}$. Having computed the required proximum mapping, we present the full algorithm in the next section.

3.4. Algorithm

We use the Primal-Dual Hybrid Gradient Algorithm (PDHGMp) [29, 90] to solve equation (3.4) and thereby model (3.1) by finding a saddle point of the associated Lagrangian

$$L(x,u,v,z;\ b_u,b_v,b_z) \ = \ \frac{1}{2}\|m\circ(y-u)\|_2^2+\lambda\|\,|v|\,\|_1+\zeta\|z\|_{\infty,1}+\frac{\nu}{2}\|x\|_2^2+\iota_{(\triangle_p)^N}(x) \\ +\left\langle \begin{pmatrix} b_u\\ b_v\\ b_z\end{pmatrix}\ ,\ \begin{pmatrix} I_N\otimes K\\ \nabla\otimes I_p\\ I_{Np}\end{pmatrix}x-\begin{pmatrix} u\\ v\\ z\end{pmatrix}\right\rangle .$$

This leads to Algorithm 3, which is guaranteed to converge to a saddle point of the Lagrangian by [28, Theorem 6.4].

Algorithm 3 PDHGMp for (3.4)

Initialization: $\tau,\sigma>0$ with $\tau\sigma\le 1/\left\|\begin{pmatrix} I_N\otimes K\\ \nabla\otimes I_p\\ I_{Np}\end{pmatrix}\right\|_2^2$,

$x^{(0)}\in\mathbb{R}^{Np},\ u^{(0)}\in\mathbb{R}^{NL},\ v^{(0)}\in\mathbb{R}^{2Np},\quad b_u^{(0)}=\bar{b}_u^{(0)}\in\mathbb{R}^{NL},\quad b_v^{(0)}=\bar{b}_v^{(0)}\in\mathbb{R}^{2Np}.$

Iterations:

For $r=0,1,\dots$

1. $\widetilde{x}^{(r)} := x^{(r)}-\tau\sigma\left(I_N\otimes K^\top \,\middle|\, \nabla^\top\otimes I_p \,\middle|\, I_{Np}\right)\begin{pmatrix}\bar{b}_u^{(r)}\\ \bar{b}_v^{(r)}\\ \bar{b}_z^{(r)}\end{pmatrix}$

   $x^{(r+1)} := \operatorname*{argmin}_x\left\{\iota_{(\triangle_p)^N}(x)+\frac{\nu}{2}\|x\|_2^2+\frac{1}{2\tau}\|x-\widetilde{x}^{(r)}\|_2^2\right\}$

2. $u^{(r+1)} := \operatorname*{argmin}_u\left\{\frac{1}{2}\|(y-u)\|_2^2+\frac{\sigma}{2}\|(I_N\otimes K)x^{(r+1)}-u+b_u^{(r)}\|_2^2\right\}$

   $v^{(r+1)} := \operatorname*{argmin}_v\left\{\lambda\|\,|v|\,\|_1+\frac{\sigma}{2}\|(\nabla\otimes I_p)x^{(r+1)}-v+b_v^{(r)}\|_2^2\right\}$

   $z^{(r+1)} := \operatorname*{argmin}_z\left\{\zeta\|z\|_{\infty,1}+\iota_{\ge 0}(z)+\frac{\sigma}{2}\|x^{(r+1)}-z+b_z^{(r)}\|_2^2\right\}$

3. $b_u^{(r+1)} := b_u^{(r)}+(I_N\otimes K)x^{(r+1)}-u^{(r+1)}$

   $b_v^{(r+1)} := b_v^{(r)}+(\nabla\otimes I_p)x^{(r+1)}-v^{(r+1)}$

   $b_z^{(r+1)} := b_z^{(r)}+x^{(r+1)}-z^{(r+1)}$

4. $\bar{b}_u^{(r+1)} := 2b_u^{(r+1)}-b_u^{(r)}$

   $\bar{b}_v^{(r+1)} := 2b_v^{(r+1)}-b_v^{(r)}$

   $\bar{b}_z^{(r+1)} := 2b_z^{(r+1)}-b_z^{(r)}$

We have described how to find the $z^{(r+1)}$-update. All other updates were already part of Algorithm 2 and have been commented on in the previous chapter. We can thus

compute all steps of Algorithm 3 efficiently and consider its numerical performance in the next section.

3.5. Numerical Results

We are interested in the improvements obtained from the inclusion of the $l_{\infty,1}$-sparsity term. First we comment on the simulated datasets and measures of performance, then we show the improvements by giving numerical results obtained with Algorithm 3 on a coarse dataset in Section 3.5.3 and on a more finely structured dataset in Section 3.5.4.

3.5.1. Simulated Datasets

We create a hypercube by randomly ordering the library and then mixing the first four endmembers with four given abundances. The abundances of all other endmembers are set to zero, i.e., endmembers 5–50 are not part of the mixture.

Our spectral library is based on the one used by Iordache, Bioucas-Dias and Plaza in [59], which is available online as part of their demonstrating code for "Minimum Volume Simplex Analysis", at `http://www.lx.it.pt/~bioucas/code.htm`. This library contains a random selection of 240 materials of the USGS library, called splib06, released in 2007 and available online at `speclab.cr.usgs.gov/spectral.lib06`. The library contains spectra of different mineral types measured at 224 bands centered at wavelengths ranging from 0.3 to 3 µm.

We have removed very similar spectra from the library by passing through the library once and removing spectrum l_j, $j = 1, \ldots, 224$, if $\cos(\measuredangle(l_i, l_j)) > .995$ for some $i < j$. This ensures that the pairwise angles between two spectra are at least $\arccos(.995) \approx 2.5°$. Still the unmixing problem is badly conditioned and simple matrix inversion could increase an input error by a factor of

$$\operatorname{cond}(K_{\text{lib}}) := \|K_{\text{lib}}\| \, \|K_{\text{lib}}^{-1}\| \approx 10^9.$$

Therefore we use a smaller library, consisting of 50 spectra chosen at random from the 166 remaining spectra.

3.5.2. Performance Discriminators

Standard performance discriminators are not a very good measure of the quality of the unmixing. Comparing the l^2-distance between the found abundances to the original abundances (or the signal-to-noise ratio) is not a good measure of how informative the solution is to the human perception: In a piecewise smooth image with lowered contrast it is often easier to recognize features of interest than in a very noisy image with full contrast.

Further, when reproducing a sparse solution, the l^2-distance alone does not weight information properly to reflect possible post-processing. If it is clear from the algorithmic

result, which of the endmembers are part of the mixture, then any noise in the abundances corresponding to other endmembers can be neglected, because a post-processing could perform a second unmixing only with the chosen endmembers. Let us therefore first consider directly the sparsity pattern of the solution. More precisely, given the computed abundance cube X^{cube}, with $\{X^{\text{cube}}_{i,j,k}\}_{i,j}$ denoting the abundances of the k-th endmember throughout the image, we compute the vectors $v_\infty(X^{\text{cube}})$ and $v_1(X^{\text{cube}})$ by

$$\left(v_\infty(X^{\text{cube}})\right)_k := \max_{i,j} X^{\text{cube}}_{i,j,k} \quad \text{and} \quad \left(v_1(X^{\text{cube}})\right)_k := \sum_{i,j} X^{\text{cube}}_{i,j,k}.$$

3.5.3. Results and Discussion for Coarse Structures

We simulate a data cube by mixing the first four spectra of the endmember library K, shown in Figure 3.1 (left) with piecewise constant abundance channels of size 24×30 similar to those depicted in Figure 2.9, now with two steps instead of four "steps" between the pure corners. In particular there are 8×10 homogeneous regions in the four corners, which contain the pure materials. We proceed further as in the previous chapter, add Gaussian noise with standard deviation of about 10% of the maximum independently to each pixel of the hypercube and finally simulate the measurement with a defective line camera, discarding part of the measured lines. To obtain the simulated instance `instDm3` we keep 30% of the lines. A few channels of the hypercube $Y^{\text{cube}} = \texttt{instDm3}$ are shown in Figure 3.1 (right).

From the hypercube Y^{cube}, reshaped into the data matrix Y and the matrix K, whose columns are the chosen 50 spectra from the USGS library (with the four endmembers among them), we obtain the abundance matrix X by solving (3.1) with Algorithm 3 and parameter $\nu = 10^{-5}$ and 30000 iterations fixed throughout this section, which took about 12 minutes. In practice a smaller number of iterations can be used, and the Lagrangian simplex constraint need not be enforced to high accuracy in each iteration.

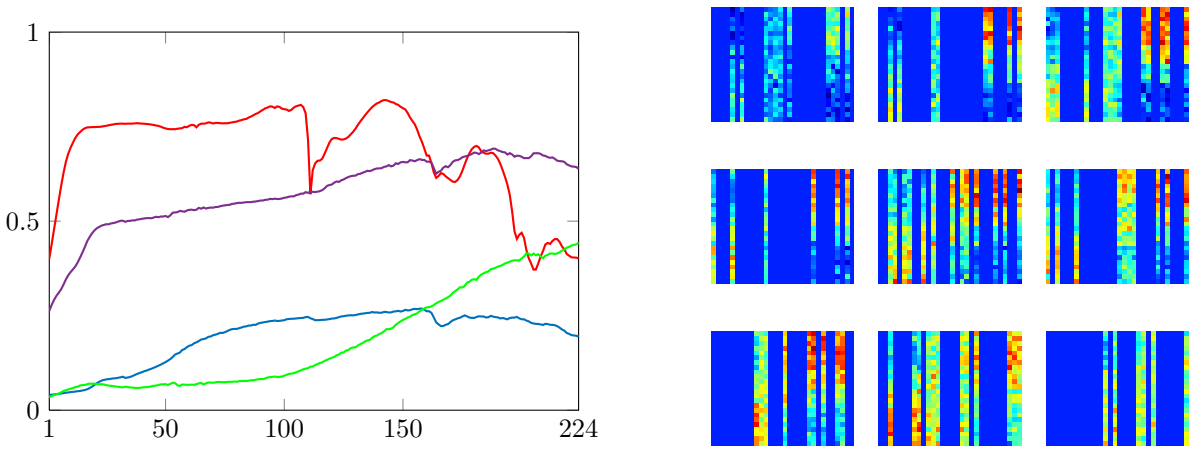

Figure 3.1.: The simulated hypercube `instDm3` with a coarse abundance structure: The four spectral endmembers used (left); Channels $1, 26, 51, \ldots, 201$ of the simulated measurement which keeps 30% of the hypercube (right).

Improvement The standard measure of the accuracy of the unmixing is the distance between the recovered abundances X and the true abundances, $\|X - X_{\text{ori}}\|_F$. It is shown for different values of λ and ζ in Table 3.1.

Table 3.1.: The abundance reconstruction error $\|X - X_{\text{ori}}\|_F$ for different values of the regularization parameters, on the coarse instance `instDm3`.

$\lambda \setminus \zeta$	0	1e-4	7.2e-4	2.16e-3	7.2e-3	0.0216	0.07	0.216	0.72	7.2
0	10.80	10.80	10.79	10.78	10.72	10.58	10.33	9.89	8.98	7.77
0.001	10.26	10.26	10.25	10.24	10.18	10.08	9.84	9.46	8.63	7.57
0.01	7.19	7.19	7.19	7.18	7.16	7.12	7.02	6.81	6.46	6.05
0.10	2.94	2.94	2.94	2.94	2.94	2.94	2.94	2.93	2.92	3.50
0.25	**2.64**	**2.64**	**2.64**	**2.64**	**2.64**	**2.64**	**2.64**	2.66	2.70	3.64

In the first row, for $\zeta = 7.2$, a result of 7.77 is achieved, which is better than the result of 10.80 without regularization, $\lambda = \zeta = 0$. Hence there is a clear improvement from using the $l_{\infty,1}$-norm. However, the effect of spatial regularization is much clearer: Changing the value of λ from 0 to 0.10, decreases the recovery error from 10.80 to 2.64. For a value of λ of 0.1, there is only a slight improvement by adding the $l_{\infty,1}$-regularization. For strongest spatial regularization, there is no further improvement by adding the $l_{\infty,1}$-regularization (last row), the quality even deteriorates in the last entries of the last row.

The Effect of Sparsity We notice that, even when $\zeta = 0$, the $l_{\infty,1}$-norm decreases for $\lambda = 0.001, 0.01$ compared to $\lambda = 0$ because on channels which are mostly zero, spatial smoothing moves outliers towards zero.

This effect can be seen very clearly in the sparsity vectors $v_1(X^{\text{cube}})$ obtained from the hypercube `instDm3` with only 30% of the data known. We plot these in Figure 3.2 for different values of λ, ζ. Let us comment on Figure 3.2.

Without the either TV- or $l_{\infty,1}$-regularization (plot 1), all components of $v_1(Xcube)$ are large, i.e. all abundance channels are used.

Both $l_{\infty,1}$-norm (plot 2) and TV-norm (plot 4) suppress the noise in most channels, while also lowering the values in the first four channels. If the regularization parameter is too large, then some coefficients in the first four channels also drop towards zero, both for $l_{\infty,1}$-norm (plot 3) and for the TV-norm (plot 6). If TV-penalization is used, (plot 4–6: '31', '38', '40'), then the influence of the $l_{\infty,1}$-norm is small, as seen by comparing the result for TV- and $l_{\infty,1}$-regularization (plot 5) to the result for TV-regularization alone (plot 4).

To summarize, we can say that for images with large homogeneous regions, the used abundances can be recovered equally well both by TV-regularization alone and by combined TV- and $l_{\infty,1}$-regularization, as long as the regularization parameters are chosen appropriately.

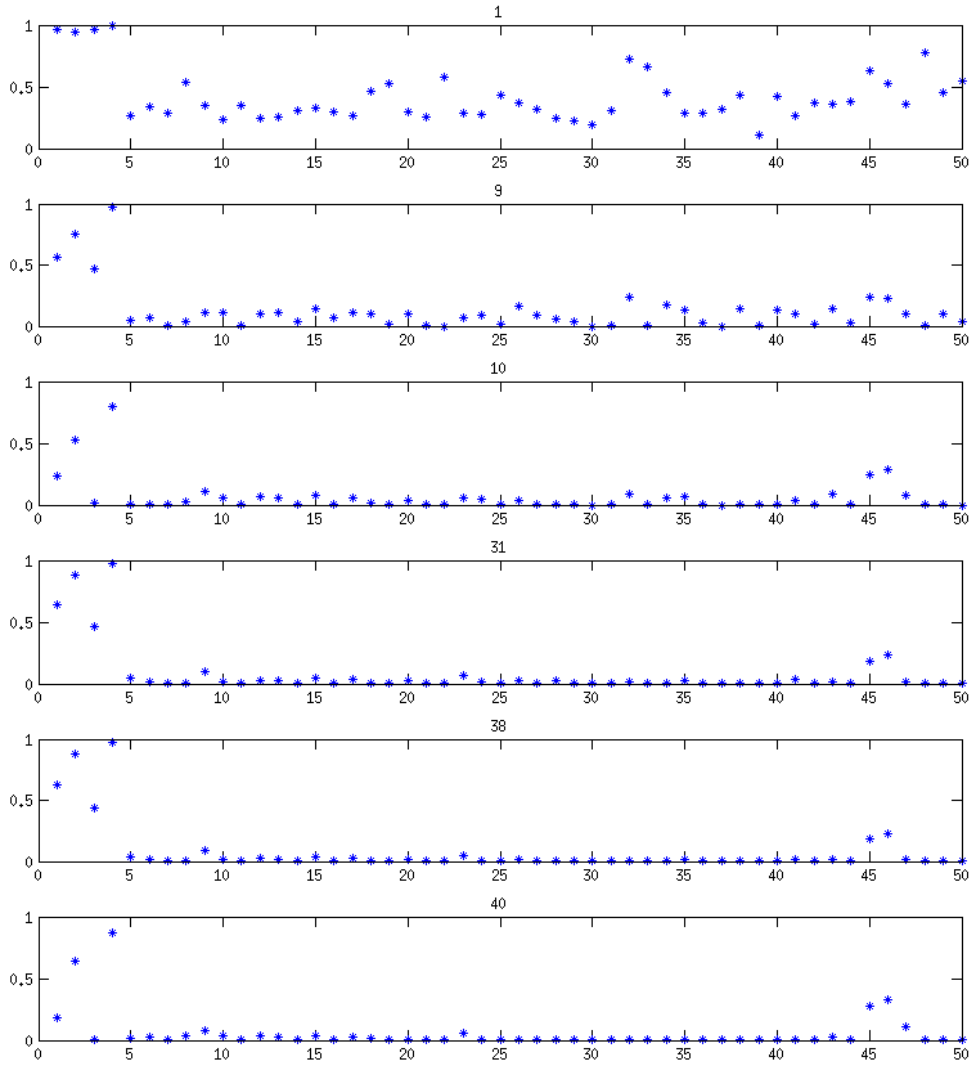

Figure 3.2.: The sparsity vector $v_1(X^{\text{cube}})$ of the abundance cube, obtained when unmixing `instDm3` for each of the following choices of regularization parameters: plots 1–3: $(\lambda, \zeta) = (0, 0), (0, .72), (0, 7.2)$ and plots 4–6: $(\lambda, \zeta) = (.01, 0),\ (.01,\ .216),\ (.01,\ .72)$.

3.5.4. Results and Discussion for Fine Structures

Improvement In the previous example TV-unmixing already worked well, due to the presence of large homogeneous regions. Let us now present a scenario, where the unmixing can be improved by adding the $l_{\infty,1}$-regularization. In general the success of spatial regularization depends on the scale separation between the relevant features of the image and the noise. There should be more room for improvement of pure TV-unmixing, when smaller features are present in the image.

As a very simple example we consider an image with finer structures in the abundances. Again we mix four randomly chosen endmembers, now with the abundances shown in Figure 3.3 (left). Figure 3.4 shows the chosen spectra (left) and the remaining spectra of the library (right). Again we add 10% noise. Figure 3.3 (right) shows four channels of the resulting hypercube.

We obtain the unmixed hypercube by minimizing model (3.1) with Algorithm 3. The following table shows the l^2-distance between the reconstructed and original hypercube for different values of the regularization parameters (λ, ζ).

Table 3.2.: The abundance reconstruction error $\|X - X_{\text{ori}}\|_F$ for different values of the regularization parameters, when unmixing from fine scale abundances.

$\lambda \setminus \zeta$	0	1e-4	7.2e-4	2.16e-3	7.2e-3	0.0216	0.07	0.216	0.72	7.2
0.00	13.32	13.16	12.64	11.83	10.44	9.85	9.43	9.56	9.74	10.48
0.00	12.55	12.30	11.96	11.14	10.21	9.68	9.30	9.47	9.67	10.44
0.01	9.99	9.91	9.78	9.54	9.06	8.81	**8.73**	9.21	9.56	10.59
0.10	13.74	13.73	13.72	13.73	13.78	13.90	14.26	15.14	15.59	16.54
0.25	18.49	18.50	18.52	18.57	18.69	18.79	18.92	18.97	18.99	19.06

TV- and $l_{\infty,1}$-regularization alone each yield similar improvements compared to the unmixing without regularization $(\lambda, \zeta) = (0, 0)$. The optimum for $l_{\infty,1}$-regularization alone with a value of 9.43 is slightly better than for TV-regularization alone with a value of 9.99. The best result is achieved for a combination of both regularizers, with a value of 8.73 for $\lambda = 0.10,\ \zeta = 0.07$.

Figure 3.5 shows the abundances obtained without regularization and with combined regularization. Figure 3.6 shows the abundances obtained from using only one of the regularization terms.

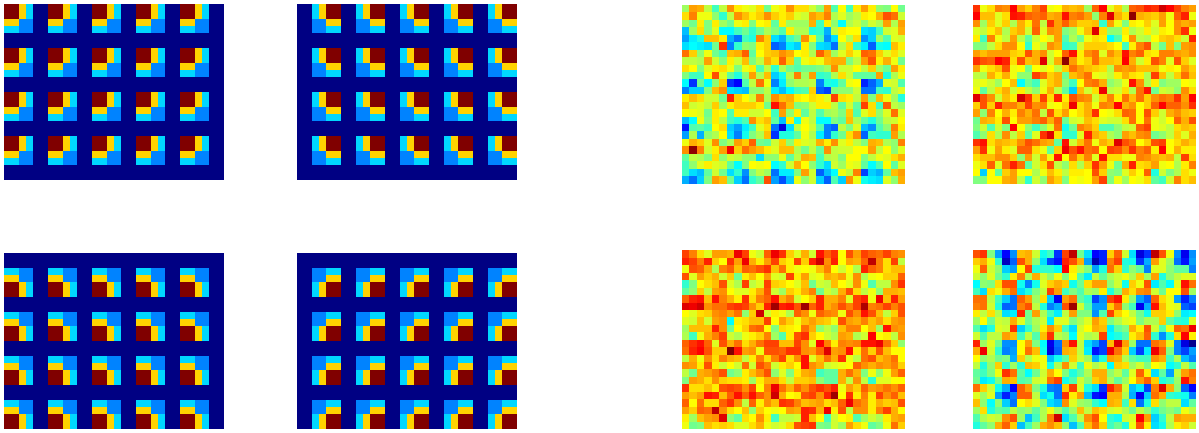

Figure 3.3.: The abundances in `instDfineB10` (left), shown are endmembers 1–4, other abundances being zero; Channels 1, 26, 55 and 201 of the corresponding noisy input hypercube (right).

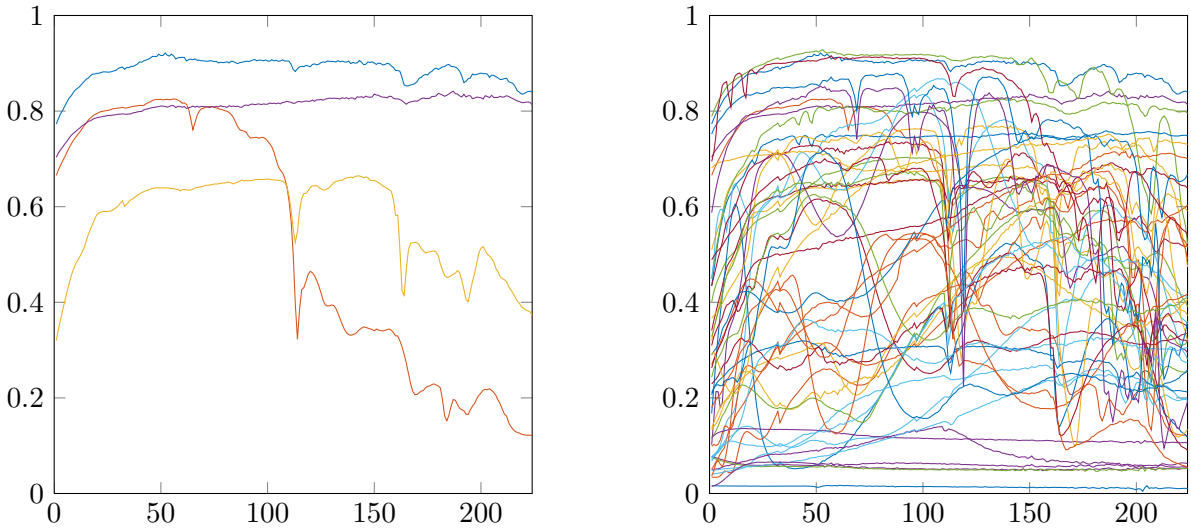

Figure 3.4.: Spectra from which the mixtures in `instDfineB10` are generated (left), and all spectra contained in the library (right).

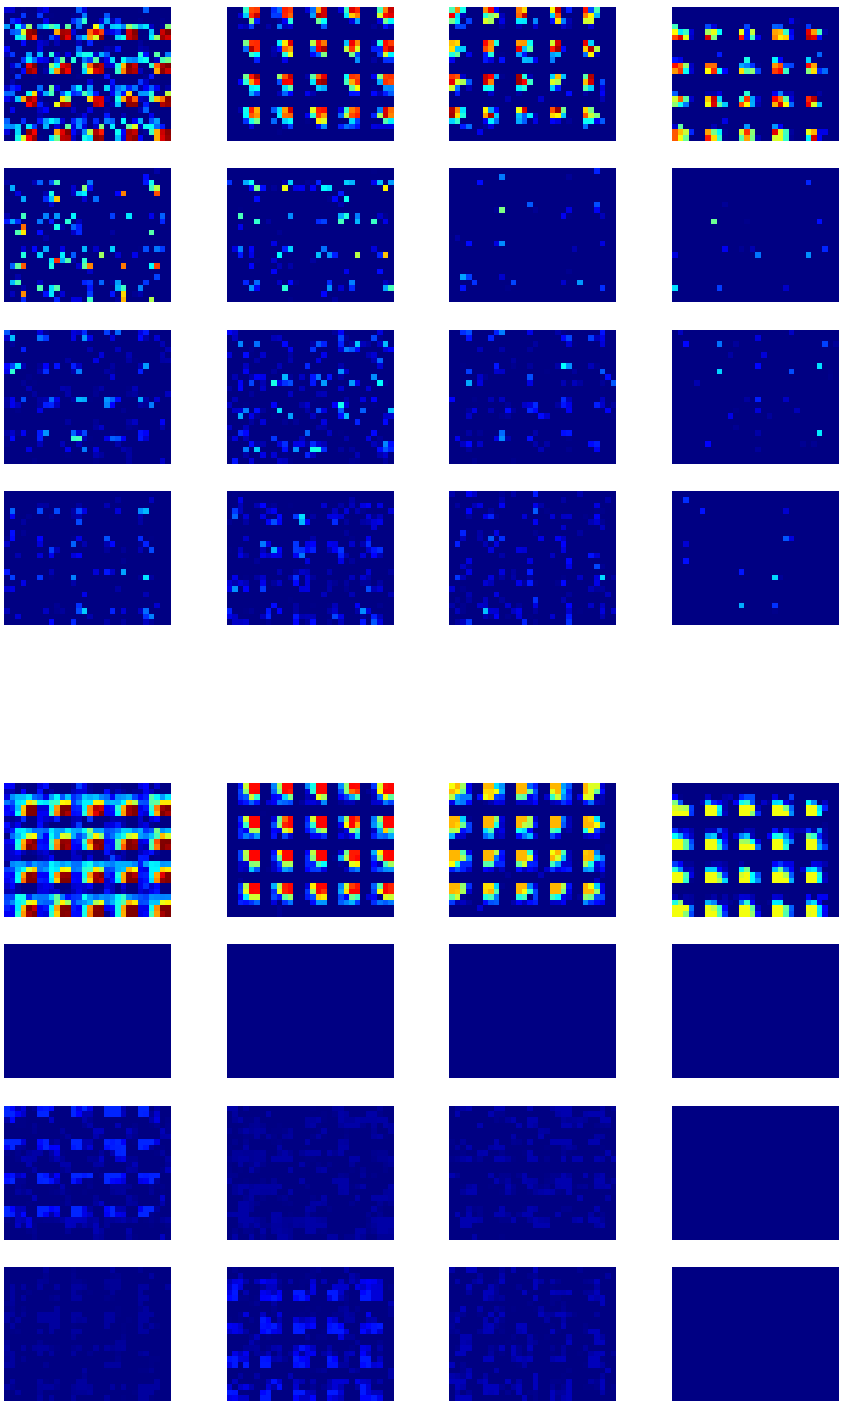

Figure 3.5.: Recovered abundances: (top) without regularization, i.e. for $(\lambda, \zeta) = (0, 0)$ and (bottom) for the best choice of parameters of $\lambda = 0.10$, $\zeta = 0.07$. In each case the same 16 channels out of the 50 recovered channels are shown.

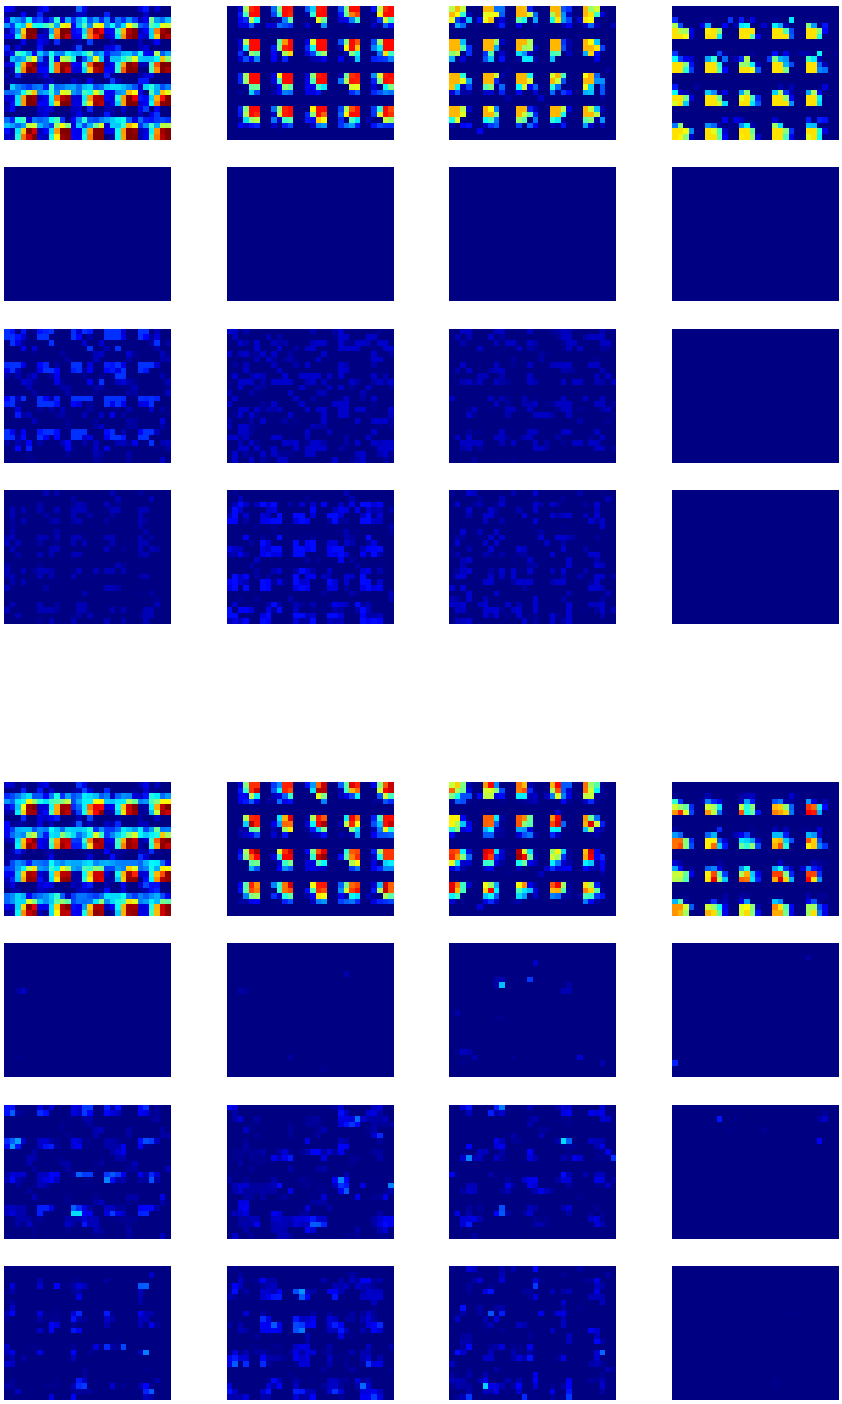

Figure 3.6.: Recovered abundances using either TV- or $l_{\infty,1}$-regularization: (top) for $l_{\infty,1}$-regularization alone, with $\zeta = 0.07$; (bottom) for TV-regularization alone with $\lambda = 0.01$. The same selection of recovered channels is shown as in Figure 3.5.

4. Unsupervised Unmixing

Outline

In Chapter 2 and Chapter 3 we assumed that spectral endmembers were known or contained in a library. There are practical scenarios where neither is the case. Consider the following instance of explorative data analysis: A researcher in biology is interested in finding patterns in the plant distribution from a hyperspectral satellite image. This requires a dimension reduction from several hundred channels (or in other applications such as mass spectroscopy even several thousand channels) to a handful of abundance channels, from which interesting regions or patterns can recognized.

Minimizing not only over the abundances $X \in \mathbb{R}_{\geq}^{p,N}$, but also over the endmembers $K \in \mathbb{R}_{\geq}^{L,p}$, we arrive at an extension of the problem of **nonnegative matrix factorization (NMF)**. Given Y, an $L \times N$ nonnegative matrix of observations, and an integer $p \in \mathbb{N}$ with $1 \leq p \leq \min(L, N)$, the NMF-Problem is the task of finding a nonnegative $L \times p$ matrix K and a nonnegative $p \times N$ matrix X such that $Y = KX$. As a factorization need not exist for $p < \min(L, N)$, the equality constraint is replaced in practical applications by the requirement that $Y \approx KX$ with respect to some similarity measure. In the variational setting (1.5), introduced at the end of Chapter 1, this can be written as minimization problem

$$\operatorname*{argmin}_{K,X} \big(\mathcal{D}(K, X; Y) + \mathcal{R}(K, X)\big)$$

where, for the NMF-problem, the regularization term $\mathcal{R}$ includes nonnegativity constraints on both K and X. Extending from the NMF-problem to the unmixing problem, we include further regularization terms in $\mathcal{R}$, which can enforce sparsity or spatial regularity.

The data fidelity $\mathcal{D}(K, X; Y)$ term measures how well the product KX approximates Y and a least-squares distance is a common choice. Under the assumption of Poisson noise, other measures are preferred. Here we use a sum over the Kullback-Leibler divergence (KL-divergence) between corresponding columns following [69, 70], which is given for

vectors Kx and y by

$$\sum_{i=1}^{L}\left(y_i \log \frac{y_i}{(Kx)_i} - y_i + (Kx)_i\right) = \left\langle \mathbb{1}_L, y \circ \log \frac{y}{Kx} - y + Kx \right\rangle.$$

In this chapter we propose a fast multiplicative algorithm to minimize a variational formulation of the unmixing problem, using the Kullback-Leibler divergence together with a simplex-constraint, and analyse its convergence.

4.1. Introduction and Related Work

Complexity In general the NMF-problem is NP-hard [105]. Complexity for fixed p has been discussed in [4]. Despite its complexity, the NMF problem formulations have been successful in a series of applications, which suggests that some real-world NMF problems might not be as difficult as the general NMF problem. Arora et al. [4] showed that there exists a subclass of nonnegative matrices, referred to as **separable**, for which the NMF problem can be solved in polynomial time. Separability requires that there exists an NMF (K, X) of the input matrix Y where the columns of K are a subset of the columns of Y. More precisely, a matrix Y is **p-separable** if there exists an index set $J \subseteq \{1, 2, \ldots, N\}$ with cardinality p and a $p \times N$ nonnegative matrix X such that $Y = Y(:, J)X$. Equivalently, the cone spanned by p columns of Y contains all columns of Y.

In hyperspectral unmixing separability is known as the **pure-pixel assumption** and in document classification as the existence of an "anchor" word for each topic.

In practice the given separable matrix will be perturbed by noise and N. Gilis and A. Vavasis [49] consider the **near-separable NMF** problem defined in [48]: "Given a noisy p-separable matrix $\tilde{Y} = Y + n$ with $Y = KX = K[I_p, X']\Pi$, where K and X' are nonnegative matrices, I_p is the $p \times p$ identity matrix, Π is a permutation matrix, and n is the noise with $\max_j \|n(:, j)\|_2 \leq \varepsilon$ for some $\varepsilon \geq 0$, find a set J of p indices such that $Y(:, J) \approx K$." Several algorithms proposed for this problem, are sensitive to the condition number of K, according to [49], where a preconditioning and a semidefinite programming approach are suggested. These include [4, 5] with several complexity results, [92] the Hottopixx algorithm suitable for large data sets, and the robust recursive algorithm in [48].

Another approach uses clustering or removal of similar columns to reduce to a subset Y' of N' columns of Y, where $N' \ll N$, from which the authors of [43] propose to select the relevant columns by seeking an $N' \times N'$ matrix T with

$$Y' \approx Y'T.$$

They impose weighted l^1-norms and row-sparsity encouraging penalization on T and finally select those columns of Y' corresponding to nonzero rows of T.

In other applications the separability assumption is violated. Remote sensing hyperspectral images may be heavily mixed. All approaches without a separability assumption only seek local minimizers [19, 51, 65, 69, 70, 74, 88].

Multiplicative Algorithms and Kullback-Leibler objective Alternating multiplicative update rules have been popularized by D. D. Lee and H. S. Seung [70] for the NMF problem. Minimizing either a least squares criterion or the column-wise Kullback-Leibler divergence between Y and KX given by

$$\mathrm{KL}\,(Y, KX) = \sum_{ij} \left(Y_{ij} \log \frac{Y_{ij}}{(KX)_{ij}} - Y_{ij} + (KX)_{ij} \right),$$

they considered the algorithmic scheme of alternating minimization with respect to K and X using the following update rules for $k = 0, 1, \ldots$, starting from arbitrary but strictly positive matrices $K^{(0)}$ and $X^{(0)}$:

$$\begin{aligned} K^{(k+1)} &= \frac{K^{(k)}}{\mathbb{1}_{L,N}(X^{(k)})^\top} \circ \frac{Y}{K^{(k)}X^{(k)}}(X^{(k)})^\top; \\ X^{(k+1)} &= \frac{X^{(k)}}{(K^{(k+1)})^\top \mathbb{1}_{L,N}} \circ (K^{(k+1)})^\top \frac{Y}{K^{(k+1)}X^{(k)}}, \end{aligned} \tag{4.1}$$

where $\mathbb{1}_{L,N}$ is an $L \times N$ matrix of ones, $X^\top$ is the transpose of X and we use the Hadamard (entrywise) product $\circ$ and division. Lee and Seung gave surrogates (for a definition see the next section) and derived the above updates in the quadratic and Kullback-Leibler case, thus showing the decrease of the respective objective functions.

