



On the Atoms of Robustness

Robust Matrix Decomposition
for Spectral Imaging

Cover painting by Magda De Bisschop

On the Atoms of Robustness

Robust Matrix Decomposition for Spectral Imaging

Roger A.R. Moens
Delft University of Technology

Under supervision of
Prof. Dr. Brandl, Dr. Ing. Van de Plas

This thesis was typeset using \LaTeX .

For the degree of

Master of Science in Systems and Control

Master of Science in Aerospace Engineering

at the Delft University of Technology



Delft Center for
Systems and Control

 **TU**Delft

The TU Delft logo features a stylized flame or torch icon above the letters 'TU', followed by the word 'Delft' in a bold, sans-serif font.

*En de lucht wordt donker
en de maan verschijnt
maar ik word niet dronken
want er is iets dat harder toeslaat
en dat maar niet verdwijnt.*

*En de maan klimt hoger
en m'n wereld kraakt
en ik word dover
ik hoor m'n eigen woorden vallen
maar niemand wordt geraakt
en ik zoek naar het woord dat alles open maakt.*

- The Scene, Open

Abstract

Modern imaging modalities across many application domains increasingly acquire a large number of very high-dimensional measurements, commonly collecting hundreds to millions of variables per spatial resolution element. That high-dimensional nature can severely challenge traditional (often Euclidean distance based) approaches to noise and dimensionality reduction. Furthermore, statistical analysis of such data is often hampered by the curse of dimensionality, concomitant with large numbers of pixels and channels, by the growing abundance of low signal-to-noise ratio (SNR) measurements, and by the detrimental effects of noise accumulation^a. It is therefore necessary to find efficient means of reducing high-dimensional (and large data size) measurement sets to a lower-dimensional representation while incurring minimal information loss. Addressing this challenge is essential (a) to enable processing of the massively multivariate measurement sets acquired by several promising imaging technologies today (commonly hundreds of gigabytes to terabytes per experiment), (b) to avoid that computational analysis becomes a bottleneck for the development of new instrumental capabilities, with hardware setups theoretically capable of yielding terabyte to petabyte imaging but currently considered impractical, and (c) to enable archiving massively multivariate measurement sets, often required to be stored for several years and containing information suitable for additional science projects. This thesis focuses on this challenge specifically. First, it explores structured and regularized matrix decomposition methods based on the ℓ_1 -norm, e.g. building on principal component pursuit (PCP)^b, to address this challenge. Second, it delivers custom implementations of these methods. Third, it develops and implements an application-driven framework for automatic setting of hyperparameters. Finally, it applies and compares these methods using several spectral imaging case studies spanning both very small scales, in molecular imaging of organic tissue, as well as very large scales, in the spectroscopic imaging of Outer space.

^aFan, J., Han, F. & Liu, H. Challenges of big data analysis. *National science review* **1**, 293–314 (2014)

^bCandès, E. J., Li, X., Ma, Y. & Wright, J. Robust principal component analysis? *Journal of the ACM (JACM)* **58**, 1–37 (2011)

Methodologically, we study the quantitative and qualitative behavior of such methods under different inexact recovery conditions to develop a notion of potential breakdown phenomena. Secondly, we propose a novel application-driven framework for hyperparameter optimization, capable of taking domain- and application-specific knowledge into account, and formulate an algorithm for solving this optimization with time complexity not larger than $\mathcal{O}(qmnr)$ on a q -fold sampled grid for a rank- r measurement matrix $M \in \mathbb{R}^{m \times n}$. Thirdly, we implement different resource-restricted PCP versions achieving 5 to 10 fold improvement in running time with respect to original in-memory implementations. And finally, we demonstrate, to our knowledge, a first, out-of-memory implementation able to process measurement matrices that would otherwise require one or more terabytes of memory.

Application-wise, we advance the capabilities of imaging mass spectrometry (IMS) and its use for exploring the molecular content of biological tissue, with different strategies: (1) the application of PCP on MALDI-(tims)TOF measurements, avoiding sparse noise skewing of the results; (2) the use of our framework to exceed traditional PCP hyperparameter setting performance for specific application goals; (3) in-memory linear scaling of our PCP implementations for MALDI-TOF, with compression ratios exceeding 10^3 while providing summarizing spatial and spectral reconstruction scores; (4) a further improvement in compression ratio over PCA, achieving a 1.73-fold increase for the same spatial reconstruction score; and (5) a first-ever robust matrix decomposition of a MALDI-timsTOF measurement set that would normally require >1 TB memory. These developments will be essential in making full ion mobility (IM) IMS, a powerful generalization of MALDI-TOF IMS that collects two spectral dimensions per pixel (or 4-mode data tensors), and that can yield petabyte measurements using standard techniques, a practical reality.

We furthermore explore what these techniques can deliver for spectral imaging at very large scales. To accomplish this, we focus on integral field spectroscopy and especially the MUSE instrument at ESO's Very Large Telescope (VLT), which explores Outer space through the spectral signatures emitted and absorbed by atoms and ions. We showcase the usability of robust matrix decomposition methods for: (1) co-adding 54 Hubble Deep Field South (HDFS) exposures, while reducing instrumental and sky background residuals, and achieving 1.2 – 2.4 fold reduction in spectral variability in comparison to median, and mean-based approaches; (2) the first concurrent identification of spatial and spectral signatures of instrumental residuals in exposures of the Hubble Ultra-Deep Field (UDF-06) influencing average spectra up to 25% for subsequent subtraction, improving the detectability of faint structures, while promising a reduction in self-induced artifacts with respect to current running median solutions; (3) exploration purposes on spiral galaxies, such as AM2210-262 in the Piscis Austrinus constellation, for subsequent measurement of galaxy kinematics, e.g. the rotation curve for spiral galaxies, and (4) enabling the extraction of line emissions from continuum spectra in a subfield of the Hubble Deep Field South (HDFS)

while concurrently and consistently increasing the SNR of these line emissions 2-fold on median and maximally up to 4-fold with respect to the reduced data. These developments will be essential in further enabling the investigation of faint and complex astrophysical signatures, one of the scientific goals of the MUSE instrument^a. Secondly, our methodology will enable achieving higher SNR values with less exposures, subsequently increasing observation time for other studies. Finally, the presented achievable compression ratios are promising for future reduction pipelines and the archiving of scientific data from novel instruments, e.g. the Extremely Large Telescope (ELT).^b

^aBacon, R. et al. The MUSE second-generation VLT instrument. In *Ground-based and Airborne Instrumentation for Astronomy III*, 7735, 773508 (2010).

^bGilmozzi, R. & Spyromilio, J. The European Extremely Large Telescope (E-ELT). *The Messenger* 127, 3 (2007).