C. De Mol and L. Lecharlier [69] have extended the multiplicative approach by additive penalization terms and have made the first claim of a convergence proof. In [69] an algorithm is given for the special instance of the problem given by blind deconvolution in incoherent optical imaging, where the original image X, the blurred image Y and the space-invariant **point spread function (PSF)** K are nonnegative light intensities. Given a blurred image Y, with $Y_{i,j} \geq 0$ for all i, j, the functional

$$\mathcal{F}(K, X) = \mathrm{KL}\,(Y, K * X) + \frac{\mu}{2}\|K\|_F^2 + \lambda\|X\|_{1,1} + \frac{\nu}{2}\|X\|_F^2 \tag{4.2}$$

is minimized over $K \in \mathbb{R}^{L,p}$ and $X \in \mathbb{R}^{p,N}$ under the constraints

$$X_{i,j} \geq 0, \quad K_{i,j} \geq 0 \quad \text{and} \quad \|K\|_{1,1} = 1, \tag{4.3}$$

with nonnegative regularization parameters λ, μ, ν. Here the **Frobenius norm** $\|A\|_F$ and the l^1-norm $\|F\|_{1,1}$ of an $m \times n$ matrix A are given, respectively, by

$$\|A\|_F^2 := \sum_{i=1}^{m}\sum_{j=1}^{n} A_{i,j}^2 \quad \text{and} \quad \|A\|_{1,1} := \sum_{i=1}^{m}\sum_{j=1}^{n} |A_{i,j}|.$$

The matrix multiplication KX has been replaced by the 2D circular convolution $K * X$ and $L = N = p =: n$, so that Y, X, K are all $n \times n$ matrices, with $X_{i,j}$ denoting the light intensity in pixel (i, j).

In the case where the PSF K is known, the algorithm (4.1) without regularization had previously been well-known in astronomy as Richardson-Lucy's one [75, 93] and in medical imaging as (EM)ML (Expectation-Maximization Maximum Likelihood).

The alternate minimization algorithm for blind deconvolution had appeared already in [46, 56]. One of the achievements in [69] is that the addition of penalty terms that prevent the algorithm from converging to the trivial solution where the PSF K is a delta-peak and the recovered image X equals the original one. The authors suggest an extension with a smoothed total-variation (TV) penalty on the image, using the separable surrogate for the TV-penalty introduced in [39]. For details on the TV-penalty and the surrogate see equations (4.55) and (4.56) in Section (4.6), where we summarize [39]. Meanwhile there exists a rich literature on blind deconvolution also known as blind deblurring. See Section 5.5.3 in [30] with comments on existance and uniqueness and models up to 2005 and the references on more recent developments in the paper by Figueiredo et al. [84]. Note also the related paper [1], where the boundary is treated nicely.

For strictly convex objective functions, uniqueness of the minimizer allows to prove convergence of the iterates, see Theorem 4.19. However, the variational formulation of NMF is no longer convex due to the similarity measure between the product KX and the input matrix Y. In fact, both $\mathcal{F}(K, X) := \mathrm{KL}\,(Y, KX)$ and $\mathcal{F}(K, X) := \frac{1}{2}\|Y - KX\|_F^2$ are **biconvex**, but not jointly convex in the pair (K, X). Recall that a function $\mathcal{F}(K, X)$ of two blocks of variables is biconvex if the functions $\mathcal{F}(\cdot, X)$ and $\mathcal{F}(K, \cdot)$ are both convex.

For non-convex objective functions convergence results are rare in the literature. A recent exception is [24], wherein J. Bolte, S. Sabach and M. Teboulle outline a convergence analysis framework building on the Kurdyka–Łojasiewicz property. However, their requirement that the partial gradients are globally Lipschitz is not fulfilled for the Kullback-Leibler divergence.

In [69] a convergence result for a multiplicative algorithm minimizing the blind deconvolution functional (4.2) has been stated without a proof. A hint is given to used Zangwill's convergence theory. When attempting to use Zangwill's convergence theory for the blind hyperspectral unmixing algorithm proposed here, it is not clear on which set the update operator is continuous. In particular surrogates do not exist at limit points. Further complications arise because the functional does not prevent rows of X and columns of K from converging to zero, which leads to limit points at which the update is not defined.

Finesso and Spreij have reinterpreted the original algorithm (4.1) from [70], lifting the involved matrices to three index tensors [45]. They conclude that limit points (K, X) of the algorithm are stationary, provided that K has no zero columns and X has no zero rows.

Least squares objective For the related problem with the least squares objective $\frac{1}{2}\|Y - KX\|_F^2$ instead of the KL-divergence, multiplicative updates have also been proposed in [70]. Lin [74] gives a nice succinct discussion of the difficulties in proving convergence and proposes a modification which has the same complexity and guarantees that limit points are stationary points.

In [51] an accelerated version of the least squares Lee-Seung algorithm is demonstrated, yet the authors conclude that neither the original nor the accelerated algorithm converge

in a reasonable number of iterations.

The authors in the survey [32] give a comparison which shows different sides: They also report a larger number of iterations for the Lee-Seung algorithm than for comparative methods, while reporting at the same time competitive CPU time until numerical convergence.

In [65] an alternating exact minimization of the least squares objective by an active set method is proposed. For stationarity of every limit point of the sequence of iterates, they refer to existing convergence theory on two block descent schemes [52], because the subproblems are solved exactly. The authors numerically demonstrate a faster convergence compared to [70] and the alternating least squares algorithm (ALS) from [19].

Contribution We follow here a suggestion by C. De Mol and consider an algorithm for the unmixing of hyperspectral images which builds on the surrogate inequality for the KL-divergence suggested by Lee and Seung [70], adds l_2-regularization as done for image deblurring in (4.2), and enforces simplex constraints by computing a Lagrange multipliers.

We analyse the convergence properties of the proposed algorithm. In particular we show that the update mappings of the algorithm do not have a continuous extension to the boundary of the positive quadrant on which the Lee-Seung algorithm is defined.

We identify a set on which the update mappings are continuous and which contains all limit points of the sequence of iterates.

The regularization terms allow us to strengthen the descent properties from [70] into strict descent. By extending these descent properties to limit points we show that limit points are fixed points and that fixed points in the interior of the positive quadrant satisfy the KKT conditions.

In the different setting where the objective function is convex, and the iterative algorithm has been obtained from minimizing a strongly convex surrogate, we prove a general convergence result and give an application.

The remainder of this chapter is structured as follows: In Section 4.2 we introduce the mathematical model and Section 4.3 recalls how optimization transfer leads to decreasing algorithms. In Section 4.4 we derive the suggested algorithm for the minimization of the extended non-convex Kullback-Leibler objective for NMF. In Section 4.5 we prove convergence of a subsequence of the iterates to a fixed point which, under additional conditions, is a stationary point. Section 4.6 discusses a stronger result for convex objective functions with a unique minimizer.

4.2. Approximate NMF with Kullback-Leibler Divergence

In this section we introduce our unmixing model, which extends the NMF-problem with the Kullback-Leibler divergence by regularization terms and simplex constraints.

We define the generalized **Kullback-Leibler divergence** between matrices Y, KX by

$$\mathrm{KL}(Y, KX) := \sum_{i=1}^{L}\sum_{j=1}^{N}\left(Y_{ij}\log\frac{Y_{ij}}{(KX)_{ij}} - Y_{ij} + (KX)_{ij}\right), \qquad \text{and}$$
$$\widetilde{\mathrm{KL}}(Y, KX) := \langle \mathbb{1}_{L,N}, KX\rangle_F - \langle Y, \log KX\rangle_F.$$

For minimization with respect to K, X only the non-constant terms of KL collected in $\widetilde{\mathrm{KL}}$ will be relevant. We wish to solve the minimization problem

$$\operatorname*{argmin}_{K,X}\ \mathcal{E}(K,X) \qquad \text{s.t.} \quad K \in \mathbb{R}^{L,p}_{\geq},\ X \in \triangle_p^N, \tag{4.4}$$

where

$$\mathcal{E}(K,X) := \widetilde{\mathrm{KL}}(Y, KX) + \frac{\mu}{2}\|K\|_F^2 + \frac{\nu}{2}\|X\|_F^2, \tag{4.5}$$

with μ, ν positive regularization parameters. Recall the definition of the p-dimensional unit simplex $\triangle_p$ introduced earlier,

$$\triangle_p := \left\{x \in \mathbb{R}^p \colon x \geq 0,\ \sum_{i=1}^{p} x_i = 1\right\}$$

and the Hadamard entry-wise product $\circ$. For two matrices A, B of the same size,

$$\langle A, B\rangle_F = \sum_{ij} A_{ij}B_{ij}$$

denotes the inner product associated to the Frobenius matrix norm $\|\cdot\|_F$. Further, $\mathbb{1}_p \in \mathbb{R}^p$ is a vector of ones and $\mathbb{1}_{r,s} \in \mathbb{R}^{r,s}$ is an $r \times s$ matrix of ones. Here and in the following, $\log\colon \mathbb{R} \to \mathbb{R}\cup\{+\infty\}$ denotes the convex extension of the usual natural logarithm $\ln\colon \mathbb{R}_> \to \mathbb{R}$, defined by

$$\log t = \begin{cases} \ln t & \text{for } t > 0, \\ -\infty & \text{for } t \leq 0. \end{cases}$$

The cost function $\mathcal{E}$ is smooth on its effective domain

$$\Omega := \operatorname{dom}\mathcal{E} = \{(K,X) \in \mathbb{R}^{L,p}_{\geq} \times \mathbb{R}^{p,N}_{\geq} \colon \mathcal{E}(K,X) < +\infty\}.$$

For $\varepsilon \geq 0$ let

$$\Omega_\varepsilon := \left\{(K,X) \in \mathbb{R}^{L,p}_{\geq} \times \mathbb{R}^{p,N}_{\geq} \colon \forall i,j\ (KX)_{i,j} > \varepsilon\right\}. \tag{4.6}$$

Then $\Omega = \Omega_0$ and $\mathcal{E}(K,X) \to \infty$ for $\min_{i,j}(KX)_{i,j} \to 0$ so that every sublevel set of $\mathcal{E}$ is a subset of Ω_ε for some $\varepsilon > 0$. Further $\|X\|_F^2 + \|K\|_F^2$ is coercive on $\mathbb{R}^{L,p} \times \mathbb{R}^{p,N}$. This has the following consequence: If an iterative algorithm started at some point $z^{(0)} \in \Omega$ produces a sequence of iterates $z^{(r)} := (K^{(r)}, X^{(r)})$ with the property that $\mathcal{E}$ is nonincreasing along the sequence $\{z^{(r)}\}_r$, then all $z^{(r)}$ are contained in a compact subset $\Pi \subset \Omega$. Indeed $\Pi \subset \Omega_\varepsilon$ for some $\varepsilon > 0$ depending on $\mathcal{E}(z^{(0)})$. A method to derive such decreasing algorithms is introduced in the next section.

4.3. Optimization Transfer

Majorization-minimization (MM), also called **optimization transfer** [68], is a technique to derive iterative algorithms which guarantee a monotone decrease of the cost function. At each iteration this technique replaces the minimization of the cost function $\mathcal{E}(z)$ by the minimization of a substitute function $\mathcal{E}_{\text{sur}}(z, \widehat{z})$, called **surrogate** [70], which depends on the current solution estimate $\widehat{z}$ and satisfies the two conditions

(i) $\mathcal{E}(\widehat{z}) = \mathcal{E}_{\text{sur}}(\widehat{z}, \widehat{z})$ for each $\widehat{z} \in \mathbb{R}^d$,

(ii) $\mathcal{E}(z) \leq \mathcal{E}_{\text{sur}}(z, \widehat{z})$ for all $z, \widehat{z} \in \mathbb{R}^d$.

The general procedure is as follows, see e.g. [39]:

- Start with an initial estimate $z^{(0)} \in \mathbb{R}^d$ of the solution.
- Obtain the $(r+1)$-th estimate as

$$z^{(r+1)} = \operatorname*{argmin}_{z \in \mathbb{R}^d} \mathcal{E}_{\text{sur}}(z, z^{(r)}). \tag{4.7}$$

Then the following lemma holds [38, 68]:

Lemma 4.1 *The sequences $\mathcal{E}(z^{(r)})$ and $\mathcal{E}_{\text{sur}}(z^{(r+1)}, z^{(r)})$ defined by recurrence (4.7) are nonincreasing.*

Proof. This is the usual proof for optimization transfer algorithms [38, 68].

$$\mathcal{E}(z^{(r+1)}) \leq \mathcal{E}_{\text{sur}}(z^{(r+1)}, z^{(r)}) \leq \mathcal{E}_{\text{sur}}(z^{(r)}, z^{(r)}) = \mathcal{E}(z^{(r)}) \leq \mathcal{E}_{\text{sur}}(z^{(r)}, z^{(r-1)}),$$

where the second inequality follows because $z^{(r+1)}$ minimizes $\mathcal{E}_{\text{sur}}(\cdot, z^{(r)})$ by (4.7). □

The objective $\mathcal{E}(K, X)$ for NMF we have defined above depends on two blocks of variables K, X. To find a local minimum in Ω, we will alternate minimization with respect to K and X, defining block-wise updates

$$X^{(r+1)} = \operatorname*{argmin}_{X} \mathcal{E}_{\text{sur},1}^{K^{(r)}}(X, X^{(r)}), \tag{4.8}$$

$$K^{(r+1)} = \operatorname*{argmin}_{K} \mathcal{E}_{\text{sur},2}^{X^{(r+1)}}(K, K^{(r)}), \tag{4.9}$$

where $\mathcal{E}_{\text{sur},1}^{K^{(r)}}$ and $\mathcal{E}_{\text{sur},2}^{X^{(r+1)}}$ are block-wise surrogates, which we define next. Define sections of Ω on which one variable is fixed,

$$\Omega_1(K) := \left\{ X \in \mathbb{R}^{p,N} \colon (K, X) \in \Omega \right\},$$
$$\Omega_2(X) := \left\{ K \in \mathbb{R}^{L,p} \colon (K, X) \in \Omega \right\}.$$

For each $K \in \mathbb{R}_{\geq}^{L,p}$ without zero rows, let $\mathcal{E}_{\mathrm{sur},1}^{K}\colon \Omega_1(K) \times \mathbb{R}_{>}^{p,N} \to \mathbb{R} \cup \{+\infty\}$ be a surrogate for $\mathcal{E}(K,\cdot)$, i.e.

$$\begin{aligned} \mathcal{E}_{\mathrm{sur},1}^{K}(\hat{X},\hat{X}) &= \mathcal{E}(K,\hat{X}) \quad \text{for all } \hat{X} \in \Omega_1(K), \\ \mathcal{E}_{\mathrm{sur},1}^{K}(X,\hat{X}) &\geq \mathcal{E}(K,\hat{X}) \quad \text{for all } X,\hat{X} \in \Omega_1(K), \end{aligned}$$

and similarly, for each $X \in \mathbb{R}_{\geq}^{p,N}$ without zero columns, let $\mathcal{E}_{\mathrm{sur},2}^{X}\colon \Omega_2(X) \times \mathbb{R}_{>}^{L,p} \to \mathbb{R} \cup \{+\infty\}$ be a surrogate for $\mathcal{E}(\cdot,X)$, i.e.

$$\begin{aligned} \mathcal{E}_{\mathrm{sur},2}^{X}(\hat{K},\hat{K}) &= \mathcal{E}(\hat{K},X) \quad \text{for all } \hat{K} \in \Omega_2(X), \\ \mathcal{E}_{\mathrm{sur},2}^{X}(K,\hat{K}) &\geq \mathcal{E}(\hat{K},X) \quad \text{for all } K,\hat{K} \in \Omega_2(X), \end{aligned} \tag{4.10}$$

Then alternating minimization (4.8), (4.9) implies monotone decrease

$$\mathcal{E}(K^{(r+1)},X^{(r+1)}) \leq \mathcal{E}(K^{(r)},X^{(r+1)}) \leq \mathcal{E}(K^{(r)},X^{(r)}). \tag{4.11}$$

Note that, for $K \in \mathbb{R}_{\geq}^{L,p}$ without zero rows, the inclusion $\mathbb{R}_{>}^{p,N} \subseteq \Omega_1(K)$ holds and for $X \in \mathbb{R}_{\geq}^{p,N}$ without zero columns, the inclusion $\mathbb{R}_{>}^{L,p} \subseteq \Omega_2(X)$ holds. The assumption that X has no zero columns will automatically be fulfilled because each iteration of the algorithm that we are going to propose, will ensure $X \in \triangle_p^N$. The property that the block variable K does not have zero rows will be kept by the algorithm because, after each iteration, $(K^{(r+1)},X^{(r+1)})$ will be in Ω due to the decrease of the objective that we are going to prove in the sequel.

4.4. NMF Algorithm via Optimization Transfer

In this section we follow a suggestion by C. De Mol, define surrogates for the given blind unmixing model (4.4) and derive Algorithm 4 below. We propose the following surrogates, extending [70].

Lemma 4.2 *Let $X \in \mathbb{R}_{\geq}^{p,N}$ have no zero columns. Then the function*

$$\mathcal{E}_{\mathrm{sur},2}^{X}(K,A) := \langle \mathbb{1}_{L,N}, KX\rangle_F - \left\langle \frac{Y}{AX}, \left(A \circ \log\frac{K}{A}\right) X \right\rangle_F - \langle Y, \log AX\rangle_F + \frac{\mu}{2}\|K\|_F^2,$$

defined on $\mathbb{R}_{\geq}^{L,p} \times \mathbb{R}_{>}^{L,p}$, is a surrogate for $\mathcal{E}(\cdot,X) - \frac{\nu}{2}\|X\|_F^2$.

Proof. In deriving the surrogate via Jensen's inequality we follow [70]. For each $(i,j) \in \{1,\ldots,L\} \times \{1,\ldots,N\}$, using that $(AX)_{ij} > 0$, we introduce a vector σ_{ij} of weights

$$\sigma_{ija} := \frac{A_{ia}X_{aj}}{(AX)_{ij}}\ .$$

By convexity of $-\log$ we obtain

$$\begin{aligned} -\log KX &= \left[-\log\left(\sum_a \sigma_{ija} \frac{K_{ia} X_{aj}}{\sigma_{ija}} \right) \right]_{ij} \\ &\leq \left[-\sum_a \sigma_{ija} \log\left(\frac{K_{ia} X_{aj}}{\sigma_{ija}} \right) \right]_{ij} \\ &= \left[-\frac{1}{(AX)_{ij}} \sum_a A_{ia} X_{aj} \left(\log \frac{K_{ia}}{A_{ia}} + \log(AX)_{ij} \right) \right]_{ij} \\ &= -\frac{1}{AX} \circ \left(A \circ \log \frac{K}{A} \right) X \;-\; \log AX \end{aligned} \tag{4.12}$$

with equality at the point $K = A$. Using this expression we obtain for $K \in \mathbb{R}^{L,p}_{\geq}$ and $A \in \mathbb{R}^{L,p}_{>}$ that

$$\begin{aligned} \mathcal{E}(K, X) &= \langle \mathbb{1}_{L,N}, KX \rangle_F - \langle Y, \log(KX) \rangle_F + \frac{\mu}{2} \|K\|_F^2 \\ &\leq \langle \mathbb{1}_{L,N}, KX \rangle_F - \left\langle \frac{Y}{AX}, \left(A \circ \log \frac{K}{A} \right) X \right\rangle_F - \langle Y, \log AX \rangle_F + \frac{\mu}{2} \|K\|_F^2 \end{aligned} \tag{4.13}$$

with equality for $K = A$, which proves the surrogate properties (4.10). □

Next we discuss the corresponding update obtained from (4.9). Define $\mathcal{A}_2 \colon \Omega \to \mathbb{R}^{L,p}_{\geq}$ by

$$\mathcal{A}_2(K, X) = \frac{1}{2\mu} \left(\sqrt{(\mathbb{1}_{L,N} X^\top) \circ (\mathbb{1}_{L,N} X^\top) + 4\mu K \circ \frac{Y}{KX}\, X^T} - \mathbb{1}_{L,N} X^\top \right). \tag{4.14}$$

Note that, if $(K, X) \in \Omega$, for some $K \in \mathbb{R}^{L,p}_{>}$, then X has no zero column. (Indeed X having no zero column is equivalent to $\Omega_1(X)$ being nonempty.) Suppose further that X has no zero rows. Then the equality $\frac{1}{2\mu}(\sqrt{B^2 + 4\mu A} - B) = \frac{2A}{\sqrt{B^2 + 4\mu A} + B}$, for real numbers $A, B > 0$, allows us to rewrite $\mathcal{A}_2$ in the form

$$\mathcal{A}_2(K, X) = \frac{2K \circ \left(\frac{Y}{KX} X^\top \right)}{\mathbb{1}_{L,N} X^\top + \sqrt{\mathbb{1}_{L,N} X^\top \circ \mathbb{1}_{L,N} X^\top + 4\mu K \circ \left(\frac{Y}{KX} X^\top \right)}}. \tag{4.15}$$

Thus, for $X \in \mathbb{R}^{p,N}_{>}$, the operator $\mathcal{A}_2(\cdot, X) \colon K^{(r)} \mapsto K^{(r+1)}$ maps $\mathbb{R}^{L,p}_{>}$ into $\mathbb{R}^{L,p}_{>}$.

Remark 4.3 *The mapping $\mathcal{A}_2$ does not have a continuous extension to $\mathbb{R}^{L,p}_{>} \times \mathbb{R}^{p,N}_{\geq}$.*

Proof. Consider $T\colon x \mapsto K \circ \left(\frac{y}{Kx} x^\top\right)$ on $\mathbb{R}^{p,N}_>$ for $p = L = 2$, $N = 1$, $y = \mathbb{1}_{L,N}$, and $K = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$. We show that the limit $x \to (0,0)^\top$ is not well defined. Indeed, consider for $s \geq 0$ the sets

$$\gamma_s := \{(x_1, x_2)^\top \in \mathbb{R}^2 \colon x_1 = s x_2\}.$$

On γ_s we then have

$$T(x) = K \circ \left(\frac{\mathbb{1}_{L,N}}{Kx} x^\top\right) = K \circ \begin{pmatrix} \frac{s}{2s+1} & \frac{1}{2s+1} \\ \frac{s}{s+2} & \frac{1}{s+2} \end{pmatrix} = \begin{pmatrix} \frac{2s}{2s+1} & \frac{1}{2s+1} \\ \frac{s}{s+2} & \frac{2}{s+2} \end{pmatrix}$$

and therefore

$$\lim_{\substack{x \to 0 \\ x \in \gamma_{\frac{1}{2}}}} T(x) \neq \lim_{\substack{x \to 0 \\ x \in \gamma_1}} T(x)$$

i.e. the limit depends on the direction along which x approaches zero. Note that T as above is one summand in $K \circ \frac{Y}{KX} X^\top$ for $N > 1$, so that $X \mapsto K \circ \frac{Y}{KX} X^\top$ is not continuous when individual columns approach zero. □

The following lemma is given to outline the later procedure. From the preceding remark, it is clear, that we should require that X has no zero column. In fact this is automatically fulfilled if $(K, X) \in \Omega$ for some $K \in \mathbb{R}^{L,p}_\geq$. The following lemma makes the additional assumption that X has no zero rows. Our later convergence analysis will build on a column-wise version to be given as Corollary 4.9, which is proven without this additional assumption.

Lemma 4.4 *Let $X \in \mathbb{R}^{L,p}_\geq$ have no zero column and no zero row, let $Y \in \mathbb{R}^{L,N}_>$ and let $\mathcal{E}^X_{\mathrm{sur},2}$ be the surrogate from Lemma 4.2. Then the following statements hold.*

(i) *For any $K^{(r)} \in \mathbb{R}^{L,p}_>$ we have $\mathcal{A}_2(K^{(r)}, X) = \operatorname{argmin}_{K \in \mathbb{R}^{L,p}_>} \mathcal{E}^X_{\mathrm{sur},2}(K, K^{(r)})$.*

(ii) *Writing $K^{(r+1)} := \mathcal{A}_2(K^{(r)}, X)$, we have $\mathcal{E}(K^{(r+1)}, X) \leq \mathcal{E}(K^{(r)}, X)$.*

Proof of Lemma 4.4. Once (i) has been shown, (ii) follows by Lemma 4.1. Therefore we now show (i). Let $A := K^{(r)} \in \mathbb{R}^{L,p}_>$. From the assumptions it follows that $AX > 0$ and thus $\varphi := \mathcal{E}^X_{\mathrm{sur},2}(\cdot, A)$ is differentiable on $\mathbb{R}^{L,p}_>$ and the gradient with respect to $\langle \cdot, \cdot \rangle_F$ is given by

$$\nabla \varphi(K) = -\frac{A}{K} \circ \left(\frac{Y}{AX} X^\top\right) + \mathbb{1}_{L,N} X^\top + \mu K.$$

The summand of $\mathcal{E}^X_{\mathrm{sur},2}(\cdot, A)$ depending on the component K_{ia} is

$$-A_{ia} \Big(\frac{Y}{AX} X^\top\Big)_{ia} \log K_{ia} + (\mathbb{1}_{L,N} X^\top)_{ia} K_{ia} + \frac{\mu}{2} K_{ia}^2.$$

As $0 < A_{ia} \big(\frac{Y}{AX} X^\top\big)_{ia}$, this summand is strictly convex and coercive for $K_{ia} \in \mathbb{R}_>$. Hence $\varphi = \mathcal{E}^X_{\mathrm{sur},2}(\cdot, A)$ attains its unique minimum at a point $K \in \mathbb{R}^{L,p}_>$, at which $\nabla \varphi$ vanishes, i.e. which satisfies

$$\mu\, K \circ K + K \circ \mathbb{1}_{L,N} X^T - A \circ \frac{Y}{AX}\, X^T = 0.$$

Note that for $\mu = 0$ the solution is given by

$$K = \frac{A \circ \left(\frac{Y}{AX}\ X^\top\right)}{\mathbb{1}_{L,N} X^\top}.$$

For $\mu \neq 0$, the quadratic equation has two solutions for each component of K, of which only one is positive. It follows that the unique minimizer K of $\mathcal{E}^X_{\mathrm{sur},2}(\cdot, A)$ in $\mathbb{R}^{L,p}_>$ is given by

$$K = \frac{1}{2\mu}\left(\sqrt{(\mathbb{1}_{L,N} X^\top) \circ (\mathbb{1}_{L,N} X^\top) + 4\mu A \circ \frac{Y}{AX}\ X^T} - \mathbb{1}_{L,N} X^\top\right),$$

which proves (i). □

Observe that the objective function $\mathcal{E}$ is separable with respect to rows of K after removing terms depending only on X. We define $e_2 \colon \mathbb{R}^{1,p}_\geq \times \mathbb{R}^{p,N}_\geq \times \mathbb{R}^{1,N}_> \to \mathbb{R} \cup \{+\infty\}$ by

$$e_2(\cdot, X, y) \colon k \mapsto \langle (kX - y \circ \log kX)^\top, \mathbb{1}_N \rangle + \frac{\mu}{2}\|k\|_2^2 \tag{4.16}$$

and have

$$\mathcal{E}(K, X) = \sum_{k=1}^{L} e_2(k_l, X, y_l) + \frac{\nu}{2}\|X\|_F^2, \tag{4.17}$$

where k_l and y_l denote the l-th rows of K, Y respectively.

Similarly, the X-dependent part of $\mathcal{E}$ can be written as a sum of terms, each depending only on one column of X. Define $e_1 \colon \mathbb{R}^{L,p}_\geq \times \mathbb{R}^{p}_\geq \times \mathbb{R}^{L}_> \to \mathbb{R} \cup \{+\infty\}$ by

$$e_1(K, \cdot, y) \colon x \mapsto \langle Kx - y \circ \log Kx, \mathbb{1}_L \rangle + \frac{\nu}{2}\|x\|_2^2 \tag{4.18}$$

so that

$$\mathcal{E}(K, X) = \sum_{j=1}^{N} e_1(K, \underline{x}(j), \underline{y}(j)) + \frac{\mu}{2}\|K\|_F^2, \tag{4.19}$$

where $\underline{x}(j)$ and $\underline{y}(j)$ denote the j-th columns of X, Y respectively.

Next we derive the X-update in a similar manner as the K-update. As the column-sum-to-one constraint needs to be taken into account, we consider one column of X at a time. The proof of the following lemma is analogous to the proof of Lemma 4.2.

Lemma 4.5 *Let $K \in \mathbb{R}^{L,p}_\geq$ have no zero rows and $y \in \mathbb{R}^L_>$. The function $e^K_{\mathrm{sur},1} \colon \mathbb{R}^p_\geq \times \mathbb{R}^p_> \to \mathbb{R} \cup \{+\infty\}$ given by*

$$e^K_{\mathrm{sur},1}(x, a) := \langle \mathbb{1}_L, Kx \rangle - \left\langle \frac{y}{Ka},\ K\left(a \circ \log \frac{x}{a}\right) + Ka \circ \log Ka \right\rangle + \frac{\nu}{2}\|x\|_2^2 \tag{4.20}$$

is a surrogate for $e_1(K, \cdot, y)$.

We next find the corresponding update and deduce the decrease of the objective from strong convexity of the surrogate. Let $C \subseteq \mathbb{R}^d$ be a convex set. Recall that a function $f\colon C \to \mathbb{R} \cup \{+\infty\}$ is **strongly convex** with modulus $\mu > 0$ if $f - \frac{\mu}{2}\|\cdot\|_2^2$ is convex.

For an equivalent characterization, recall the usual order relation on the set of positive semidefinite $d \times d$ matrices defined by $B \succeq A$ if and only if $B - A$ is **positive semidefinite**, i.e. $x^\top (B - A)x \geq 0$ for all $x \in \mathbb{R}^d$. Then the following characterization holds.

Lemma 4.6 *Let C be an open convex subset of $\mathbb{R}^d$ and $\mu > 0$ and let $f\colon C \to \mathbb{R} \cup \{+\infty\}$ be twice differentiable on C. Then f is strongly convex with modulus μ if and only if*

$$\operatorname{Hess}_f(x) \succeq \mu \mathrm{I}.$$

Proof. Recall that a function $g\colon C \to \mathbb{R} \cup \{+\infty\}$ is convex if and only if $\operatorname{Hess}_f(x)$ is positive semidefinite for all $x \in C$, see [21, Proposition 1.2.6]. Applying this to $g = f - \frac{\mu}{2}\|\cdot\|_2^2$ we obtain that $f\colon C \to \mathbb{R} \cup \{+\infty\}$ is strongly convex with modulus $\mu > 0$ if and only if, for all $x \in C$,

$$0 \preceq \operatorname{Hess}_g(x) = \operatorname{Hess}_g(x) - \mu \mathrm{I}. \qquad \square$$

Strong convexity implies the following lower bound on the growth. For $x, y \in \mathbb{R}^d$ let $[x, y]$ denote the line segment between x and y.

Lemma 4.7 *Let $C \subset \mathbb{R}^d$ be an open convex set and $f\colon C \to \mathbb{R} \cup \{+\infty\}$ a proper lsc function which is strongly convex with parameter $\mu > 0$. Let $x, y \in C$ and assume that*

(i) *f is twice differentiable in a neighbourhood of $[x, y]$ and*

(ii) *$f(x) \leq f(z)$ for all points $z \in [x, y]$.*

Then

$$f(y) - f(x) \geq \frac{\mu}{2}\|y - x\|_2^2. \tag{4.21}$$

Proof. We are going to apply Taylor's theorem to the function $\varphi\colon [0, 1] \to \mathbb{R} \cup \{+\infty\}$ defined by $\varphi(t) := f\big(x + t(y - x)\big)$. Note that $\varphi'(t) = \big\langle \nabla f\big(x + t(y - x)\big),\, y - x\big\rangle$ and $\varphi''(t) = \langle y - x, H_t(y - x)\rangle$, where $H_t := \operatorname{Hess}_f\big(x + t(y - x)\big) \in \mathbb{R}^{d,d}$ for $t \in [0, 1]$.

By Taylor's theorem there is a $\xi \in (0, 1)$ with $\varphi(1) - \varphi(0) = \varphi'(0) + \frac{1}{2}\varphi''(\xi)$. From assumption (ii) it follows that $\varphi'(0) \geq 0$ and thus

$$f(y) - f(x) \geq \tfrac{1}{2}\langle y - x, H_\xi(y - x)\rangle.$$

Now strong convexity of f implies that $\operatorname{Hess}_f\big(x + \xi(y - x)\big) \succeq \mu \mathrm{I}$ and we obtain (4.21). $\square$

Next, before using strong convexity of the surrogate, we define the domain for the update mapping of columns of X, in agreement with $\Omega_1(K)$ and similarly for rows of K. Let

$$\omega_1(K) := \big\{x \in \mathbb{R}^p_{\geq}\colon Kx > 0\big\}, \quad \text{for } K \in \mathbb{R}^{L,p}_{\geq},$$
$$\omega_2(X) := \big\{k \in \mathbb{R}^{1,p}_{\geq}\colon kX > 0\big\}, \quad \text{for } X \in \mathbb{R}^{p,N}_{\geq}.$$

Note that $\omega_1(K)$ is empty if K has zero rows and $\omega_2(X)$ is empty if X has zero columns.

Theorem 4.8 *Let $K \in \mathbb{R}_{\geq}^{L,p}$ have no zero rows, let $y \in \mathbb{R}_{>}^{L}$. Let the surrogate $e_{\mathrm{sur},1}^{K}$ for $e_1(K,\cdot,y)$ be defined by (4.20). For $\nu > 0$ define the operator $\mathcal{A}_1'(K,\cdot,y)\colon \omega_1(K) \to \mathbb{R}_{\geq}^{p}$ by*

$$\mathcal{A}_1'(K,x,y) := \frac{1}{2\nu}\left(\sqrt{(K^\top \mathbb{1}_L + l\mathbb{1}_p)^2 + 4\nu x \circ K^\top\Big(\frac{y}{Kx}\Big)} - (K^\top \mathbb{1}_L + l\mathbb{1}_p)\right), \tag{4.22}$$

where $l \in \mathbb{R}$ is the solution of

$$2\nu + (\|K\|_{1,1} + pl) - \mathbb{1}_p^\top \sqrt{(K^\top \mathbb{1}_L + l\mathbb{1}_p)^2 + 4\nu x \circ K^\top\Big(\frac{y}{Kx}\Big)} = 0.$$

Then the following statements hold:

(i) *For $x^{(r)} \in \mathbb{R}_{>}^{p}$, the function $e_{\mathrm{sur},1}^{K}(\cdot, x^{(r)})$ is strongly convex with modulus ν.*

(ii) *For $x^{(r)} \in \mathbb{R}_{>}^{p}$, the unique minimizer of $e_{\mathrm{sur},1}^{K}(\cdot, x^{(r)})$ in $\triangle_p$ is given by $\mathcal{A}_1'(K, x^{(r)}, y)$.*

(iii) *For $x^{(r)} \in \omega_1(K)$ and $x^{(r+1)} := \mathcal{A}_1'(K, x^{(r)}, y)$, we have*

$$e_1(K, x^{(r)}, y) - e_1(K, x^{(r+1)}, y) \geq \frac{\nu}{2}\|x^{(r+1)} - x^{(r)}\|_2^2.$$

If $\nu = 0$ and K does not have zero columns, define the operator $\mathcal{A}_1'(K,\cdot,y)\colon \mathbb{R}_{>}^{p} \to \mathbb{R}_{\geq}^{p}$ by

$$\mathcal{A}_1'(K,x,y) := \frac{x \circ K^\top \frac{y}{Kx}}{K^\top \mathbb{1}_L + l\mathbb{1}_p}, \tag{4.23}$$

where $l \in \mathbb{R}$ is the solution of

$$\sum_{q=1}^{p} \frac{x_q\big(K^\top \frac{y}{Kx}\big)_q}{\big(K^\top \mathbb{1}_L\big)_q + l} = 1. \tag{4.24}$$

Then, for $x^{(r)} \in \mathbb{R}_{>}^{p}$, a minimizer of $e_{\mathrm{sur},1}^{K}(\cdot, x^{(r)})$ in $\triangle_p$ is given by $\mathcal{A}_1'(K, x^{(r)}, y)$ and for $x^{(r+1)} := \mathcal{A}_1'(K, x^{(r)}, y)$ we have

$$e_1(K, x^{(r+1)}, y) \leq e_1(K, x^{(r)}, y).$$

Proof of Theorem 4.8. (i) Let $\varphi := e_{\mathrm{sur},1}^{K}(\cdot, a)$. Then

$$\mathrm{Hess}_\varphi(x) = \mathrm{diag}\Big(\frac{a}{x \circ x} \circ K^\top\Big(\frac{y}{Ka}\Big)\Big) + \nu \mathrm{I} \succeq \nu \mathrm{I}.$$

and it follows from Lemma 4.6 that φ is strongly convex with modulus ν.

(ii) The gradient of on $\mathbb{R}_{>}^{p}$ is given by

$$\nabla\varphi(x) = K^\top \mathbb{1}_L - \frac{a}{x} \circ K^\top\Big(\frac{y}{Ka}\Big) + \nu x. \tag{4.25}$$

In order to find $\operatorname{argmin}_{\triangle_p} \varphi$ we introduce the Lagrangian for the constraint $\mathbb{1}_p^\top x = 1$. Let $\mathcal{L}\colon \mathbb{R}^p \times \mathbb{R}$ be given by

$$\mathcal{L}(x,l) := e^K_{\mathrm{sur},1}(x,a) + l(\mathbb{1}_p^\top x - 1) \tag{4.26}$$

We find the minimizer as in the proof of Lemma 4.4. The summand of $e^K_{\mathrm{sur},1}(x,a)$ depending on x_i is given by

$$\varphi^i(x_i) := (K^\top \mathbb{1})_i \, x_i - a_i \Big(K^\top \frac{y}{Ka}\Big)_i \log \frac{x_i}{a_i} + \frac{\nu}{2} x_i^2. \tag{4.27}$$

Suppose first that **K does not have zero columns**. Then $K^\top \mathbb{1}_L > 0$ and φ^i is coercive for $x_q \in \mathbb{R}_>$, for each $i = 1, \dots, p$. Hence the minimizer x of φ over $\triangle_p$ belongs to $\triangle_p \cap \mathbb{R}^p_> = \{x \in \mathbb{R}^p \colon \mathbb{1}^\top x = 1\} \cap \mathbb{R}^p_>$ and satisfies

$$0 = \nabla_x \mathcal{L}(x,l) = \nabla \varphi(x) + l\mathbb{1}_p = K^\top \mathbb{1}_L - \frac{a}{x} \circ K^\top \Big(\frac{y}{Ka}\Big) + \nu x + l\mathbb{1}_p \tag{4.28}$$

or equivalently

$$\nu x \circ x + x \circ (K^\top \mathbb{1}_L + l\mathbb{1}_p) - a \circ K^\top \Big(\frac{y}{Ka}\Big) = 0. \tag{4.29}$$

Case 1: $\nu > 0$

For $\nu \neq 0$, this quadratic equation has one positive solution in every component and for $l \in \mathbb{R}$ the unique minimizer $x^l \in \mathbb{R}^p$ of $\mathcal{L}(\cdot, l)$ is given by

$$x^l := \frac{1}{2\nu} \left(\sqrt{(K^\top \mathbb{1}_L + l\mathbb{1}_p)^2 + 4\nu a \circ K^\top \Big(\frac{y}{Ka}\Big)} - (K^\top \mathbb{1}_L + l\mathbb{1}_p) \right) \tag{4.30}$$

$$= \frac{2a \circ K^\top \left(\frac{y}{Ka}\right)}{K^\top \mathbb{1}_L + l\mathbb{1}_p + \sqrt{(K^\top \mathbb{1}_L + l\mathbb{1}_p)^2 + 4\nu a \circ K^\top \left(\frac{y}{Ka}\right)}}. \tag{4.31}$$

Next we show that there is a value $l \in \mathbb{R}$ for which x^l is in $\triangle_p$. From (4.31) we see $x^l \in \mathbb{R}^p_>$. Now $\mathbb{1}_p^\top x^l = 1$ is equivalent to $2\nu \mathbb{1}_p^\top x^l = 2\nu$, i.e.

$$2\nu = \mathbb{1}_p^\top \sqrt{(K^\top \mathbb{1}_L + l\mathbb{1}_p)^2 + 4\nu a \circ K^\top \Big(\frac{y}{Ka}\Big)} - (\mathbb{1}_p^\top K^\top \mathbb{1}_L + pl) =: F(l). \tag{4.32}$$

Defining, for $q = 1, \dots, p$ the quantities

$$\beta_q := (K^\top \mathbb{1}_L)_q, \quad \alpha_q := 4\nu a_q \Big(K^\top \frac{y}{Ka}\Big)_q \quad \text{and } f_q(l) := \sqrt{(\beta_q + l)^2 + \alpha_q} - (\beta_q + l)$$

we have

$$F(l) = \sum_{q=1}^{p} f_q(l).$$

From $K \geq 0$ having no zero columns, we deduce that $0 < \alpha_q$, $0 < \beta_q$ and that $f_q'(l) = (\beta_q + l)/\sqrt{(\beta_q + l)^2 + \alpha_q} - 1 < 0$ and $f_q''(l) = \alpha \left(\sqrt{(\beta_q + l)^2 + \alpha_q}\right)^{-3} > 0$. It follows

that f_q is convex, continuous and strictly decreasing on $\mathbb{R}$ with $\lim_{t\to-\infty} f_q(t) = +\infty$ and $\lim_{t\to+\infty} f_q(t) = 0$ for $q = 1, \dots, p$. The same then holds for $F\colon \mathbb{R} \to \mathbb{R}_>$, $F(l) = \sum_{q=1}^{p} f_q(l)$. Therefore (4.32) has a unique solution $l \in \mathbb{R}$.
Case 2: $\nu = 0$
For $\nu = 0$, the solution x^l of equation (4.29) is

$$x^l = \frac{a \circ K^\top \frac{y}{Ka}}{K^\top \mathbb{1}_L + l\mathbb{1}_p} \tag{4.33}$$

and l should be chosen so that

$$1 = \mathbb{1}_p^\top x^l = \sum_{q=1}^{p} \frac{a_q \left(K^\top \frac{y}{Ka}\right)_q}{\left(K^\top \mathbb{1}_L\right)_q + l} =: F(l). \tag{4.34}$$

Let $\gamma := \min_{q\in\{1,\dots,p\}} [K^\top \mathbb{1}_L]_q$. As F is continuous on $(-\gamma, +\infty)$ and strictly decreasing with $\lim_{t\searrow-\gamma} F(t) = +\infty$, $\lim_{t\to\infty} F(t) = 0$, there is a unique solution l of (4.34).

In both cases, the vector x^l, corresponding to the solution $l \in \mathbb{R}$ of (4.32) and (4.34) respectively, is a stationary point of $\varphi = e^K_{\mathrm{sur},1}(\cdot, a)$ on $\triangle_p \cap \mathbb{R}^p_>$. By strong convexity, φ has a unique minimzer and this minimizer is a stationary point and thus equals x^l.

Suppose now that **K has zero columns** and let $\nu > 0$. Note that this cannot occur in Algorithm 4 and is only relevant for the convergence analysis later on. Let $J \subset \{1, \dots, L\}$ be the set of indices with $k_q = 0$. For $q \in J$ the function φ^q from (4.27) becomes $\varphi^q(x_q) = \frac{\nu}{2} x_q^2$ so the x_q-dependent term of $\mathcal{L}(x, l)$ is

$$\frac{\nu}{2} x_q^2 + l x_q,$$

which is still coercive for $x_q \to +\infty$, but no longer for $x \to 0$. For a local minimizer x of $\mathcal{L}(x, l)$ we can thus only say $x_q \in [0, +\infty)$, in particular $x_q = 0$ is possible for every zero column k_q of K.

If $x_q > 0$ then optimality requires $0 = \partial_{x_q} \mathcal{L}(x, l) = \nu x_q + l$, which yields $x_q = -\frac{l}{\nu}$ and $l < 0$. Conversely, for $x_q = 0$, optimality yields $0 \le \partial_{x_q} \mathcal{L}(x, l) = l$. Together we have for every $q \in J$ that

$$\begin{aligned} l \ge 0 \quad &\Leftrightarrow \quad x_q = 0, \\ l < 0 \quad &\Leftrightarrow \quad x_q = -\frac{l}{\nu}. \end{aligned}$$

This agrees with formula (4.30), which for $k_q = 0$ reads

$$(x^l)_q = \frac{1}{2\nu}(\sqrt{l^2} - l) = \begin{cases} -\frac{l}{\nu} & \text{for } l < 0, \\ 0 & \text{for } l \ge 0. \end{cases}$$

We summarize that formula (4.30) remains valid for all $q \in \{1, \dots, p\}$.

It remains to show the existence of a Lagrange multiplier l for which $\mathbb{1}_p^\top x^l = 1$. For $q \in J$ we have

$$\alpha_q = 0, \quad \beta_q = 0, \quad f_q(l) = \sqrt{l^2} - l = \begin{cases} -2l & \text{for } l < 0, \\ 0 & \text{for } l \geq 0. \end{cases}$$

Thus, f_q is still convex continuous and decreasing in $\mathbb{R}$ if $q \in J$. For indices $j \notin J$, the previous argument for nonzero columns of K remains valid and guarantees at least one j for which f_j is strictly decreasing. It follows that the function F is convex continuous and strictly decreasing and it still holds that (4.32) has a unique solution $l \in \mathbb{R}$, as required.

For later reference we note, also if K has zero columns, it still holds that $x_q > 0$ implies

$$0 = \partial_{x_q} \mathcal{L}(x^l, l) \tag{4.35}$$

for x^l given by (4.30) .

(iii) If $x^{(r)} \in \mathbb{R}^p_>$, then the claim follows from Lemma 4.7 for $\nu > 0$ because φ is strongly convex with modulus ν and $C = \triangle_p$ is convex. For $\nu = 0$ the claimed decrease of the objective is just the surrogate property.

If $x := x^{(r),\varepsilon} \in \omega_1(K) \setminus \mathbb{R}^p_>$, we make no claim for $\nu = 0$. To show that the strict decrease still holds in the case $\nu > 0$, we use continuity and an approximation argument. Let $x^{(r),\varepsilon} := x + \varepsilon \mathbb{1}_p$. Then

$$x^{(r),\varepsilon} \in \mathbb{R}^p_>, \quad x^{(r),\varepsilon} \to x^{(r)} \text{ for } \varepsilon \to 0.$$

and by continuity of $\mathcal{A}'$ also

$$\mathcal{A}'(K, x^{(r),\varepsilon}, y) =: w^\varepsilon \quad \longrightarrow \quad x^{(r+1)} := \mathcal{A}'(K, x^{(r)}, y). \tag{4.36}$$

As $x^{(r),\varepsilon} \in \mathbb{R}^p_>$ and the surrogate $e^K_{\mathrm{sur},1}(\cdot, x^{(r),\varepsilon})$ is still defined, by the previous argument, we get

$$e_1(K, x^{(r),\varepsilon}, y) - e_1(K, w^\varepsilon, y) \geq \|w^\varepsilon - x^{(r),\varepsilon}\|_2^2.$$

Passing to the limit $\varepsilon \to 0$ with the help of (4.36) we obtain

$$e_1(K, x^{(r)}, y) - e_1(K, x^{(r+1)}, y) \geq \|x^{(r+1)} - x^{(r)}\|_2^2. \qquad \square$$

Analogously, we introduce the operator $\mathcal{A}'_2(k, X, y')$ obtained as the restriction of $\mathcal{A}_2(K, X)$ to corresponding rows k, y' of K, Y, and state that it decreases the objective e_2.

Corollary 4.9 *Suppose that $X \in \mathbb{R}^{p,N}_\geq$ has no zero columns, let $y' \in \mathbb{R}^{1,N}_>$ and $\mu > 0$. Define the operator $\mathcal{A}'_2(\cdot, X, y') \colon \omega_2(X) \to \mathbb{R}^{1,p}_\geq$ by*

$$\mathcal{A}'_2(k, X, y') := \frac{1}{2\mu} \left(\sqrt{\mathbb{1}_N^\top X^\top \circ \mathbb{1}_N^\top X^\top + 4\mu k \circ \left(\frac{y'}{kX} X^\top \right)} - \mathbb{1}_N^\top X^\top \right). \tag{4.37}$$

Then, for $k^{(r)} \in \omega_1(X)$ and $k^{(r+1)} := \mathcal{A}'_2(k^{(r)}, X, y')$, we have

$$e_2(k^{(r)}, X, y') - e_2(k^{(r+1)}, X, y') \geq \frac{\mu}{2} \|k^{(r+1)} - k^{(r)}\|_2^2.$$

Proof of Corollary 4.9. The K-update was obtained directly in Lemma 4.2 and Lemma 4.4, because there were no row-constraints to be taken into account. Of course we could have proceeded row-wise in an analogous way to the column-wise procedure for the X-update. In this way, one obtains the update (4.15) row-wise, which for given X,y' and k is just (4.37). The decrease property follows by repeating the argument used to prove Theorem 4.8 (i)–(iii). □

The convergence analysis requires that the operator $\mathcal{A}_1$ continuously extends to the set of limit points of iterates and that this extension still leads to a decrease of the energy. The set of limit points might contain pairs (K, X) with K having zero columns. Therefore we restrict the later convergence analysis to the case $\nu > 0$, in which continuity of $\mathcal{A}_1$ and decrease of e_2 have been shown also for points (K, X) with K having zero columns. Note that a continuous extension of the update to $\mathcal{A}_1'$, defined for $\nu = 0$ by the equations (4.33) and (4.34), from

$$Z_{1,\text{cols}} \times \mathbb{R}^p_> \times \mathbb{R}^L_> \quad \text{to} \quad Z_1 \times \mathbb{R}^p_> \times \mathbb{R}^L_>,$$

where

$$\begin{aligned} Z_1 \quad &:= \{K \in \mathbb{R}^{L,p}_{\geq} \colon K \text{ has no zero rows}\}, \\ Z_{1,\text{cols}} &:= \{K \in \mathbb{R}^{L,p}_{\geq} \colon K \text{ has no zero rows and no zero columns}\}, \end{aligned}$$

is not trivial. Indeed, the following naive extension α of $\mathcal{A}_1'$ does not lead to a decrease in the energy function.

Example 4.10 *For a given $K \in Z_1$ and $y \in \mathbb{R}^p_>$ let $J := \{q \colon k_q = 0\} \subset \{1, \dots, p\}$. Define the operator $\alpha(K, \cdot, y) \colon \mathbb{R}^p_> \to \mathbb{R}^p_{\geq}$ by*

$$\alpha(x)_j = \begin{cases} \frac{x_j (K^\top \frac{y}{Kx})_j}{(K^\top \mathbb{1}_L)_j + l}, & j \notin J \\ 0, & j \in J, \end{cases}$$

where l is the root of

$$\sum_{j \notin J} \frac{x_j \left(K^\top \frac{y}{Kx}\right)_j}{\left(K^\top \mathbb{1}_L\right)_j + l} = 1.$$

Choose some $k_1, k_2, y_1, y_2 \in \mathbb{R}_>$ satisfying

$$1 > \frac{y_1 + y_2}{k_1 + k_2} =: t^*, \tag{4.38}$$

and define

$$K = \begin{pmatrix} k_1 & 0 \\ k_2 & 0 \end{pmatrix}, \quad x^* = \begin{pmatrix} t^* \\ 1 - t^* \end{pmatrix}, \quad y = \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}.$$

Then, introducing $x^{**} =: \alpha(K, x^*, y)$, we have

$$e_1(K, x^{**}, y) > e_1(K, x^*, y).$$

Note that (4.38) is satisfied for $y_1 = y_2 = \frac{1}{2}$, $k_1 = k_2 = 1$.

Proof. Let $x_t = (t, 1-t)^\top$, $t \in (0,1)$. We then have

$$\begin{aligned} e_1(K, x_t, y) &= \langle Kx_t - y \circ \log Kx_t, \mathbb{1}_2 \rangle \\ &= \langle tk - y \circ \log tk, \mathbb{1}_2 \rangle \\ &= tk_1 - y_1 \log(tk_1) + tk_2 - y_2 \log(tk_2) \\ &= t(k_1 + k_2) - (y_1 + y_2) \log t - y_1 \log k_1 - y_2 \log k_2. \end{aligned}$$

The function

$$f(t) := t(k_1 + k_2) - (y_1 + y_2) \log t$$

is strongly convex and coercive on $\mathbb{R}_>$ and its unique minimizer is the root t^* of $f'(t) = k_1 + k_2 - (y_1 + y_2)$. Assuming (4.38), it follows that

$$f(t^*) < f(1).$$

Using that

$$e_1(K, x_1, y) - e_1(K, x_{t^*}, y) = f(1) - f(t^*)$$

the claim follows if $x_1 = \alpha(K, x_{t^*}, y)$. Now $J = \{2\}$, $\alpha(x)_2 = 0$ and the columns sum condition $\alpha(x)_1 + \alpha(x)_2 = 1$ implies $\alpha(x) = 1$, i.e., $x_1 = \alpha(K, x_{t^*}, y)$, as required. □

Repeating the argument from Lemma 4.8 for $x = \underline{x}(j)^{(r)}$, $y = \underline{y}(j)$, $j = 1, \ldots, N$ where $\underline{x}(j)$, $\underline{y}(j)$ denote the j-th columns of $X^{(r)}, Y$ respectively, we obtain in each column a column update $\underline{x}(j)^{(r+1)}$. Collecting the corresponding Lagrange multipliers l_j into a vector $\mathbf{l} = (l_1, \ldots, l_N)^\top$, the full update $X^{\mathbf{l}} = (\underline{x}(1)^{(r+1)}, \ldots, \underline{x}(N)^{(r+1)})$ can be written

$$X^{(r+1)}(\mathbf{l}) := \frac{1}{2\nu} \left(\sqrt{(K^\top \mathbb{1}_{L,N} + \mathbb{1}_p \mathbf{l}^\top)^2 + 4\nu X^{(r)} \circ K^\top \Big(\frac{Y}{KX^{(r)}} \Big)} - (K^\top \mathbb{1}_{L,N} + \mathbb{1}_p \mathbf{l}^\top) \right) \tag{4.39}$$

where the jth component l_j of $\mathbf{l} \in \mathbb{R}^N$ is defined by (4.32) with $a = \underline{x}(j)^{(r)}$, i.e. l_j is the root of

$$2\nu = \mathbb{1}_p^\top \sqrt{(K^\top \mathbb{1}_L + l_j \mathbb{1}_p)^2 + 4\nu \underline{x}(j)^{(r)} \circ K^\top \Big(\frac{\underline{y}(j)}{K\underline{x}(j)^{(r)}} \Big)} - (\mathbb{1}_p^\top K^\top \mathbb{1}_L + p l_j). \tag{4.40}$$

By convexity of the right-hand side, we can find l_j with a few Newton iterations.

Let $\mathcal{A}_1 \colon \Omega \to \mathbb{R}^{p,N}_{\geq}$ denote the mapping $(K, X^{(r)}) \mapsto X^{(r+1)}$ defined by equations (4.39) and (4.40). For $(K, X^{(r)}) \in \Omega$ and $K \in \mathbb{R}^{L,p}_>$ equation (4.39) can be written

$$X^{(r+1)}(\mathbf{l}) = \frac{2X^{(r)} \circ K^\top \Big(\frac{Y}{KX^{(r)}} \Big)}{K^\top \mathbb{1}_{L,N} + \mathbb{1}_p \mathbf{l}^\top + \sqrt{(K^\top \mathbb{1}_{L,N} + \mathbb{1}_p \mathbf{l}^\top)^2 + 4\nu X^{(r)} \circ K^\top \Big(\frac{Y}{KX^{(r)}} \Big)}},$$

and the operator $\mathcal{A}_1(K, \cdot) \colon X^{(r)} \mapsto X^{(r+1)}$ maps $\mathbb{R}^{p,N}_>$ into $\mathbb{R}^{p,N}_>$.

Alternating the updates (4.39) and (4.15) we obtain Algorithm 4.

Algorithm 4 Multiplicative Algorithm for Blind Unmixing (4.5), adapted from [69]

Input: Given are the data matrix $Y \in \mathbb{R}_{>}^{L,N}$ and $p \in \mathbb{N}$, as well as regularization parameters $\nu \geq 0$, $\mu > 0$.
Initialization: Initialize $z^{(0)} = \left(K^{(0)}, X^{(0)}\right) \in \mathbb{R}_{>}^{L,p} \times \mathbb{R}_{>}^{p,N}$ with arbitrary strictly positive values, satisfying the column normalization constraints

$$\sum_{q=1}^{p} X_{qj}^{(0)} = 1 \qquad \text{for } j = 1, \ldots, N.$$

At each iteration: (until stopping criterion is reached)

i) Introduce a vector $\mathbf{l} = (l_1, \ldots, l_N)^\top$ of Lagrange parameters, with l_j corresponding to the constraint that the jth column of $X^{(r+1)}$ should sum to one, and determine its entries as follows:

a) if $\nu \neq 0$, for $j = 1, \ldots, N$ determine l_j by solving equation (4.40), i.e.

$$2\nu + \mathbb{1}_p^\top \left(K^{(r)\top} \mathbb{1}_L + l_j \mathbb{1}_p\right) - \mathbb{1}_p^\top \sqrt{(K^{(r)\top} \mathbb{1}_L + l_j \mathbb{1}_p)^2 + 4\nu\, a \circ K^{(r)\top}\left(\frac{y}{K^{(r)} a}\right)} = 0,$$

where a denotes the jth column of $X^{(r)}$ and y denotes the jth column of Y.

b) if $\nu = 0$, for $j = 1, \ldots, N$ determine l_j by solving equation (4.34), i.e.

$$1 = \sum_{i=1}^{p} \frac{a_i \left[K^{(r)\top} \frac{y}{K^{(r)} a}\right]_i}{\left[K^{(r)\top} \mathbb{1}_L\right]_i + l_j};$$

ii) compute $A^{(r)} = X^{(r)} \circ K^{(r)\top}\left(\frac{Y}{K^{(r)} X^{(r)}}\right)$,

iii) compute $B^{(r)} = K^{(r)\top} \mathbb{1}_{LN} + \mathbb{1}_p \mathbf{l}^\top$,

iv) set $X^{(r+1)} = \dfrac{2A^{(r)}}{B^{(r)} + \sqrt{B^{(r)} \circ B^{(r)} + 4\nu A^{(r)}}}$,

v) compute $C = K^{(r+1)} \circ \left(\dfrac{Y}{K^{(r)} X^{(r+1)}}\, X^{(r+1)\top}\right)$,

vi) compute $D^{(r+1)} = \mathbb{1}_{LN} X^{(r+1)\top}$,

vii) set $K^{(r+1)} = \dfrac{2C^{(r+1)}}{D^{(r+1)} + \sqrt{D^{(r+1)} \circ D^{(r+1)} + 4\mu C^{(r+1)}}}$.

We stated that the Lagrange multipliers can be found with a few Newton iterations and briefly summarize the underlying results. The **Newton scheme** to solve $f(x) = 0$ for a differentiable function $f\colon \mathbb{R} \to \mathbb{R}$ is the iteration defined by

$$x^{(k+1)} = x^{(k)} - \frac{f(x^{(k)})}{f'(x^{(k)})}, \qquad k = 0, 1, \ldots, \tag{4.41}$$

for some starting point $x^{(0)} \in \mathbb{R}$. The right-hand side of (4.41) is the intersection point of the tangent

$$f(x^{(k)}) + f'(x^{(k)})(x - x^{(k)})$$

to the graph of f in the point $(x^{(k)}, f(x^{(k)}))$ with the x-axis.

For convex functions the Newton scheme converges globally, compare [54, Satz 18.3].

Theorem 4.11 *Let $\mathcal{I} \subset \mathbb{R}$ be an interval and $f\colon \mathcal{I} \to \mathbb{R}$ differentiable strictly decreasing and convex with a root $\hat{x} \in \mathcal{I}$. Then the Newton iterations converge for all $x^{(0)} \in \mathcal{I}$ with $x^{(0)} \leq \hat{x}$ monotonically to $\hat{x}$. If $\mathcal{I} = \mathbb{R}$ and $x^{(0)} \in \mathbb{R}$ is chosen arbitrarily, then convergence is monotone after the first iteration.*

The same results holds we replace the assumption that f be "strictly decreasing" and $x_0 \leq \hat{x}$ by "strictly increasing" and $x_0 \geq \hat{x}$.

Proof. The first part of the proof can be found in [54, Satz 18.3] and is given here for the convenience of the reader. Assuming that $x^{(k)} \leq \hat{x}$ for some $k \geq 0$, we prove the induction step

$$x^{(k)} \leq x^{(k+1)} \leq \hat{x}. \tag{4.42}$$

With f being differentiable, convexity implies that the graph of f is everywhere above the tangent to f in the point $x^{(k)}$, i.e.

$$f(x) \geq f(x^{(k)}) + f'(x^{(k)})(x - x^{(k)}) =: t(x).$$

For $x = x^{(k+1)}$ the right-hand side of this equation vanishes by the update rule (4.41). Therefore $f(x^{(k+1)})$ is nonnegative and, by strict monotonicity of f, thus $x^{(k+1)} \leq \hat{x}$. On the other hand (4.41) implies $x^{(k+1)} \geq x^{(k)}$, using that $f'(x^{(k)})$ is negative by assumption, and $f(x^{(k)})$ is nonnegative by the induction hypothesis. Convergence now follows from monotonicity and boundedness of the iterations. By (4.41) the limit is a root of f. This proves the first statement.

Assuming $\mathcal{I} = \mathbb{R}$ and $x^{(0)} \geq \hat{x}$, we show that $x^{(1)} \leq \hat{x}$. Indeed, for any starting point $x^{(0)} \in \mathbb{R}$, convexity implies that the tangent $t(x)$ in $x^{(0)}$ is everywhere below the graph of f so that $t(x^{(0)}) = f(x^{(0)}) \geq 0 \geq t(\hat{x})$ and t has a root $x^{(1)}$ in $[x^{(0)}, \hat{x}]$, as required.

To replace strictly decreasing by strictly increasing, consider $g(x) := f(-x)$. □

We say that a sequence $x^{(k)}$ **converges quadratically** to $\hat{x}$ if it converges to $\hat{x}$ and

$$\|x^{(k+1)} - \hat{x}\| \leq C\|x^{(k)} - \hat{x}\|^2$$

for some $C > 0$ and all sufficiently large $k \in \mathbb{N}$. Informally this means that "the number of correct digits doubles in every iteration". We quote the following well known result on the convergence on the quadratic convergence of the Newton iterates [54, Theorem 19.1]:

Theorem 4.12 *Let $\|\cdot\|$ and $\|\cdot\|_{\mathrm{op}}$ be compatible norms in $\mathbb{R}^n$ and $\mathbb{R}^{n,n}$, and let $\Omega \subset \mathbb{R}^n$ be open. Let $F\colon \Omega \to \mathbb{R}^n$ be continuously differentiable with a root $\hat{x} \in \Omega$. Let the Jacobi matrix $F'(x)$ be invertible for all x in a ball $\mathcal{U} \subset \Omega$ around $\hat{x}$ and assume that*

$$\|F'(x)^{-1}\big(F'(y) - F'(x)\big)\|_{\mathrm{op}} \le L\|y - x\| \tag{4.43}$$

for all $x, y \in \mathcal{U}$ with some constant $L > 0$. Then the Newton iterates converge locally quadratically to $\hat{x}$.

The next remark details that these results apply and gives suitable starting points.

Remark 4.13 1. *While for the case $\nu > 0$, any starting point can be chosen by the second statement of Theorem 4.11, this is not the case for $\nu = 0$, where too large choices of the initial estimate $l^{(0)}$ could lead to a first iterate $l^{(1)}$ outside the the range of definition $(-\gamma, +\infty)$.*

First we give starting points $l^{(0)}$ for which Theorem 4.11 guarantees convergence for the cases $\nu > 0$ and $\nu = 0$. Recall the definitions of α_q, β_q and f_q, $q = 1, \dots, p$ from the proof of Theorem 4.8. For $\nu > 0$ we estimate

$$f_q(l) = \sqrt{(\beta_q + l)^2 + \alpha_q} - (\beta_q + l) \ge -(\max_q \beta_q + l).$$

Choosing $l^{(0)} := -\frac{2\nu}{p} - \max_q \beta_q$ ensures that the right-hand side is at least $\frac{2\nu}{p}$ so that $F(l^{(0)})$ is at least 2ν and $l^{(0)}$ is smaller than the desired root $\hat{l}$ of $F(l) = 2\nu$.

For $\nu = 0$, define $\alpha'_q := a_q(K^\top \frac{y}{Ka})_q$. Let $r \in \{1, \dots, p\}$ be the index of a smallest entry of the vector $K^\top \mathbb{1}_L$, i.e. an index with $(K^\top \mathbb{1}_L)_r = \gamma$. Estimating F from (4.32) by

$$F(l) = \sum_{q=1}^{p} \frac{\alpha'_q}{(K^\top \mathbb{1}_L)_q + l} \ge \frac{\alpha'_r}{(K^\top \mathbb{1}_L)_r + l}$$

we see that $l^{(0)} := \alpha'_r - \gamma$ yields $F(l^{(0)}) \ge 1$ and thus $l^{(0)}$ is smaller than the root of $F(l^{(0)}) - 1$.

Let us summarize. In the case $\nu > 0$, we choose

$$l^{(0)} := -\frac{2\nu}{p} - \max_q \beta_q$$

and in the case $\nu = 0$, we choose

$$l^{(0)} := \alpha'_r - \gamma \quad \text{for } r \in \operatorname*{argmin}_{r \in \{1,\dots,p\}} (K^\top \mathbb{1}_L)_r.$$

With these choices, Theorem 4.11 guarantees that the Newton iterations with starting point $l^{(0)}$ converge monotonically to roots of (4.32) and (4.34) respectively.

2. *We see next that quadratic convergence holds. Observe that both for $\nu > 0$ and $\nu = 0$ the function F' is M-Lipschitz on $[l^{(0)}, +\infty)$ for some $M > 0$, because F'' is bounded*

on $[l^{(0)}, +\infty)$. *Next,* $|F'|$ *is bounded below by* $\varepsilon > 0$ *on a neighbourhood* U *of* $\hat{l}$ *if we set* $\varepsilon := \frac{1}{2}|F'(\hat{l})|$, *where* $\hat{l}$ *is the desired root.*

Together, for all $l_1, l_2 \in [l^{(0)}, +\infty) \cap U$, *we have*

$$|F'(l_1)^{-1}\big(F'(l_2) - F'(l_1)\big)| \leq \frac{M}{\varepsilon}|l_1 - l_2|, \tag{4.44}$$

so F *satisfies condition* (4.43). *By Theorem 4.12, the Newton iterates converge locally quadratically both for* $\nu > 0$ *and* $\nu = 0$.

Indeed, looking at the proof of [54, Satz 19.1], we see that is is sufficient to require (4.43) *on line segments between successive iterates to obtain that for all* k *the inequality*

$$\|l^{(k+1)} - \hat{l}\|_2 \leq \frac{M}{2\varepsilon}\|l^{(k)} - \hat{l}\|_2^2 \tag{4.45}$$

holds. By 1. *we have* $l^{(k)} \in [l^{(0)}, \hat{l}]$ *for all* k, *and* (4.43) *is satisfied on* $[l^{(0)}, \hat{l}]$. *Therefore* (4.45) *holds for all* k.

4.5. Convergence of the NMF Algorithm

In this section we prove that the sequence of iterates produced by Algorithm 4 has a convergent subsequence and that the every limit point of a subsequence is a fixed point. Further every fixed point in the interior of the positive quadrant satisfies the KKT-conditions.

Recall the partial updates $\mathcal{A}_1 \colon (K^{(r)}, X^{(r)}) \mapsto X^{(k+1)}$ and $\mathcal{A}_2 \colon (K^{(r)}, X^{(k+1)}) \mapsto K^{(k+1)}$. Writing $\pi_1 \colon (K, X) \mapsto K$ for the projection and similarly for π_2, the concatenation

$$\mathcal{A} = (\mathcal{A}_2, \pi_2) \circ (\pi_1, \mathcal{A}_1) \colon (K^{(r)}, X^{(r)}) \mapsto (X^{(k+1)}, K^{(k+1)}) \tag{4.46}$$

defines one iteration of Algorithm 4. We note continuity.

Lemma 4.14 *The partial updates* $\mathcal{A}_1$ *and* $\mathcal{A}_2$ *are continuous on* Ω *and hence* $\mathcal{A} = \mathcal{A}_2 \circ \mathcal{A}_1$ *is continuous on* Ω.

Proof. As the function $t \mapsto \frac{1}{t}$ is continuous on $(0, +\infty)$, the mapping $(K, X) \mapsto \frac{1}{KX}$ is continuous on Ω and it follows from (4.14) that the mapping $\mathcal{A}_2 \colon (K^{(r)}, X) \mapsto K^{(r+1)}$ is continuous.

Next we show that $\mathcal{A}_1$ is continuous. The Lagrange multipliers l_j, $j = 1, \ldots, N$ are defined by (4.40) or (4.34), which are both polynomials in l_j and therefore depend continuously [55] on the coefficients of these polynomials, which in turn depend continuously on $X \circ (K^{(r)})^\top\big(\frac{Y}{K^{(r)}X}\big)$ and $(K^{(r)})^\top \mathbb{1}_n$. Then continuity of the mapping $\mathcal{A}_1$ follows as for $\mathcal{A}_2$. □

Of course continuity holds also row- and column-wise for the row- and column-wise update operators $\mathcal{A}_1', \mathcal{A}_2'$ defined by (4.22) and (4.37).

Corollary 4.15 *Let $L, p, N \in \mathbb{N}$. Then the following assertions hold.*

(i) *For $y \in \mathbb{R}_{>}^{L}$ and $K \in \mathbb{R}_{\geq}^{L,p}$ the operator $\mathcal{A}_1'(K, \cdot, y)$ is continuous on the set $\omega_1(K)$.*

(ii) *For $y \in \mathbb{R}_{>}^{1,N}$ and $X \in \mathbb{R}_{\geq}^{p,N}$ the operator $\mathcal{A}_2'(\cdot, X, y)$ is continuous on the set $\omega_2(X)$.*

Before giving the convergence proof, we restate the immediate implication of Lemma 4.8 that the cost functions are decreasing with a strict decrease except at fixed points. Recall the summands e_1, e_2 of the objective defined by (4.18) and (4.16).

Lemma 4.16 *Let $\nu > 0$, $X \in \mathbb{R}_{\geq}^{p,N}$ and $y \in \mathbb{R}_{>}^{1,N}$. For any $k^* \in \omega_2(X)$ and $k^{**} := \mathcal{A}_2'(k^*, X, y)$ the following statements hold:*

(i) $e_2(k^{**}, X, y) \leq e_2(k^*, X, y)$.

(ii) *If $e_2(k^{**}, X, y) = e_2(k^*, X, y)$ then $k^{**} = k^*$.*

Similarly, let $\mu > 0$, let $K \in \mathbb{R}_{\geq}^{L,p}$ and $y \in \mathbb{R}_{>}^{L}$. For $x^* \in \omega_1(K)$ and $x^{**} := \mathcal{A}_1'(K, x^*, y)$ the following statements hold:

(iii) $e_1(K, x^{**}, y) \leq e_1(K, x^*, y)$.

(iv) *If $e_1(K, x^{**}, y) = e_1(K, x^*, y)$ then $x^{**} = x^*$.*

Now we turn to convergence. For brevity we denote the iterates generated by Algorithm 4 by $z^{(r)} := (K^{(r)}, X^{(r)})$, $r = 1, 2, \dots$.

Theorem 4.17 *Let $z^{(0)} \in \mathbb{R}_{>}^{L,p} \times \mathbb{R}_{>}^{p,N}$ and $\nu > 0$. The sequence of iterates $z^{(r)}$ generated by Algorithm 4 with starting point $z^{(0)}$ has a convergent subsequence. Furthermore, every limit point of a subsequence is a fixed point of $\mathcal{A}$.*

Proof. 1) For a starting point $z^{(0)} \in \mathbb{R}_{>}^{L,p} \times \mathbb{R}_{>}^{p,N}$ we have $\mathcal{E}(z^{(0)}) < +\infty$. By (4.11), the energy is decreasing under the updates (4.8) and (4.9). Additionally $\mathcal{E}$ is bounded below by zero and therefore

$$\lim_{r \to \infty} \mathcal{E}(z^{(r)}) = \mathcal{E}^*,$$

for some $\mathcal{E}^* \in \mathbb{R}$.

By the decrease all iterates $z^{(r)} = (K^{(r)}, X^{(r)})$ are contained in the $\mathcal{E}(z^{(0)})$-sublevel set of $\mathcal{E}$, which we denote by Π. Recall that coercivity implies that for some $\varepsilon > 0$, we have $\Pi \subseteq \Omega_\varepsilon$. Note further that Π is closed and bounded and thus compact.

2) Let $z^* = (K^*, X^*)$ be an accumulation point of $\{z^{(r)}\}_r$ and r_j a subsequence with $z^* = \lim_{j \to \infty} z^{(r_j)}$. By continuity of $\mathcal{E}$, we have $\mathcal{E}(z^*) = \mathcal{E}^*$ and if z^{**} is any other accumulation point, then

$$\mathcal{E}(z^{**}) = \mathcal{E}^* = \mathcal{E}(z^*). \tag{4.47}$$

3) Consider the sequence of next iterates $w^j := \mathcal{A}(z^{r_j})$ where $\mathcal{A}$ is the update from (4.46). Let $\{j_k\}_k \subset \mathbb{N}$ be a convergent subsequence with $\lim_{k \to \infty} w^{j_k} =: z^{**} = (K^{**}, X^{**}) \in \Pi \subseteq \Omega_\varepsilon$. By continuity of $\mathcal{A}$, we can pass to the limit $k \to \infty$ in $w^{j_k} = \mathcal{A}(z^{r_{j_k}})$ and obtain

$$z^{**} = \mathcal{A}(z^*).$$

As the objective $\mathcal{E}$ is decreasing under the substeps $\mathcal{A}_1, \mathcal{A}_2$, in the same way as (4.47) was derived, it also follows that

$$\mathcal{E}(K^{**}, X^{**}) = \mathcal{E}(K^*, X^{**}) = \mathcal{E}(K^*, X^*). \tag{4.48}$$

4) We aim to show that $z^{**} = z^*$. To this end we now show that the first equality in (4.48) implies $K^{**} = K^*$, by considering all rows of K^* individually.

By (4.17) we have

$$\mathcal{E}(K^*, X^*) = \sum_{l=1}^{L} e_2(k_l^*, X^*, y_l) + \frac{\nu}{2}\|X^*\|_F^2,$$

where k_l^*, k_l^{**} and y_l denote the l-th rows of K^*, K^{**} and Y respectively, $l = 1, \dots, L$. Subtracting the corresponding equation for $\mathcal{E}(K^{**}, X^*)$ yields

$$\mathcal{E}(K^*, X^*) - \mathcal{E}(K^{**}, X^*) = \sum_{l=1}^{L} \Big(e_2(k_l^*, X^*, y_l) - e_2(k_l^{**}, X^*, y_l) \Big). \tag{4.49}$$

Observing that $(K^*, X^*) \in \Pi \subseteq \Omega_\varepsilon$ implies $k_l^* \in \omega_2(X)$, $l = 1, \dots, L$, and recalling from Corollary 4.9 that the update operator $\mathcal{A}_2 \colon (K^*, X^*) \mapsto (K^{**}, X^*)$ is row-wise represented by $\mathcal{A}_2'$, i.e.,

$$k_l^{**} = \mathcal{A}_2'(k_l^*, X^*, y_l), \quad l = 1, \dots, L,$$

we deduce with Lemma 4.16 that

$$e_2(k_l^{**}, X^*, y_l) \leq e_2(k_l^*, X^*, y_l), \quad l = 1, \dots, L. \tag{4.50}$$

Hence each summand on the right-hand side of (4.49) is nonnegative. By the first equation in (4.48) the sum is zero, hence

$$e_2(k_l^{**}, X^*, y_l) = e_2(k_l^*, X^*, y_l), \quad l = 1, \dots, L. \tag{4.51}$$

Using again Lemma 4.16, we obtain

$$k_l^{**} = k_l^*, \quad l = 1, \dots, L. \tag{4.52}$$

5) Similarly we conclude from the second equality in (4.48) that $X^{**} = X^*$, by splitting $\mathcal{E}(K^*, X^{**})$ and $\mathcal{E}(K^*, X^*)$ into sums corresponding to columns of X^*, X^{**}. Together with 4) we have shown $z^{**} = z^*$ and z^* is a fixed point of $\mathcal{A}$. □

Still, fixed points in the interior of the positive quadrant satisfy the KKT-conditions for minimization problem (4.4).

Theorem 4.18 *Let $Y \in \mathbb{R}_>^{L,N}$ and $\nu > 0$. Let $z = (K, X) \in \mathbb{R}_>^{L,p} \times \mathbb{R}_>^{p,N}$ be a fixed point of $\mathcal{A}$. Then z satisfies the KKT-conditions for the minimization of $\mathcal{E}$ over $\mathbb{R}_\geq^{L,p} \times \triangle_p^N$, which, on $\mathbb{R}_>^{L,p} \times \mathbb{R}_>^{p,N}$ become*

(i) $\nabla_K \mathcal{E}(K, X) = 0$,

(ii) $\nabla_{\underline{x}(j)} \mathcal{E}(K, X) \in span(\mathbb{1}_p), \quad j = 1, \ldots, N$,

where $\underline{x}(j)$ denotes the j-th column of X.