Contents

Abstract	ix
I Context	1
1 Introduction	2
1.1 Problem Statement	2
1.2 Goals and Objectives	3
1.3 Significance	5
1.4 Thesis Outline	5
2 A Mathematical Primer	8
2.1 Introduction	8
2.2 Basic Notations	9
2.3 Vectors and Vector Spaces	10
2.4 Matrices and its Algebra	14
2.5 Summary	23
3 Application Domains	25
3.1 Introduction	25
3.2 Imaging Mass Spectrometry: MALDI	25
3.3 Integral field Spectroscopy: MUSE	31
3.4 Summary	34
4 Literature Review	37
4.1 Introduction	37
4.2 Extended Linear Mixture Model	43
4.3 Methods	45
4.4 Summary	80

II	Research	93
5	A Comparative Study under Inexact Recovery Conditions	94
5.1	Introduction	94
5.2	Experimental Design	95
5.3	Quantitative Performance Analysis	103
5.4	Qualitative Performance Analysis	109
5.5	Conclusions	111
6	An Application-Driven Framework for Parameter Optimization	114
6.1	Introduction	114
6.2	Experimental Design	116
6.3	Case Studies	124
6.4	Conclusions	128
7	Acceleration and Scaling of PCP	132
7.1	Introduction	132
7.2	Experimental Design	133
7.3	Acceleration Performance Analysis	142
7.4	Scaling Performance Analysis	146
7.5	Conclusions	150
8	Case Studies I: Imaging Mass Spectrometry	156
8.1	Introduction	156
8.2	Experimental Design	158
8.3	Case Study A: Parameter Setting	164
8.4	Case Study B: A Comparative Study	173
8.5	Case Study C: Scale-Dependent Comparison	180
8.6	Case Study D: Terabyte Scale	187
8.7	Conclusions	192
9	Case Studies II: Integral-Field Spectroscopy	195
9.1	Introduction	195
9.2	Experimental Design	196
9.3	Case Study A: Co-Addition	201
9.4	Case Study B: Artifact Identification	210
9.5	Case Study C: Component Evaluation	217
9.6	Case Study D: Line Emission Extraction	225
9.7	Conclusions	232
10	Conclusions	235
10.1	Methodology	235
10.2	Application Domains	239
10.3	Contributions	243
10.4	Recommendations	244

10.5 Collaborative Perspective	245
III Appendices	248
Appendices	249
A Additional Materials	249
A.1 A Comparative Study under Inexact Recovery Conditions	250
A.2 An Application-Driven Framework for Parameter Optimization	252
A.3 Acceleration and Scaling of PCP	254
A.4 Case Studies II: Intregal-Field Spectroscopy	258
B HannibalSpecter	262
C Oak Ridge National Lab: Research Proposal	263

Part I

Context

Chapter 1

Introduction

After this chapter you will be able to

- assess this research thesis' position within the broader spectrum of dimensionality reduction methods
- review our goals and objectives, and corresponding research questions
- understand the areas of impact for this thesis
- maneuver yourself through this thesis via the outline

1.1 Problem Statement

Today, imaging modalities used in astronomy and analytical chemistry are providing profound insights into the earliest times of our Universe¹, as well as into the molecules constituting our human tissue². Concomitantly, these modalities deliver massively multivariate measurement sets, running into the hundreds of gigabytes to terabytes and beyond, making direct human interpretation often impractical, if not impossible. Their high-dimensional nature induces challenges, such as the curse of dimensionality and the abundance of low signal-to-noise ratio (S/N) measurements, which ultimately tend to hamper efficient statistical analysis on such data sets³. Furthermore, innovative instrumental principles, such as those used in the METIS instrument⁴ for astronomy, and technological advancements, such as MALDI timsTOF⁵ for analytical chemistry, are pushing the limits of spatial and spectral resolutions, as well as instrumental sensitivity and detection coverage even further. The latter specifically emphasizes the importance of tackling the computational challenges related to this high-dimensional nature in order to keep facilitating scientific discovery.

In the recent past, simple approaches based on mean and median, and more advanced weighted schemes exploiting mathematical structure of systems, such as the Kalman filter⁶, have been successfully applied for noise reduction purposes. Furthermore, the use of matrix decomposition methodology has been proven valuable for both noise and dimensionality reduction through its ability to exploit signal structure on larger scales. One example of this type of methodology is the ever-popular principal component analysis (PCA). Unfortunately, high-dimensionality renders many popular Euclidean distance-based methods, such as PCA, sometimes ineffectual³. Concurrently, simple outlier removal mechanisms, often performed manually in the past, have become practically tedious, if not impossible due to the large number of measurements to process.

While these methodological problems have (partially) been resolved in recent research, e.g. a PCA approach robust against sparse outliers⁷, these solutions have not yet seeped through into many application domains. This thesis seeks to fulfill part of this effort.

1.2 Goals and Objectives

The aim of this thesis consists of the development of novel strategies to efficiently reduce high-dimensional (and large data size) spectral imaging measurement sets to a lower-dimensional representation while incurring minimal information loss. Addressing this challenge is (a) essential to enable processing of the massively multivariate measurement sets acquired by several promising imaging technologies today (commonly hundreds of gigabytes to terabytes per experiment), and (b) to avoid that computational analysis becomes a bottleneck for the development of new instrumental capabilities, with hardware setups theoretically capable of yielding terabyte to petabyte imaging but currently considered impractical. This thesis focuses on this challenge specifically, explores structured and regularized matrix decomposition methods based on the ℓ_1 -norm, e.g. building on principal component pursuit (PCP)⁷, to address it, and applies and compares these methods using several spectral imaging case studies spanning both very small scales, in molecular imaging of organic tissue, as well as very large scales, in the spectroscopic imaging of Outer space.

Our first objective consists of translating some of the theory, related to structured and regularized matrix decomposition methods, into practice. This includes providing insights into break down phenomena and the tuning of its corresponding hyperparameters. Finally, we review the acceleration and scalability of the developed strategies. Those findings, each answering a particular subquestion, will contribute to the formulation of an answer on research question A (see below).

Our second objective consists of applying and comparing these methods using several spectral imaging case studies spanning both very small scales, in molecular imaging of organic tissue, as well as very large scales, in the spectroscopic

imaging of Outer space. For these application domains we collaborate with the

A. Can matrix decomposition methods based on the ℓ_1 -robustness principles efficiently reduce massively multivariate and high-dimensional spectral imaging measurement matrices to a lower-dimensional representation while incurring minimal information loss?

A.1. How do these methods perform with respect to classical dimensionality reduction methods under inexact recovery conditions: (a) quantitatively, i.e. in terms of efficiency, compression ratio and information loss; (b) qualitatively, i.e. what break down phenomena appear in the atoms, dictionaries, and matrix terms?

A.2. Can a framework be built to efficiently tune the regularization parameters of the above-mentioned methods, but customized for a specific application?

A.3. Can a memory and compute resources restricted version of such approaches be developed, and how does it perform with respect to the original batch version that assumes the entire data set can be held in memory?

B. How can the application-driven methodology we have developed, be used to address application domain-specific challenges?

B.1. In molecular imaging, and specifically imaging mass spectrometry high-dimensionality and *massive data set footprints* are the most pressing bottleneck. How can we apply the developed methodology to measurement matrices of MALDI TOF and MALDI timsTOF nature, respectively of mouse kidney and mouse pup origin, specifically for *dimensionality reduction and compression purposes*?

B.2. In astronomical imaging, and specifically integral field spectroscopy *low signal-to-noise ratios and instrumental artifacts* are the most pressing bottleneck. How can we apply the developed methodology to measurement matrices of the MUSE instrument picturing a spiral galaxy (AM2210-262), the Hubble Ultra Deep-Field (UDF-06), and different exposures of the Hubble Deep Field South (HDFS), *for noise reduction purposes*?

Mass Spectrometry Research Center (MSRC), Vanderbilt University, in Nashville, Tennessee (USA) for imaging mass spectrometry, and the Leibniz Institute (AIP), in Potsdam (GER) for integral field spectroscopy. Those findings, each answer-

ing a particular subquestion, will contribute to the formulation of an answer on research question B (see below).

1.3 Significance

The significance of this thesis primarily resides in the application domains and what these methods enable there.

Imaging Mass Spectrometry For imaging mass spectrometry, these developments will be essential in making full ion mobility (IM) IMS, a powerful generalization of MALDI-TOF IMS that collects two spectral dimensions per pixel (or 4-mode data tensors), and that can yield petabyte measurements using standard techniques, a practical reality.

Integral field spectroscopy For integral field spectroscopy, these developments will be essential in further enabling the investigation of faint and complex astrophysical signatures, one of the scientific goals of the MUSE instrument⁸. Secondly, our methodology will enable achieving higher SNR values with less exposures, subsequently increasing observation time for other studies. Finally, the presented achievable compression ratios are promising for future reduction pipelines and the archiving of scientific data from novel instruments, e.g. the Extremely Large Telescope (ELT)⁹.

1.4 Thesis Outline

Our thesis is structured in two parts: (a) a contextual part, and (b) a research part. For students and researchers without extensive background knowledge, but interested in the methodological aspects, we advise perusing the mathematical primer in chapter 2 and the literature review in chapter 4. For researchers interested in the methodological advancements of this thesis we advise chapter 5, chapter 6, and chapter 7. If the reader's interest lies in the practical applications in the domain of astronomical imaging, we advise chapter 9, and for applications in the domain of imaging mass spectrometry, we advise chapter 8. An short introduction of both application domains is provided in chapter 3.

Part I - Context: This part mainly focuses on the context of our research.

Chapter 1: Introduction.

Chapter 2: We describe the mathematical basis, useful for the methodological chapters of the thesis.

Chapter 3: We describe the application domains in which we apply the developed methodology. This chapter is useful as a primer for readers of chapter 8 and/or 9.

Chapter 4: Introduces the available literature, and shows a mismatch between methodology and practical application in the aforementioned domains. This chapter gives a methodological overview of matrix decomposition approaches and robust versions of those.

Part II - Research: This part addresses the research questions stated above, with each chapter addressing one subquestion.

Chapter 5: We translate the methodology to practical use and inspect its behavior in imperfect situations, as is the case in practical settings. This chapter addresses subquestion A.1.

Chapter 6: We address the parameter setting problem and provide a framework to set parameters optimally according to the needs of a specific application. This chapter addresses subquestion A.2.

Chapter 7: We address acceleration and scaling problems, i.e. we seek to make the underlying algorithms faster and practically scalable to the terabyte data size range. This chapter addresses subquestion A.3.

Chapter 8: We provide case studies for imaging mass spectrometry. This chapter addresses subquestion B.1.

Chapter 9: We provide case studies for integral field spectroscopy. This chapter addresses subquestion B.2.

Chapter 10: Conclusions.

Bibliography

- [1] Wisotzki, L. *et al.* Nearly all the sky is covered by lyman- α emission around high-redshift galaxies. *Nature* **562**, 229–232 (2018).
- [2] Consortium, H. *et al.* The human body at cellular resolution: the nih human biomolecular atlas program. *Nature* **574**, 187 (2019).
- [3] Fan, J., Han, F. & Liu, H. Challenges of big data analysis. *National science review* **1**, 293–314 (2014).
- [4] Brandl, B. R. *et al.* Instrument concept and science case for the mid-ir e-elt imager and spectrograph metis. In *Ground-based and Airborne Instrumentation for Astronomy III*, vol. 7735, 77352G (International Society for Optics and Photonics, 2010).
- [5] Spraggins, J. M. *et al.* High-performance molecular imaging with maldi trapped ion-mobility time-of-flight (timstof) mass spectrometry. *Analytical chemistry* **91**, 14552–14560 (2019).
- [6] Kalman, R. E. A new approach to linear filtering and prediction problems (1960).
- [7] Candès, E. J., Li, X., Ma, Y. & Wright, J. Robust principal component analysis? *Journal of the ACM (JACM)* **58**, 1–37 (2011).
- [8] Bacon, R. *et al.* The muse second-generation vlt instrument. In *Ground-based and Airborne Instrumentation for Astronomy III*, vol. 7735, 773508 (International Society for Optics and Photonics, 2010).
- [9] Gilmozzi, R. & Spyromilio, J. The european extremely large telescope (e-elt). *The Messenger* **127**, 3 (2007).

Chapter 2

A Mathematical Primer

After this chapter you will be able to

- understand the notation and basics of scalars, vectors, matrices used elsewhere in this thesis;
- follow the basics of matrix decomposition;
- understand the essentials of convex optimization;

2.1 Introduction

In this chapter we review basic notation and provide the necessary notions for our literature study and research. Those basics consists of a short recap on linear algebra and convex optimization. In a first section, we review vectors and matrices, but also norms and matrix decomposition. The second section is focused on convex optimization, handling convex relaxations, convex-cardinality problems, and shortly touches upon the low-rank recovery problem. Finally, we provide, where possible, the geometric interpretation of the concepts to help guide intuition.

“It is not easy to visualize a 10-dimensional plane in 11-dimensional space. It is harder to see eleven of those planes intersecting at a single point in that space – but somehow it is almost possible. With three planes in three dimensions it can certainly be done. Then linear algebra moves the problem into four dimensions, or eleven dimensions, where the intuition has to imagine the geometry (and gets it right)” - Strang¹

2.2 Basic Notations

Throughout this thesis, we will maintain the same notation as the book *Convex Optimization*², with some small differences. The set of real numbers is denoted by \mathbb{R} , the set of real n -vectors is denoted by \mathbb{R}^n , and the set of real m by n matrices is denoted by $\mathbb{R}^{m \times n}$. We use the following notations for a

- scalar: $x \in \mathbb{R}$,
- vector: $\mathbf{x} \in \mathbb{R}^n$,
- matrix: $A \in \mathbb{R}^{m \times n}$.

A vector \mathbf{x} is considered as a column vector and defined as

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}. \quad (2.1)$$

A row vector will be denoted by \mathbf{x}^\top . A vector is always boldfaced, and its entries, i.e. scalars, are denoted by x_i . Finally, a matrix is represented as

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix} = [\mathbf{a}_1 \quad \mathbf{a}_2 \quad \cdots \quad \mathbf{a}_n]. \quad (2.2)$$

Similar to the vectors, a matrix will be strictly denoted by capital letters and its entries, i.e. scalars, are denoted by a_{ij} . Note that for our research we are only focused on the domain of real numbers.

Diagonal The diagonal of a square matrix $A \in \mathbb{R}^{n \times n}$ can be represented as:

$$\mathbf{d} = \mathbf{diag}(A), \quad (2.3)$$

where $\mathbf{d} \in \mathbb{R}^n$ and contains all diagonal elements of A , i.e. a_{ii} . However, the same notation is used to denote the diagonal matrix formed by the vector elements on the diagonal:

$$B = \mathbf{diag}(\mathbf{d}). \quad (2.4)$$

Note that $A = B$ only if A was already a diagonal matrix, otherwise $A \neq B$.