Proof. From the definition of $\mathcal{E}$ in (4.5) we compute

$$\nabla_K \mathcal{E}(K, X) = \mathbb{1}_{L,N} X^\top - \frac{Y}{KX} X^\top + \mu K,$$
$$\nabla_X \mathcal{E}(K, X) = K^\top \mathbb{1}_{L,N} - K^\top \frac{Y}{KX} + \nu X.$$

As z is a fixed point of $\mathcal{A}$, we have $\mathcal{A}_1(K, X) = X$ and $\mathcal{A}_2(K, X) = K$. On the one hand, from $\mathcal{A}_2(K, X) = K$ it follows by Lemma 4.4 that K is a minimizer of $\varphi = \mathcal{E}^X_{\mathrm{sur},2}(\cdot, A)$ over $\mathbb{R}_>^{L,p}$ for $A = K$ and satisfies

$$0 = \nabla\varphi(K) = -\frac{A}{K} \circ \Big(\frac{Y}{AX} X^\top\Big) + \mathbb{1}_{L,N} X^\top + \mu K$$

with $A = K$. This proves (i).

On the other hand, from $\mathcal{A}_1(K, X) = X$ it follows first of all for each $j \in \{1, \ldots, N\}$ that the j-th columns $\underline{x}(j)$ of X belongs to $\triangle_p$. Further $\underline{x}(j)$ minimizes the Lagrangian for the minimization over $\triangle_p$ of the surrogate corresponding to the j-th column, defined in (4.26) by

$$\mathcal{L}(x, l) = e^K_{\mathrm{sur},1} K(x, a) + l(\mathbb{1}_p^\top x - 1),$$

where now the old column is $a = \underline{x}(j)$. Details on the latter statement are in (4.28) and (4.35).

From (4.25) we have

$$\nabla_x e^K_{\mathrm{sur},1}(x, a) = K^\top \mathbb{1}_L - \frac{a}{x} K^\top \Big(\frac{y}{Ka}\Big) + \nu x.$$

In particular

$$\nabla_x e^K_{\mathrm{sur},1}(\underline{x}(j), \underline{x}(j)) = \nabla_{\underline{x}(j)} \mathcal{E}(K, X).$$

With $\underline{x}(j)$ minimizing $\mathcal{L}(\cdot, l)$ for $a = \underline{x}(j)$, it follows that

$$\begin{aligned} 0 = \nabla_x \mathcal{L}(\underline{x}(j), l) &= \nabla_x e^K_{\mathrm{sur},1}(\underline{x}(j), \underline{x}(j)) + l\mathbb{1}_p \\ &= \nabla_{\underline{x}(j)} \mathcal{E}(K, X) + l\mathbb{1}_p, \quad j = 1, \ldots, N, \end{aligned}$$

which is (ii). □

4.6. Convergence of Optimization Transfer Algorithms under Uniqueness

If we consider the simpler case where the iterative mapping stems from a strongly convex surrogate and the energy has a unique minimizer, we get the following much stronger result.

Theorem 4.19 *Let $\Omega \subset \mathbb{R}^d$ be closed, let $\mathcal{E}\colon \Omega \to \mathbb{R} \cup \{+\infty\}$ be bounded below, let $\mathcal{E}_{\text{sur}}$ be a surrogate for $\mathcal{E}$ and let $\mathcal{A}\colon \Omega \to \Omega$ denote an iterative update mapping which satisfies the following assumptions:*

(i) *$\mathcal{A}$ is continuous.*

(ii) *$\mathcal{A}$ has a unique fixed point.*

(iii) *$\mathcal{A}(x) = \operatorname{argmin}_\Omega \mathcal{E}_{\text{sur}}(\cdot, x)$ for every $x \in \Omega$.*

(iv) *$\operatorname{Hess}_{\mathcal{E}_{\text{sur}}(\cdot,x)}$ is uniformly bounded below by $\xi \mathrm{I}$, for all $x \in \Omega$, some $\xi > 0$.*

Then, if the sequence of iterates is contained in a compact set, it converges.

The above will typically be applied replacing (ii) with

(ii') Each fixed point of $\mathcal{A}$ is a minimizer of $\mathcal{E}$ and $\mathcal{E}$ is strictly convex.

Proof. Let $\{x^{(k)}\}_k$ be the sequence of iterates generated by $\mathcal{A}$ from a starting point $x^{(0)} \in \Omega$ via the update $x^{(k+1)} := \mathcal{A}(x^{(k)})$, $k = 0, 1, 2, \ldots$. Assumption (iii) and Lemma 4.1 imply

$$\mathcal{E}(x^{(k+1)}) \leq \mathcal{E}(x^{(k)}), \quad k = 0, 1, 2, \ldots,$$

and with $\mathcal{E}$ being bounded below, it follows that $\mathcal{E}(x^{(k)})$ is monotone decreasing to some limit $\mathcal{E}^* \in \mathbb{R}$. From the lower bound on the Hessian of the surrogate in assumption (iv) and the surrogate properties it follows that

$$\begin{aligned} \mathcal{E}(x^{(k)}) - \mathcal{E}(x^{(k+1)}) &\geq \mathcal{E}_{\text{sur}}(x^{(k)}, x^{(k)}) - \mathcal{E}_{\text{sur}}(x^{(k+1)}, x^{(k)}) \\ &\geq \xi \, \|x^{(k+1)} - x^{(k)}\|_2^2. \end{aligned} \tag{4.53}$$

With $\mathcal{E}(x^{(k)}) \searrow \mathcal{E}^*$ this implies

$$\|x^{(k+1)} - x^{(k)}\|_2 \leq \xi^{-\frac{1}{2}} \left(\mathcal{E}(x^{(k)}) - \mathcal{E}(x^{(k+1)})\right)^{\frac{1}{2}} \to 0 \tag{4.54}$$

for $k \to \infty$.

As the sequence of iterates is contained in a compact set, it has a convergent subsequence. Let $x^{(k_j)}$ be a convergent subsequence with limit $\tilde{x}$. By the triangle inequality,

$$\|\mathcal{A}(\tilde{x}) - \tilde{x}\|_2 \leq \|\mathcal{A}(\tilde{x}) - \mathcal{A}(x^{(k_j)})\|_2 + \|\mathcal{A}(x^{(k_j)}) - x^{(k_j)}\|_2 + \|x^{(k_j)} - \tilde{x}\|_2.$$

Taking the limit $j \to \infty$, the first term on the right-hand side goes to zero by continuity of $\mathcal{A}$. The second term goes to zero by (4.54). Thus all three terms on the right-hand side converge to zero as $j \to \infty$, and $\tilde{x}$ must be a fixed point of $\mathcal{A}$. By assumption (ii), the fixed point of $\mathcal{A}$ is unique, hence it follows that the whole sequence converges,

$$x^{(k)} \to \tilde{x} := Fix\,\mathcal{A}$$

for $k \to \infty$.

By the decrease property (4.53), every minimizer of $\mathcal{E}$ is a fixed point of $\mathcal{A}$. Hence $\mathcal{E}$ has a unique minimizer and this minimizer is given by $\tilde{x}$. We have thus shown that the whole sequence of iterates converges to the minimizer of $\mathcal{E}$. □

The assumptions in Theorem 4.19 are realistic in some settings and have been used in [39]. M. Defrise, C. Vanhove and X. Lin [39] consider reconstruction from CT data with TV-regularization, propose a block-iterative algorithm and prove its convergence to a regularized solution.

They consider a linear measurement system $y = Ax$, with a matrix $A \in \mathbb{R}^{MP,N}$. The goal is to reconstruct the image $x \in \mathbb{R}^N$, discretized on a grid of N voxels, from a set of M projections measured on a detector with P pixels. The measured data is written as a vector of M projections $y = (y_1^\top, y_2^\top, \ldots, y_M^\top)^\top$, with $y_j \in \mathbb{R}^P$, for $j = 1, \ldots, M$.

Accordingly they define a vector $\sigma = (\sigma_1^\top, \sigma_2^\top, \ldots, \sigma_M^\top)^\top \in \mathbb{R}^{MP}$ with σ_j denoting the estimated variances of measurement y_j. Assuming a Gaussian distribution of the measurements they wish to find the minimizer of the penalized weighted least squares cost function

$$\mathcal{E}(x) = \|\tfrac{1}{\sigma} \circ (y - Ax)\|_2^2 + \beta P(x),$$

where the first term is the data fidelity term and the second term ist the penalty weighted by the regularization parameter $\beta > 0$. Their penalty is a smooth discrete approximation of the TV-penalty

$$P(x) = \sum_{j=1}^{N} |\nabla_i x|_\varepsilon, \tag{4.55}$$

with the notation

$$|\nabla_i x|_\varepsilon := \sqrt{\varepsilon^2 + \sum_{j \in B_i} (x_i - x_j)^2}$$

for the discrete approximation of the magnitude of the gradient of the image at voxel i, calculated using a neighbourhood $B_i \subset \{1, \ldots, N\}$ of voxel i and a small value $\varepsilon > 0$. Typically in 2D, with $i = (i_x, i_y)$ one might take $B_{(i_x,i_y)} = \big\{(i_x + 1, i_y), (i_x, i_y + 1)\big\}$, with the two neighbours of (i_x, i_y) giving the finite difference estimate of the image gradient components along the x-axis and y-axis at (i_x, i_y). Let $\overline{B}_j$ denote the adjoint neighbourhood of $\overline{B}_j$, defined by $k \in \overline{B}_j$ iff $j \in B_k$. Then, for the stated 2D example, the adjoint neighbourhood is $\overline{B}_{i_x,i_y} = \big\{(i_x - 1, i_y), (i_x, i_y - 1)\big\}$.

They define the surrogate for the TV-penalty by

$$P_{\text{sur}}(x, \hat{x}) = \sum_{i=1}^{N} \frac{\varepsilon^2 + \sum_{j \in B_i} \{(x_i - x_j)(\hat{x}_i - \hat{x}_j) + (\hat{x}_j - x_j)^2 + (\hat{x}_i - x_i)^2\}}{2|\nabla_i \hat{x}|_\varepsilon}, \tag{4.56}$$

which has the form

$$P_{\text{sur}}(x, \hat{x}) = \sum_{i=1}^{N} p_i(\hat{x})(x_i - z_i(\hat{x}))^2 \quad + \text{terms independent of } x$$

for suitable vectors $p, z \in \mathbb{R}^N$ depending only on $\hat{x}$.

Their surrogate iteration is started at some initial estimate $x^{(0)} \in \mathbb{R}^N$ and obtains the $(k+1)$-th estimate as

$$x^{(k+1)} = \operatorname*{argmin}_{x \in \mathbb{R}^N} \mathcal{E}_{\text{sur}}(x, x^{(k)})$$

with

$$\mathcal{E}_{\text{sur}}(x, \hat{x}) = \|\tfrac{1}{\sigma} \circ (y - Ax)\|_2^2 + \beta P_{\text{sur}}(x, \hat{x}). \tag{4.57}$$

To solve this quadratic minimization problem they use the regularized fixed-block algorithm [42] as inner iteration.

Showing first that the $x^{(k)}$ are bounded, $\|x^{(k)}\|_2 \leq C$, for some constant C, they can deduce that $\beta p_i(x^{(k)}) \geq \xi$ for all i, k. Assume in the following that $\sigma_j = \mathbb{1}_p$ for all j. The general case follows by rescaling A and the data. Then the Hessian of $\mathcal{E}_{\text{sur}}$ satisfies

$$\frac{1}{2} \operatorname{Hess}_{\mathcal{E}_{\text{sur}}(\cdot, x^{(k)})} = A^\top A + \beta \operatorname{diag}\left((p_i(x))_{i=1}^N\right) \succeq \xi \mathrm{I},$$

which gives assumtion (iv) of Theorem 4.19. In their case, assumption (i) follows immediately from the formula for the update mapping as minimum of the quadratic surrogate and (ii) is deduced similar to (ii') by showing that every fixed point is a stationary point of $\mathcal{E}$, [39, Lemma 3], and appealing to uniqueness. To prove compactness of the iterates [39, Lemma 1], they require that the constant image is not in the null space of A and that pairs of pixels are linked by a sequence of neighbourhoods. Then, for $\beta > 0$, $\varepsilon > 0$, convergence follows as in Theorem 4.19, see [39, Theorem 1].

Part II.

Moreau-Yosida Envelopes in Hadamard Spaces

5. Prerequisites

Outline

In this chapter, we introduce basic notation from convex analysis in Banach and Hadamard spaces, including weak convergence, Γ-convergence and Mosco convergence. In the last section we recall two growth bounds needed in the sequel.

5.1. Banach Spaces

For normed spaces X, Y, let us denote by $\mathcal{L}(X;Y)$ the continuous linear functions from X into Y and write $\langle \cdot\,,\,\cdot \rangle : X \times X^* \to \mathbb{R}$ for the dual pairing. For a real normed space X, we call $X^* := \mathcal{L}(X, \mathbb{R})$ its **dual space**. Setting

$$\langle x^*\,,\, J_X x \rangle_{X^*} := \langle x\,,\, x^* \rangle_X \qquad \text{for } x \in X,\, x^* \in X^*$$

defines an isometric mapping $J_X \in \mathcal{L}(X; X^{**})$, where $X^{**} := (X^*)^* = \mathcal{L}(X^*; \mathbb{R})$ is the **second dual** of X. A Banach space X is called **reflexive** if and only if the above isometry J_X is surjective. For a good introduction and further properties see [2].

A Banach space (resp. its norm) is called **uniformly convex** if for every $\varepsilon > 0$ there is a $\delta > 0$ such that, for all points x, y in the unit ball, $\|x-y\| \geq \varepsilon$ implies $\|\frac{1}{2}(x+y)\| \leq 1-\delta$. The Milman–Pettis theorem says that a uniformly convex space is reflexive. Particular examples are the L^p-spaces for $1 < p < +\infty$.

The **domain** $\operatorname{dom} f$ of a function $f\colon X \to \mathbb{R} \cup \{+\infty\}$ is given by $\operatorname{dom} f := \{x \in X : f(x) < +\infty\}$ and f is **proper** if $\operatorname{dom} f \neq \emptyset$. A function $f\colon X \to \mathbb{R} \cup \{+\infty\}$ is called **lower semicontinuous** (lsc) if the level sets $\{x \in X\colon f(x) \leq \alpha\}$ are closed for all $\alpha \in \mathbb{R}$. A function $f\colon X \to \mathbb{R} \cup \{+\infty\}$ is **convex** if for all $x, y \in \operatorname{dom} f$ and all $t \in [0,1]$ the relation

$$f\left(tx + (1-t)y\right) \leq tf(x) + (1-t)f(y)$$

holds. For a proper convex lsc function $f\colon X \to \mathbb{R} \cup \{+\infty\}$, the **proximal mapping** $\operatorname{Prox}_{\lambda f}\colon X \to X$ is defined by

$$\operatorname{Prox}_{\lambda f}(x) = \operatorname*{argmin}_{y \in X} \Big\{ f(y) + \tfrac{1}{2\lambda}\, \mathrm{d}(x,y)^2 \Big\}, \tag{5.1}$$

where $\mathrm{d}(x,y) := \|x-y\|$. For brevity, instead of $\mathrm{Prox}_{\lambda f}$ we will also write $J_{\lambda f}$ or just J_λ if f is fixed. Indeed, the minimizer of the right-hand side exists and is unique. For $\lambda > 0$, the **Moreau-Yosida envelope** of f is given by

$$f_\lambda(x) := \inf_{y\in X}\Big\{ f(y) + \tfrac{1}{2\lambda}\,\mathrm{d}(x,y)^2 \Big\}. \tag{5.2}$$

Let X be a Banach space. A sequence $\{x_n\}_n \subseteq X$ **converges weakly** to $x \in X$, written $x_n \overset{w}{\to} x$, if and only if

$$\langle x_n\,,\,u\rangle \to \langle x\,,\,u\rangle \quad \text{for all } u \in X^*.$$

We call a function **weakly lower semicontinuous** (**weakly lsc**), if it is lower semicontinuous w.r.t. the weak convergence of sequences. Hence a weakly lsc function is automatically lower semicontinuous. It is difficult in general to prove weak lower semicontinuity. However, if $f\colon X \to \mathbb{R}$ is a convex function, then f is weakly lsc if and only if f is strongly lsc. For a proof of this result on finite-valued functions see [9, Theorem 2.1.2].

A Banach space has the **Kadets-Klee property**, also called **Radon-Riesz property**, if for every sequence $\{x_n\}_n$ it holds that

$$x_n \overset{w}{\to} x \text{ and } \|x_n\| \to \|x\| \quad \Leftrightarrow \quad x_n \to x. \tag{5.3}$$

In particular, Hilbert spaces have the Kadets-Klee property.

Let X be a reflexive Banach space and $f\colon X \to \mathbb{R}\cup\{+\infty\}$ a proper, convex, lsc function. The **subdifferential** ∂f of f is the possibly multivalued operator from X into X^* whose graph is equal to

$$\begin{aligned}\partial f :=& \{(x,\xi)\in X\times X^*\colon f(u) \geq f(x) + \langle \xi\,,\ u-x\rangle,\ \text{for all } u\in X\} \\ =& \{(x,\xi)\in X\times X^*\colon f(x) + f^*(x) = \langle \xi, x\rangle\}\,,\end{aligned} \tag{5.4}$$

see e.g. (3.46) in [8]. Note that we use the same symbol for the operator and its graph. In order to prove a subdifferential sum rule in infinite-dimensional manifolds, we will need the follwing result from [16, Corollary 16.38].

Theorem 5.1 *Let* H *be a real Hilbert space, let* $f_1, f_2\colon \mathrm{H} \to \mathbb{R}\cup\{+\infty\}$ *be proper convex lsc functions such that one of the following holds:*

(i) $\operatorname{dom} f_1 \cap \operatorname{int}\operatorname{dom} f_2 \neq \emptyset$.

(ii) $\operatorname{dom} f_2 = \mathrm{H}$.

Then $\partial(f_1+f_2) = \partial f_1 + \partial f_2$.

5.2. Γ-Convergence and Mosco Convergence

In this section we introduce Γ-convergence in a metric space $(\mathcal{X}, \mathrm{d})$, and Mosco convergence in a Banach space $(X, \|\cdot\|)$. Although we need only the definition of Mosco convergence, in service of the unacquainted reader, we give a motivation via the Young-Fenchel transform.

Γ-convergence (E. De Giorgi, 1975) has been called the "commonly-recognized notion of convergence for variational problems" by the author of the introductory textbook [26]. For a family of mathematical problems, depending on a parameter, one can often foresee a limiting behaviour, which may be captured by replacing the family with a new, often simpler problem, in which the parameter has disappeared. Examples range from materials with microstructure and homogenization, over discretization to convex relaxations of non-convex problems.

To capture the behaviour it is usually sufficient to consider a sequence of parameter values and we index over $n \in \mathbb{N}$ in the following. Rephrased in variational terms, for a family of minimization problems depending on a parameter,

$$\min\bigl\{f_n(x)\colon x \in \mathcal{X}_n\bigr\},$$

the asymptotic behaviour of the solutions is described by defining a limit problem (no longer depending on n),

$$\min\bigl\{f(x)\colon x \in \mathcal{X}\bigr\}.$$

Γ-convergence implies that minimizers of the limit function f are related to approximate minimizers of the functions f_n. Before we state this fundamental result, we need two definitions.

Definition 5.2 *A sequence $\{f_n\}_n\colon \mathcal{X} \to \mathbb{R} \cup \{+\infty\}$ is said to* **Γ-converge** *to $f\colon \mathcal{X} \to \mathbb{R} \cup \{+\infty\}$, as $n \to \infty$, abbreviated $f_n \xrightarrow{\Gamma} f$, if and only if, for each $x \in \mathcal{X}$, we have*

(i) *$f(x) \le \liminf_{n\to\infty} f_n(x_n)$, whenever $x_n \to x$, and*

(ii) *there is a sequence $\{y_n\} \subseteq \mathcal{X}$ such that $y_n \to x$ and $f_n(y_n) \to f(x)$.*

Weierstrass' theorem asserts that a lower semicontinuous function attains its minimum if it is coercive. We define equi-coercivity [26, Definition 1.19]:

Definition 5.3 *A function $f\colon \mathcal{X} \to \mathbb{R} \cup \{+\infty\}$ is* **coercive** *if for all $t \in \mathbb{R}$ the set $\{f \le t\}$ is precompact. A function $f\colon \mathcal{X} \to \mathbb{R} \cup \{+\infty\}$ is* **mildly coercive** *if there is a non-empty compact set $K \subseteq \mathcal{X}$ such that $\inf_{\mathcal{X}} f = \inf_K f$. A sequence $\{f_n\}_n$ is* **equi-mildly coercive** *if there is a non-empty compact set $K \subseteq \mathcal{X}$ such that $\inf_{\mathcal{X}} f_n = \inf_K f_n$, for all n.*

We quote further [26, Remark 1.20]: If f is coercive then it is mildly coercive. An example of a non-coercive, mildly coercive function is given by any periodic function $f\colon \mathbb{R} \to \mathbb{R}$.

We have the following fundamental result.

Theorem 5.4 *Let $\{f_n\}_n \colon \mathcal{X} \to \mathbb{R} \cup \{+\infty\}$, be an equi-mildly coercive sequence of functions Γ-converging to a proper function $f \colon \mathcal{X} \to \mathbb{R} \cup \{+\infty\}$. Let us call a sequence $\{x_n\}_n$* ***minimizing*** *for $\{f_n\}_n$, if $\lim_n \big(f_n(x_n) - \inf_{\mathcal{X}} f_n\big) = 0$. Then, the following statements hold:*

(i) $\lim\limits_{n\to\infty} \inf\limits_{\mathcal{X}} f_n = \min\limits_{\mathcal{X}} f$.

(ii) *Every cluster point of a minimizing sequence for $\{f_n\}_n$ is a minimizer of f.*

(iii) *Every minimizer of f is cluster point of a minimizing sequence.*

A proof in metric spaces can be found in [26, Theorem 1.21].

Remark 5.5 *Theorem 5.4 also holds in general topological spaces. We will need part (ii) in the general setting, proven e.g. in [8, Theorem 1.10]. The definition of τ-epi convergence in the latter [8, Definition 1.9], is equivalent to Γ-convergence by [26, Proposition 1.18].*

A sequence $\{f_n\}_n$ of functions defined on a Banach space, **Γ(seq.-w)-converges** to f, if Definition 5.2 holds with convergence replaced by weak convergence, i.e. writing $x_n \overset{w}{\to} x$ in part (i) and $y_n \overset{w}{\to} x$ in part (ii) therein.

The Γ-limit of a sequence of convex functions is convex. We will further need the following elementary lemma, which is a special case of [8, Corollary 1.18]:

Lemma 5.6 *Let $\{a_{n,k}\}_{k,n\in\mathbb{N}}$ be a double indexed family of numbers in $\mathbb{R} \cup \{+\infty\}$ with $\lim_{n\to\infty} \lim_{k\to\infty} a_{n,k} = A$. Then there is an nondecreasing mapping $n \colon \mathbb{N} \to \mathbb{N}$ with*

$$\lim_{k\to\infty} a_{n(k),k} = A.$$

Proof. There exists an increasing sequence $\{k_n\}_n \subset \mathbb{N}$ such that for all n and all $k \geq k_n$ one has $|a_{n,k} - \lim_{p\to\infty} a_{n,p}| < \frac{1}{n}$. Let further $k_0 := 1$ and define a nondecreasing mapping $n \colon \mathbb{N} \to \mathbb{N}$ by

$$n(l) := \begin{cases} 1 & \text{if } 1 \leq l < k_1, \\ j & \text{if } k_j \leq l < k_{j+1} \text{ for } j = 1, 2, \dots . \end{cases}$$

Then $l \geq k_j = k_{n(l)}$ for $l \in [k_j, k_{j+1})$ and thus for all $l \geq k_1$ we have

$$\big|a_{n(l),l} - \lim_{k\to\infty} a_{n(l),k}\big| \leq \tfrac{1}{n(l)}.$$

From $n(l) \to \infty$ as $l \to \infty$ it follows that

$$\big|\lim_{k\to\infty} a_{n(l),k} - A\,\big| \to 0,$$

as $l \to \infty$. Together we obtain for $l \to \infty$ that

$$\big|a_{n(l),l} - A\,\big| \leq \big|a_{n(l),l} - \lim_{k\to\infty} a_{n(l),k}\big| + \big|\lim_{k\to\infty} a_{n(l),k} - A\,\big| \to 0. \qquad \square$$

Clearly the proof can be adapted if the limits are replaced by lim inf, and for related statements see also [8, Lemma 1.15–1.17].

Remark 5.7 *Note that a topological formulation of* Γ*-convergence agrees with the sequential definition in Definition 5.2 whenever every point has a countable neighbourhood basis, that is, whenever for every point* $x \in \mathcal{X}$ *there is a sequence* $N_1, N_2, \ldots$ *of open neighbourhoods of* x *such that for every neighbourhood* N *of* x *there is an integer* j *with* $N_j \subseteq N$*. Such a space is called* ***first countable****. In particular, metric spaces are first countable.*

We would like to add a note of caution: For example the weak topology is not metrizable and a topological definition of Γ*-convergence for the weak topology is indeed different from the sequential definition adopted here, see the sketch of an example in [11, p. 270].*

A stronger notion of convergence in infinite-dimensional Banach spaces than the notion of Γ-convergence has been introduced by U. Mosco [83]. We briefly motivate the definition: Denote by $\Gamma_0(X)$ the family of all proper convex lsc functions on a reflexive Banach space X. The Young-Fenchel transform f^* of $f \in \Gamma_0(X)$ is defined on X^* by

$$f^*(x^*) = \sup_{x \in X} \left(\langle x^*, x \rangle - f(x) \right), \quad x^* \in X^*,$$

it satisfies $f^* \in \Gamma_0(X^*)$ and is involutory, i.e. $f^{**} = f$ provided we identify X^{**} with X. Let $\{f_n\}_n$ be a sequence of proper convex lsc functions on a separable reflexive Banach space X, which satisfies the following equi-coerciveness assumption: For every sequence $\{x_n\}_n \subset X$, $n \in \mathbb{N}$ we have

$$\sup_{n \in \mathbb{N}} f_n(x_n) < +\infty \quad \longrightarrow \quad \sup_{n \in \mathbb{N}} \|x_n\| < +\infty.$$

Then the following statements are equivalent [8, Theorem 3.11]:

(i) $f_n \overset{\Gamma(\text{seq.-w})}{\longrightarrow} f$, with respect to weak convergence of sequences on X.

(ii) $f_n^* \xrightarrow{\Gamma} f^*$, with respect to strong convergence on X^*.

The norm topology and weak topology are exchanged under conjugation. The following definition by U. Mosco, [83], combines both:

Definition 5.8 *A sequence* $\{f_n\}_n \colon X \to \mathbb{R} \cup \{+\infty\}$ ***converges*** *to* $f \colon X \to \mathbb{R} \cup \{+\infty\}$ ***in the sense of Mosco****, as* $n \to \infty$*, abbreviated* $f_n \xrightarrow{M} f$*, if and only if, for each* $x \in X$*, we have*

(i) $f(x) \leq \liminf_{n \to \infty} f_n(x_n)$*, whenever* $x_n \overset{w}{\to} x$*, and*

(ii) *there is a sequence* $\{y_n\} \subseteq X$ *such that* $y_n \to x$ *and* $f_n(y_n) \to f(x)$*.*

Now we can state [8, Theorem 3.18], on the space $\Gamma_0(X)$ of proper convex lsc functions defined on a reflexive Banach space X.

Theorem 5.9 *Let $\{f_n\}_n \subseteq \Gamma_0(X)$ and $f \in \Gamma_0(X)$. One then has the equivalence of the following statements.*

(i) $f^n \xrightarrow{M} f$ *in* $\Gamma_0(X)$

(ii) $(f^n)^* \xrightarrow{M} f^*$ *in* $\Gamma_0(X^*)$.

In other words, the Fenchel transformation is bicontinuous if we equip $\Gamma_0(X)$ and $\Gamma_0(X^*)$ with the topology of Mosco convergence.

5.3. Hadamard Spaces

In this section we collect prerequisites in Hadamard spaces. We start with basic definitions, then introduce weak convergence and related properties, and finally recall two results on the growth of convex functions.

5.3.1. Preliminaries

We start with the basic notation in Hadamard spaces. A complete metric space $(\mathcal{H}, \mathrm{d})$ is called a **Hadamard space** if every two points $x, y \in \mathcal{H}$ are connected by a geodesic and the following condition holds true

$$\mathrm{d}(x,v)^2 + \mathrm{d}(y,w)^2 \leq \mathrm{d}(x,w)^2 + \mathrm{d}(y,v)^2 + \mathrm{d}(x,y)^2 + \mathrm{d}(v,w)^2, \tag{5.5}$$

for all $x, y, v, w \in \mathcal{H}$. Inequality (5.5) implies that Hadamard spaces have nonpositive curvature. Recall that a curve $\gamma\colon [0,1] \to \mathcal{H}$ is called a **geodesic** if for all $t_1, t_2 \in [0,1]$ the relation

$$\mathrm{d}\big(\gamma(t_1), \gamma(t_2)\big) = |t_1 - t_2|\ \mathrm{d}\big(\gamma(0), \gamma(1)\big)$$

holds true. In Hadamard spaces, geodesics are unique. For $x, y \in \mathcal{H}$, we denote by $\gamma_{\widehat{x,y}}$ the geodesic starting at x and reaching y at time $t = 1$.

A set $\mathcal{C} \subseteq \mathcal{H}$ is **convex**, if for all $x, y \in \mathcal{H}$ the geodesic $\gamma_{\widehat{x,y}}$ lies in $\mathcal{C}$. The intersection of an arbitrary family of convex closed sets is itself a convex closed set. For a closed convex subset $\mathcal{C} \subset \mathcal{H}$ and $x \in \mathcal{H}$, the **metric projection** of x to $\mathcal{C}$ is the point

$$\mathcal{P}_{\mathcal{C}}(x) := \operatorname*{argmin}_{c \in \mathcal{C}} \mathrm{d}(x,c).$$

Usually we write $\mathcal{P}_{\mathcal{C}}\, x$ instead of $\mathcal{P}_{\mathcal{C}}(x)$. The metric projection exists and is unique, see [11, Theorem 2.1.12].

The definition of proper and lsc functions carries over from the Banach space setting. A function $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ is called **convex** if for all $x, y \in \mathcal{H}$ the function $f \circ \gamma_{\widehat{x,y}}$ is convex, i.e., we have

$$f\big(\gamma_{\widehat{x,y}}(t)\big) \leq t f\big(\gamma_{\widehat{x,y}}(0)\big) + (1-t) f\big(\gamma_{\widehat{x,y}}(1)\big), \text{ for all } t \in [0,1] \tag{5.6}$$

and **concave** if we have the opposite sign in (5.6). The function f is **strictly convex** if the strict inequality holds for all $0 < t < 1$, and **strongly convex** with parameter $\kappa > 0$ if for all $x, y \in \mathcal{H}$ and all $t \in [0,1]$ we have

$$f\big(\gamma_{\widehat{x,y}}(t)\big) \leq tf\big(\gamma_{\widehat{x,y}}(0)\big) + (1-t)f\big(\gamma_{\widehat{x,y}}(1)\big) - \kappa t(1-t)\,\mathrm{d}\big(\gamma_{\widehat{x,y}}(0), \gamma_{\widehat{x,y}}(1)\big).$$

The squared distance $\mathrm{d}^2(\cdot, y)\colon \mathcal{H} \to \mathbb{R}_{\geq 0}$ in a Hadamard space is strongly convex with $\kappa = 1$. Every proper strongly convex lsc functions $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ admits a unique minimizer, see, e.g., [11, Proposition 2.2.17].

For a proper convex lsc function $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ and $\lambda > 0$, the definition of the **Moreau-Yosida envelope** $f_\lambda\colon \mathcal{H} \to (-\infty, +\infty)$ and the **proximal mapping**, $\mathrm{Prox}_{\lambda f}\colon \mathcal{H} \to \mathcal{H}$ carries over from Banach spaces, with the inner product distance replaced by the Hadamard distance. By the strong convexity of $\mathrm{d}(\cdot, x)^2$, the function $f + \frac{1}{2\lambda}\,\mathrm{d}(\cdot, x)^2$ is also strongly convex so that again a unique minimizer exists. The proximal mapping $\mathrm{Prox}_{\lambda f}$ of a proper convex function f is continuous, see [11, Theorem 2.2.22]. The Moreau-Yosida envelope of a proper convex lsc function possesses the semigroup property

$$(f_\lambda)_\mu(x) = f_{\lambda+\mu}(x), \quad \lambda, \mu > 0.$$

Moreover we have for all $x \in \mathcal{H}$ that

$$\lim_{\lambda \to 0} f_\lambda(x) = f(x), \tag{5.7}$$

see [8, Theorem 2.46]. We will also need the following auxiliary lemma.

Lemma 5.10 *Let $(\mathcal{H}, d)$ be a Hadamard space and let $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ be a convex lsc function. Then, abbreviating $J_{\lambda f} := \mathrm{Prox}_{\lambda f}$, for all $x, y \in \mathcal{H}$ and $\lambda > 0$, we have*

$$f\left(J_{\lambda f}x\right) + \tfrac{1}{2\lambda} d\left(x, J_{\lambda f}x\right)^2 + \tfrac{1}{2\lambda} d\left(J_{\lambda f}x, y\right)^2 \leq f(y) + \tfrac{1}{2\lambda} d(x,y)^2. \tag{5.8}$$

A proof can be found in [11, Lemma 2.2.23].

5.3.2. Weak Convergence

To discuss the relation between the convergence of sequences $\{f_n\}_n$ and those of their Moreau-Yosida envelopes $\{f_{n,\lambda}\}_n$ we will need the convergence notion in the sense of Mosco in Hadamard spaces see e.g. in [11, p. 103]. To this end let us introduce weak convergence in Hadamard spaces. For a bounded sequence $\{x_n\}_n$ of points $x_n \in \mathcal{H}$, we define the function $\omega\colon \mathcal{H} \to [0, +\infty)$ by

$$\omega(x;\, \{x_n\}_n) := \limsup_{n \to \infty} \mathrm{d}(x, x_n)^2. \tag{5.9}$$

This function is strongly convex and therefore has a unique minimzer, which is called the **asymptotic center** of $\{x_n\}_n$, see [11, p. 58]. A sequence $\{x_n\}_n$ **weakly converges** to a point $x \in \mathcal{H}$ if it is bounded and x is the asymptotic center of each subsequence of $\{x_n\}_n$. We write $x_n \overset{w}{\to} x$. Note that $x_n \to x$ implies $x_n \overset{w}{\to} x$. Now Definition 5.8 of Mosco convergence carries over to Hadamard spaces.

For the weak limit we collect the following properties proved in [11, Propositions 3.1.2, 3.1.6, 3.1.3, Lemma 3.2.1].

Theorem 5.11 *Let $\mathcal{H}$ be a Hadamard space and $\{x_n\}_n$ a sequence of points $x_n \in \mathcal{H}$.*

(i) *If $\{x_n\}_n$ is bounded, then it has a weakly convergent subsequence.*

(ii) *Let $x \in \mathcal{H}$. We have $x_n \to x$ if and only if both $x_n \overset{w}{\to} x$ and, for some $y \in \mathcal{H}$, $d(x_n, y) \to d(x, y)$ (Kadets-Klee property).*

(iii) *Let $\{x_n\}_n$ be bounded. Then $x_n \overset{w}{\to} x$ if and only if $\mathcal{P}_\gamma x_n \to x$ as $n \to \infty$ for every geodesic $\gamma\colon [0,1] \to \mathcal{H}$ with $x \in \gamma$.*

(iv) *If $\{x_n\}_n$ is contained in a closed convex set $\mathcal{C} \subseteq \mathcal{H}$ and converges weakly to a point $x \in \mathcal{H}$, then $x \in \mathcal{C}$.*

The properties (i) and (iv) together imply that closed bounded convex sets in a Hadamard space are compact with respect to weak convergence of sequences.

A function $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ is **weakly lsc** if it is lsc with respect to weak convergence, i.e. for all $x \in \operatorname{dom} f$ and every sequence $x_n \overset{w}{\to} x$ we have $\liminf_{n\to\infty} f(x_n) \geq f(x)$. As before in Banach spaces, we give a convenient sufficient condition [11, Lemma 3.2.3]:

Lemma 5.12 *If $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ is a convex lsc function then f is weakly lsc.*

5.3.3. Growth Bounds for Convex Functions

To conclude this chapter we introduce two growth bounds following [11]: Lemma 5.15 holds for any convex function in a Hadamard space, and Lemma 5.16 holds uniformly along a Mosco convergent sequence of convex functions.

While closed and bounded sets are not compact in general Hadamard spaces, the following result known for compact sets in Euclidean space holds under the additional assumption of convexity [11, Prop. 2.1.16]:

Lemma 5.13 *Let $\mathcal{H}$ be a Hadamard space. If $\{\mathcal{C}_i\}_{i\in I}$ is a nested non-increasing family of closed bounded convex sets in $\mathcal{H}$, where I is an arbitrary directed set, then $\cap_{i\in I}\mathcal{C}_i \neq \emptyset$.*

The proof in [11, Prop. 2.1.16] fixes one point and uses negative curvature to deduce that the sequence of its projections to the sets $\mathcal{C}_i$ is Cauchy. Alternatively, the result follows naturally from the remark on compactness after Theorem 5.11. Next, we quote two lemmas from [11, Lemma. 2.2.12, Lemma 2.2.13]:

Lemma 5.14 *Let $\mathcal{H}$ be a Hadamard space and $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ be a convex lsc function. Then f is bounded from below on bounded sets.*

Proof. Let $\mathcal{C} \subseteq \mathcal{H}$ be bounded. Enlarging $\mathcal{C}$ if necessary, we may assume that $\mathcal{C}$ is closed and convex. If $\inf_{\mathcal{C}} f = -\infty$, then the sets $S_k := \{x \in \mathcal{C}\colon f(x) \leq -k\}$ for $k \in \mathbb{N}$ are all nonempty. Since each S_k is closed, convex and bounded, Lemma 5.13 yields a point $z \in \cap_{k\in\mathbb{N}} S_k$. By lower semicontinuity, $f(z) = -\infty$, which is a contradiction. □

Lemma 5.15 *Let $\mathcal{H}$ be a Hadamard space and let $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$, be a convex lsc function. For each $x_0 \in \mathcal{H}$ there exist constants $\alpha, \beta \in \mathbb{R}$ such that*

$$f(y) \geq -\alpha\, \mathrm{d}(y, x_0) - \beta,$$

for all $y \in \mathcal{H}$.

We omit the proof of Lemma 5.15 in [11, Lemma. 2.2.13], because it is a simpler version of the proof of Lemma 5.16, see Remark 5.17.