Trace The trace of a matrix is defined as the sum of its diagonal:

$$\mathbf{tr}(A) = \sum_{i=1}^n a_{ii}, \quad (2.5)$$

with $A \in \mathbb{R}^{n \times n}$.

Vectorization The vectorization of a matrix is a linear transformation, converting a matrix into a column vector

$$\mathbf{vec}(A) = \begin{bmatrix} a_{11} \\ \vdots \\ a_{m1} \\ a_{12} \\ \cdots \\ a_{m2} \\ \cdots \\ a_{1n} \\ \cdots \\ a_{mn} \end{bmatrix}, \quad (2.6)$$

with $A \in \mathbb{R}^{m \times n}$.

2.3 Vectors and Vector Spaces

We maintain the same definitions as the book *Linear algebra and its applications*¹. Hence, a vector space is defined as the space \mathbb{R}^n , consisting of all column vectors with n components and with the components being real numbers. Within these vector spaces two operations are possible:

1. addition of two vectors;
2. multiplication of vectors by scalars.

Secondly, a *basis* for a vector space is a set of vectors and has two properties:

1. it is linearly independent;
2. it spans the space.

Finally, a subspace of a vector space is a nonempty subset that satisfies two requirements:

1. the addition of any vectors \mathbf{x} and \mathbf{y} in the subspace, lies in the subspace;
2. the multiplication of any vector \mathbf{x} in the subspace and any scalar α , lies in the subspace.

2.3.1 Multiplication

Two types of multiplication are of interest here: the inner and outer product.

Inner Product For two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, their inner product is defined as

$$\mathbf{x}^\top \mathbf{y} = x_1 y_1 + x_2 y_2 + \cdots + x_n y_n. \quad (2.7)$$

Note that only for vectors of the same length an inner product can be taken.

Outer Product For two vectors $\mathbf{x} \in \mathbb{R}^n, \mathbf{y} \in \mathbb{R}^m$, their outer product is defined as

$$\mathbf{xy}^\top = \begin{bmatrix} x_1 y_1 & x_1 y_2 & \cdots & x_1 y_m \\ x_2 y_1 & x_2 y_2 & \cdots & x_2 y_m \\ \vdots & \vdots & \ddots & \vdots \\ x_n y_1 & x_n y_2 & \cdots & x_n y_m \end{bmatrix} \quad (2.8)$$

Note that in the case that $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$: $\text{tr}(\mathbf{xy}^\top) = \mathbf{x}^\top \mathbf{y}$.

2.3.2 Orthogonality

Two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ are said to be orthogonal if

$$\mathbf{x}^\top \mathbf{y} = 0. \quad (2.9)$$

2.3.3 Norms

For this section, we follow the norm definitions as given in the book *Matrix Computations*³. A vector norm on \mathbb{R}^n is a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ that satisfies the following properties:

1. nonnegativity and mapping of the additive identity:
 $f(\mathbf{x}) \geq 0, f(\mathbf{x}) = 0$ iff $\mathbf{x} = \mathbf{0}$,
2. triangle inequality:
 $f(\mathbf{x} + \mathbf{y}) \leq f(\mathbf{x}) + f(\mathbf{y})$,
3. homogeneity:
 $f(\alpha \mathbf{x}) = |\alpha| f(\mathbf{x})$

for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and $\alpha \in \mathbb{R}$. A common class of vector norms are the ℓ_p -norms. They are denoted as $\|\cdot\|_p$, and defined for $p \geq 1$ as

$$\|\mathbf{x}\|_p = \left(\sum_{i=1}^n |x_i|^p \right)^{\frac{1}{p}} \quad (2.10)$$

Example 2.1: some commonly used ℓ_p -norms on $\mathbf{x} \in \mathbb{R}^n$ are:

- $\|\mathbf{x}\|_1 = \sum_{i=1}^n |x_i|$ (a.k.a. Manhattan norm)
- $\|\mathbf{x}\|_2 = \sqrt{\sum_{i=1}^n x_i^2} = \sqrt{\mathbf{x}^\top \mathbf{x}}$ (a.k.a. Euclidean norm)
- $\|\mathbf{x}\|_\infty = \max_i |x_i|$ (a.k.a. Chebyshev norm)

Finally, we will also make use of the notation $\|\mathbf{x}\|_0 = \#\{i : x_i \neq 0\}$ for the ℓ_0 -norm. It reflects the cardinality of a vector, i.e. the number of non-zero vector entries. Note that it is not an actual norm as it does not satisfy the conditions as specified previously, e.g. the homogeneity condition does not hold. Norms are often used as measure for the size of an error or the amplitude of a signal in a certain direction.

“A norm on a vector space plays the same role as absolute value: it furnishes a distance measure. More precisely, \mathbb{R}^n together with a norm on \mathbb{R}^n defines a metric space rendering the familiar notations of neighborhood, open sets, convergence, and continuity.” - Golub et al.³

2.3.4 Orthogonal Projection

A final, interesting operation is the orthogonal projection. In the vector case it is defined as

$$\mathbf{w} = \frac{\mathbf{v}^\top \mathbf{u}}{\mathbf{u}^\top \mathbf{u}} \mathbf{u}, \quad (2.11)$$

where \mathbf{w} is the projection, $\mathbf{v} \in \mathbb{R}^n$ is the projected vector and $\mathbf{u} \in \mathbb{R}^n$ is the vector on which is projected. In the Euclidean sense, the orthogonal projection of a vector can be considered as the vector closest to \mathbf{v} residing in the subspace spanned by \mathbf{u} , i.e. closest vector property⁴.

2.3.5 Geometrical Interpretation

Inner Product The inner product of two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ can alternatively be calculated as

$$\mathbf{a}^\top \mathbf{b} = \|\mathbf{a}\|_2 \|\mathbf{b}\|_2 \cos \theta, \quad (2.12)$$

where θ describes the angle between both vectors, looking from an Euclidean point of view. See **Fig. 2.1** for a visual interpretation in \mathbb{R}^2 . Note that in the

case of orthogonal vectors, i.e. $\theta = 90^\circ$, the inner product becomes zero as $\cos(90^\circ) = 0$. Finally, we observe that when $\|b\|_2 = 1$, we obtain the length of the orthogonal projection of a in the direction of b , i.e. $\|a\|_2 \cos \theta$.




Figure 2.1 | Geometrical representation of the inner product in \mathbb{R}^2 .

Unit Norm Spheres The unit norm sphere, i.e. when $\|\cdot\|_p = 1$, can be best visualized in a two or three dimensional space, e.g. $x \in \mathbb{R}^2$ or \mathbb{R}^3 . The two dimensional case is given in **Fig. 2.2**. Observe that for $\|x\|_0 = 1$ the origin is not included. Note that when we change unity to any arbitrary non-negative number, we obtain a scaled version of the norm sphere, i.e. inflated or deflated.




Figure 2.2 | Geometric interpretation of unit norm sphere ($\|\cdot\|_p = 1$) in a two-dimensional case, i.e. $x \in \mathbb{R}^2$.

2.4 Matrices and its Algebra

We introduce the reader to the concept of two important subspaces, i.e. the column space and the row space, and the concept of rank. Therefore, we maintain the notation and definitions as given in *Matrix Computations*³.