Mosco convergence of a sequence of convex functions to a proper function implies that the growth bound from Lemma 5.15 holds with the same constants for all functions in the sequence.

Lemma 5.16 *Let $\mathcal{H}$ be a Hadamard space and let $f_n\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$, $n \in \mathbb{N}$ be proper, convex, lsc functions. with $f_n \xrightarrow{M} f$, as $n \to \infty$, for some proper function $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$. Then, for each $x_0 \in \mathcal{H}$, there exist constants $\alpha, \beta \in \mathbb{R}$ such that*

$$f_n(y) \geq -\alpha\, \mathrm{d}(y, x_0) - \beta.$$

for all $y \in \mathcal{H}$ and all $n \in \mathbb{N}$.

The following argument is part of the proof of [11, Theorem 5.2.4] and given here, with small modifications, for the convenience of the reader.

Proof. Suppose the claim does not hold, that is, for each $k \in \mathbb{N}$ there exists $n_k \in \mathbb{N}$ and $x_k \in \mathcal{H}$ such that

$$f_{n_k}(x_k) < -k[\mathrm{d}(x_k, x_0) + 1]. \tag{5.10}$$

Without loss of generality we may assume $n_k \to \infty$ as $k \to \infty$, because otherwise, a subsequence of the indices is constant, $n_{k_l} = j$, for some index j, and (5.10) implies $\liminf_{l\to\infty} f^j(x_{k_l}) \leq -\infty$, contradicting Lemma 5.15.

If $\{x_k\}$ were bounded, then there would be some $\hat{x} \in \mathcal{H}$ and a subsequence $\{x_k\}$, still denoted $\{x_k\}$ such that $x_k \overset{w}{\to} \hat{x}$. By the Mosco convergence of the sequence $\{f_n\}$ this would imply

$$f(\hat{x}) \leq \liminf_{k\to\infty} f_{n_k}(x_k) \leq -\limsup_{k\to\infty} k[\mathrm{d}(x_k, x_0) + 1] \leq -\infty, \tag{5.11}$$

which is impossible.

Hence, we may assume that $\{x_k\}$ is unbounded. Choose $y_0 \in \operatorname{dom} f$ and find $y_k \to y_0$ such that $f_{n_k}(y_k) \to f(y_0)$. Define

$$z_k := (1 - t_k)y_k + t_k x_k, \qquad \text{with} \quad t_k := \frac{1}{\sqrt{k}\, \mathrm{d}(x_k, y_k)};$$

recall that this means $z_k := \gamma(t_k)$, where $\gamma\colon [0,1] \to \mathcal{H}$ is the geodesic with $\gamma(0) = y_k$, $\gamma(1) = x_k$.

Then $z_k \to y_0$. By convexity

$$\begin{aligned} f_{n_k}(z_k) &\leq (1-t_k) f_{n_k}(y_k) + t_k f_{n_k}(x_k) \\ &\leq (1-t_k) f_{n_k}(y_k) - t_k k[\mathrm{d}(x_k, x_0) + 1] \\ &\leq (1-t_k) f_{n_k}(y_k) - \sqrt{k}\, \frac{\mathrm{d}(x_k, x_0) + 1}{\mathrm{d}(x_k, y_k)}. \end{aligned}$$

Hence, using that $f_{n_k}(y_k) \to f(y_0) \in \mathbb{R}$, we get

$$f(y_0) \leq \liminf_{k\to\infty} f_{n_k}(z_k) \leq -\infty,$$

which is not possible, either. □

Remark 5.17 *The proof of Lemma 5.15 is simpler, see [11, p. 39]. The final contradiction follows directly from lower semi-continuity of f, where Mosco convergence is used in Lemma 5.16. Before, instead of inequality* (5.11)*, we argue as follows: If the sequence $\{x_n\}$ were bounded, then the analogue to* (5.10)*, $f(x_k) < -k[\mathrm{d}(x_k, x_0) + 1]$ would contradict Lemma 5.14.*

6. Moreau-Yosida Envelopes in Hadamard Spaces

Outline

Proximal mappings and Moreau-Yosida envelopes of convex functions play a central role in convex analysis. In particular, they appeared in various minimization algorithms which have recently found application in image processing and machine learning. For overviews, see for instance [16, 28, 87].

In this chapter we study Moreau-Yosida envelopes and some implications of negative curvature. In Section 6.1 we begin with an introduction to trigonometry in constant negative curvature, and thereafter introduce nonzero curvature bounds via comparison geometry.

In Section 6.2 we consider the implications of negative curvature on differentiability properties of Moreau-Yosida envelopes and prove several properties of independent interest: Under the assumption of a lower bound on the curvature we show that weak convergence in Hadamard manifolds implies weak convergence in the tangent space. Finally we prove a formula for the gradient of Moreau-Yosida envelopes in terms of the proximal mapping on Hadamard manifolds.

In Section 6.3 we first use these properties to prove the equivalence between pointwise convergence of the Moreau-Yosida envelopes and Mosco convergence in Hadamard manifolds. Then we prove the same result in the more general setting of Hadamard spaces. This work in Hadamard spaces from Section 6.3.3 has been continued in [15].

6.1. Introduction to Hadamard Manifolds

In this section we consider Hadamard spaces $(\mathcal{H}, \mathrm{d})$ which are at the same time connected Riemannian manifolds: we deal with Hadamard manifolds, for which we use the same notation $(\mathcal{H}, \mathrm{d})$.

We start with an introduction and basic properties in Section 6.1.1 and Section 6.1.2. In Section 6.1.3 we present different hyperbolic model spaces and trigonometry. In Section 6.1.4 we prove that Hadamard manifolds with bounded sectional curvatures are geodesic spaces with the same curvature bounds. Our exposition builds upon the equivalence of two curvature definitions, in Theorem 6.10, and the Cartan-Alexandrov-Topogonov Comparison Theorem. We use Corollary 6.12 later in Section 6.2.3.

6.1.1. Relevance of Hadamard Manifolds

We restrict to Hadamard manifolds for several reasons. On the one hand, in the study of convex functions, the assumption of nonpositive curvature is natural. Indeed, on complete Riemannian manifolds with finite volume, there is no nontrivial continuous convex function [108]. The authors of [36] generalize the result to strictly monotone vector fields. While for extended valued functions nontrivial examples exist, other problems occur without the assumption of nonpositive curvature:

Consider e.g. the two-dimensional sphere $\mathcal{S}^2 \subset \mathbb{R}^3$ with its intrinsic metric. The ball $\mathrm{B}_1((1,0,0)^\top) =: C$, contained in the upper hemisphere, is a closed convex set and the distance of the south pole $(-1,0,0)^\top$ to C is minimized at every point of the boundary of C. Moreover, the projection $\mathcal{P}_C$ to C is discontinuous at the south pole and the squared distance $\frac{1}{2}\,\mathrm{d}_C(\cdot)^2$ from C is neither convex nor differentiable in a neighbourhood of the south pole. Note that $\mathcal{P}_C$ and $\frac{1}{2}\,\mathrm{d}_C(\cdot)^2$ are, respectively, the proximal mapping and Moreau-Yosida envelope of ι_C.

Nonpositive curvature implies many nice properties, which include the uniqueness of geodesics and of the proximal mapping to convex sets as well as differentiability of the squared distance to a point. Recent extensions of classical algorithms from convex analysis to Hadamard spaces include the proximal point algorithm [10, 44], the proximal point algorithm with Bregman distances [86] and Douglas Rachford Splitting [17].

On the other hand Hadamard manifolds naturally arise in many interesting applications. Let us give two examples below.

DT-MRI By magnetic resonance tomography one can obtain a diffusion tensor image (DT-MRI), in which each pixel is the 3×3 symmetric positive definite (spd) matrix representing the diffusivity of the measured tissue at the corresponding location. The manifold $\mathcal{P}(3)$ of 3×3 spd matrices together with the affine invariant metric

$$g_M(u,v) := \operatorname{tr}\left(x^{-1}ux^{-1}v\right), \quad u, v \in T_x\mathcal{P}(3)$$

is a Riemannian manifold of non-constant curvature. The algorithm in [17] is suitable for the restoration of images with values in symmetric Hadamard manifolds and convergence

has been proved for manifolds with constant sectional curvature. Figure 6.1 shows the denoising result[1] for a planar section of the Camino dataset[2] [33] obtained in [17].

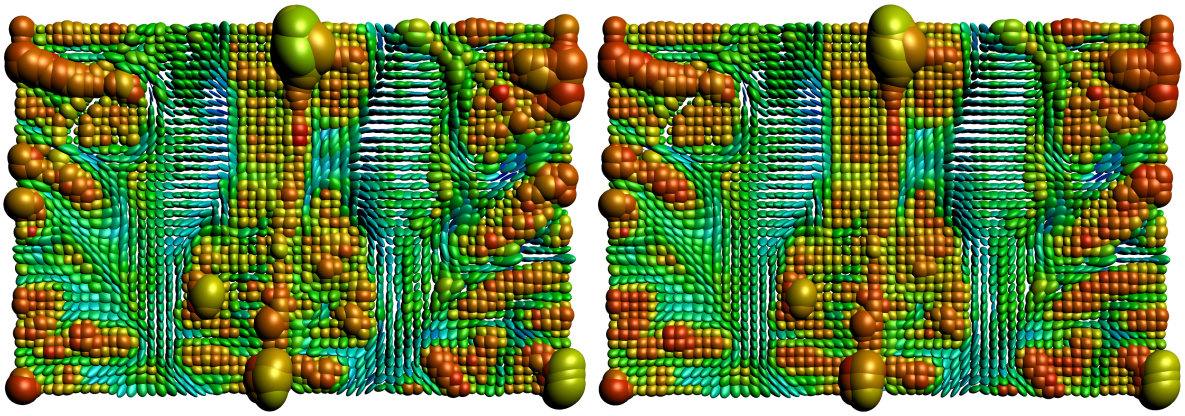

Figure 6.1.: Original data from the Camino set (left) and denoising (right). We visualize an spd matrix $A \in \mathcal{P}(3)$ by the isosurface ellipsoid $\{x \in \mathbb{R}^3 \colon x^\top A x = 1\}$, whose major axis lengths are the eigenvalues of A. The colouring corresponds to the anisotropy index from [79]. We see the successful denoising especially in the more homogeneous regions.

Retina Images We consider the coupled denoising of multiple shots of the same object. The m shots are given as m similar gray-value images $\{g^k_{i,j}\}_{(i,j)\in\mathcal{G}}$, for $k = 1, \dots, m$. Here $\mathcal{G}$ denotes the domain grid of the image, e.g. $\mathcal{G} := \{1, \dots, n_1\} \times \{1, \dots, n_2\}$ for an $n_1 \times n_2$ image. We assume that the gray values $g^k_{i,j}$, $k = 1, \dots, m$ at the same pixel (i, j) are realizations of a univariate Gaussian random variable with distribution $\mathcal{N}(\mu_{i,j}, \sigma_{i,j})$. We estimate the mean and standard deviation by the maximum likelihood estimators

$$\hat{\mu}_{i,j} = \frac{1}{m}\sum_{k=1}^{m} g^k_{i,j}, \quad \hat{\sigma}_{i,j} = \sqrt{\frac{1}{m}\sum_{k=1}^{m}(g^k_{i,j} - \hat{\mu}_{i,j})^2}.$$

One is interested in the underlying image of means $\{\mu_{i,j}\}_{(i,j)\in\mathcal{G}}$. Instead of attempting to denoise $\{\hat{\mu}_{i,j}\}_{(i,j)\in G}$, the approach in [17] uses knowledge of the sample variances in the denoising process. To this end, the image of pairs $\{(\hat{\mu}_{i,j}, \hat{\sigma}_{i,j})\}_{(i,j)\in\mathcal{G}}$ is regarded as image $f \colon \mathcal{G} \to \mathcal{N}$ with values in the Riemannian manifold $\mathcal{N}$ of univariate nondegenerate Gaussian probability distributions parametrized by the mean and standard deviation, equipped with the Fisher metric

$$g_F(u, v) := \frac{u_1 v_1 + 2u_2 v_2}{\sigma^2}, \quad \text{for } u, v \in T_{(\mu,\sigma)}\mathcal{N}. \tag{6.1}$$

[1] Courtesy of J. Persch and R. Bergmann
[2] see http://cmic.cs.ucl.ac.uk/camino

Let us briefly recall the origin of this expression: A distance measure for probability distributions with density function $f(x;\theta)$ parametrized by $\theta = (\theta_1, \dots, \theta_n)$ is given by the Fischer information matrix [35, Chapter 11, p. 397],

$$F(\theta) := \left(\int f(x;\theta) \, \frac{\partial \ln f(x;\theta)}{\partial \theta_i} \frac{\partial \ln f(x;\theta)}{\partial \theta_j} \, dx \right)_{i,j=1}^{n}.$$

In the case of univariate Gaussian densities this reduces to

$$F(\mu, \sigma) = \begin{pmatrix} \frac{1}{\sigma^2} & 0 \\ 0 & \frac{2}{\sigma^2} \end{pmatrix}$$

and $g_F(u, v) = u^\top F v$ yields (6.1). Note further that $(\mathcal{N}, g_F)$ is isometric to the rescaling $(\mathbb{H}^2, 2g_H)$ of the hyperbolic plane, introduced in Section 6.1.3 below, via the isomorphism $\pi \colon \mathcal{N} \to \mathbb{H}^2$ given by

$$\pi(\mu, \sigma) = \left(\tfrac{\mu}{\sqrt{2}}, \sigma\right), \qquad \pi^{-1}(x_1, x_2) = (\sqrt{2}\, x_1, x_2),$$

see also [34].

In this setting the Parallel Douglas Rachford Algorithm [17] has been shown to converge. The algorithm is applied in [17] to denoise retina data, kindly provided by J. Angulo [3]: A sequence of $m = 20$ images $(g_{i,j}^k)_{i,j=1}^{384}$ of size 384×384 has been taken of the same retina with a CCD (coupled charged-device) camera. Due to the short exposure time, the noise level is high. The top row of Figure 6.2 shows the first and last image of the sequence. From these images we obtain the image $f \colon \mathcal{G} \to \mathcal{N}$ with $(f_{i,j})_{i,j=1}^{384} = ((\hat{\mu}_{i,j}, \hat{\sigma}_{i,j}))_{i,j=1}^{384}$ depicted in the middle row of Figure 6.2, which is still noisy and also blurred due to the motion of the eye. The bottom row of Figure 6.2 shows the denoising result obtained in [17].

6.1.2. Basic Properties and Definitions

Recall that in this section $(\mathcal{H}, \mathrm{d})$ denotes a Hadamard manifold. Let $T_x\mathcal{H}$ be the **tangent space** at x on $\mathcal{H}$, which is assumed to be a Hilbert space for the rest of this section. By $\langle \cdot, \cdot \rangle_x$ we denote the corresponding Riemannian metric. We will often skip the index x when it is clear from the context which tangential space is meant. By $\exp_x \colon T_x\mathcal{H} \to \mathcal{H}$ we denote the **exponential map** at x, i.e., $\exp_x \xi = \gamma_{x,\xi}(1) = y$, where $\gamma_{x,\xi}$ is the geodesic starting at x in direction ξ, i.e., $\gamma_{x,\xi}(0) = x$ and $\dot{\gamma}_{x,\xi}(0) = \xi$. Its inverse, the **logarithmic map** $\log_x \colon \mathcal{H} \to T_x\mathcal{H}$ is given by $\log_x y = \xi$, where $\gamma_{x,\xi}(1) = y$. Let $\mathcal{T}_{xy}\, \xi \in T_y\mathcal{H}$ denote the **parallel transported** (via Levi-Civita connection) tangent vector $\xi \in T_x\mathcal{H}$ along the unique geodesic from x to y. Note that parallel transport preserves the Riemannian metric. A Riemannian manifold is **geodesically complete** if the domain of every geodesic $\gamma_{\overset{\frown}{ab}} \colon [0,1] \to \mathcal{H}$ can be extended to all of $\mathbb{R}$.

By ∇f we denote the **Riemannian gradient** of a differentiable function $f \colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$, i.e., $\langle \nabla f(x), \xi \rangle = \mathrm{D} f_x[\xi]$ for all $\xi \in T_x\mathcal{H}$, where $\mathrm{D}_x f$ is the derivative of f at $x \in \mathcal{H}$. In particular we have for $\mathrm{d}(\cdot, x)^2 : \mathcal{H} \to \mathbb{R}_{\geq 0}$ that

$$\nabla \, \mathrm{d}(\cdot, x)^2(y) = -2 \log_y x. \tag{6.2}$$

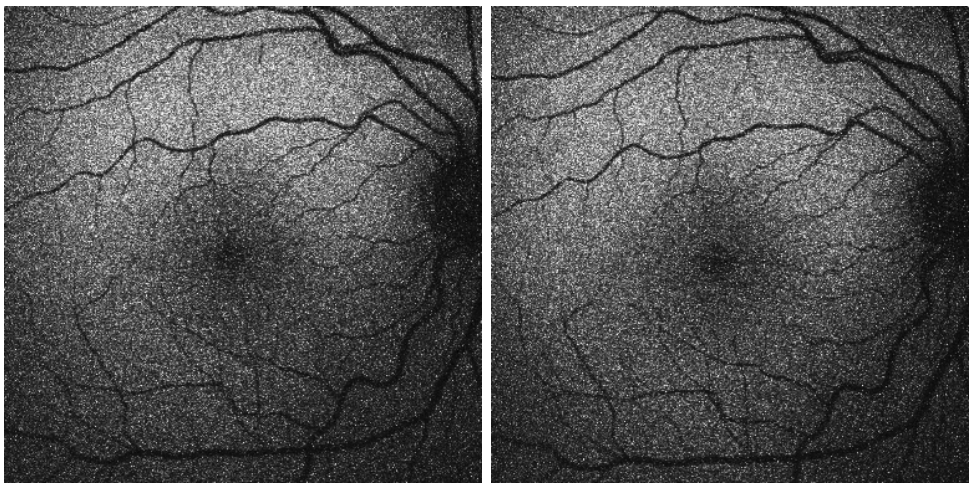

(a) First image $g^{(1)}$ (b) Last image $g^{(20)}$

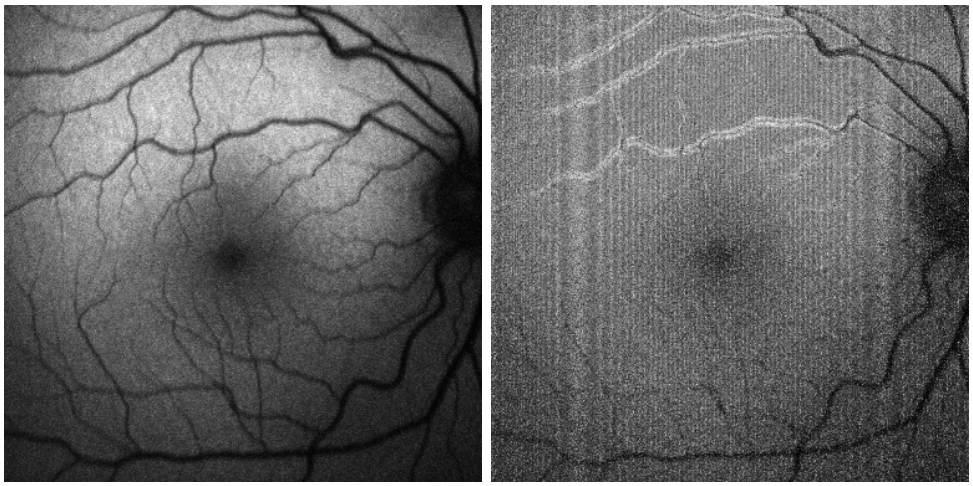

(c) Original mean (d) Original standard deviation

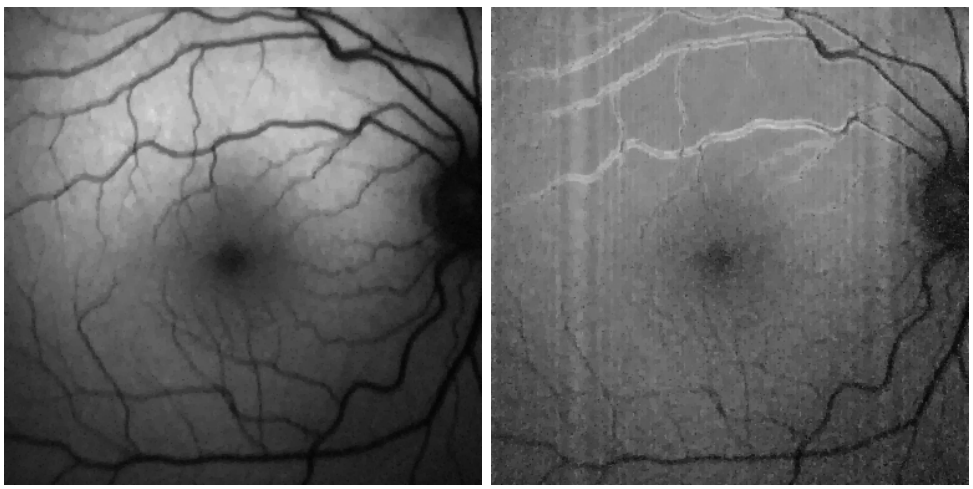

(e) Restored mean (f) Restored standard deviation

Figure 6.2.: Denoising of the retina data from [3] with the Douglas Rachford Algorithm in [17]. The restored image keeps the main features, e.g. veins in the mean and their movement in the variance in an area around them.

The **subdifferential** of $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ at $x \in \operatorname{dom} f$ is defined by

$$\partial f(x) := \big\{\xi \in T_x\mathcal{H} : f(y) \geq f(x) + \langle \xi, \log_x y\rangle \text{ for all } y \in \operatorname{dom} f\big\}. \tag{6.3}$$

We agree that $\partial f(x) = \emptyset$ if $x \notin \operatorname{dom} f$. If the Riemannian gradient $\nabla f(x)$ of f in $x \in \mathcal{H}$ exists, then $\partial f(x) = \big\{\nabla f(x)\big\}$. Further we see immediately from the definition that $x \in \mathcal{H}$ is a global minimizer of a proper f if and only if

$$0 \in \partial f(x). \tag{6.4}$$

The following theorem has been proved in [73] for general finite-dimensional complete connected Riemannian manifolds with an additional assumption on $\operatorname{dom} f \cap \operatorname{dom} g$ which is always fulfilled for finite-dimensional Hadamard spaces.

Theorem 6.1 *Let $\mathcal{H}$ be a finite-dimensional Hadamard manifold and $f, g\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ proper convex functions. Let $x \in \operatorname{dom} f \cap \mathit{int}(\operatorname{dom} g)$. Then we have the subdifferential sum rule*

$$\partial(f + g)(x) = \partial f(x) + \partial g(x).$$

Let $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ be proper and convex. The **directional derivative** of f at a point $x \in \operatorname{dom} f$ in the direction $v \in T_x\mathcal{H}$ is defined by

$$f'(x; v) := \lim_{t \to 0+} \frac{f(\exp_x tv) - f(x)}{t}.$$

We agree that $f'(x; v) = +\infty$, if $x \notin \operatorname{dom} f$.

6.1.3. Hyperbolic Model Spaces

The hyperbolic space $\mathbb{H}^d$ in dimension d is the connected Riemannian manifold of constant sectional curvature -1. Several models are equivalent [71, Proposition 3.5]:

Proposition 6.2 *The following Riemannian manifolds are mutually isometric.*

(i) *The **hyperboloid model** $(\mathbb{H}^d_M, g_M)$ of dimension d is defined as the "upper sheet" $\{x_{d+1} > 0\}$ of the two sheeted hyperboloid in $\mathbb{R}^{d+1}$,*

$$\mathbb{H}^d_M := \Big\{x \in \mathbb{R}^{d+1} \colon \langle x, x\rangle_M = -x_{d+1}^2 + \sum_{i=1}^{d} x_i^2 = -1,\ x_{d+1} > 0\Big\} \subseteq \mathbb{R}^{d+1},$$

*together with the **Minkowski inner product** $g_M := \langle \cdot, \cdot \rangle_M$. given by metric $\langle x, y\rangle_M := -x_{d+1}\, y_{d+1} + \sum_{i=1}^{d} x_i y_i$.*

(ii) *The **Poincaré ball** $(\mathbb{H}^d_B, g_B)$ is the unit ball $\mathbb{H}^d_B := \mathrm{B}^d \subseteq \mathbb{R}^d$, together with the metric*

$$g_B(u, v) := \frac{4\langle u, v\rangle}{\big(1 - \|x\|^2\big)^2}, \quad u, v \in T_x\mathrm{B}^d.$$

(iii) *The **Poincaré half-space** is the upper half-space* $\mathbb{H}^d_P := \{x \in \mathbb{R}^d \colon x_d > 0\}$ *with the metric*

$$g_P(u,v) := \frac{\langle u, v\rangle}{x_d^2}, \quad u, v \in T_x\mathbb{H}^d_P.$$

Most important to us will be the **hyperbolic plane** $\mathbb{H} := \mathbb{H}^2$, in dimension 2, represented by any of the above models together with its intrinsic distance $\mathrm{d}_{\mathbb{H}}$ derived from the Riemannian metric. The equivalence is shown via the following isometries [71, Proposition 3.5]: The isometry $\pi_1 \colon \mathbb{H}^d_M \to \mathbb{H}^d_B$ is defined by

$$\pi_1(x) := \frac{1}{1+x_{d+1}}\,\tilde{x}, \quad \pi_1^{-1}(y) = \frac{1}{1-\|y\|^2}\begin{pmatrix} 2y \\ 1+\|y\|^2 \end{pmatrix},$$

where $x = (x_1, \dots, x_d, x_{d+1})^\top = (\tilde{x}^\top, x_{d+1})^\top$. The isometry $\pi_2 \colon \mathbb{H}^d_B \to \mathbb{H}^d_P$ is defined by

$$\pi_2(x) := \tfrac{1}{\|\tilde{x}\|^2+(x_d-1)^2}\begin{pmatrix} 2\tilde{x} \\ 1-\|x\|^2-x_d^2 \end{pmatrix}, \quad \pi_2^{-1}(x) = \tfrac{1}{\|\tilde{y}\|^2+(y_d+1)^2}\begin{pmatrix} 2\tilde{y} \\ \|\tilde{y}\|^2+y_d^2-1 \end{pmatrix},$$

where $x = (x_1, \dots, x_{d-1}, x_d)^\top = (\tilde{x}^\top, x_d)^\top$. For more models and explicit formulas of distance, exponential and logarithmic map and geodesics in the hyperboloid model see also [17, Appendix B].

Hyperbolic Trigonometry

In a triangle $\triangle ABC$ we denote the interior angles at the vertices A, B, C by α, β, γ and the lengths of opposite sides by a, b, c. Analogously, in $\triangle\overline{ABC}$ we write $\overline{a}, \overline{b}, \overline{c}$ and $\overline{\alpha}, \overline{\beta}, \overline{\gamma}$. Recall the basic cosine and sine rules, see e.g. [60, Section III 5.1, 5.2].

Theorem 6.3 *Let $\triangle ABC$ be a triangle in the hyperbolic plane. Then the following relations are satisfied.*

(i) *Sine Law:* $\dfrac{\sinh a}{\sin\alpha} = \dfrac{\sinh b}{\sin\beta} = \dfrac{\sinh c}{\sin\gamma},$

(ii) *Cosine Law I:* $\cosh a = \cosh b\,\cosh c - \sinh b\,\sinh c\,\cos\alpha,$

(iii) *Cosine Law II:* $\cosh a = \dfrac{\cos\beta\,\cos\gamma + \cos\alpha}{\sin\beta\,\sin\gamma}.$

Lemma 6.4 *Let $\triangle ABC$ be a triangle in the hyperbolic plane with $\gamma = \frac{\pi}{2}$. Then*

$$\cos\alpha = \frac{\tanh b}{\tanh c}. \tag{6.5}$$

Proof. From the Cosine Law I, Theorem 6.3(ii), applied for side c, we obtain $\cosh c = \cosh a\,\cosh b$. Substituting into the Cosine Law I, applied for side a, gives

$$\cos\alpha\,\sinh b\,\sinh c = \cosh b\,\cosh c - \frac{\cosh c}{\cosh b} = \frac{\cosh c}{\cosh b}\left(\cosh^2 b - 1\right).$$

Using $\cosh^2 b - 1 = \sinh^2 b$ and reordering we arrive at (6.5). □

The following lemma can be considered as the limiting case of the Cosine Law II for $\alpha = 0$:

Lemma 6.5 *Let $\triangle ABC$ be a triangle in which the vertices B and C are ordinary points in the hyperbolic plane while A is a point at infinity. Then*

$$\cosh a \, \sin\beta \, \sin\gamma = \cos\beta \, \cos\gamma + 1.$$

In particular, if $\gamma = \frac{\pi}{2}$ then

$$\cosh a \, \sin\beta = 1, \quad \sinh a \, \tan\beta = 1, \quad \textit{and} \quad \tanh a \, \sec\beta = 1,$$

where $\sec\beta = 1/\cos\beta$.

For a proof see e.g. [60, Section III 6.1].

6.1.4. Curvature Bounds in Geodesic Spaces

Curvature bounds in geodesic spaces are defined by comparison. Recall our definition of the hyperbolic plane $(\mathbb{H}^2, \mathrm{d}_{\mathbb{H}})$ in Section 6.1.3. For the sphere $\mathcal{S}^2 := \{x \in \mathbb{R}^3 \colon \|x\| = 1\}$, let $\mathrm{d}_{\mathcal{S}}$ denote the intrinsic distance obtained when equipping the tangent space with the Riemannian metric given by the standard scalar product in $\mathbb{R}^3$, which e.g. implies that antipodal points have distance π. For each $\kappa \in \mathbb{R}$ we define a **model space:**

(i) If $\kappa = 0$ then $\mathbb{M}_0 := \mathbb{R}^2$ with $\mathrm{d}_\kappa(x,y) := \|x - y\|$.

(ii) If $\kappa > 0$ then $\mathbb{M}_\kappa := \mathcal{S}^2$ with distance $\mathrm{d}_\kappa(x,y) := \frac{1}{\sqrt{\kappa}} \, \mathrm{d}_{\mathcal{S}}(x,y)$.

(iii) If $\kappa < 0$ then $\mathbb{M}_\kappa := \mathbb{H}^2$ with $\mathrm{d}_\kappa(x,y) := \frac{1}{\sqrt{-\kappa}} \, \mathrm{d}_{\mathbb{H}}(x,y)$.

As a rescaling of the metric is a conformal (angle-preserving) mapping, by the definition of $\mathbb{M}_\kappa$ the standard trigonometric identities from $\mathcal{S}^2, \mathbb{R}^2$ and $\mathbb{H}^2$ have corresponding versions in $\mathbb{M}_\kappa$. The Cosine Law I from Theorem 6.3 has the following generalization to spaces of constant curvature.

Theorem 6.6 *Let $\triangle xyz$ be a triangle in $\mathbb{M}_\kappa$ with $a = \mathrm{d}(y,z)$, $b = \mathrm{d}(z,x)$, $c = \mathrm{d}(x,y)$ and inner angle γ at the point z. One of the following is true.*

(i) *If $\kappa < 0$ then* $\cosh(\sqrt{-\kappa}c) = \cosh(\sqrt{-\kappa}\,a)\cosh(\sqrt{-\kappa}\,b) - \sinh(\sqrt{-\kappa}\,a)\sinh(\sqrt{-\kappa}\,b)\cos\gamma.$

(ii) *If $\kappa = 0$ then* $c^2 = a^2 + b^2 - 2ab\cos\gamma.$

(iii) *If $\kappa > 0$ then* $\cos(\sqrt{\kappa}\,c) = \cos(\sqrt{\kappa}\,a)\cos(\sqrt{\kappa}\,b) + \sin(\sqrt{\kappa}\,a)\sin(\sqrt{\kappa}\,b)\cos\gamma.$

In $\mathbb{M}_\kappa$, for $\kappa < 0$, we are going to need the following rescaled versions of Lemma 6.4 and Lemma 6.5. In a triangle $\triangle ABC \subset \mathbb{M}_\kappa$, $\kappa < 0$, with $\gamma = \frac{\pi}{2}$ we have

$$\cos\beta = \frac{\tanh(\sqrt{-\kappa}\,a)}{\tanh(\sqrt{-\kappa}\,c)}. \tag{6.6}$$

Further, if A is a vertex at infinity and the opposite side has length a, the third angle $\beta =: \Pi_\kappa(a)$ satisfies

$$\cos\Pi_\kappa(a) = \tanh(\sqrt{-\kappa}\,a). \tag{6.7}$$

Given points x, y, z in a complete geodesic space $\mathcal{X}$, we say that three points $\overline{x}, \overline{y}, \overline{z} \in \mathbb{M}_\kappa$ form a **comparison triangle** for $\triangle xyz$, if $\mathrm{d}_\kappa(\overline{x}, \overline{y}) = \mathrm{d}(x, y)$, $\mathrm{d}_\kappa(\overline{y}, \overline{z}) = \mathrm{d}(y, z)$ and $\mathrm{d}_\kappa(\overline{z}, \overline{x}) = \mathrm{d}(z, x)$, see Figure 6.3.

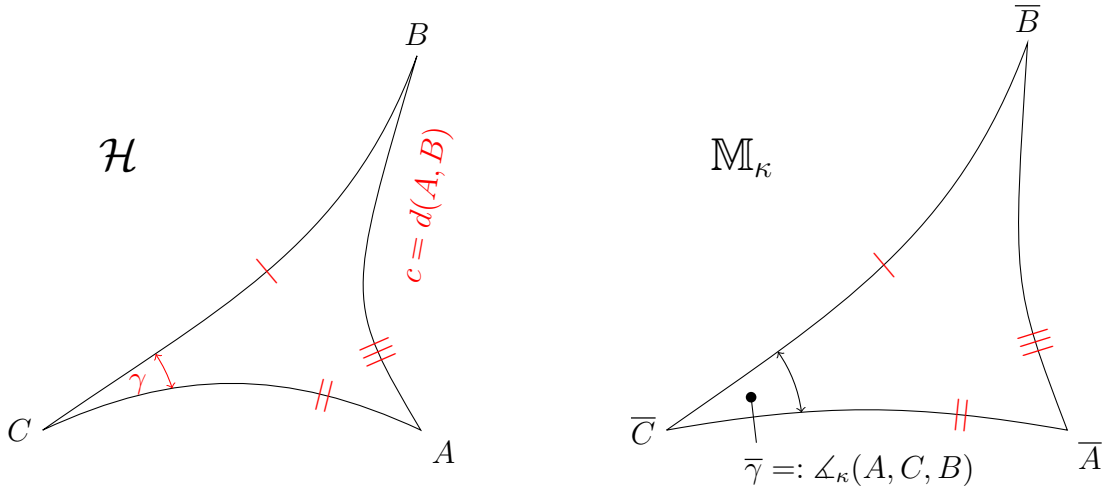

Figure 6.3.: A triangle $\triangle ABC$ in a Hadamard manifold $\mathcal{H}$ (left), and its comparison triangle $\triangle\overline{ABC}$ in $\mathbb{M}_\kappa$ (right); the comparison angle $\measuredangle_\kappa(A, C, B)$, for the angle $\measuredangle ACB$ in $\mathcal{H}$ is defined to be the angle $\measuredangle\overline{ACB}$ in $\mathbb{M}_\kappa$.

Note that $\mathbb{M}_\kappa$ is bounded (i.e. has finite diameter) if $\kappa > 0$, and not bounded if $\kappa \le 0$. Denote the diameter of $\mathbb{M}_\kappa$ by R_κ, i.e.,

$$R_\kappa = \begin{cases} \pi/\sqrt{\kappa}, & \kappa > 0 \\ +\infty, & \kappa \le 0. \end{cases}$$

Comparison triangles always exist for $\kappa \le 0$ and for $\kappa > 0$ if the triangle is not too large:

Lemma 6.7 *Let $a, b, c \in \mathbb{R}_{\ge 0}$ with $a \le b+c$, $b \le c+a$, $c \le a+b$ and $\kappa(a+b+c)^2 \le (2\pi)^2$, i.e., $a + b + c \le 2R_\kappa$, which is a void condition for $\kappa \le 0$. Then there are $\overline{x}, \overline{y}, \overline{z} \in \mathbb{M}_\kappa$ with $\mathrm{d}_\kappa(\overline{x}, \overline{y}) = a$, $\mathrm{d}_\kappa(\overline{y}, \overline{z}) = b$ and $\mathrm{d}_\kappa(\overline{z}, \overline{x}) = c$.*

Note that, for three points $A, B, C \in \mathbb{R}^2$, by the cosine law, the angle between the sements $[C, A]$ and $[C, B]$ equals

$$\measuredangle ACB = \arccos \frac{\mathrm{d}(C,A)^2 + \mathrm{d}(C,B)^2 - \mathrm{d}(A,B)}{2\,\mathrm{d}(C,A)\,\mathrm{d}(C,B)},$$

where $\mathrm{d}(A,B) = \|A - B\|$. Let (X, d) be a geodesic space. We define the **angle** between two geodesics $\gamma\colon [0,1] \to \mathcal{X}$ and $\eta\colon [0,1] \to \mathcal{X}$ with $\gamma_0 = \eta_0 =: C$ by

$$\measuredangle(\gamma,\eta) := \limsup_{s,t\to+0} \arccos \frac{\mathrm{d}(C,\gamma_s)^2 + \mathrm{d}(C,\eta_t)^2 - \mathrm{d}(\gamma_s,\eta_t)^2}{2\,\mathrm{d}(C,\gamma_s)\,\mathrm{d}(C,\eta_t)}, \tag{6.8}$$

i.e. $\measuredangle(\gamma,\eta) = \limsup_{s,t\to+0} \measuredangle\overline{\gamma}_s\overline{C}\overline{\eta}_t$, where $\triangle\overline{C}\overline{\gamma}_s\overline{\eta}_t$ is a comparison triangle for $\triangle C\gamma_s\eta_t$ in $\mathbb{R}^2$.

In Hadamard spaces, the functions $s \mapsto \measuredangle\overline{\gamma}_s\overline{C}\overline{\eta}_t$ and $t \mapsto \measuredangle\overline{\gamma}_s\overline{C}\overline{\eta}_t$ are both monotone, so that the lim sup in (6.8) is a limit and $\measuredangle(\gamma,\eta) = \lim_{t\to+0} \measuredangle\overline{\gamma}_t\overline{C}\overline{\eta}_t$. Angles satisfy a triangle inequality [27, Theorem 3.6.34] and, if measured between two segments in a Hadamard space, angles depend continuously on the endpoints [11, Proposition 1.2.8].