Column Space The column space of a matrix $A \in \mathbb{R}^{m \times n}$ is a vector space spanned by the columns of A and is a subspace of \mathbb{R}^m . It is sometimes also called the range (**ran**) of the matrix A and defined as

$$\mathbf{ran}(A) = \{y \in \mathbb{R}^m : y = Ax \text{ for some } x \in \mathbb{R}^n\}. \quad (2.13)$$

Row Space Similar to the column space of a matrix, the row space is a vector space spanned by the rows of A and is a subspace of \mathbb{R}^n . It is defined as

$$\mathbf{ran}(A^T) = \{x \in \mathbb{R}^n : x = A^T y \text{ for some } y \in \mathbb{R}^m\}. \quad (2.14)$$

Rank A final important property related to a matrix $A \in \mathbb{R}^{m \times n}$ is its rank. It can be defined as

$$\mathbf{rank}(A) = \mathbf{dim}(\mathbf{ran}(A)), \quad (2.15)$$

where **dim** is the dimension of the subspace and reflects its “degrees of freedom”. In the context of matrix compression, minimizing the number of degrees of freedom is key as less so-called components will be required to describe the underlying matrix adequately. Finally, note that one should be vigilant when computing the rank, e.g. in Python⁵, as it might differ from the actual rank due to computational errors.

2.4.1 Norms

Equivalent to vector norm definition, there exist a definition for matrix norms. In particular, a function $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ that satisfies the following properties:

1. nonnegativity and mapping of the additive identity:
 $f(A) \geq 0, f(A) = 0 \text{ iff } A = 0,$
2. triangle inequality:
 $f(A + B) \leq f(A) + f(B),$
3. homogeneity:
 $f(\alpha A) = |\alpha| f(A)$

for $A, B \in \mathbb{R}^{m \times n}$ and $\alpha \in \mathbb{R}$ defines a matrix norm¹. A particular class of matrix norms are related to the ℓ_p -induced norms and are defined as

$$\|A\|_{\ell_p} = \sup_{\mathbf{x} \neq \mathbf{0}} \frac{\|A\mathbf{x}\|_p}{\|\mathbf{x}\|_p} = \max_{\|\mathbf{x}\|_p=1} \|A\mathbf{x}\|_p. \quad (2.16)$$

Example 2.2: some commonly used matrix ℓ_p -induced norms on $A \in \mathbb{R}^{m \times n}$ are:

- $\|A\|_{\ell_1} = \max_j \sum_{i=1}^m |a_{ij}|$
- $\|A\|_{\ell_2} = \sqrt{\max_i \lambda_i(A^T A)} = \sqrt{\lambda_{\max}(A^T A)} = \sigma_1(A)$
- $\|A\|_{\ell_\infty} = \max_i \sum_{j=1}^n |a_{ij}|$

where λ_i are the eigenvalues of matrix $A^T A$, and σ_1 is the largest singular value (see the Singular Value Decomposition in section 2.4.4).

A second class of norms are so-called p-q-norms and defined as

$$\|A\|_{p,q} = \left(\sum_{j=1}^m \left(\sum_{i=1}^n |a_{ij}|^p \right)^{q/p} \right)^{1/q}. \quad (2.17)$$

Finally, three other important norms are the Frobenius norm ($\|\cdot\|_F$), nuclear norm ($\|\cdot\|_*$) and zero-norm or cardinality ($\|\cdot\|_0$), they are respectively defined as

$$\begin{aligned} \|A\|_F &= \sqrt{\sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2} = \sqrt{\text{tr}(AA^T)} = \sqrt{\sum_{i=1}^r \sigma_i^2(A)}, \\ \|A\|_* &= \sum_{i=1}^r \sigma_i(A), \\ \|A\|_0 &= \#\{(i, j) : a_{ij} \neq 0\}, \end{aligned} \quad (2.18)$$

where r is the matrix rank and σ_i are the singular values related to matrix $A \in \mathbb{R}^{m \times n}$ (see the Singular Value Decomposition in section 2.4.4). Note that the Frobenius norm can be interpreted as the ℓ_2 -norm on the singular values, and the nuclear norm as the ℓ_1 -norm on the singular values. Finally, the zero-norm or cardinality of a matrix A , is defined similar to the zero-norm of a vector, and once more does not satisfy the qualifications for a matrix norm.

2.4.2 Orthogonal Projection

Similar to the vector case we can define a projection for matrices. This projection operator P_V projects a matrix onto the subspace spanned by the columns or rows of matrix $V \in \mathbb{R}^{m \times n}$, we provide its definition for the column space of V , such that

$$P_V = V(V^T V)^{-1} V^T. \quad (2.19)$$

Pre-multiplying a to be projected matrix with the projection operator P_V provides the projection.

2.4.3 Geometric Interpretation

ℓ_p -induced Norms An interpretation of the matrix-vector is the matrix being a linear transformation on the original vector. Since we defined our ℓ_p -induced norms in equation 2.16 as the maximum of a unit norm vector under a transformation, we can visualize this transformation by transforming the actual unit norm sphere and defining it as the length of largest vector on the transformed sphere. An example is given in **Fig. 2.3** for the matrix

$$A = \begin{bmatrix} 1 & 2 \\ 0 & 2 \end{bmatrix}. \quad (2.20)$$




Figure 2.3 | Geometric interpretation of matrix norms induced by vector norms in two-dimensional case ($x \in \mathbb{R}^2$). The blue lines give the corresponding vector norms, where the red contours give the image after multiplying with A as given in Eq. 2.20. The dashed black line gives the maximal distance, i.e. the norm. Reproduced after Example 3.1 [pg. 19]⁴.

2.4.4 Factorization

Factorization or decomposition of a matrix is a common principle with matrix related operations and applications. An analogy can be seen with factoring numbers, where one wants to find the constituents of certain numbers in a canonical form for further analysis. We will call these factors of the decomposition the factor matrices.

Example 2.3: finding the zeros of a third-order equation is one of the most simple examples of factorization:

$$\begin{aligned}x^3 - 6x^2 + 11x - 6 &= 0 \\(x - 1)(x - 2)(x - 3) &= 0\end{aligned}$$

as we can see, the factored equation gives more insight in our original problem.

The same simple principle holds for matrix factorization. By exploiting structure one can use it for efficiently solving linear equations, reducing or simplifying calculations, collaborative filtering, noise filtering, e.g. by imposing a certain canonical matrix form for the factor matrices, unmixing of spectral signatures and discovery of latent features underlying the interactions between two or more different entities.

“Therefore, in order to better understand the behavior and properties of a complex system, a natural approach is to decompose the system into its simpler components.” - Chandrasekaran et al.⁶

In this section, we introduce two important basic decompositions: the QR decomposition and singular value decomposition (a.k.a. SVD). We use those basic algorithms in the later chapter as they form the basis for many modern algorithms.

2.4.4.1 QR

We use the definition as stated in *Linear algebra and its applications*¹: every m by n matrix A with linearly independent columns can be factored into

$$A = QR \tag{2.21}$$

The columns of Q are orthonormal, and R is upper triangular and invertible. To achieve this decomposition, different implementations exist based on Householder, block Householder and Givens transformations, all with their own specific conditions and properties.

“One algorithmic idea in numerical linear algebra is more important than all the others: QR factorization.” - Trefethen et al.⁴

Nevertheless, a simple interpretation of the QR decomposition is possible via householder reflections, which consists of reflecting vectors such that one of its elements becomes zero, i.e. this is what we see with R being upper triangular. If we interpret a matrix as a concatenation of vectors and assume $\mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{R}^n$ and $\mathbf{A} \in \mathbb{R}^{n \times 3}$, then

$$\mathbf{A} = \mathbf{Q}\mathbf{R}$$

$$\begin{bmatrix} \vdots & \vdots & \vdots \\ \mathbf{a} & \mathbf{b} & \mathbf{c} \\ \vdots & \vdots & \vdots \end{bmatrix} = \begin{bmatrix} \vdots & \vdots & \vdots \\ \mathbf{q}_1 & \mathbf{q}_2 & \mathbf{q}_3 \\ \vdots & \vdots & \vdots \end{bmatrix} \begin{bmatrix} \mathbf{q}_1^\top \mathbf{a} & \mathbf{q}_1^\top \mathbf{b} & \mathbf{q}_1^\top \mathbf{c} \\ 0 & \mathbf{q}_2^\top \mathbf{b} & \mathbf{q}_2^\top \mathbf{c} \\ 0 & 0 & \mathbf{q}_3^\top \mathbf{c} \end{bmatrix} \quad (2.22)$$

Our matrix Q is orthonormal and as such its vectors span the full range of A. This implies that if we rotate or reflect its vectors (depending on $\det(\mathbf{Q}) = \pm 1$), \mathbf{a} , \mathbf{b} and \mathbf{c} , onto this new basis, e.g. $\mathbf{q}_1^\top \mathbf{a}$, we obtain the rotated vectors defined in our new orthonormal system. Finally, it can be associated with Legendre polynomials in continuous space⁴.

Uniqueness The uniqueness of the QR decomposition is guaranteed and stated in *Numerical linear algebra*⁴ under Theorem 7.2: each $\mathbf{A} \in \mathbb{C}^{m \times n}$ ($m \geq n$) of full rank has a unique reduced QR factorization $\mathbf{A} = \hat{\mathbf{Q}}\hat{\mathbf{R}}$ with $r_{jj} > 0$. The $\hat{\cdot}$ represents the reduced factors.

Computational Complexity The computational complexity of the Householder orthogonalization is $2mn^2 - 2/3n^3$ flops, and can be accessed e.g. via the DGE-QRF routine in LAPACK.

2.4.4.2 Singular Value Decomposition

Building upon the QR decomposition is the singular value decomposition (SVD), which plays a prominent role in the characterization of many matrix “nearness problems” and data analysis³. We state its definition as given in theorem 2.4.1 of *Matrix Computations*³: if A is a real m-by-n matrix, then there exist orthogonal matrices

$$\begin{aligned} \mathbf{U} &= [\mathbf{u}_1, \dots, \mathbf{u}_m] \in \mathbb{R}^{m \times m} \text{ and} \\ \mathbf{V} &= [\mathbf{v}_1, \dots, \mathbf{v}_m] \in \mathbb{R}^{n \times n} \end{aligned} \quad (2.23)$$

such that

$$\mathbf{U}^\top \mathbf{A} \mathbf{V} = \mathbf{\Sigma} = \mathbf{diag}(\sigma_1, \dots, \sigma_p) \in \mathbb{R}^{m \times n}, \quad p = \min\{m, n\} \quad (2.24)$$

where $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_p \geq 0$ and are called the singular values. Observe that we apply two transformation matrices on the left and right side of the matrix containing the singular values to reconstruct our original matrix A. They represent rotations on both column and row space of A.

“The SVD makes it possible for us to say that every matrix is diagonal – if only one uses the proper bases for the domain and range spaces.” - Trefethen et al.⁴

An interesting property of the SVD is its characterization for both its null space and range by the orthonormal bases as

$$\begin{aligned}\mathbf{null}(A) &= \mathbf{span}\{v_{r+1}, \dots, v_n\}, \\ \mathbf{ran}(A) &= \mathbf{span}\{u_1, \dots, u_r\},\end{aligned}\tag{2.25}$$

where $r = \mathbf{rank}(A)$. Furthermore, the rank of a diagonal matrix Σ reduces to the cardinality of its diagonal as

$$\begin{aligned}\mathbf{rank}(A) &= \mathbf{rank}(\Sigma), \\ &= \mathbf{rank}(\mathbf{diag}(\sigma)), \\ &= \mathbf{card}(\sigma),\end{aligned}\tag{2.26}$$

where $\sigma \in \mathbb{R}^p$ with $p = \min\{m, n\}$. Secondly, and important when handling large matrices, is the concept of a thin SVD. As earlier defined, the matrix rank is always smaller or equal to the smallest dimension of its column and row space, i.e. $\mathbf{rank}(A) \leq \min\{m, n\}$. Hence, when our matrix A is rank-deficient, we don't need to store the full U , and Σ matrices, but can reduce it to a thin version: if $A = U\Sigma V^T \in \mathbb{R}^{m \times n}$ and $m \geq n$,

$$A = U_1 \Sigma_1 V^T,\tag{2.27}$$

where

$$\begin{aligned}U_1 &= [u_1, \dots, u_n] \in \mathbb{R}^{m \times n}, \text{ and} \\ \Sigma_1 &= [\sigma_1, \dots, \sigma_n] \in \mathbb{R}^{n \times n}.\end{aligned}\tag{2.28}$$

Thirdly, an interesting property of the SVD and its form is that we can construct the original matrix $A \in \mathbb{R}^{m \times n}$ and $\mathbf{rank}(A) = r$ as a sum of scaled outer products

$$A = \sum_{i=1}^r \sigma_i u_i v_i^T.\tag{2.29}$$

The latter can be seen as the decomposition of matrix A in rank-one matrices. And as such, its decomposition represents the minimal number of rank-one matrices to reconstruct A . Finally, an important theorem is the so called Eckhart-Young theorem stating that if $k < r (= \mathbf{rank}(A))$ and

$$A_k = \sum_{i=1}^k \sigma_i u_i v_i^T,\tag{2.30}$$

then

$$\begin{aligned} \min. \quad & \|A - B\|_2 = \|A - A_k\|_2 = \sigma_{k+1}, \\ \text{subj. } & \mathbf{rank}(B) = k. \end{aligned} \quad (2.31)$$

Hence, we observe that the SVD can optimally identify nearby matrices of lesser rank, where our metric is defined by the 2-norm or Frobenius norm. Its relation to QR decomposition is that an initial QR factorization is commonly performed on the matrix $A \in \mathbb{R}^{m \times n}$ when $m \gg n$. Thereafter, the SVD decomposition is applied on R to increase efficiency, this drastically improves the required number of floating point operations⁷. Furthermore, it is also utilized by randomized SVD methods. The process is given as

$$\begin{aligned} A &= QR, \\ &= Q(\mathbf{U}\Sigma\mathbf{V}^T), \\ &= (Q\mathbf{U})\Sigma\mathbf{V}^T. \end{aligned} \quad (2.32)$$

Secondly, the QR decomposition is sometimes also used within the SVD routine itself, however, in an iterative version. Different SVD implementations and methods exists and are implemented in different linear algebra packages. We advise the review *The singular value decomposition: Anatomy of optimizing an algorithm for extreme scale*⁸ for the interested reader. An simple overview of the packages and some of the SVD methods is given below:

1. EISPACK: via eigenvalue decomposition
2. LINPACK
3. (Sca)LAPACK:
 - based on blocked Householder transformations,
 - based on divide and conquer implementations
4. ARPACK and PROPACK: for sparse and structured matrices
5. Bisection and inverse iterations
6. Multiple Relatively Robust Representations
7. PLASMA, MAGMA and DPLASMA: two stage reductions
8. Jacobi methods

Finally note that for a non-negative matrix A , the first right and left singular vector are non-negative. This is referred to as the Perron-Frobenius vector⁹.