Remark 6.8 *i) The authors of [27] call* (6.8) *an upper angle and reserve the term angle for the case when the* lim sup *is a limit. This holds for shortest paths in lengths spaces with a curvature bound [27, Remark 3.6.31].*

ii) Note that, for $\mathcal{X} = \mathbb{R}^2$ and general paths, [27, Proposition 3.3.27] states that the lim sup *in* (6.8) *is a limit if and only both paths are differentiable at* 0.

Further, for any $A, B, C \in \mathcal{X}$ and $\kappa \in \mathbb{R}$, the **comparison angle** in $\mathbb{M}_\kappa$ for the angle $\measuredangle ACB$ is defined to be the Riemannian angle $\measuredangle_\kappa\overline{ACB}$ at C in a comparison triangle $\triangle\overline{ABC}$ in $\mathbb{M}_\kappa$, and is denoted by $\measuredangle_\kappa(A,C,B)$, see Figure 6.3 (right).

Recall our definition of Hadamard spaces, i.e. geodesic spaces of nonpositive curvature. Other characterizations are possible. The following theorem summarizes [11, Theorem 1.3.3] and the remarks after [11, Definition 1.2.1].

Theorem 6.9 *Let $\mathcal{X}$ be a geodesic space. The following assertions are equivalent:*

(i) *For all $x, y, u, v \in \mathcal{X}$ we have*

$$\mathrm{d}(x,u)^2 + \mathrm{d}(y,v)^2 \le \mathrm{d}(x,y)^2 + \mathrm{d}(y,u)^2 + \mathrm{d}(u,v)^2 + \mathrm{d}(v,x)^2.$$

(ii) *The angle between the sides of every geodesic triangle in $\mathcal{X}$ is no greater than the angle between the corresponding sides of its comparison triangle in $\mathbb{R}^2$.*

(iii) *For every triangle $\triangle xyz$ in $\mathcal{X}$ and any point p on a geodesic from y to z, we have*

$$\mathrm{d}(x,p) \le \|\overline{x} - \overline{p}\|,$$

where $\triangle\overline{xyz}$ is a comparison triangle in $\mathbb{R}^2$ and $\overline{p}$ is the point on the line from $\overline{y}$ to $\overline{z}$ with $\mathrm{d}(\overline{p},\overline{x}) = \|p - x\|$.

(iv) *The distance function is* 1*-convex, i.e. for every geodesic* $\gamma\colon [0,1] \to \mathcal{X}$*, every point* $p \in \mathcal{X}$ *and all* $t \in [0,1]$ *we have*

$$\mathrm{d}(p,\gamma_t)^2 \leq (1-t)\,\mathrm{d}(p,\gamma_0)^2 + t\,\mathrm{d}(p,\gamma_1)^2 - t(1-t)\,\mathrm{d}(\gamma_0,\gamma_1)^2.$$

The equivalence between (iii) and (iv) is an elementary calculation with the Euclidean inner product in $\mathbb{R}^2$. For further equivalences and the proof we refer to [11, Theorem 1.3.2, Theorem 1.3.3]. Let us state that one may replace the inequality in (i) by

$$\mathrm{d}(x,u)^2 + \mathrm{d}(y,v)^2 \leq \mathrm{d}(x,y)^2 + \mathrm{d}(u,v)^2 + 2\,\mathrm{d}(x,v)\,\mathrm{d}(y,u).$$

and that in (iv) is equivalent to the same statement allowing just $t = \frac{1}{2}$.

The equivalence between (ii) and (iii) holds with any model space in place of $\mathbb{R}^2$. The excellent introductory text [27, Theorem 4.3.5] gives a proof for $\kappa = 0$, and outlines its extension to the following general statement in [27, Exercise 4.6.3].

Theorem 6.10 *Let* $\mathcal{X}$ *be a geodesic space and* $\kappa \in \mathbb{R}$*. The following are equivalent:*

(i) *Every point of* $\mathcal{X}$ *has a neighbourhood* U *such that for every triangle* $\triangle ABC$ *contained in* U *with* $a+b+c \leq 2R_\kappa$ *and any point* $D \in [A,C]$ *the inequality*

$$\mathrm{d}(B,D) \leq \mathrm{d}_\kappa(\overline{B},\overline{D})$$

holds, where $\triangle\overline{ABC}$ *is a comparison triangle in* $\mathbb{M}_\kappa$ *and* $\overline{D} \in [\overline{AC}]$ *is the point such that* $\mathrm{d}(\overline{A},\overline{D}) = \mathrm{d}(A,D)$*.*

(ii) *Every point of* $\mathcal{X}$ *has a neighbourhood* U *such that for every triangle* $\triangle ABC$ *contained in* U *with* $a+b+c \leq 2R_\kappa$ *we have*

$$\measuredangle ACB \leq \measuredangle_\kappa(A,C,B).$$

Further, (i) *and* (ii) *are also equivalent if we reverse both inequalities.*

A geodesic space has **curvature bounded above** by κ if any of the conditions in Theorem 6.10 holds, and **curvature bounded below** by κ if the reverse inequalities hold. Geodesic spaces with curvature bounds are called **Alexandrov spaces** and those with curvature bounded above by κ are also referred to as $\mathrm{CAT}(\kappa)$.

In the remainder of this section we show that Hadamard manifolds with bounded sectional curvatures are geodesic spaces with the same curvature bounds.

Relation to Riemannian Curvature Riemannian manifolds are length spaces and curvature bounds as introduced above for geodesic spaces are related to bounds on the Gaussian curvature. We prepare the statement in the setting of Hadamard manifolds, which are geodesic spaces, and for $\kappa \leq 0$, implying that comparison triangles exist.

Let A, B, C be three points in a Hadamard manifold $\mathcal{H}$ and let $\widetilde{A}, \widetilde{B}, \widetilde{C}$ be three points in $\mathbb{M}_\kappa$, $\kappa \leq 0$ such that $\mathrm{d}(\widetilde{B},\widetilde{C}) = \mathrm{d}(B,C)$, $\mathrm{d}(\widetilde{A},\widetilde{C}) = \mathrm{d}(A,C)$ and the angle between

the shortest paths $\gamma_{\widehat{C,A}}$ and $\gamma_{\widehat{C,B}}$ at C is equal to that between the shortest paths $\gamma_{\widehat{\widetilde{C},\widetilde{A}}}$ and $\gamma_{\widehat{\widetilde{C},\widetilde{B}}}$ at $\widetilde{C}$, see Figure 6.4. We call $\triangle\widetilde{A}\widetilde{B}\widetilde{C}$ an **angle comparison triangle** and quote a Cartan-Alexandrov-Topogonov Comparison Theorem from [27, Theorem 6.5.6]:

Theorem 6.11 *Let $\triangle ABC$ be a triangle in a Hadamard manifold $\mathcal{H}$ and let $\triangle\widetilde{A}\widetilde{B}\widetilde{C} \subset \mathbb{M}_\kappa$ be as above. Then*

(i) $\mathrm{d}(A,B) \geq \mathrm{d}(\widetilde{A},\widetilde{B})$ *provided that the Gaussian curvature in $\mathcal{H}$ satisfies $\kappa \geq K(P)$ for all $P \in \mathcal{H}$ and*

(ii) $\mathrm{d}(A,B) \leq \mathrm{d}(\widetilde{A},\widetilde{B})$ *provided that the Gaussian curvature in $\mathcal{H}$ satisfies $\kappa \leq K(P)$ for all $P \in \mathcal{H}$.*

Theorem 6.11 has been proved in [27, Theorem 6.5.6] in the more general setting with $\mathcal{H}$ replaced by a neighbourhood U in a Riemannian manifold such that there is only one shortest path between any two points in U. Note that by [27, Lemma 5.2.11] every point in a typical (see there for details) Riemannian manifold possesses such a neighbourhood.

As immediate consequence we obtain an angle comparison property for Hadamard manifolds (more generally Riemannian manifolds) with curvature bounds.

Corollary 6.12 *Let $\triangle ABC$ be a triangle in a Hadamard manifold $\mathcal{H}$. Then the angle γ at the point C satisfies*

(i) $\gamma \leq \measuredangle_\kappa(A,C,B)$ *provided $K(P) \leq \kappa$ for all $P \in \mathcal{H}$ and*

(ii) $\gamma \geq \measuredangle_\kappa(A,C,B)$ *provided $K(P) \geq \kappa$ for all $P \in \mathcal{H}$.*

Proof. Let $\triangle\widetilde{A}\widetilde{B}\widetilde{C} \subset \mathbb{M}_\kappa$ be an angle comparison triangle as above and let $\triangle\overline{ABC} \subset \mathbb{M}_\kappa$ be a comparison triangle for $\triangle ABC$. Then $\widetilde{a} = \overline{a} = a$, $\widetilde{b} = \overline{b} = b$ and the triangles $\triangle\overline{ABC}$ and $\triangle\widetilde{A}\widetilde{B}\widetilde{C}$ are nonisomorphic only if $\overline{\gamma} \neq \widetilde{\gamma}$. By the Cosine Law I the length $c(\gamma)$ of the third side in a triangle in $\mathbb{M}_\kappa$ having two sides of length a and b with angle γ between them is given by

$$c(\gamma) = \tfrac{1}{\sqrt{-\kappa}} \cosh^{-1}\left(\cosh(\sqrt{-\kappa}\,a)\,\cosh(\sqrt{-\kappa}\,b) - \sinh(\sqrt{-\kappa}\,a)\,\sinh(\sqrt{-\kappa}\,b)\,\cos\gamma\right).$$

Note that $c(\gamma)$ is monotone increasing for $\gamma \in [0,\pi]$. Therefore, having $\widetilde{c} \leq c = \overline{c}$ from case (i) of Theorem 6.11 (respectively $\geq$ in case (ii)) immediately implies $\widetilde{\gamma} \leq \overline{\gamma}$ (respectively $\geq$). By definition $\measuredangle_\kappa(A,C,B) = \overline{\gamma}$, by construction $\widetilde{\gamma} = \gamma$ and the proof is complete. □

Hence Hadamard manifolds with bounded sectional curvatures are geodesic spaces with the same curvature bounds. More generally, Riemannian manifolds are length spaces with the same curvature bounds and this follows by the same line of argument from [27, Theorem 6.5.6].

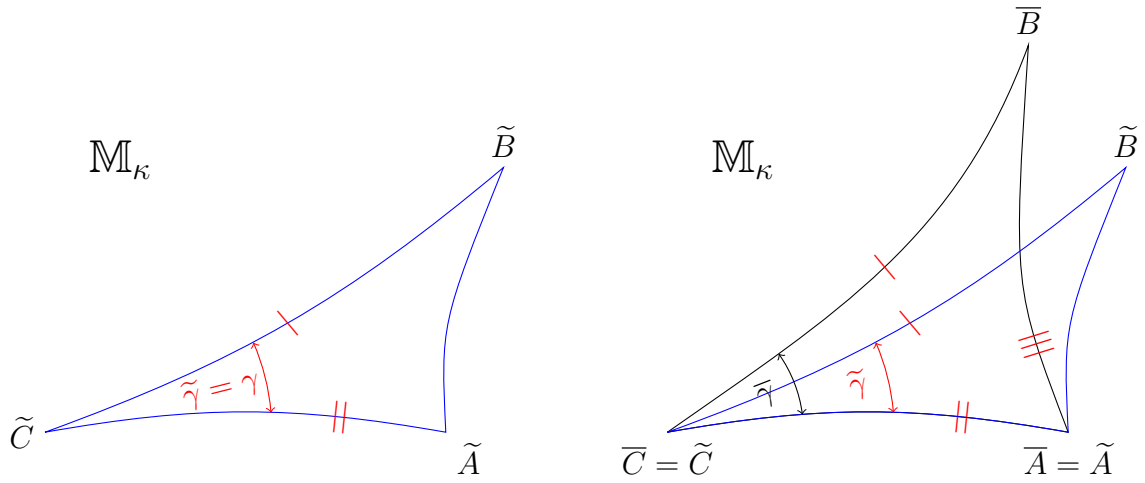

Figure 6.4.: Angle comparison triangle $\triangle\widetilde{A}\widetilde{B}\widetilde{C}$ in $\mathbb{M}_\kappa$ for the triangle $\triangle ABC$ from Figure 6.3, (left). If the Gaussian curvature at every point of $\mathcal{H}$ is at most κ, the Cartan-Alexandrov-Topogonov Comparison Theorem states that the opposite side $\widetilde{A}\widetilde{B}$ of an angle comparison triangle in $\mathbb{M}_\kappa$ is shorter than the opposite side $\overline{AB}$ of a comparison triangle. The comparison angle $\overline{\gamma} = \measuredangle_\kappa(A, C, B)$ in $\mathbb{M}_\kappa$ must therefore be larger than $\widetilde{\gamma} = \gamma$, see Corollary 6.12, (right).

6.2. Useful Properties and Moreau-Yosida Envelopes

In this section we consider differential properties of Moreau-Yosida envelopes and along the way obtain several results which are of interest in their own. The main result of this section is the differentiability of Moreau-Yosida envelopes and an analytic expression for the gradient in terms of the proximal mapping, see Theorem 6.22. We are going to show the usefulness of this result by proving with its help the extension of a result by H. Attouch in Chapter 6.3.

6.2.1. Subdifferential Sum Rule

In Section 6.2.5 we are going to obtain an analytical expression for ∇f_λ. To this end we have to apply a subdifferential sum rule which is to the best of our knowledge only known for finite-dimensional manifolds, i.e. for manifolds with Euclidean tangent spaces, see Theorem 6.1. In this section we are going to prove the generalization to infinite-dimensional Hadamard manifolds.

Finite convex functions in Euclidean space are locally Lipschitz. Let us prove the infinite-dimensional generalization of this simple fact.

Lemma 6.13 *Let $\mathcal{H}$ be a geodesically complete Riemannian manifold. Let $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ be convex. Suppose that there are $x_0 \in \mathcal{H}$, $R > 0$ and $m, M \in \mathbb{R}$ with*

$$m \leq f(x) \leq M$$

for all $x \in \mathrm{B}_{2R}(x_0)$. Then f is Lipschitz continuous on $\mathrm{B}_R(x_0)$. More precisely,

$$|f(b) - f(a)| \leq \frac{M-m}{R}\,\mathrm{d}(a,b) \tag{6.9}$$

for all $a, b \in \mathrm{B}_R(x_0)$.

Proof. For given points $a, b \in \mathrm{B}_R(x_0)$ with $b \neq a$ consider the geodesic $\gamma_{\widehat{ab}}\colon [0,1] \to \mathcal{H}$. By geodesic completeness the domain of $\gamma_{\widehat{ab}}$ can be extended to all of $\mathbb{R}$. Let $c := \gamma_{\widehat{ab}}\big(1 + \frac{R}{\mathrm{d}(a,b)}\big)$. By the constant speed property we have $\mathrm{d}(a,c) = \mathrm{d}(a,b) + R$ and thus $c \in \mathrm{B}_{2R}(x_0)$. Further $\gamma_{\widehat{ac}}$ is just a reparameterization of $\gamma_{\widehat{ab}}$ and

$$b = \gamma_{\widehat{ab}}(1) = \gamma_{\widehat{ac}}\left(\frac{\mathrm{d}(a,b)}{\mathrm{d}(a,c)}\right).$$

From convexity of f it follows that

$$f \circ \gamma_{\widehat{ac}}\left(\frac{\mathrm{d}(a,b)}{\mathrm{d}(a,c)}\right) \leq \left(1 - \frac{\mathrm{d}(a,b)}{\mathrm{d}(a,c)}\right) f \circ \gamma_{\widehat{ac}}(0) + \frac{\mathrm{d}(a,b)}{\mathrm{d}(a,c)} f \circ \gamma_{\widehat{ac}}(1).$$

Reordering and using the given bounds, this implies

$$f(b) - f(a) \leq \frac{\mathrm{d}(a,b)}{\mathrm{d}(a,c)}\big(f(c) - f(a)\big) \leq \frac{\mathrm{d}(a,b)}{R}(M-m).$$

Exchanging the roles of a and b, we obtain (6.9). □

Corollary 6.14 *Let $\mathcal{H}$ be a Hadamard manifold modelled on a Hilbert space and $x \in \mathcal{H}$. Let $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ be a proper convex lsc function which is bounded in a neighbourhood of x. Then $0 \in \operatorname{int}\big(\operatorname{dom} f'(x;\cdot)\big)$.*

Proof. By Lemma 6.13, there exists $\varepsilon > 0$ such that the function f is Lipschitz on $\mathrm{B}_\varepsilon(x)$, with some constant $L > 0$. Note that convexity of f implies monotonicity of the map $t \mapsto \frac{1}{t}\big(f(\exp_x(tv)) - f(x)\big)$ on $[0,1]$ so that the limit $t \to +0$ exists in $[-\infty, +\infty]$. For $v \in \mathrm{B}_\varepsilon(0) \subseteq T_x\mathcal{H}$ and $t \in (0,1)$ it follows that

$$|f'(x;v)| = \left|\lim_{t\to+0} \frac{f(\exp_x(tv)) - f(x)}{t}\right| \leq \tfrac{1}{t}\,L\|tv\| = L\|v\|,$$

which shows that $|f'(x;\cdot)|$ is bounded by εL on $\mathrm{B}_\varepsilon(0)$. □

The following representation has been shown in [73, Prop. 3.8 (ii)] for finite-dimensional Riemannian manifolds, with the proof remaining unchanged for manifolds modelled on Hilbert space.

Lemma 6.15 *Let $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ be a proper convex function on a Riemannian manifold, and let $x \in \operatorname{dom} f$. Then we have the subdifferential representation*

$$\partial f(x) = \big\{w \in T_x\mathcal{H}\colon \langle w, v\rangle \leq f'(x;v) \quad \textit{for all } v \in T_x\mathcal{H}\big\}. \tag{6.10}$$

With this we can extend Theorem 6.1 to the infinite-dimensional case.

Theorem 6.16 *Let $\mathcal{H}$ be a Hadamard manifold modelled on a Hilbert space. Let $f_1, f_2 \colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ be proper convex lsc functions, let $x \in \operatorname{dom} f_1 \cap \operatorname{dom} f_2$ and let f_2 be bounded in a neighbourhood of x. Then*

$$\partial\big(f_1 + f_2\big)(x) = \partial f_1(x) + \partial f_2(x). \tag{6.11}$$

Proof of Theorem 6.16. We extend the proof in [73, Prop. 4.3] to the infinite-dimensional case. First note that the subdifferential of a proper convex function f at $x \in \mathcal{H}$ can be written

$$\partial f(x) = \partial f'(x;\cdot)(0). \tag{6.12}$$

To see this, note that both sides are empty if $x \notin \operatorname{dom} f$. Let now $x \in \operatorname{dom} f$. By the definition of the subdifferential, $w \in \partial f'(x;\cdot)(0)$ if for each $u \in T_0(T_x\mathcal{H})$ and each geodesic γ from 0 to u,

$$f'(x;u) \geq f'(x;0) + \langle w, \gamma'(0)\rangle.$$

As the unique geodesic γ in $T_0(T_x\mathcal{H})$ from 0 to u is given by $\gamma(t) = tu$, with $\gamma'(0) = u$, this is equivalent to $f'(x;u) \geq \langle w, u\rangle$ for all $u \in T_x\mathcal{H}$, i.e., the subdifferential representation (6.10), and hence to $w \in \partial f(x)$.

Returning to the proof, it follows from (6.12), for $i = 1, 2$, that

$$\partial f_i(x) = \partial f_i'(x;\cdot)(0) \quad \text{and} \quad \partial(f_1 + f_2)(x) = \partial\big[\, f_1'(x;\cdot) + f_2'(x;\cdot)\big](0).$$

By Corollary 6.14, we have $0 \in \operatorname{int}(\operatorname{dom} f_2'(x;\cdot))$ and hence, by the classical Moreau-Rockafellar Theorem in linear spaces, Theorem 5.1, it follows that

$$\partial\big[\, f_1'(x;\cdot) + f_2'(x;\cdot)\big](0) = \partial f_1'(x;\cdot)(0) + \partial f_2'(x;\cdot)(0) = \partial f_1(x) + \partial f_2(x).$$

Together with the previous equation, this proves the subdifferential sum rule (6.11). □

6.2.2. Equivalent Curvature Definitions

We give an equivalent curvature definition (K. T. Sturm, Lecture at the University of Bonn, 2012). This section is included as reference and neither the statement nor the equivalence will be used anywhere else in this thesis.

For a locally bounded Borel function $f \colon \mathcal{X} \to \mathbb{R}$, we say that $u \colon \mathcal{X} \to \mathbb{R} \cup \{+\infty\}$ satisfies

$$\mathrm{D}^2 u \geq f \quad \text{on } \mathcal{X} \text{ along geodesics}$$

if and only if for every geodesic $\gamma \colon [0,1] \to \mathcal{X}$ and all $t \in (0,1)$ we have

$$u(\gamma_t) \leq (1-t)u(\gamma_0) + tu(\gamma_1) - |\dot\gamma|^2 \int_0^1 g_{s,t}\, f(\gamma_s)\, ds,$$

where g denotes the Green function on $(0,1)$,

$$g_{s,t} = \begin{cases} (1-s)t, & \text{if } 0 \leq t \leq s \leq 1 \\ s(1-t), & \text{if } 1 \geq t > s \geq 0. \end{cases}$$

Note that $\int_0^1 g_{s,t}\,ds = \frac{1}{2}t(1-t)$. This is a reasonable definition because it can be shown that u satisfies $\mathrm{D}^2 u \geq f$ along geodesics if and only if, for all geodesics γ with $u(\gamma_0) < +\infty$, $u(\gamma_1) < +\infty$ the function $t \mapsto -u(\gamma_t)$ is lsc on $[0,1]$, continuous on $(0,1)$ and satisfies $\partial_t^2 u(\gamma_t) \geq f(\gamma_t)|\dot{\gamma}|^2$ in distributional sense on $(0,1)$.

With the notation just introduced we have the language to speak about curvature bounds.

Theorem 6.17 *For any $\kappa \in \mathbb{R}$ and any geodesic space $(\mathcal{X}, \mathrm{d})$ the following are equivalent:*

(i) *For all $z \in \mathcal{X}$ we have along geodesics*

$$\begin{cases} \mathrm{D}^2 \frac{cosh(\sqrt{-\kappa}\,\mathrm{d}(z,\cdot))}{\kappa} \geq -\cosh(\sqrt{-\kappa}\,\mathrm{d}(z,\cdot)), & \text{if } \kappa < 0, \\ \mathrm{D}^2\, \mathrm{d}^2(z,\cdot)/2 \leq 1, & \text{if } \kappa = 0, \\ \mathrm{D}^2 \frac{\cos(\sqrt{\kappa}\,\mathrm{d}(z,\cdot))}{\kappa} \geq -\cos(\sqrt{\kappa}\,\mathrm{d}(z,\cdot)), & \text{if } \kappa > 0, \end{cases}$$

in $\mathcal{X}_\kappa(z) := \left\{ x \in \mathcal{X} \colon \kappa\, \mathrm{d}^2(z,x) < \left(\frac{\pi}{2}\right)^2 \right\}$.

(ii) *For every geodesic triangle (α, β, γ) with perimeter $L(\alpha) + L(\beta) + L(\gamma) \leq \frac{2\pi}{\sqrt{\kappa}}$ and corresponding comparison triangle $(\overline{\alpha}, \overline{\beta}, \overline{\gamma})$ in $\mathbb{M}_\kappa$ and all $s, t \in [0,1]$ we have*

$$\mathrm{d}(\alpha_s, \beta_t) \geq \mathrm{d}_\kappa(\overline{\alpha}_s, \overline{\beta}_t).$$

In (ii), a triple of geodesics $(\overline{\alpha}, \overline{\beta}, \overline{\gamma})$ is a geodesic triangle if $\alpha_1 = \beta_0$, $\beta_1 = \gamma_0$ and $\gamma_1 = \alpha_0$. A geodesic triangle in $\mathbb{M}_\kappa$ is a comparison triangle if the geodesics have the same lengths.

6.2.3. Weak Convergence in the Tangent Space

In finite-dimensional Hadamard manifolds we have that

$$x_n \to x \quad \Rightarrow \quad \log_x x_n \to 0.$$

In this section we prove that the same is true with respect to weak convergence in infinite-dimensional Hadamard manifolds whose curvature is locally bounded below. We are going to say that $\mathcal{H}$ fulfills property (A), if for every sequence $\{x_n\}_n$ the following implication (A) holds true:

$$\text{(A)} \qquad x_n \overset{w}{\to} x \quad \Rightarrow \quad \log_x x_n \overset{w}{\to} 0.$$

Theorem 6.18 *Let $\mathcal{H}$ be a Hadamard manifold with sectional curvatures bounded below by some constant $\kappa < 0$. Then $\mathcal{H}$ has property (A).*

Proof. Given $A \in T_x\mathcal{H} \setminus \{0\}$ we need to show $\langle \log_x x_n, A\rangle \to 0$. Define the geodesic $\gamma\colon \mathbb{R} \to \mathcal{H}$ by

$$\gamma(t) = \exp_x(tA).$$

As weakly convergent sequences are by definition bounded, there exists $R > 0$ with $\mathrm{d}(x, x_n) \le R$ for all n. Further, we know by Theorem 5.11(iii) that

$$z_n := \mathcal{P}_\gamma\, x_n \to x \tag{6.13}$$

for $n \to \infty$. Let $n_1 < n_2 < \cdots < n_k < \ldots$ be an enumeration of the indices n with $z_n = x$. Then $\log_x x_{n_k} \perp A$ for all k. We may therefore restrict attention to the remaining indices, assuming from now that $z_n \neq x$ for all n. Then $A/\|A\| = \pm \log_x z_n / \|\log_x z_n\|$ and the claim is equivalent to

$$\begin{aligned} 0 &= \lim_{n\to\infty} \langle \log_x x_n, \tfrac{A}{\|A\|}\rangle \\ &= \lim_{n\to\infty} \|\log_x x_n\| \cos \measuredangle x_n x z_n, = \lim_{n\to\infty} c_n \cos\beta_n \end{aligned} \tag{6.14}$$

where $\beta_n := \measuredangle x_n x z_n$ denotes the inner angle at the point x and we abbreviate from now on $c_n := \mathrm{d}(x, x_n)$, $a_n := \mathrm{d}(x, z_n)$.

Let $\triangle \overline{x}_n \overline{x}\, \overline{z}_n$ be a comparison triangle for $\triangle x_n x z_n$ in $\mathbb{M}_\kappa$, i.e. $\mathrm{d}(\overline{x}, \overline{x}_n) = \mathrm{d}(x, x_n) = c_n$, $\mathrm{d}(\overline{x}_n, \overline{z}_n) = \mathrm{d}(x_n, z_n)$ and $\mathrm{d}(\overline{x}, \overline{z}_n) = \mathrm{d}(x, z_n) = a_n$. By Corollary 6.12 the comparison angle $\overline{\beta}_n := \measuredangle \overline{x}_n \overline{x}\, \overline{z}_n$ in the model space of more negative curvature is smaller, that is

$$0 \le \overline{\beta}_n \le \beta_n \le \tfrac{\pi}{2} \quad \text{and} \quad 0 \le \measuredangle \overline{x}_n \overline{z}_n \overline{x} \le \tfrac{\pi}{2}. \tag{6.15}$$

For each n, we argue differently depending on how $\overline{\beta}_n$ compares to the angle of parallelism $\Pi(a_n)$ in the model space $\mathbb{M}_\kappa$.

Case 1: $\overline{\beta}_n \ge \Pi(a_n)$.
Then $\frac{\pi}{2} \ge \beta_n \ge \overline{\beta}_n \ge \Pi(a_n) > 0$. By (6.7) we know $\cos\Pi(a_n) = \tanh(\sqrt{-\kappa}\, a_n)$ and obtain with $a_n \to 0$ from (6.13) that

$$c_n \cos\beta_n \le c_n \cos\Pi(a_n) = c_n \tanh(\sqrt{-\kappa}\, a_n) \le R \tanh(\sqrt{-\kappa} a_n) \to 0$$

as $n \to \infty$.

Case 2: $\overline{\beta}_n < \Pi(a_n)$.
Then there is a point $\overline{x}'_n$ on the geodesic ray from $\overline{x}$ through $\overline{x}_n$ with $\measuredangle \overline{x}'_n \overline{z}_n \overline{x} = \frac{\pi}{2}$, which satisfies $c_n = \mathrm{d}(x, x_n) < \mathrm{d}(\overline{x}, \overline{x}'_n)$.

Preparing to estimate more carefully than in the previous case, consider the function $t \mapsto \frac{t}{\tanh t}$. Using that $t \le \sinh t$ on $[0, +\infty)$, its derivative is

$$t \mapsto \frac{1}{\tanh t} - \frac{t}{\sinh^2 t} \ge \frac{1}{\tanh t} - \frac{1}{\sinh t} \ge 0.$$

Therefore the function $t \mapsto \frac{t}{\tanh t}$ is increasing on $[0, R]$ and bounded above by $\frac{R}{\tanh R}$ on $[0, R]$. The cos-relation (6.6) in $\triangle \overline{x}'_n \overline{x} \overline{z}_n$ yields

$$\cos\overline{\beta}_n = \frac{\tanh(\sqrt{-\kappa}\, a_n)}{\tanh(\sqrt{-\kappa}\ \mathrm{d}(\overline{x}, \overline{x}'_n))} \le \frac{\tanh(\sqrt{-\kappa}\, a_n)}{\tanh(\sqrt{-\kappa}\, c_n)}.$$

Using (6.15) we conclude

$$\begin{aligned} c_n \cos\beta_n \le c_n \cos\overline{\beta}_n &\le \frac{1}{\sqrt{-\kappa}} \frac{\sqrt{-\kappa}\, c_n}{\tanh(\sqrt{-\kappa}\, c_n)} \tanh(\sqrt{-\kappa}\, a_n) \\ &\le \frac{1}{\sqrt{-\kappa}} \frac{R}{\tanh R} \tanh(\sqrt{-\kappa}\, a_n) \end{aligned}$$

and (6.14) follows. □

6.2.4. Parameter Dependence of the Moreau-Yosida Envelope

The following lemma has been proved for reflexive Banach spaces with strictly convex primal and dual norms fulfilling the Kadets-Klee property in [8, Lemma 3.27]. Here is the generalization to Hadamard spaces.

Lemma 6.19 *Let $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ be a proper convex lsc function and $x \in \mathcal{H}$. Then the function $\lambda \mapsto \lambda f_\lambda(x)$ is a concave, continuously differentiable function on $(0, +\infty)$ with derivative*

$$\frac{\mathrm{d}}{\mathrm{d}\lambda} \lambda f_\lambda(x) = f(J_\lambda x), \quad J_\lambda x := \operatorname{Prox}_{\lambda f}(x). \tag{6.16}$$

Proof. 1. Let us start by proving that the function

$$\lambda \mapsto f(J_\lambda x) \tag{6.17}$$

is continuous on $(0, +\infty)$. For a different proof via the resolvent identity [11, Prop. 2.2.24], see [11, Prop. 2.2.26].

Suppose that $\lambda > 0$ and that $\{\lambda_n\}_n$ is a sequence of real numbers with $\lambda_n \to \lambda$ for $n \to \infty$. Define a sequence of functions $\{G_n\}_n$ and the function G by

$$G_n(u) := f(u) + \tfrac{1}{2\lambda_n} \mathrm{d}(x,u)^2, \quad G(u) := f(u) + \tfrac{1}{2\lambda} \mathrm{d}(x,u)^2.$$

We make the claim that the sequence $\{G_n\}_n$ Γ-converges to G in the topology defined by the weak convergence of sequences. Indeed, to show the lim inf-inequality, note that every weakly convergent sequence, $u_n \overset{w}{\to} u$, is bounded, i.e. for some $L > 0$ we have $\mathrm{d}(x, u_n) \le L$, and hence

$$\liminf_{n\to\infty} \left(f(u_n) + \tfrac{1}{2\lambda_n} \mathrm{d}(x,u_n)^2 \right) \ge \liminf_{n\to\infty} G(u_n) - \lim_{n\to\infty} \left| \tfrac{1}{2\lambda} - \tfrac{1}{2\lambda_n} \right| L^2 \ge G(u),$$

the last inequality coming from weak lower semicontinuity of the convex lsc function G. To prove the lim sup-inequality we take the constant sequence.

Having shown the claim, note further that the sequence $\{G_n\}_n$ is equi-mildly coercive. To see this, we argue as follows: Let $M > 0$ be an upper bound for the sequence $\{\lambda_n\}_n$. By [11, Lemma. 2.2.13] there are constants $\alpha, \beta \in \mathbb{R}$ such that $f(y) \ge -\alpha\, \mathrm{d}(y,x) - \beta$ for all $y \in \mathcal{H}$. Thus

$$\begin{aligned} G_n(u) &\ge \tfrac{1}{2\lambda_n} \mathrm{d}(x,u)^2 - \alpha\, \mathrm{d}(x,u) - \beta \\ &\ge \tfrac{1}{2M} \mathrm{d}(x,u)^2 - \alpha\, \mathrm{d}(x,u) - \beta. \end{aligned}$$

Let us chose $t_0 > 0$ such that, for all $t \geq t_0$,

$$\tfrac{1}{2M}t^2 - \alpha t - \beta \geq f(x) + 1.$$

Then $\mathrm{d}(x,u) \geq t$ implies $G_n(u) > f(x) = G_n(x)$ and hence, for all n,

$$\operatorname{argmin} G_n \subseteq \overline{B_t(x)}. \tag{6.18}$$

For every sequence contained in $\overline{B_t(x)}$, Theorem 5.11(i) guarantees the existence of a weakly convergent subsequence. Further $\overline{B_t(x)}$ is closed and convex, so that by Theorem 5.11(iv) every limit of such a subsequence belongs to $\overline{B_t(x)}$. In other words, $\overline{B_t(x)}$ is compact with respect to the weak topology and we have shown by (6.18) that the sequence $\{G_n\}_n$ is equi-mildly coercive in the weak topology.

As the sets of minimizers are given by $\operatorname{argmin} G = \{J_\lambda x\}$ and $\operatorname{argmin} G_n = \{J_{\lambda_n} x\}$, respectively, it follows from Theorem 5.4 and Remark 5.5 that

$$J_{\lambda_n} x \overset{w}{\to} J_\lambda x. \tag{6.19}$$

The theorem further tells us that

$$\inf_{u\in\mathcal{H}} G_n(u) \to \inf_{u\in\mathcal{H}} G(u),$$

that is

$$f(J_{\lambda_n}x) + \tfrac{1}{2\lambda_n}\,\mathrm{d}(x, J_{\lambda_n}x)^2 \longrightarrow f(J_\lambda x) + \tfrac{1}{2\lambda}\,\mathrm{d}(x, J_\lambda x)^2, \qquad \text{as } n \to \infty. \tag{6.20}$$

From (6.19) and the lower semicontinuity w.r.t. weak convergence in $\mathcal{H}$ of the convex functions f and $\mathrm{d}(\cdot, x)^2$, we get

$$f(J_\lambda x) \leq \liminf_n f(J_{\lambda_n} x), \tag{6.21}$$

$$\tfrac{1}{2\lambda}\,\mathrm{d}(x, J_\lambda x)^2 \leq \liminf_n \tfrac{1}{2\lambda_n}\,\mathrm{d}(x, J_{\lambda_n} x)^2. \tag{6.22}$$

Combining (6.21) and (6.22),

$$\begin{aligned} f(J_\lambda x) + \tfrac{1}{2\lambda}\,\mathrm{d}(x, J_\lambda x) &\leq \liminf_n f(J_{\lambda_n} x) + \tfrac{1}{2\lambda}\,\mathrm{d}(x, J_\lambda x)^2 \\ &\leq \limsup_n f(J_{\lambda_n} x) + \tfrac{1}{2\lambda}\,\mathrm{d}(x, J_\lambda x)^2 \\ &\leq \limsup_n f(J_{\lambda_n} x) + \liminf_n \tfrac{1}{2\lambda_n}\,\mathrm{d}(x, J_{\lambda_n} x)^2 \\ &\leq \limsup_n \Big(f(J_{\lambda_n} x) + \tfrac{1}{2\lambda_n}\,\mathrm{d}(x, J_{\lambda_n} x)^2 \Big). \end{aligned}$$

By (6.20) both sides are equal, hence all inequalities are equalities and we have

$$f(J_{\lambda_n} x) \to f(J_\lambda x),$$

which proves part 1.

Then also

$$\mathrm{d}(x, J_{\lambda_n} x) \to \mathrm{d}(x, J_\lambda x).$$

As a side-remark, let us note that, combined with (6.19), by Theorem 5.11(iii), we obtain for all $x \in \mathcal{H}$ that the mapping $\lambda \mapsto J_\lambda x$ is continuous from $(0, +\infty)$ into $\mathcal{H}$.

2. Let us now conclude the proof. From the equality

$$\lambda\, f_\lambda(x) = \inf_{u \in \mathcal{H}} \big\{ \lambda f(u) + \tfrac{1}{2}\, \mathrm{d}(x, u)^2 \big\}$$

it follows that $\lambda \mapsto \lambda\, f_\lambda(x)$ is a concave function. Given $\lambda_1, \lambda_2 > 0$, by definition of f_λ,

$$\begin{aligned} \lambda_1\, f_{\lambda_1}(x) - \lambda_2 f_{\lambda_2}(x) &\leq \Big(\lambda_1 f(J_{\lambda_2} x) + \tfrac{1}{2}\, \mathrm{d}(x, J_{\lambda_2} x])^2 \Big) - \Big(\lambda_2 f(J_{\lambda_2} x) + \tfrac{1}{2}\, \mathrm{d}(x, J_{\lambda_2} x])^2 \Big) \\ &= (\lambda_1 - \lambda_2) f(J_{\lambda_2} x). \end{aligned}$$

Exchanging λ_1 and λ_2 we obtain

$$(\lambda_1 - \lambda_2) f(J_{\lambda_1} x) \leq \lambda_1\, f_{(\lambda_1)}(x) - \lambda_2\, f_{(\lambda_2)}(x) \leq (\lambda_1 - \lambda_2) f(J_{\lambda_2} x). \tag{6.23}$$

Dividing by $\lambda_1 - \lambda_2$ (first with $\lambda_1 > \lambda_2$, then with $\lambda_1 < \lambda_2$), we let λ_2 tend to λ_1. Using the continuity proven in part 1, we thus obtain the derivative formula (6.16). □

We will use the above lemme in combination with the following result which can be found e.g. in [8, Lemma 3.28].