Uniqueness Its uniqueness is defined in *Numerical linear algebra*⁴: every matrix $A \in \mathbb{C}^{m \times n}$ has a singular value decomposition. Furthermore, the singular values $\{\sigma_j\}$ are uniquely determined, and, if A is square and the σ_j are distinct, the left and right singular vectors $\{u_j\}$ and $\{v_j\}$ are uniquely determined up to complex signs, i.e. complex scalar factors of absolute value 1.

Computational Complexity The computational complexity is typically $6mn^2 + 16n^2$ for $m \geq n$ for the SVD (DGESVD) making use of the QR iteration, and $6mn^2 + 8n^3$ for SVD making use of D&C (DGESD) for $m \geq n$.

2.4.4.3 Geometric Interpretation

QR Factorization We will show the principle geometric interpretation for the Householder QR factorization. First of all, constructing our orthogonal Q matrix consist of multiplication of multiple reflections, as such

$$Q = Q_n \cdots Q_2 Q_1. \quad (2.33)$$

Furthermore, each unitary matrix Q_k is chosen to be of the form

$$Q_k = \begin{bmatrix} I & 0 \\ 0 & F \end{bmatrix}, \quad (2.34)$$

where $I \in \mathbb{R}^{(k-1) \times (k-1)}$ and F is a unitary matrix $\in \mathbb{R}^{(m-k+1) \times (m-k+1)}$, called the Householder reflector. The basic idea of this reflector is to reflect the space across the hyperplane H . The latter is visualized in **Fig. 2.4**.




Figure 2.4 | Geometrical interpretation of the Householder reflections. The vector x is projected through a Householder reflection on the Euclidean axis, in both negative (red) and positive (blue) direction. For stability purposes the negative direction is preferred. Reproduced after Figure 10.2 [pg. 72]⁴.

When the reflector is applied, every point on one side of the hyperplane H is mapped to its mirror image on the other side. Hence, we can define F as

$$F = I - 2 \frac{vv^T}{v^T v} \quad (2.35)$$

where v is defined by $\pm \|x\| e_1 - x$. However, for numerical stability the point farthest from x is chosen⁴. Hence in this case $-\|x\| e_1 - x$ would be the preferred choice.

Singular Value Decomposition The singular value decomposition can be geometrically interpreted via two routes. One where $U\Sigma$ are seen as an principal semiaxes of A , and V is seen as a preimage of those principal semiaxes. A second interpretation is seeing the singular values of a matrix A as the length of the semiaxes of a hyperellipsoid E that can be defined as³:

$$E = \{Ax : \|x\|_2 = 1\}. \quad (2.36)$$

We choose the interpretation of the latter. We can decompose A into three matrices U (orthogonal), Σ (diagonal) and V (orthogonal). An orthogonal matrix is a linear transformation and hence in Euclidean space can be seen as a rotation or reflection^a. A diagonal matrix can be seen as a scaling operation. Hence, we perform a rotation, scaling and final rotation on our vector x . This can be seen in **Fig. 2.5**.




Figure 2.5 | Geometrical interpretation of the singular value decomposition of a 2×2 matrix. In the left image a rotation of the principal semiaxes can be seen, after which a scaling and second rotation is applied towards the right image. Reproduced after Figure 4.1 [pg. 26]⁴.

^athis can be defined by the sign of its determinant

2.5 Summary

In this chapter, we have introduced the notation that we maintain throughout this thesis. More specifically the basic notations for scalars, vectors and matrices were given along with some basic notions. In second instance, we introduced vectors and vectors spaces and their vector norms. This enabled the transition to matrices, where we discussed the importance of two important subspaces, introduced different types of matrix norms, and finally discussed two common factorization methods: the QR decomposition and the singular value decomposition. Especially, the latter two are of importance for later chapters.

Bibliography

- [1] Strang, G. Linear algebra and its applications, thomson learning. *Inc., London* (1988).
- [2] Boyd, S. & Vandenberghe, L. *Convex optimization* (Cambridge university press, 2004).
- [3] Golub, G. H. & Van Loan, C. F. *Matrix computations*, vol. 3 (JHU press, 2012).
- [4] Trefethen, L. N. & Bau III, D. *Numerical linear algebra*, vol. 50 (Siam, 1997).
- [5] Van Rossum, G. & Drake, F. L. *Python 3 Reference Manual* (CreateSpace, Scotts Valley, CA, 2009).
- [6] Chandrasekaran, V., Sanghavi, S., Parrilo, P. A. & Willsky, A. S. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization* **21**, 572–596 (2011).
- [7] Chan, T. F. An improved algorithm for computing the singular value decomposition. *ACM Transactions on Mathematical Software (TOMS)* **8**, 72–83 (1982).
- [8] Dongarra, J. *et al.* The singular value decomposition: Anatomy of optimizing an algorithm for extreme scale. *SIAM review* **60**, 808–865 (2018).
- [9] Cichocki, A., Zdunek, R., Phan, A. H. & Amari, S.-i. *Nonnegative matrix and tensor factorizations: applications to exploratory multi-way data analysis and blind source separation* (John Wiley & Sons, 2009).

Chapter 3

Application Domains

After this chapter you will be able to

- follow the basics of imaging mass spectrometry
- understand the challenges of MALDI-IMS based molecular imaging
- follow the basics of the MUSE instrument
- understand the challenges of the MUSE instrument

3.1 Introduction

In this chapter, we briefly introduce the two application domains that will be used for case studies of the developed strategies. We include for both the basic instrumentation principles, notions of the data processing of the raw data acquired, and challenges faced in these instrumental settings.

3.2 Imaging Mass Spectrometry: MALDI

In this section, we describe imaging mass spectrometry (IMS) and more specifically the matrix-assisted laser desorption/ionization type of IMS, referred to as MALDI-IMS. First, we detail its importance in a global setting. Second, we discuss the sample preparation, the instrumental principles, and the processing of the raw data. Finally, we discuss some of its challenges.