Lemma 6.20 *Let $\{h_n\}_n$ be a sequence of concave, differentiable functions $h_n : \mathbb{R} \to \mathbb{R}$ which converge pointwise to a concave, continuously differentiable function h. Then, for all $\lambda > 0$, it holds that*

$$h_n'(\lambda) \to h'(\lambda), \quad \text{as } n \to \infty.$$

Moreover, the convergence is uniform on every compact subset of $(0, +\infty)$.

6.2.5. Gradients of Moreau-Yosida Envelope

In this section we prove an expression for the gradient of Moreau-Yosida envelopes in terms of the proximum mapping, for functions defined on a finite- or infinite-dimensional Hadamard manifold $\mathcal{H}$.

Conbining the subdifferential sum rule from Section 6.2.1, Theorem 6.16, with (6.2) and (6.4), we immediately obtain:

Lemma 6.21 (Derivative at Proximum) *Let $f \colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ be a proper convex lsc function and $p_x := \mathrm{Prox}_{\lambda f}(x)$, then*

$$\tfrac{1}{\lambda} \log_{p_x}(x) \in \partial f(p_x). \tag{6.24}$$

Starting from Lemma 6.21, we prove the following expression.

Theorem 6.22 *Let $f\colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ be a proper convex lsc function defined on a Hadamard manifold $\mathcal{H}$. Then, for all $\lambda > 0$, the Moreau-Yosida envelope f_λ is differentiable with Riemannian gradient*

$$\nabla f_\lambda(x) = -\tfrac{1}{\lambda} \log_x(J_\lambda x)\,, \quad J_\lambda x := \operatorname{Prox}_{\lambda f}(x).$$

Proof. We have to prove for every $x \in \mathcal{H}$ and $y \to x$ that

$$\big|f_\lambda(y) - f_\lambda(x) - \langle -\tfrac{1}{\lambda} \log_x(J_\lambda x), \log_x y\rangle\big| = o(\mathrm{d}(x,y))\,. \tag{6.25}$$

Recall that we write $f \in o(t)$ as $t \to 0$ if $f(t)/t \to 0$ for $t \to 0$.

We estimate $f_\lambda(y) - f_\lambda(x) - \langle -\frac{1}{\lambda} \log_x(J_\lambda x), \log_x y\rangle$ from above and below.

1. Lemma 6.21 guarantees $\frac{1}{\lambda} \log_{J_\lambda y}(y) \in \partial f(J_\lambda y)$ and by the convexity of f we obtain

$$f(J_\lambda x) - f(J_\lambda y) \geq \langle \log_{J_\lambda y}(J_\lambda x), \tfrac{1}{\lambda} \log_{J_\lambda y}(y)\rangle. \tag{6.26}$$

We know that $-\log_y(J_\lambda x)$ is in the subgradient of the convex function $\frac{1}{2}\,\mathrm{d}(\cdot, J_\lambda x)^2$ at y so that

$$\tfrac{1}{2}\,\mathrm{d}(x, J_\lambda x)^2 - \tfrac{1}{2}\,\mathrm{d}(y, J_\lambda x)^2 \geq \langle \log_y x, -\log_y(J_\lambda x)\rangle,$$

and similarly as $-\log_{J_\lambda y} y$ is in the subgradient of $\frac{1}{2}\,\mathrm{d}(\cdot, y)^2$ at $J_\lambda y$ we have

$$\tfrac{1}{2}\,\mathrm{d}(J_\lambda x, y)^2 - \tfrac{1}{2}\,\mathrm{d}(J_\lambda y, y)^2 \geq \langle \log_{J_\lambda y}(J_\lambda x), -\log_{J_\lambda y}(y)\rangle.$$

Adding the last two estimates gives

$$\tfrac{1}{2}\,\mathrm{d}(x, J_\lambda x)^2 - \tfrac{1}{2}\,\mathrm{d}(y, J_\lambda y)^2 \geq -\langle \log_y x, \log_y(J_\lambda x)\rangle - \langle \log_{J_\lambda y}(J_\lambda x), \log_{J_\lambda y}(y)\rangle. \tag{6.27}$$

Multiplying by $\frac{1}{\lambda}$ and adding (6.26), we obtain after rearranging

$$\begin{aligned} f_\lambda(y) - f_\lambda(x) &\leq \tfrac{1}{\lambda}\langle \log_y x, \log_y(J_\lambda x)\rangle \\ &= \tfrac{1}{\lambda}\langle -\log_x y, \mathcal{T}_{yx}\big(\log_y(J_\lambda x)\big)\rangle. \end{aligned} \tag{6.28}$$

Consequently we get

$$f_\lambda(y) - f_\lambda(x) - \langle -\tfrac{1}{\lambda} \log_x(J_\lambda x), \log_x y\rangle \leq -\tfrac{1}{\lambda}\langle \mathcal{T}_{yx}\big(\log_y(J_\lambda x)\big) - \log_x(J_\lambda x), \log_x y\rangle. \tag{6.29}$$

As the parallel transport is an isometry, we see that

$$\|\mathcal{T}_{yx} \log_y(J_\lambda x) - \log_x(J_\lambda x)\| = \|-\log_{J_\lambda x}(y) + \log_{J_\lambda x}(x)\| \leq \varepsilon_x\big(\mathrm{d}(x,y)\big)$$

where $\varepsilon_x\colon \mathbb{R}_+ \to \mathbb{R}_+$ with $\varepsilon_x(t) \to 0$ as $t \to 0$ is the modulus of continuity of the continuous function $\log_{J_\lambda x}$ at x. We conclude that the right-hand side in (6.29) can be estimated as

$$\begin{aligned} f_\lambda(y) - f_\lambda(x) - \langle -\tfrac{1}{\lambda} \log_x(J_\lambda x), \log_x y\rangle &\leq \tfrac{1}{\lambda}\,\mathrm{d}(x,y)\,\varepsilon_x\big(\mathrm{d}(x,y)\big) \\ &\in o\big(\mathrm{d}(x,y)\big) \end{aligned} \tag{6.30}$$

as $y \to x$.

2. Exchanging x and y in (6.28) we obtain

$$-\tfrac{1}{\lambda}\langle \log_x(J_\lambda y), \log_x y\rangle \leq f_\lambda(y) - f_\lambda(x)$$

and hence

$$\begin{aligned} f_\lambda(y) - f_\lambda(x) - \langle -\tfrac{1}{\lambda}\log_x(J_\lambda x), \log_x y\rangle &\geq -\tfrac{1}{\lambda}\langle \log_x(J_\lambda y) - \log_x(J_\lambda x), \log_x y\rangle \\ &\geq -\tfrac{1}{\lambda}\,\mathrm{d}(x,y)\,\kappa_x\big(\mathrm{d}(x,y)\big) \end{aligned} \tag{6.31}$$

where $\kappa_x\colon \mathbb{R}_+ \to \mathbb{R}_+$ with $\kappa_x(t) \to 0$ as $t \to 0$ is the modulus of continuity at x of $\log_x \circ J_\lambda$, where both $\log_x$ and the proximal mapping J_λ are continuous. Thus the right-hand side of (6.31) is in $o\big(\mathrm{d}(x,y)\big)$. This completes the proof. □

6.3. Convergence of Functions and their Moreau-Yosida Envelopes

In this section we establish a relation between Moreau-Yosida envelopes and Mosco convergence, using the differential properties of Moreau-Yosida envelopes proven in the previous section. Specifically, a well known result of H. Attouch says that Mosco convergence of a sequence of convex lower semicontinuous functions on a Hilbert space is completely characterized by the pointwise convergence of their Moreau-Yosida envelopes [7, Theorem 1.2]. H. Attouch later on extended this result into a certain class of Banach spaces, see Theorem 6.24 below.

The following theorem was proved by H. Attouch, see [7, Theorem 1.2] and [8, Theorem 3.26].

Theorem 6.23 *Let H be a Hilbert space and let $f_n\colon H \to \mathbb{R}\cup\{+\infty\}$, $n \in \mathbb{N}$, be proper convex lower semicontinuous functions. Then the following statements are equivalent:*

i) $\{f_n\}_n$ converges to a function $f\colon H \to \mathbb{R}\cup\{+\infty\}$ in the sense of Mosco, $f_n \xrightarrow{M} f$.

ii) The sequence of Moreau-Yosida envelopes $\{f_{n,\lambda}\}_n$ of $\{f_n\}_n$ converge pointwise to the Moreau-Yosida envelope f_λ of f for all $\lambda > 0$.

In this section we show the usefulness of the properties of Moreau-Yosida envelopes derived in the previous section by generalizing with their help Theorem 6.23 to Hadamard Manifolds. More precisely, while it has already been known that the Mosco convergence of a sequence of convex lower semicontinuous functions on a Hadamard space implies the pointwise convergence of the corresponding Moreau-Yosida envelopes, the converse implication was an open question. We now fill this gap. In Section 6.3.2 we use the differential properties of Moreau-Yosida envelopes to give a proof in Hadamard manifolds. Thereafter, in Section 6.3.3 we give an independent proof in the general setting of Hadamard spaces.

Related Work Note that both Γ- and Mosco convergences have already been used in the framework of Hadamard spaces. In [61], J. Jost studied harmonic mappings with metric space targets and as a tool he introduced Γ-convergence on Hadamard spaces. He also defined the Mosco convergence by saying that a sequence of convex lsc functions on a Hadamard space Mosco converges if their Moreau-Yosida envelopes converge pointwise [61]. In [66], K. Kuwae and T. Shioya studied both Γ- and Mosco convergence in Hadamard space in depth and obtained numerous generalizations. They already gave the standard definition of the Mosco convergence used in this paper (relying on the notion of weak convergence) and right after their Definition 5.7 in [66] they note "Jost's definition of Mosco convergence... seems unfitting in view of Mosco's original definition." By our main result it follows that both definitions are equivalent.

In [66, Proposition 5.12], the authors prove that the Mosco convergence of *nonnegative* convex lsc functions on a Hadamard space implies the pointwise convergence of their Moreau-Yosida envelopes. This result was later proved in [13, Theorem 4.1] without the nonnegativity assumption. The inverse implication was left open; see [11, Question 5.2.5]. We now answer this question in the positive. As a corollary of our main result we obtain that the Mosco convergence of convex closed sets is equivalent to the Frolík-Wijsman convergence.

In [13, 66] the Mosco convergence of functions on Hadamard spaces was studied in connection with *gradient flows*. In particular, it was shown in [13] that the Mosco convergence of convex lsc functions on a Hadamard space implies the pointwise convergence of the associated gradient flow semigroups. Interestingly, apart from applications of Hadamard space gradient flows into harmonic mappings theory, see e.g., [61], [102, Section 8], there have been also other motivations. Most remarkably, gradient flows of a convex function on a Hadamard space appear as an important tool in Kähler geometry in connection with Donaldson's conjecture on Calabi flows [18, 101]. It has also similarly inspired new developments in Riemannian geometry [53]. Finally, in [14], a gradient flow of a convex continuous function was used to construct a Lipschitz retraction in a Hadamard space. For discrete-time gradient flows of convex functions in Hadamard spaces and their applications, see [11, 12].

6.3.1. Previous Results

Let us begin our exposition with two previous extensions of Theorem 6.23. On the one hand, H. Attouch proved an extension to a certain class of Banach spaces, Theorem 6.24. On the other hand, M. Bačák has shown a partial extension to Hadamard spaces, Theorem 6.25. The backward direction has still been open in Hadamard spaces, and we fill the gap.

Extension to Banach Spaces In order to derive the subdifferential of the Moreau-Yosida envelope in Section 6.2.5 we made use of the gradient of the squared distance function. A natural minimal condition on a reflexive Banach space in which one can hope for f_λ to be differentiable, is therefore that norm and dual norm are Frechét differentiable except at the origin. This class of reflexive Banach spaces is characterized

by the following definition, see [8, p. 305] for a detailed statement and the bibliography in [98]. We define

$$\mathcal{C} = \left\{ \big(X, \|\cdot\|\big) : \begin{array}{l} X \text{ is a reflexive Banach space, } \|\cdot\|_X \text{ and } \|\cdot\|_{X'} \text{ are} \\ \text{strictly convex and satisfy (5.3)} \end{array} \right\}. \tag{6.32}$$

H. Attouch [8, Theorem 3.26] proved the following extension of Theorem 6.23.

Theorem 6.24 *Let X be a reflexive Banach space belonging to the class $\mathcal{C}$. Let $\{f_n\}_n$ be a sequence of proper convex lsc functions $f_n \colon X \to \mathbb{R} \cup \{+\infty\}$ and let $f \colon X \to \mathbb{R} \cup \{+\infty\}$ be a proper convex lsc function. Then the following statements are equivalent:*

(i) $f_n \xrightarrow{M} f$ *as* $n \to \infty$ *i.e. the sequence* $\{f_n\}_n$ *Mosco converges to* f.

(ii) $f_{n,\lambda}(x) \to f_\lambda(x)$ *as* $n \to \infty$, *for all* $\lambda > 0$ *and all* $x \in X$.

(iii) $\begin{cases} \operatorname{Prox}_{\lambda f_n} x \to \operatorname{Prox}_{\lambda f} x \quad \text{for all } \lambda > 0 \text{ and all } x \in X \quad \text{and} \\ \exists (u_n, v_n) \in \partial f_n,\ \exists (u, v) \in \partial f \text{ s.t. } u_n \xrightarrow{X} u,\ v_n \xrightarrow{X^*} v,\ f_n(u_n) \to f(u). \end{cases}$

Here ∂f is the graph of the subdifferential defined by (5.4).

Partial Extension to Hadamard Spaces Implication (i) $\Rightarrow$ (ii) in Theorem 6.23 has been generalized to Hadamard spaces in [11, Theorem 5.2.4]:

Theorem 6.25 *Let $\mathcal{H}$ be a Hadamard space, let $\{f_n\}_n$ be a sequence of proper convex lsc functions $f_n \colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ and let $f \colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ be a proper function. If $f_n \xrightarrow{M} f$ as $n \to \infty$ then*

$$\lim_{n\to\infty} f_{n,\lambda}(x) = f_\lambda(x), \tag{6.33}$$

and

$$\lim_{n\to\infty} \operatorname{Prox}_{\lambda f_n}(x) = \operatorname{Prox}_{\lambda f}(x), \tag{6.34}$$

for every $\lambda \in (0, +\infty)$ and $x \in \mathcal{H}$.

For the convenience of the reader, let us recall the proof. The argument around (6.36) has been slightly modified from the proof in [11, Theorem 5.2.4].

Proof. In the following we are going to abbreviate $J_\lambda^n := \operatorname{Prox}_{\lambda f_n}$ and $J_\lambda := \operatorname{Prox}_{\lambda f}$. Note that the Mosco limit is automatically convex and lsc, so that J_λ is well-defined.

1. Given $x_0 \in \mathcal{H}$, by Lemma 5.16, there exist $\alpha, \beta \in \mathbb{R}$ such that

$$f_n(y) \geq -\alpha\, \mathrm{d}(y, x_0) - \beta, \tag{6.35}$$

for every $y \in \mathcal{H}$ and $n \in \mathbb{N}$. Choose a sequence $\{u_n\}_n \subset \mathcal{H}$ such that $u_n \to x_0$ and $f_n(u_n) \to f(x_0)$. From the definition of $J_\lambda^n x$, we have

$$f_n(u_n) + \tfrac{1}{2\lambda}\, \mathrm{d}(x, u_n)^2 \geq f_n(J_\lambda^n x) + \tfrac{1}{2\lambda}\, \mathrm{d}(x, J_\lambda^n x)^2.$$

By (6.35) with $y = J_\lambda^n x$ and the triangle inequality, further

$$\begin{aligned}\tfrac{1}{2\lambda}\,\mathrm{d}(x, J_\lambda^n x)^2 &\le f_n(u_n) + \tfrac{1}{2\lambda}\,\mathrm{d}(x, u_n)^2 + \alpha\,\mathrm{d}(J_\lambda^n x, x_0) + \beta \\ &\le f_n(u_n) + \tfrac{1}{2\lambda}\,\mathrm{d}(x, u_n)^2 + \alpha\,\mathrm{d}(x, J_\lambda^n x) + \alpha\,\mathrm{d}(x, x_0) + \beta.\end{aligned}$$

Since the term $\mathrm{d}(x, J_\lambda^n x)$ occurs quadratically on the left-hand side and linearly on the right-hand side, it follows that the sequence $\{J_\lambda^n x\}_n$ is bounded.

2. Boundedness of the sequence $\{J_\lambda^n x\}_n$ implies the existence of a weak cluster point $c \in \mathcal{H}$. Let $\{n_k\}_k \subset \mathbb{N}$ be a subsequence of indices with $J_\lambda^{n_k} x \overset{w}{\to} c$ in $\mathcal{H}$, for $k \to \infty$.

The following is slightly modified from the proof in [11, Theorem 5.2.4]. We believe that, for the last inequality [11, eq. 5.2.26] to hold, we need to restrict ourselves to the subsequence $\{n_k\}_k$ first, as follows.

As $f_n \xrightarrow{M} f$, there is a sequence $\{y_n\}_n \subset \mathcal{H}$ such that $y_n \to J_\lambda x$ and $f_n(y_n) \to f(J_\lambda x)$. Then, restricting first to the subsequence, and using in this order, the definition of $f_{n_k,\lambda}$, that y_n is a recovery sequence, the definition of $J_\lambda x$ tested with c and finally Mosco convergence and weak lower semicontinuity of $\mathrm{d}(x,\cdot)^2$, we obtain

$$\begin{aligned}\limsup_{k\to\infty} f_{n_k,\lambda}(x) &\le \limsup_{k\to\infty}\Big(f_{n_k}(y_{n_k}) + \tfrac{1}{2\lambda}\,\mathrm{d}(x, y_{n_k})^2\Big) \\ &= f(J_\lambda x) + \tfrac{1}{2\lambda}\,\mathrm{d}(x, J_\lambda x)^2 \\ &\le f(c) + \tfrac{1}{2\lambda}\,\mathrm{d}(x, c)^2 \\ &\le \liminf_{k\to\infty}\Big(f_{n_k}(J_\lambda^{n_k} x) + \tfrac{1}{2\lambda}\,\mathrm{d}(x, J_\lambda^{n_k} x)^2\Big). \qquad (6.36)\end{aligned}$$

Thus all inequalities are equalities, which gives $J_\lambda x = c$, by the uniqueness of $J_\lambda x$. As c was arbitrary, this implies $J_\lambda^n x \overset{w}{\to} J_\lambda x$, as $n \to \infty$.

Now the chain of equalities holds with limits taken along the whole sequence. For real sequences $\{a_n\}_n$, $\{b_n\}_n$ one has $\liminf a_n + \limsup b_n \le \limsup(a_n + b_n)$. Using this estimate and reordering the first line of (6.36) yields

$$\begin{aligned}\limsup_{n\to\infty} \tfrac{1}{2\lambda}\,\mathrm{d}(x, J_\lambda^n x)^2 &\le -\liminf_{n\to\infty} f_n(J_\lambda^n x) + \limsup_{n\to\infty} f_n(y_n) + \limsup_{n\to\infty} \tfrac{1}{2\lambda}\,\mathrm{d}(x, y_n)^2 \\ &\le -\liminf_{n\to\infty} f_n(J_\lambda^n x) + f(J_\lambda x) + \tfrac{1}{2\lambda}\,\mathrm{d}(x, J_\lambda x)^2 \\ &\le \tfrac{1}{2\lambda}\,\mathrm{d}(x, J_\lambda x)^2 \\ &\le \tfrac{1}{2\lambda}\liminf_{n\to\infty} \mathrm{d}(x, J_\lambda^n x)^2. \qquad (6.37)\end{aligned}$$

By Theorem 5.11 (iii), the weak convergence of the $J_\lambda^n x$, from equation (6.36), and convergence of the "distance to a point", from equation (6.37), together imply strong convergence

$$J_\lambda^n x \to J_\lambda x, \quad \text{as } n \to \infty,$$

which proves (6.34). Finally, inequality (6.36) with f_λ in the second line, immediately gives (6.33), and the proof is complete. □

Remark 6.26 (Simpler proof in $\mathbb{R}^n$) *If the functions are defined not on a general Hadamard space but on $\mathbb{R}^d$, Mosco convergence agrees with Γ-convergence and Theorem 6.25 follows from Theorem 5.4, applied for each $\lambda > 0$ and $x \in \mathbb{R}^d$ as follows: Define functions g, g^n by*

$$g^n(\cdot) = f_n(\cdot) + \tfrac{1}{2\lambda}\,\mathrm{d}(\cdot, x)^2, \quad n = 1, 2, \dots$$
$$g(\cdot) = f(\cdot) + \tfrac{1}{2\lambda}\,\mathrm{d}(\cdot, x)^2.$$

Then it follows from equation (6.35) that the g_n are equi-mildly coercive. Minima and minimizers of the g_n are $f_{n,\lambda}(x)$ and $J_\lambda^n x$, and the unique minimizer of g is $J_\lambda x$, hence Theorem 5.4 and Remark 5.5 yield that $f_{n,\lambda}(x) \to f_\lambda(x)$ and $J_\lambda^n x \to J_\lambda x$, respectively, as $n \to \infty$.

Remark 6.27 *This argument does not work for Mosco convergence in the infinite-dimensional case, exactly for the reason which made Mosco convergence an interesting concept: It is an intermediate notion between Γ-convergence w.r.t. weak convergence of sequences and Γ-convergence w.r.t. strong convergence of sequences.*

In an infinite-dimensional setting Mosco convergence implies that the g^n Γ-converge w.r.t. strong convergence of sequences, but bounded sets are not compact w.r.t. strong convergence; on the other hand, while bounded sets are weakly sequentially compact, Mosco convergence does not allow us to conclude that the g^n Γ-converge w.r.t. weak convergence of sequences.

6.3.2. Generalization to Hadamard Manifolds using Gradients of Moreau-Yosida Envelopes

In the previous section we derived several properties of Moreau-Yosida envelopes. In this section we show their usefulness in generalizing Theorem 6.23 to Hadamard manifolds. First we present the line of argument in the simpler setting of finite-dimensional manifolds, see Theorem 6.29, then give the proof in infinite-dimensional manifolds, see Theorem 6.30.

We will further need the following auxiliary lemma, which holds more generally for metric spaces. The Lipschitz constants of the envelopes in the last part are obtained in similar way, as for a single function in [8, Theorem 2.64].

Lemma 6.28 *Let $\mathcal{H}$ be a Hadamard space and $\{f_n\}_n$ a sequence of functions $f_n \colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$. Suppose that there exist $0 < \lambda_1 < \lambda_2$, such that for $\lambda \in \{\lambda_1, \lambda_2\}$ the Moreau-Yosida envelopes fulfil $f_{n,\lambda} \to f_\lambda$ pointwise as $n \to \infty$, for some proper function $f \colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$. Then, given $x_0 \in \mathcal{H}$, there exists $n_0 \in \mathbb{N}$ such that the following relations hold:*

(i) *The functions f_n have a common minorizing quadratic, i.e., there exists $r > 0$ such that for all $x \in \mathcal{H}$ and all $n \geq n_0$ we have*

$$f_n(x) \geq -r\,\mathrm{d}(x, x_0)^2 - r.$$

(ii) *There exists $R_0 > 0$ and a sequence $\{x_n\}_n \subset B_{R_0}(x_0)$ in the ball $B_{R_0}(x_0)$ of radius R_0 around x_0, such that for all $n \geq n_0$*

$$f_n(x_n) \leq L$$

with some constant $L \in \mathbb{R}$.

(iii) *For each $\lambda > 0$, the sequence of Moreau-Yosida envelopes $\{f_{n,\lambda}\}_n$ is locally equi-bounded, i.e. for all $R > 0$, there exists $K_R \in \mathbb{R}$ such that for all $x \in B_R(x_0)$ and all $n \geq n_0$,*

$$f_{n,\lambda}(x) \leq K_R.$$

(iv) *The sequence of Moreau-Yosida envelopes $\{f_{n,\lambda}\}_n$ is locally equi-Lipschitz, i.e., there is a function $C\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, bounded on bounded subsets, such that for all $\lambda \in (0, \frac{1}{8r}]$, all $x, \tilde{x} \in B_R(x_0)$ and all $n \geq n_0$ it holds*

$$|f_{n,\lambda}(x) - f_{n,\lambda}(\tilde{x})| < \tfrac{1}{\lambda}\,\mathrm{d}(x, \tilde{x})\, C(R).$$

Moreover, $C(R)$ depends only and in a continuous way on R, R_0, r and $f_{\lambda_1}(x_0)$.

Proof. (i) Since f is proper, the function f_λ is everywhere finite. By assumption, there is some n_0 such that for all $n \geq n_0$ and $\lambda \in \{\lambda_1, \lambda_2\}$,

$$|f_{n,\lambda}(x_0) - f_\lambda(x_0)| \leq 1. \tag{6.38}$$

Since

$$f_{n,\lambda}(x_0) = \inf_y \Big\{ f_n(y) + \tfrac{1}{2\lambda}\,\mathrm{d}(y, x_0)^2 \Big\}$$

we obtain for $\lambda \in \{\lambda_1, \lambda_2\}$, $y \in \mathcal{H}$ and $n \geq n_0$,

$$\begin{aligned} f_n(y) &\geq f_{n,\lambda}(x_0) - \tfrac{1}{2\lambda}\,\mathrm{d}(y, x_0)^2 \\ &\geq f_\lambda(x_0) - 1 - \tfrac{1}{2\lambda}\,\mathrm{d}(y, x_0)^2 \end{aligned} \tag{6.39}$$

which gives assertion (i) for $r := \max\big(\frac{1}{2\lambda}, 1 - f_\lambda(x_0)\big)$.

(ii) Let $n \geq n_0$, where n_0 is defined as in (i) and $x_n := \mathrm{Prox}_{\lambda_1 f_n}(x_0)$. By (6.38) we obtain

$$f_{\lambda_1}(x_0) + 1 \geq f_{n,\lambda_1}(x_0) = f_n(x_n) + \tfrac{1}{2\lambda_1}\,\mathrm{d}(x_n, x_0)^2, \tag{6.40}$$

and from (6.39) with $y := x_n$ and $\lambda := \lambda_2$ further

$$1 - f_{\lambda_2}(x_0) \geq -\tfrac{1}{2\lambda_2}\,\mathrm{d}(x_n, x_0)^2 - f_n(x_n).$$

Summing up these inequalities gives

$$f_{\lambda_1}(x_0) - f_{\lambda_2}(x_0) + 2 \geq \Big(\tfrac{1}{2\lambda_1} - \tfrac{1}{2\lambda_2}\Big)\,\mathrm{d}(x_n, x_0)^2. \tag{6.41}$$

Consequently, $\{x_n\}_n \subset B_{R_0}(x_0)$ for $R_0 > \Big((f_{\lambda_1}(x_0) - f_{\lambda_2}(x_0) + 2)/(\frac{1}{2\lambda_1} - \frac{1}{2\lambda_2})\Big)^{\frac{1}{2}}$ and $n \geq n_0$. Since by (6.40) also $f_n(x_n) \leq f_{\lambda_1}(x_0) + 1 =: L$, we obtain (ii).

(iii) Given $R > 0$, for every $n \geq n_0$, and $x \in B_R(x_0)$, we have with $x_n := \operatorname{Prox}_{\lambda_1 f_n}(x_0)$ that

$$\mathrm{d}(x, x_n) \leq \mathrm{d}(x, x_0) + \mathrm{d}(x_0, x_n) \leq R + R_0.$$

Hence, for all $\lambda > 0$,

$$f_{n,\lambda}(x) \leq f_n(x_n) + \tfrac{1}{2\lambda}\,\mathrm{d}(x_n, x)^2 \leq L + \tfrac{1}{2\lambda}(R + R_0)^2 =: K_R,$$

as required for (iii).

(iv) For $n \geq n_0$ arbitrarily fixed, let $g := f_n$. For $0 < \lambda \leq \frac{1}{8r}$, $x \in \mathcal{H}$ and $\varepsilon > 0$, let $p = p(x, \lambda, \varepsilon, n)$ denote a point such that

$$g_\lambda(x) \leq g(p) + \tfrac{1}{2\lambda}\,\mathrm{d}(p, x)^2 \leq g_\lambda(x) + \varepsilon. \tag{6.42}$$

By the quadratic minorization property in (i) we conclude

$$-r\big(\mathrm{d}(p, x_0)^2 + 1\big) + \tfrac{1}{2\lambda}\,\mathrm{d}(p, x)^2 \leq g_\lambda(x) + \varepsilon.$$

Transforming equivalently and inserting (6.42) we get

$$\begin{aligned}\mathrm{d}(p, x)^2 &\leq 2\lambda r\ \mathrm{d}(p, x_0)^2 + 2\lambda(g_\lambda(x) + \varepsilon + r)\\ &\leq 4\lambda r\ \mathrm{d}(p, x)^2 + 4\lambda r\ \mathrm{d}(x, x_0)^2 + 2\lambda(g_\lambda(x) + \varepsilon + r)\end{aligned}$$

and then, since $1 - 4\lambda r \leq \frac{1}{2}$,

$$\mathrm{d}(p, x)^2 \leq 8\lambda r\ \mathrm{d}(x, x_0)^2 + 4\lambda(g_\lambda(x) + \varepsilon + r). \tag{6.43}$$

Consider two points $x, \tilde{x} \in \mathcal{H}$ and corresponding points p and $\tilde{p}$ satisfying (6.42). Then it follows

$$\begin{aligned}g_\lambda(\tilde{x}) &\leq g(p) + \tfrac{1}{2\lambda}\,\mathrm{d}(p, \tilde{x})^2\\ &\leq g_\lambda(x) + \varepsilon + \tfrac{1}{2\lambda}\left(\mathrm{d}(p, \tilde{x})^2 - \mathrm{d}(p, x)^2\right).\end{aligned}$$

Together with

$$\mathrm{d}(p, \tilde{x})^2 - \mathrm{d}(p, x)^2 \leq \mathrm{d}(x, \tilde{x})\,(2\,\mathrm{d}(p, x) + \mathrm{d}(\tilde{x}, x)) \leq \mathrm{d}(x, \tilde{x})\left(\mathrm{d}(p, x)^2 + 1 + \mathrm{d}(\tilde{x}, x)\right)$$

and (6.43) this becomes

$$\begin{aligned}g_\lambda(\tilde{x}) - g_\lambda(x) &\leq \varepsilon + \tfrac{1}{2\lambda}\ \mathrm{d}(x, \tilde{x})\left(\mathrm{d}(p, x)^2 + 1 + \mathrm{d}(\tilde{x}, x)\right)\\ &\leq \varepsilon + \tfrac{1}{2\lambda}\ \mathrm{d}(x, \tilde{x})\left(8\lambda r\ \mathrm{d}(x, x_0)^2 + 4\lambda(g_\lambda(x) + r + \varepsilon) + 1 + \mathrm{d}(\tilde{x}, x)\right).\end{aligned}$$

Letting $\varepsilon \to 0$ and exchanging the roles of $x, \tilde{x}$, we obtain, for $0 < \lambda \leq \frac{1}{8r}$ and $x, \tilde{x} \in B_R(x_0)$ by (iii), that

$$\begin{aligned}|g_\lambda(\tilde{x}) - g_\lambda(x)| \ &\leq \tfrac{1}{\lambda}\ \mathrm{d}(x, \tilde{x}) \sup_{x, \tilde{x} \in B_R(x_0)} \big\{4\lambda r\ \max\left(\mathrm{d}(x, x_0)^2,\ \mathrm{d}(\tilde{x}, x_0)^2\right)\\ &\qquad +2\lambda \max\left(g_\lambda(x),\ g_\lambda(\tilde{x})\right) + 2\lambda r + \tfrac{1}{2} + \tfrac{1}{2}\,\mathrm{d}(\tilde{x}, x)\big\}\\ &\leq \tfrac{1}{\lambda}\ \mathrm{d}(x, \tilde{x})\ \underbrace{\left(\tfrac{1}{2}R^2 + 2\lambda K_R + \tfrac{3}{4} + R\right)}_{C(R)}.\end{aligned}$$

This completes the proof. □

Main Result in Finite Dimensional Hadamard Manifolds

Before give the main result in infinite-dimensional Hadamard manifolds, let us consider the simpler setting of finite-dimensional Hadamard manifolds. In finite-dimensional Hadamard manifolds weak and strong convergence agree and the sequence $\{f_n\}_n$ Mosco converges to f if and only if it Γ-converges to f.

Theorem 6.29 *Let $\mathcal{H}$ be a finite-dimensional Hadamard manifold. Let $\{f_n\}_n$ be a sequence of proper convex lsc functions $f_n \colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$. For all $\lambda > 0$, let the sequence of Moreau-Yosida envelopes $f_{n,\lambda}$ converge pointwise to the Moreau-Yosida envelope f_λ of some proper convex lsc function $f \colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$. Then $f_n \xrightarrow{\Gamma} f$ as $n \to \infty$.*

Proof. We abbreviate the proximal mappings of f and f_n by $J_\lambda := \operatorname{Prox}_{\lambda f}$ and $J_\lambda^n := \operatorname{Prox}_{\lambda f_n}$. By Proposition 6.22, the envelopes f_λ and $f_{n,\lambda}$ are differentiable. By convexity, for all $\lambda > 0$, all $n \in \mathbb{N}$ and all $x, y \in \mathcal{H}$ we have the subdifferential inequality

$$f_{n,\lambda}(y) - f_{n,\lambda}(x) \geq \langle \nabla f_{n,\lambda}(x), \log_x y \rangle. \tag{6.44}$$

1. First we show for all $\lambda > 0$ and all $x \in \mathcal{H}$ that

$$\nabla f_{n,\lambda}(x) \to \nabla f_\lambda(x) \tag{6.45}$$

as $n \to \infty$. Indeed, from (6.44) and the equi-local Lipschitz property of the $\{f_{n,\lambda}\}_n$, proven in Theorem 6.28 (iv), it follows for all $\lambda > 0$ and all $x \in \mathcal{H}$ that the sequence $\{\nabla f_{n,\lambda}(x)\}_n$ is bounded. As closed bounded sets in $\mathbb{R}^N$ are compact, any subsequence has a further subsequence which converges strongly,

$$\nabla f_{n_k,\lambda}(x) \to \xi \qquad \text{as } k \to \infty,$$

for some $\xi \in T_x\mathcal{H}$. Passing to the limit in (6.44) using the convergence of $\{f_{n,\lambda}(x)\}_n$ we obtain that

$$f_\lambda(y) \geq f_\lambda(x) + \langle \xi, \log_x y \rangle,$$

for all $\lambda > 0$ and all $x, y \in \mathcal{H}$. Hence $\xi \in \partial f_\lambda(x)$ and as f_λ is differentiable, we get $\xi = \nabla f_\lambda(x)$. Therefore $\nabla f_{n,\lambda}(x) \to \nabla f_\lambda(x)$ in $T_x\mathcal{H}$.
2. Following [8, p. 317], we show Γ-convergence

$$f_n \xrightarrow{\Gamma} f.$$

Let us note that by Lemma 6.28 (i), for every fixed $x \in \mathcal{H}$, there is an $r > 0$ such that for all n large enough and all $y \in \mathcal{H}$

$$f_n(y) \geq -r d(y, x)^2 - r. \tag{6.46}$$

Now we verify both properties of Γ-convergence.

2.1 *Limsup Inequality.* From (5.7) we know that f_λ increases to f as $\lambda \to 0$. Together with the assumption, that $f_{n,\lambda}(x) \to f_\lambda(x)$, for all $x \in \mathcal{H}$ as $n \to \infty$ we obtain

$$f(x) = \lim_{\lambda \to 0} f_\lambda(x) = \lim_{\lambda \to 0} \lim_{n \to \infty} f_{n,\lambda}(x).$$

By Lemma 5.6 (diagonalization lemma), there exists a sequence $\{\lambda_n\}_n$ with $\lim_{n\to\infty}\lambda_n = 0$ such that

$$\begin{aligned} f(x) &= \lim_{n\to\infty} f_{n,\lambda_n}(x) \\ &= \lim_{n\to\infty}\Big(f_n\left(J^n_{\lambda_n}(x)\right) + \tfrac{1}{2\lambda_n} d\left(x, J^n_{\lambda_n}(x)\right)^2\Big). \end{aligned} \tag{6.47}$$

Taking $x_n := J^n_{\lambda_n}(x)$, we obtain

$$f(x) \geq \limsup_{n\to\infty} f_n(x_n).$$

It remains to show that, in the case $f(x) < +\infty$, we have $x_n \to x$ as $n \to \infty$. Indeed, inserting Lemma 6.28(i) into (6.47), we have

$$f(x) \geq \limsup_{n\to\infty}\Big(\big(\tfrac{1}{2\lambda_n} - r\big)d(x_n, x)^2 - r\Big)$$

and because $\left(\frac{1}{2\lambda_n} - r\right) \to +\infty$, this is only possible if $d(x_n, x) \to 0$ as $n \to \infty$. (In the case $f(x) = +\infty$, there is nothing to prove.)
2. *Liminf Inequality.* Suppose that $y_n \to x$ as $n \to \infty$. Then $\log_x y_n \to 0$. From the subdifferential inequality (6.44) we have

$$\begin{aligned} f_n(y_n) &\geq f_{n,\lambda}(y_n) \\ &\geq f_{n,\lambda}(x) + \langle \nabla f_{n,\lambda}(x), \log_x y_n\rangle. \end{aligned}$$

Using boundedness the sequence $\{\nabla f_{n,\lambda}(x)\}_n$, we take the limit $n \to \infty$ and get

$$\begin{aligned} \liminf_{n\to\infty} f_n(y_n) &\geq \lim_{n\to\infty}\big(f_{n,\lambda}(x) + \langle \nabla f_{n,\lambda}(x), \log_x y_n\rangle\big) \\ &= f_\lambda(x). \end{aligned}$$

Taking the supremum over $\lambda > 0$, we finally obtain

$$\liminf_{n\to\infty} f_n(y_n) \geq f(x). \qquad \square$$

Main Result on Hilbert Manifolds

Theorem 6.30 *Let $\mathcal{H}$ be a Hadamard manifold satisfying property (A). Let $\{f_n\}_n$ be a sequence of proper convex lsc functions $f_n\colon \mathcal{H} \to \mathbb{R}\cup\{+\infty\}$. For all $\lambda > 0$, let the sequence of Moreau-Yosida envelopes $f_{n,\lambda}$ converge pointwise to the Moreau-Yosida envelope f_λ of some proper function $f\colon \mathcal{H} \to \mathbb{R}\cup\{+\infty\}$. Then $f_n \xrightarrow{M} f$ as $n \to \infty$.*

Recall from Section 6.2.3 that a sufficient condition for (A) is a lower bound on the curvature. Note also that (A) holds for finite-dimensional Hadamard manifolds. In this case Mosco convergence can be replaced by Γ-convergence and part 1 of the following proof simplifies, as we have seen above. Basically the proof follows along the lines of the proof in [8, Theorem 3.26].

Proof. 1. First we prove, for all $\lambda > 0$ and all $x \in \mathcal{H}$, that the sequence $\{\nabla f_{n,\lambda}(x)\}_n$ converges to $\nabla f_\lambda(x)$.

1.1. First note that f is proper convex lsc. Indeed, for all $\lambda > 0$ and all n, the envelopes $f_{n,\lambda}$ are convex and continuous. Thus $f = \sup_{\lambda>0} f_\lambda$ is a convex lsc function. It follows further that $J_\lambda x$ is well-defined and from Theorem 6.22 that f is differentiable.
1.2. We start by showing weak convergence, i.e., that for all $\lambda > 0$ and all $x \in \mathcal{H}$ it holds

$$\nabla f_{n,\lambda}(x) \xrightarrow{w} \nabla f_\lambda(x). \tag{6.48}$$

By definition of the subdifferential (6.3) and the equi-local Lipschitz property of $f_{n,\lambda}$ proven in Lemma 6.28(iv) we have for an arbitrary fixed $y \in \mathcal{H}$ that

$$\langle \nabla f_{n,\lambda}(x), \log_x y\rangle \le f_{n,\lambda}(y) - f_{n,\lambda}(x) \le \tfrac{1}{\lambda}\,\mathrm{d}(x,y)\,C\big(\mathrm{d}(x,y)\big).$$

Since $y \in \mathcal{H}$ can be arbitrarily chosen, this implies that for every $\zeta \in T_x\mathcal{H}$, the sequence $\{\langle \nabla f_{n,\lambda}(x),\ \zeta\rangle\}_n$ is bounded. Taking $\zeta_n := \nabla f_{n,\lambda}(x)/\|\nabla f_{n,\lambda}(x)\|$, we obtain that the sequence $\{\nabla f_{n,\lambda}(x)\}_n$ is bounded in $T_x\mathcal{H}$. The tangent space $T_x\mathcal{H}$ is a Hilbert space so that the unit ball in $T_x\mathcal{H}$ is weakly sequentially compact. Thus there exists a weakly convergent subsequence $\nabla f_{n^k,\lambda}(x)$

$$\nabla f_{n^k,\lambda}(x) \xrightarrow{w} \xi, \quad \text{as } k \to \infty$$

for some $\xi \in T_x\mathcal{H}$. Passing to the limit $k \to \infty$ in the subdifferential equation

$$f_{n_k,\lambda}(y) \ge f_{n_k,\lambda}(x) + \langle \nabla f_{n_k,\lambda}(x), \log_x y\rangle$$

we obtain by the convergence of $\{f_{n,\lambda}(x)\}_n$ that

$$f_\lambda(y) \ge f_\lambda(x) + \langle \xi, \log_x y\rangle,$$

for all $\lambda > 0$ and all $x, y \in \mathcal{H}$. Hence $\xi \in \partial f_\lambda(x)$ and as f_λ is differentiable, we get $\xi = \nabla f_\lambda(x)$. Therefore $\nabla f_{n,\lambda}(x) \xrightarrow{w} \nabla f_\lambda(x)$ in $T_x\mathcal{H}$.

1.3. Next we prove the strong convergence. As, in a Hilbert space, weak convergence and convergence of the norm imply strong convergence, we need only show

$$\|\nabla f_{n,\lambda}(x)\| \to \|\nabla f_\lambda(x)\|$$

for all $x \in \mathcal{H}$. By Theorem 6.22 we know that

$$\nabla f_{n,\lambda}(x) = -\tfrac{1}{\lambda}\log_x(J^n_\lambda x), \quad J^n_\lambda x := \mathrm{Prox}_{\lambda f_n}(x)$$

Hence $\mathrm{d}(x, J^n_\lambda x) = \|\log_x(J^n_\lambda x)\| = \lambda\|\nabla f_{n,\lambda}(x)\|$ and

$$f_{n,\lambda}(x) = f_n(J^n_\lambda x) + \tfrac{1}{2\lambda}\,\mathrm{d}(x, J^n_\lambda x)^2 = f_n(J^n_\lambda x) + \tfrac{\lambda}{2}\|\nabla f_{n,\lambda}(x)\|^2.$$

By assumption, $f_{n,\lambda}(x) \to f_\lambda(x)$ so that it remains to show for all $\lambda > 0$ and all $x \in \mathcal{H}$ that

$$f_n(J^n_\lambda x) \to f(J_\lambda x). \tag{6.49}$$

To this end, let us introduce, for each x, the two sequences of real valued functions $\{h_n\}_n$ and $\{F_n\}_n$ of the variable $\lambda > 0$:

$$h_n(\lambda) := \lambda\, f_{n,\lambda}(x), \quad F_n(\lambda) := f_n(J_\lambda^n x).$$

By Lemma 6.19 we have

$$\tfrac{\mathrm{d}}{\mathrm{d}\lambda} h_n(\lambda) = F_n(\lambda).$$

for $n \in \mathbb{N}$. By assumption

$$h_n(\lambda) \to h(\lambda) := \lambda\, f_\lambda(x)$$

as $n \to \infty$. By Lemma 6.20 the pointwise convergence of $\{h_n\}_n$ implies the pointwise convergence of $\{h_n'\}_n = \{F_n\}_n$ to $h' = F$. In other words, we have shown (6.55).

2. Now we show Mosco convergence

$$f_n \xrightarrow{M} f.$$

Let us note that by Lemma 6.28(i), for every fixed $x \in \mathcal{H}$, there is an $r > 0$ such that for all n large enough and all $y \in \mathcal{H}$

$$f_n(y) \geq -r d(y,x)^2 - r. \tag{6.50}$$

Let us now verify both properties of Mosco convergence.

2.1 *Limsup Inequality.* From 5.7 we know that f_λ increases to f as $\lambda \to 0$. Together with the assumption, that $f_{n,\lambda}(x) \to f_\lambda(x)$, for all $x \in \mathcal{H}$ as $n \to \infty$ we obtain

$$f(x) = \lim_{\lambda\to 0} f_\lambda(x) = \lim_{\lambda\to 0} \lim_{n\to\infty} f_{n,\lambda}(x).$$

By the diagonalization lemma, see [8, Lemma 1.18], there exists a sequence $\{\lambda_n\}_n$ with $\lim_{n\to\infty} \lambda_n = 0$ such that

$$\begin{aligned} f(x) &= \lim_{n\to\infty} f_{n,\lambda_n}(x) \\ &= \lim_{n\to\infty} \left(f_n\big(J_{\lambda_n}^n(x)\big) + \tfrac{1}{2\lambda_n}\, \mathrm{d}\big(x, J_{\lambda_n}^n(x)\big)^2 \right). \end{aligned} \tag{6.51}$$

Taking $x_n := J_{\lambda_n}^n(x)$, we obtain

$$f(x) \geq \limsup_n f_n(x_n).$$

It remains to show that, in the case $f(x) < +\infty$, we have $x_n \to x$ as $n \to \infty$. Indeed, inserting Lemma 6.28(i) into (6.51), we have

$$f(x) \geq \limsup_{n\to\infty} \left(\big(\tfrac{1}{2\lambda_n} - r\big)\, \mathrm{d}(x_n, x)^2 - r \right)$$

and because $\big(\frac{1}{2\lambda_n} - r\big) \to +\infty$, this is only possible if $\mathrm{d}(x_n, x) \to 0$ as $n \to \infty$. (In the case $f(x) = +\infty$, there is nothing to prove.)

2.2 *Liminf Inequality.* Suppose that $x_n \overset{w}{\to} x$ as $n \to \infty$. Using strong convergence of the sequence $\{\nabla f_{n,\lambda}(x)\}_n$ from part 1 of the proof, we want to pass to the limit in

$$\begin{aligned} f_n(x_n) &\geq f_{n,\lambda}(x_n) \\ &\geq f_{n,\lambda}(x) + \langle \nabla f_{n,\lambda}(x), \log_x x_n \rangle. \end{aligned}$$

By assumption (A) the weak convergence $x_n \overset{w}{\to} x$ in $\mathcal{H}$ implies $\log_x x_n \overset{w}{\to} 0$ in $T_x\mathcal{H}$. Together with $\nabla f_{n,\lambda}(x) \to \nabla f_\lambda(x)$ from part 1 of the proof this implies the convergence $\langle \nabla f_{n,\lambda}(x), \log_x x_n \rangle \to 0$ as $n \to \infty$ and we have

$$\begin{aligned} \liminf_{n\to\infty} f_n(x_n) &\geq \lim_{n\to\infty} \big(f_{n,\lambda}(x) + \langle \nabla f_{n,\lambda}(x), \log_x x_n \rangle \big) \\ &= f_\lambda(x). \end{aligned}$$

Taking the supremum over $\lambda > 0$, we finally obtain

$$\liminf_{n\to\infty} f_n(x_n) \geq f(x). \qquad \square$$

6.3.3. Generalization to Hadamard Spaces

In this section we prove the inverse implication to Theorem 6.25, and thus generalize Theorem 6.23 to Hadamard spaces. The proof given has originated from joint work with M. Bačák, submitted for publication [15]. It is independent of the proof in Hadamard manifolds and the only result used from the previous section is Lemma 6.28 (i). Hadamard manifolds being Hadamard spaces, Theorem 6.30 from the previous section is contained in the following theorem.

Theorem 6.31 *Let $\mathcal{H}$ be a Hadamard space, $\{f_n\}_n$ a sequence of proper convex lsc functions $f_n \colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$, and $f \colon \mathcal{H} \to \mathbb{R} \cup \{+\infty\}$ a proper convex lsc function. Assume that for each $\lambda > 0$ the sequence of Moreau-Yosida envelopes $\{f_{n,\lambda}\}_n$ converges pointwise to the Moreau-Yosida envelope f_λ. Then $f_n \xrightarrow{M} f$ as $n \to \infty$.*

Proof. Observe that $f(x) \geq f_\lambda(x) \geq f\left(J_\lambda x\right)$. For $x \in \overline{\operatorname{dom} f}$, it holds by [11, Proposition 2.2.26] that $\lim_{\lambda\to+0} J_\lambda x = x$ so that the lower semicontinuity of f implies

$$f(x) = \lim_{\lambda\to+0} f_\lambda(x) = \lim_{\lambda\to+0} f\left(J_\lambda x\right). \tag{6.52}$$

1. *(Limsup Inequality)* Let us show that, given $x \in \mathcal{H}$, there exists a sequence $y_n \to x$ with $\limsup_{n\to\infty} f_n(y_n) \leq f(x)$. If $f(x) = +\infty$, then there is nothing to prove. Assume therefore $x \in \operatorname{dom} f$. Together with the assumption that $f_{n,\lambda}(x) \to f_\lambda(x)$ for all $x \in \mathcal{H}$ as $n \to \infty$, we obtain

$$f(x) = \lim_{\lambda\to+0} f_\lambda(x) = \lim_{\lambda\to+0} \lim_{n\to\infty} f_{n,\lambda}(x).$$

By the diagonalization lemma, see [8, Lemma 1.18], there exists a sequence $\{\lambda_n\}_n$ with $\lim_{n\to\infty}\lambda_n = 0$ such that

$$\begin{aligned} f(x) &= \lim_{n\to\infty} f_{n,\lambda_n}(x) \\ &= \lim_{n\to\infty} \left(f_n\left(J^n_{\lambda_n}x\right) + \tfrac{1}{2\lambda_n}\,\mathrm{d}\big(x, J^n_{\lambda_n}x\big)^2 \right), \end{aligned} \tag{6.53}$$

where $J^n_\lambda = J_{\lambda f_n}$. Hence $f(x) \geq \limsup_{n\to\infty} f_n\left(J^n_{\lambda_n}x\right)$. We put $y_n := J^n_{\lambda_n}x$ and show that $y_n \to x$. Indeed, inserting Lemma 6.28(i) into (6.53), we have

$$f(x) \geq \limsup_{n\to\infty} \left(\left(\tfrac{1}{2\lambda_n} - r\right) \mathrm{d}(x, y_n)^2 - r \right)$$

and we can conclude that $y_n \to x$.

2. Let us show that $J^n_\lambda x \to J_\lambda x$. From the previous step, we know that there exists a sequence $y_n \to J_\lambda x$ with $\limsup_{n\to\infty} f_n(y_n) \leq f\left(J_\lambda x\right)$. Then we obtain

$$f_\lambda(x) = f\left(J_\lambda x\right) + \tfrac{1}{2\lambda} d\left(x, J_\lambda x\right)^2 \geq \limsup_{n\to\infty} \left(f_n(y_n) + \tfrac{1}{2\lambda} d\left(x, y_n\right)^2 \right)$$

and by (5.8) further

$$f_\lambda(x) \geq \limsup_{n\to\infty} \left(f_{n,\lambda}(x) + \tfrac{1}{2\lambda} d\left(J^n_\lambda x, y_n\right)^2 \right) = f_\lambda(x) + \limsup_{n\to\infty} \tfrac{1}{2\lambda} d\left(J^n_\lambda x, y_n\right)^2 .$$

Hence we conclude $J^n_\lambda x \to J_\lambda x$.

3. *(Liminf Inequality)* Let $x_n \overset{w}{\to} x$. We have to prove $\liminf_{n\to\infty} f_n(x_n) \geq f(x)$. By definition of the Moreau-Yosida envelope and (5.8) we have

$$\begin{aligned} f_n(x_n) &\geq f_n\left(J^n_\lambda x_n\right) + \tfrac{1}{2\lambda} d\left(x_n, J^n_\lambda x_n\right)^2 \\ &\geq f_n\left(J^n_\lambda x\right) + \tfrac{1}{2\lambda} d\left(x, J^n_\lambda x\right)^2 + \tfrac{1}{2\lambda} d\left(J^n_\lambda x_n, J^n_\lambda x\right)^2 \\ &\quad + \tfrac{1}{2\lambda} d\left(x_n, J^n_\lambda x_n\right)^2 - \tfrac{1}{2\lambda} d\left(x, J^n_\lambda x_n\right)^2 . \end{aligned}$$

By the nonpositive curvature inequality in (5.5) we obtain

$$f_n(x_n) \geq f_n\left(J^n_\lambda x\right) + \tfrac{1}{2\lambda} d\left(J^n_\lambda x, x_n\right)^2 - \tfrac{1}{2\lambda} d\left(x, x_n\right)^2 . \tag{6.54}$$

Let us show that $f_n\left(J^n_\lambda x\right)$ converges as $n \to \infty$. Consider

$$\begin{aligned} f_{n,\lambda}(x) &= f_n(J^n_\lambda x) + \tfrac{1}{2\lambda}\,\mathrm{d}(x, J^n_\lambda x)^2, \\ f_\lambda(x) &= f(J_\lambda x) \;\; + \tfrac{1}{2\lambda}\,\mathrm{d}(x, J_\lambda x)^2. \end{aligned}$$

By assumption we have $f_{n,\lambda}(x) \to f_\lambda(x)$, and by Step 2 also $J^n_\lambda x \to J_\lambda x$ as $n \to \infty$. This implies

$$f_n(J^n_\lambda x) \to f(J_\lambda x). \tag{6.55}$$

By the definition of the weak limit of $\{x_n\}_n$, for every subsequence $n_k \to \infty$, we have

$$\limsup_{k\to\infty} \mathrm{d}(J_\lambda x, x_{n_k})^2 \geq \limsup_{k\to\infty} \mathrm{d}(x, x_{n_k})^2. \tag{6.56}$$

By the triangle inequality we obtain $\mathrm{d}(J_\lambda^{n_k}x, x_{n_k}) \geq |\,\mathrm{d}(J_\lambda^{n_k}x, J_\lambda x) - \mathrm{d}(J_\lambda x, x_{n_k})|$. Using Step 2 and (6.56) results in

$$\begin{aligned}\limsup_{k\to\infty} \mathrm{d}(J_\lambda^{n_k}x, x_{n_k})^2 &\geq \limsup_{k\to\infty} \left(\mathrm{d}(J_\lambda^{n_k}x, J_\lambda x) - \mathrm{d}(J_\lambda x, x_{n_k})\right)^2 \\ &= \limsup_{k\to\infty} \mathrm{d}(J_\lambda x, x_{n_k})^2 \\ &\geq \limsup_{k\to\infty} \mathrm{d}(x, x_{n_k})^2.\end{aligned}$$

Rearranging, we get

$$\begin{aligned}0 &\leq \limsup_{k\to\infty} \mathrm{d}(J_\lambda^{n_k}x, x_{n_k})^2 + \liminf_{k\to\infty}(-\mathrm{d}(x, x_{n_k})^2) \\ &\leq \limsup_{k\to\infty} \left(\mathrm{d}(J_\lambda^{n_k}x, x_{n_k})^2 - \mathrm{d}(x, x_{n_k})^2\right)\end{aligned}$$

and, as the subsequence was arbitrary,

$$\liminf_{n\to\infty} \left(\mathrm{d}(J_\lambda^{n}x, x_{n})^2 - \mathrm{d}(x, x_{n})^2\right) \geq 0.$$

Returning to (6.54), the previous equation and (6.55) yield

$$\liminf_{n\to\infty} f_n(x_n) \geq f(J_\lambda x).$$

If $x \in \overline{\operatorname{dom} f}$, then from (6.52) we obtain

$$\liminf_{n\to\infty} f_n(x_n) \geq f(x).$$

For $x \notin \overline{\operatorname{dom} f}$ we can repeat the above conclusions for the finite continuous convex functions $g_n := f_{n,\mu}$ and $g = f_\mu$ for some fixed $\mu > 0$ instead of f_n and f. Note that the assumptions are fulfilled by the semigroup property of the Moreau-Yosida envelopes. Finally we let $\mu \to +0$ and invoke (5.7). This concludes the proof. □

Recall that a sequence of convex closed sets $C_n \subset \mathcal{H}$ converges to a convex closed set $C \subset \mathcal{H}$ in the sense of Frolík-Wijsman if the respective distance functions converge pointwise; that is, if $\mathrm{d}(x, C_n) \to \mathrm{d}(x, C)$ for each $x \in \mathcal{H}$. This concept originated in [47, 107]. On the other hand, a sequence of convex closed sets $C_n \subset \mathcal{H}$ converges to a convex closed set $C \subset \mathcal{H}$ in the sense of Mosco if the indicator functions ι_{C_n} converge in the sense of Mosco to the indicator function ι_C. The following is a direct consequence of our main result.

Corollary 6.32 (Frolík-Wijsman convergence) *A sequence of convex closed sets $C_n \subset \mathcal{H}$ converges to a convex closed set $C \subset \mathcal{H}$ in the sense of Frolík-Wijsman if and only if it converges to C in the sense of Mosco.*

Proof. Observe that the Moreau-Yosida envelope of ι_C with $\lambda = \frac{1}{2}$ is precisely the distance function squared $\mathrm{d}(\cdot, C)^2$ and apply Theorems 6.25 and 6.31. □

Bibliography

[1] M. S. C. Almeida and M. Figueiredo. Deconvolving images with unknown boundaries using the alternating direction method of multipliers. *IEEE Transactions on Image Processing*, 22(8):3074–3086, 2013.

[2] H. W. Alt. *Lineare Funktionalanalysis*. Springer, Berlin, 6th edition, 2012.

[3] J. Angulo and S. Velasco-Forero. *Geometric Theory of Information*, chapter Morphological Processing of Univariate Gaussian Distribution-Valued Images Based on Poincaré Upper-Half Plane Representation, pages 331–366. Springer, 2014.

[4] S. Arora, R. Ge, R. Kannan and A. Moitra. Computing a nonnegative matrix factorization – provably. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC '12, pages 145–162, New York, NY, 2012.

[5] S. Arora, R. Ge and A. Moitra. Learning topic models – going beyond SVD. In *IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1–10. IEEE, 2012.

[6] M. Asbach, D. Mauruschat and B. Plinke. Understanding multi-spectral images of wood particles with matrix factorization. In *Proceedings of the Optical Characterization of Materials (OCM) Conference 2013*, pages 191–202. KIT Scientific Publishing, 2013.

[7] H. Attouch. Familles d'opérateurs maximaux monotones et mesurabilité. *Annali di Matematica Pura ed Applicata (4)*, 120:35–111, 1979.

[8] H. Attouch. *Variational Convergence for Functions and Operators*. Pitman, Boston, MA, 1984.

[9] G. Aubert and P. Kornprobst. *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations*, volume 147. Springer, New York, 2002.

[10] M. Bačák. The proximal point algorithm in metric spaces. *Israel Journal of Mathematics*, 194(2):689–701, 2013.

[11] M. Bačák. *Convex Analysis and Optimization in Hadamard Spaces*, volume 22 of *De Gruyter Series in Nonlinear Analysis and Applications*. De Gruyter, Berlin, 2014.

[12] M. Bačák. Computing medians and means in Hadamard spaces. *SIAM J. Optim.*, 24(3):1542–1566, 2014.

[13] M. Bačák. Convergence of nonlinear semigroups under nonpositive curvature. *Trans. Amer. Math. Soc.*, 367(6):3929–3953, 2015.

[14] M. Bačák and L. Kovalev. Lipschitz retractions in Hadamard spaces via gradient flow semigroups. *Preprint arXiv:1603.01836*, 2016.

[15] M. Bačák, M. J. Montag and G. Steidl. Convergence of functions and their Moreau-Yosida envelopes on Hadamard spaces. *Preprint arXiv: 1604.08047*, 2016.

[16] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, New York, 1st edition, 2011.

[17] R. Bergmann, J. Persch and G. Steidl. Parallel Douglas–Rachford algorithm for restoring images with values in symmetric Hadamard spaces. *SIAM Journal on Imaging Sciences*, 9(3):901–937, 2016.

[18] R. J. Berman, T. Darvas and C. H. Lu. Convexity of the extended K-energy and the large time behaviour of the weak Calabi flow. *Preprint arXiv:1510.01260*, 2015.

[19] M. W. Berry, M. Browne, A. N. Langville, V. P. Pauca and R. J. Plemmons. Algorithms and applications for approximate nonnegative matrix factorization. *Computational Statistics & Data Analysis*, 52(1):155–173, 2007.

[20] M. Bertalmio, A. Bertozzi and G. Sapiro. Navier-Stokes, fluid dynamics, and image and video inpainting. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 355–362, 2001.

[21] D. P. Bertsekas, A. Nedić and A. E. Ozdaglar. *Convex analysis and optimization*. Athena Scientific, Belmont, MA, 2003.

[22] J. Bioucas-Dias and J. Nascimento. Hyperspectral subspace identification. *IEEE Transactions on Geoscience and Remote Sensing*, 46(8):2435–2445, 2008.

[23] J. Bioucas-Dias, A. Plaza, N. Dobigeon, M. Parente, Q. Du, P. Gader and J. Chanussot. Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 5(2):354–379, 2012.

[24] J. Bolte, S. Sabach and M. Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1):459–494, 2014.

[25] S. Boyd, N. Parikh, E. Chu, B. Peleato and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

[26] A. Braides. Γ*-convergence for Beginners*. Oxford University Press, Oxford, 2002.

[27] D. Burago, Y. Burago and S. Ivanov. *A course in metric geometry*, volume 33. American Mathematical Society, Providence, RI, 2001.

[28] M. Burger, A. Sawatzky and G. Steidl. First order algorithms in variational image processing. In *Operator Splittings and Alternating Direction Methods*. Springer, 2017.

[29] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1): 120–145, 2011.

[30] T. F. Chan and J. Shen. *Image processing and analysis: variational, PDE, wavelet, and stochastic methods*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2005.

[31] A. Chen. The inpainting of hyperspectral images: A survey and adaptation to hyperspectral data. In *Proceedings SPIE, Image and Signal Processing for Remote Sensing XVIII*, volume 8537, page 85371K, 2012.

[32] M. Chu, F. Diele, R. Plemmons and S. Ragni. Optimality, computation, and interpretation of nonnegative matrix factorizations. *Preprint*, 2004.

[33] P. A. Cook, Y. Bai, S. Nedjati-Gilani, K. K. Seunarine, M. G. Hall, G. J. Parker and D. C. Alexander. Camino: Open-source diffusion-MRI reconstruction and processing. In *Proceedings of the International Society for Magnetic Resonance in Medicine 14*, page 2759, Seattle, WA, 2006.

[34] S. I. Costa, S. A. Santos and J. E. Strapasson. Fisher information distance: A geometrical reading. *Discrete Applied Mathematics*, 197:59–69, 2015.

[35] T. M. Cover and J. A. Thomas. *Elements of information theory*. Wiley, New York, 2nd edition, 2006.

[36] J. X. Cruz Neto, I. D. Melo and P. A. Sousa. Non-existence of strictly monotone vector fields on certain Riemannian manifolds. *Acta Mathematica Hungarica*, 146 (1):240–246, 2015.

[37] G. Dal Maso. *An Introduction to* Γ*-Convergence*. Progress in Nonlinear Differential Equations and Their Applications 8. Birkhäuser, Boston, MA, 1993.

[38] I. Daubechies, M. Defrise and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57(11):1413–1457, 2004.

[39] M. Defrise, C. Vanhove and X. Liu. An algorithm for total variation regularization in high-dimensional linear problems. *Inverse Problems*, 27(6):065002, 2011.

[40] K. Degraux, V. Cambareri, L. Jacques, B. Geelen, C. Blanch and G. Lafruit. Generalized inpainting method for hyperspectral image acquisition. In *IEEE International Conference on Image Processing*, pages 315–319, 2015. Extended preprint arXiv:1502.01853.

[41] L. Delchambre. Weighted principal component analysis: a weighted covariance eigendecomposition approach. *Monthly Notices of the Royal Astronomical Society*, 446(4):3545–3555, 2015.

[42] P. Eggermont, G. Herman and A. Lent. Iterative algorithms for large partitioned linear systems, with applications to image reconstruction. *Linear Algebra and its Applications*, 40:37–67, 1981.

[43] E. Esser, M. Möller, S. Osher, G. Sapiro and J. Xin. A convex model for nonnegative matrix factorization and dimensionality reduction on physical space. *IEEE Transactions on Image Processing*, 21(7):3239–3252, 2012.

[44] O. P. Ferreira and P. R. Oliveira. Proximal point algorithm on Riemannian manifolds. *Optimization*, 51(2):257–270, 2002.

[45] L. Finesso and P. Spreij. Nonnegative matrix factorization and I-divergence alternating minimization. *Linear Algebra and its Applications*, 416(2):270–287, 2006.

[46] D. A. Fish, J. G. Walker, A. M. Brinicombe and E. R. Pike. Blind deconvolution by means of the Richardson–Lucy algorithm. *Journal of Optical Society of America A*, 12(1):58–65, 1995.

[47] Z. Frolík. Concerning topological convergence of sets. *Czechoslovak Mathematical Journal*, 10(85):168–180, 1960.

[48] N. Gillis and S. A. Vavasis. Fast and robust recursive algorithms for separable nonnegative matrix factorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(4):698–714, 2014.

[49] N. Gillis and S. A. Vavasis. Semidefinite programming based preconditioning for more robust near-separable nonnegative matrix factorization. *SIAM Journal on Optimization*, 25(1):677–698, 2015.

[50] T. Goldstein, M. Li, X. Yuan, E. Esser and R. Baraniuk. Adaptive primal-dual hybrid gradient methods for saddle-point problems. *Preprint arXiv: 1305.0546*, 2013.

[51] E. F. Gonzalez and Y. Zhang. Accelerating the Lee-Seung algorithm for nonnegative matrix factorization. Technical report, Department of Computational and Applied Mathematics, Rice University, Houston, TX, 2005.

[52] L. Grippo and M. Sciandrone. On the convergence of the block nonlinear Gauss–Seidel method under convex constraints. *Operations Research Letters*, 26 (3):127–136, 2000.

[53] M. J. Gursky and J. Streets. A formal riemannian structure on conformal classes and the inverse gauss curvature flow. *Preprint arXiv: 1507.04781*, 2015.

[54] M. Hanke-Bourgeois. *Grundlagen der numerischen Mathematik und des wissenschaftlichen Rechnens*. Vieweg+Teubner, Wiesbaden, 3rd edition, 2009.

[55] G. Harris and C. Martin. Shorter notes: The roots of a polynomial vary continuously as a function of the coefficients. *Proceedings of the American Mathematical Society*, 100(2):390–392, 1987.

[56] T. J. Holmes. Blind deconvolution of quantum-limited incoherent imagery: maximum-likelihood approach. *Journal of Optical Society of America A*, 9(7):1052–1061, 1992.

[57] A. Hyvärinen, J. Karhunen and E. Oja. *Independent component analysis*. John Wiley & Sons, 2004.

[58] A. Hyvärinen and E. Oja. Independent component analysis: algorithms and applications. *Neural Networks*, 13(4–5):411–430, 2000.

[59] M.-D. Iordache, J. M. Bioucas-Dias and A. Plaza. Total variation spatial regularization for sparse hyperspectral unmixing. *IEEE Transactions on Geoscience and Remote Sensing*, 50(11):4484–4502, 2012.

[60] B. Iversen. *Hyperbolic geometry*. Cambridge University Press, 1 edition, 1992.

[61] J. Jost. Nonlinear Dirichlet forms. In *New directions in Dirichlet forms*, volume 8 of *AMS/IP Studies in Advanced Mathematics*, pages 1–47. American Mathematical Society, Providence, RI, 1998.

[62] M. Karas and R. Krüger. Ion formation in MALDI: The cluster ionization mechanism. *Chemical Reviews*, 103(2):427–440, 2003.

[63] M. Karas, D. Bachmann, U. Bahr and F. Hillenkamp. Matrix-assisted ultraviolet laser desorption of non-volatile compounds. *International Journal of Mass Spectrometry and Ion Processes*, 78:53–68, 1987.

[64] S. L. Keeling and K. Kunisch. Robust l1 approaches to computing the geometric median and principal and independent components. *Journal of Mathematical Imaging and Vision*, 56(1):99–124, 2016.

[65] H. Kim and H. Park. Nonnegative matrix factorization based on alternating nonnegativity constrained least squares and active set method. *SIAM Journal on Matrix Analysis and Applications*, 30(2):713–730, 2008.

[66] K. Kuwae and T. Shioya. Variational convergence over metric spaces. *Trans. Amer. Math. Soc.*, 360(1):35–75, 2008.

[67] D. Landgrebe. Hyperspectral image data analysis. *IEEE Signal Processing Magazine*, 19(1):17–28, 2002.

[68] K. Lange, D. R. Hunter and I. Yang. Optimization transfer using surrogate objective functions. *Journal of Computational and Graphical Statistics*, 9(1):1–20, 2000.

[69] L. Lecharlier and C. De Mol. Regularized blind deconvolution with Poisson data. *Journal of Physics: Conference Series*, 464(1):012003, 2013.

[70] D. D. Lee and H. S. Seung. Algorithms for non-negative matrix factorization. In *Advances in Neural Information Processing Systems 13*, pages 556–562. MIT Press, 2001.

[71] J. M. Lee. *Riemannian Manifolds. An Introduction to Curvature*. Springer, New York, 1997.

[72] R. Leitner, H. Mairer and A. Kercek. Real-time classification of polymers with NIR spectral imaging and blob analysis. *Real-Time Imaging*, 9(4):245–251, 2003.

[73] C. Li, B. S. Mordukhovich, J. Wang and J.-C. Yao. Weak sharp minima on Riemannian manifolds. *SIAM Journal on Optimization*, 21(4):1523–1560, 2011.

[74] C. J. Lin. On the convergence of multiplicative update algorithms for nonnegative matrix factorization. *IEEE Transactions on Neural Networks*, 18(6):1589–1596, 2007.

[75] L. B. Lucy. An iterative technique for the rectification of observed distributions. *The Astronomical Journal*, 79:745, 1974.

[76] J. Marshall and J. Oberwinkler. Ultraviolet vision: The colourful world of the mantis shrimp. *Nature*, 401(6756):873–874, 1999.

[77] R. Méndez-Rial, M. Calvino-Cancela and J. Martín-Herrero. Accurate implementation of anisotropic diffusion in the hypercube. *Geoscience and Remote Sensing Letters, IEEE*, 7(4):870–874, 2010.

[78] R. Méndez-Rial, M. Calvino-Cancela and J. Martín-Herrero. Anisotropic inpainting of the hypercube. *Geoscience and Remote Sensing Letters, IEEE*, 9(2):214–218, 2012.

[79] M. Moakher and P. G. Batchelor. Symmetric positive-definite matrices: From geometry to applications and visualization. In *Visualization and Processing of Tensor Fields*, pages 285–298. Springer, Berlin, 2006.

[80] G. Moerkotte, M. J. Montag, A. Repetti and G. Steidl. Proximal operator of quotient functions with application to a feasibility problem in query optimization. *Journal of Computational and Applied Mathematics*, 285:243–255, 2015.

[81] M. J. Montag and H. Stephani. Hyperspectral unmixing from incomplete and noisy data. *Journal of Imaging*, 2(1):7, 2016.

[82] J. Moreau. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société Mathématique de France*, 93:273–299, 1965.

[83] U. Mosco. Convergence of convex sets and of solutions of variational inequalities. *Advances in Mathematics*, 3(4):510–585, 1969.

[84] J. P. Oliveira, M. A. T. Figueiredo and J. M. Bioucas-Dias. Parametric blur estimation for blind restoration of natural images: Linear motion and out-of-focus. *IEEE Transactions on Image Processing*, 23(1):466–477, 2014.

[85] P. Paclík, R. Leitner and R. P. Duin. A study on design of object sorting algorithms in the industrial application using hyperspectral imaging. *Journal of Real-Time Image Processing*, 1(2):101–108, 2006.

[86] E. A. Papa Quiroz. An extension of the proximal point algorithm with Bregman distances on Hadamard manifolds. *Journal of Global Optimization*, 56(1):43–59, 2013.

[87] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231, 2013.

[88] V. P. Pauca, J. Piper and R. J. Plemmons. Nonnegative matrix factorization for spectral data analysis. *Linear Algebra and its Applications*, 416(1):29–47, 2006.

[89] J. F. Pinto da Costa, H. Alonso and L. Roque. A weighted principal component analysis and its application to gene expression data. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 8(1):246–252, 2011.

[90] T. Pock, A. Chambolle, D. Cremers and H. Bischof. A convex relaxation approach for computing minimal partitions. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 810–817, 2009.

[91] B. Rasti, M. O. Ulfarsson and J. R. Sveinsson. Hyperspectral subspace identification using SURE. *IEEE Geoscience and Remote Sensing Letters*, 12(12):2481–2485, 2015.

[92] B. Recht, C. Re, J. Tropp and V. Bittorf. Factoring nonnegative matrices with linear programs. In *Advances in Neural Information Processing Systems 25*, pages 1214–1222. Curran Associates, 2012.

[93] W. H. Richardson. Bayesian-based iterative method of image restoration. *Journal of Optical Society of America A*, 62(1):55–59, 1972.

[94] L. I. Rudin, S. Osher and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.

[95] J. J. Settle and N. A. Drake. Linear mixing and the estimation of ground cover proportions. *International Journal of Remote Sensing*, 14(6):1159–1177, 1993.

[96] S. Setzer, G. Steidl and T. Teuber. Deblurring Poissonian images by split Bregman techniques. *Journal of Visual Communication and Image Representation*, 21:193–199, 2010.

[97] B. Shafei and G. Steidl. Segmentation of images with separating layers by fuzzy c-means and convex optimization. *Journal of Visual Communication and Image Representation*, 23(4):611–621, 2012.

[98] Y. Sonntag. *Convergence au sens de Mosco; théorie et applications à l'approximation des solutions d'inéquations. Thèse d'Etat.* PhD thesis, Université de Provence, Marseille, 1982.

[99] G. Steidl and T. Teuber. Removing multiplicative noise by Douglas-Rachford splitting methods. *Journal of Mathematical Imaging and Vision*, 36(2):168–184, 2010.

[100] H. Stephani. *Automatic Segmentation and Clustering of Spectral Terahertz Data.* PhD thesis, Technische Universität Kaiserslautern and Johannes Kepler Universität Linz, 2012.

[101] J. Streets. The consistency and convergence of K-energy minimizing movements. *Transactions of the American Mathematical Society*, 368(7):5075–5091, 2016.

[102] K.-T. Sturm. Nonlinear Markov operators associated with symmetric Markov kernels and energy minimizing maps between singular spaces. *Calculus of Variations and Partial Differential Equations*, 12(4):317–357, 2001.

[103] J. Theiler and B. Wohlberg. Detection of unknown gas-phase chemical plumes in hyperspectral imagery. In *Proceedings SPIE, Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XIX*, volume 8743, page 874315, 2013.

[104] M. Valiollahzadeh and W. Yin. Hyperspectral data reconstruction combining spatial and spectral sparsity. Technical Report TR10-29, Rice University, Houston, TX, 2010.

[105] S. A. Vavasis. On the complexity of nonnegative matrix factorization. *SIAM Journal on Optimization*, 20(3):1364–1377, 2009.

[106] F. Wang, W. Cao and Z. Xu. Convergence of multi-block Bregman ADMM for nonconvex composite problems. *Preprint arXiv: 1505.03063*, 2015.

[107] R. A. Wijsman. Convergence of sequences of convex sets, cones and functions. II. *Transactions of the American Mathematical Society*, 123:32–45, 1966.

[108] S.-T. Yau. Non-existence of continuous convex functions on certain Riemannian manifolds. *Mathematische Annalen*, 207(4):269–270, 1974.

[109] W. I. Zangwill. *Nonlinear Programming: A Unified Approach*. Prentice-Hall, 1969.

Scientific Career

6/2005	*University entrance qualification (Abitur)* Peter-Wust-Gymnasium Wittlich
10/2005 – 7/2008	*Undergraduate studies in Mathematics* King's College, University of Cambridge, United Kingdom
7/2008	*Bachelor of Arts in Mathematics*
10/2008 – 6/2013	*Graduate studies in Mathematics* University of Bonn
6/2013	*Master of Science in Mathematics*
8/2013 – today	*PhD scholarship, Fraunhofer ITWM Kaiserslautern* Department of Mathematics, University of Kaiserslautern

Wissenschaftlicher Werdegang

6/2005	*Abitur* Peter-Wust-Gymnasium Wittlich
10/2005 – 7/2008	*Studium der Mathematik* King's College, Universität Cambridge, Vereinigtes Königreich
7/2008	*Bachelor of Arts in Mathematik*
10/2008 – 6/2013	*Masterstudiengang Mathematik* Rheinische Friedrich-Wilhelms-Universität Bonn
6/2013	*Master of Science in Mathematik*
8/2013 – heute	*Promotionsstipendium, Fraunhofer ITWM Kaiserslautern* Fachbereich Mathematik, Technische Universität Kaiserslautern