3.2.1 Introduction

In analytical chemistry, mass spectrometry has become a central technology to measure and map the presence and abundance of a broad array of molecular species in a given sample. When mass spectrometry is used to map the molecular content at different spatial locations throughout a sample, we typically refer to it as IMS. The particular type of imaging mass spectrometry, specifically suited for mapping biological systems, is MALDI-IMS. *What makes MALDI-IMS particularly powerful as a molecular imaging modality is that it can concurrently map hundreds of molecular species' spatial distributions throughout a given tissue sample in a single experiment. Furthermore, it differentiates itself from many other molecular imaging modalities (often microscopy based) by its ability to measure molecular presence without the need for prior tagging of the molecules of interest. This makes MALDI-IMS untargeted and particularly well suited to explore hereto unknown parts of biomolecular pathways or to discover previously unknown players in e.g. disease patho-mechanisms. Its ability to analyze complex mixtures ranging from small drug compounds and metabolites to lipids and very large proteins without the need for prior tagging makes it very versatile and has powered a rapid expansion of IMS-based applications in recent years, ranging from pathology to material science, and from drug metabolism to forensics*¹. In broad terms, an IMS experiment provides for all detectable ions within a given molecular mass range (so for a multitude of molecular species), their spatial distribution and ion intensity, which indirectly reports location-specific abundance of those species. This mapping of molecular content takes place across a prepared sample surface, which is often an organic tissue section. These tissue sections are often biological samples, which can originate e.g. from plants or fruits², mammalian tissue³ from model organisms, or even human tissue⁴, but also from for example bacterial biofilms. As a result, MALDI-IMS is used in an increasingly wide array of application domains for practical studies that include topics such as proteomics, pathology, drug metabolism, and bacterial infection.

3.2.2 Sample Preparation

A first step for most samples is cryo-sectioning, where the material, in a frozen state, is sectioned. The sectioning step is pictured in **Fig. 3.1**. After sectioning, the slices are mounted on conductive glass slides and dried, as seen in **Fig. 3.2**. After drying some additional solutions can be applied to the samples, e.g. to wash off salts or other undesired compounds. After these steps, and often a final step before measuring, a matrix solution, a crystalline layer, is sprayed onto the samples. The latter is used to facilitate non-destructive ionization of organic molecular species and enhances the intensity responses from the samples. At this stage, the samples are ready to be measured by mass spectrometry. Admittedly, this is only a rough sketch of the sample preparation process for MALDI-IMS. A large number of intermediate steps and concomitant variables, e.g. section thick-

ness, spray thickness, process temperature, process duration and solution types, are not discussed here since they are not central to the scope of the thesis, but they do play an important role in the success of IMS experiments and add to the data and meta-data provided by the mass spectrometer. The nuances related to these matters require expert knowledge and extensive wet-lab experience to do them full justice. An example of such wet-lab and sample preparation complexity can be observed in **Fig. 3.2** where three different tissue sections from the same tissue block are mounted next to each other on a glass slide: the left and right sections e.g. show cutting artifacts.




Figure 3.1 | Rat brain mounted in microtome to perform transversal sectioning. Source: author's visit at Mass Spectrometry Research Center, Vanderbilt University (TN, US).




Figure 3.2 | Transversal rat brain slices mounted on conductive glass slide. Source: author's visit at Mass Spectrometry Research Center, Vanderbilt University (TN, US).

3.2.3 Instrumental Principles

Since 2019, a novel trapped ion-mobility time-of-flight (*timsTOF*) mass spectrometer instrument has been gaining prominence for imaging mass spectrometry experiments⁵. Using MALDI as an ionization source, it generates ions from sample molecules by means of laser irradiation, conducted at distinct pixel locations on a tissue section. The instrument has two main modes: *qTOF* and *TIMS*. The first, *qTOF*, is the abbreviation for quadrupole time-of-flight mass spectrometer and returns the sample-provided intensities for a range of mass-to-charge (m/z) features when the laser is fired on a distinct spatial pixel. The detector of the instrument is based on counting statistics, such that individual values within a mass spectrum can be considered Poisson distributed in terms of ion intensity counts. Consecutively, an ion intensity image showing the spatial distribution of a specific molecular species or specific m/z value is a spatial map of ion intensity values reported across all tissue locations that were measured. Besides this mode, there is also a *TIMS* mode, which provides beside m/z based separation of the sample molecules also ion mobility based separation, a second spectral axis expressed in terms of drift time. As such, it adds an extra spectral dimension,

namely drift time ($1/K_0$, where K_0 is the ion-mobility), to the existing spectral dimension probed by mass spectrometry, namely m/z , and these two spectral dimensions are added to the two spatial dimensions along which a 2-D IMS experiment takes place. The drift time of the second spectral axis reflects the relative collisional cross sections of the ions and as such can differentiate molecule species with similar m/z values, but with different 3-dimensional shapes. Besides these different modes, all measurements can be collected in a so called positive and negative ion mode (referring to the polarity of the electro-magnetic fields used in the ion optics of the mass spectrometer), providing the possibility to respectively measure both protonated or deprotonated molecule species in a sample. A schematic visualization of the instrument and its different ion optic parts is given in **Fig. 3.3**.




Figure 3.3 | Prototype MALDI timsTOF mass spectrometer. (A) The developed imaging platform couples a MALDI source capable of high-spatial-resolution and throughput imaging to an ESI timsTOF Pro. (B) Ion-mobility separations are performed in a dual Tims funnel in the first vacuum stage of the instrument. Tims resolution is tunable by adjusting the EFG scan time, and throughput is maximized by performing accumulation and analysis in parallel on sequential pixels. Original picture and caption from **Figure 1**⁵.

3.2.4 Preprocessing

After data acquisition, it is common to preprocess by means of baseline subtraction (removing slow-moving technical variation from mass spectra), intensity normalization (projecting ion intensity scales from different mass spectra to a common consensus scale that allows direct comparison of ion intensity values across spectra), m/z alignment (projecting m/z axes from different mass spectra to a common consensus axis that allows direct comparison of ion intensity values across spectra), but also to bin m/z features and re-sample ion intensity values to increase SNR, or to perform peak-picking to reduce the dimensionality by means of domain-specific feature selection. The latter consists of picking peaks of interest within full profile descriptions of mass spectral measurements, usually retaining only the peaks with most intensity and highest signal (effectively eliminating the ability to estimate the signal noise floor and potentially removing low-abundant m/z features from downstream analysis). Furthermore, in case of

the TIMS mode, most of these preprocessing steps along the m/z axis will yield an equivalent preprocessing step along the ion-mobility dimension, leading to even complexer preprocessing pipelines and potential feature loss along the way.

3.2.5 Challenges

Traditionally, MALDI data challenges consist primarily of the need to remove non-biological variation from the data, i.e. noise related to e.g. sample preparation and instrumental conditions, and the need to effectively deal with the often impractical dimensionality of the measurement sets. With the arrival of MALDI-timsTOF⁵ and the additional spectral (ion mobility) axis it can add to IMS experiments, the massive dimensional nature of such molecular imaging data has grown into the primary impediment holding back this area of exploratory molecular imaging.

3.2.5.1 Dimensionality

One way of illustrating the dimensionality challenges involved in IMS and its generalization, ion mobility (IM) IMS, is to quote the data dimensions for a concrete molecular imaging example. An example experiment mapping the molecular content of a sagittal mouse pup tissue section measures a data tensor of which the first two modes, the spatial axes, together yield 164,808 individual pixels or spatial measurement points. If standard IMS is applied using the instrument depicted in Fig. 3.3, collecting just a mass spectrum per pixel, this will amount to acquiring intensity values for 219,140 individual m/z bins per pixel. If we store these values to disk using a standard 4 bytes per intensity value, this amounts to a raw IMS measurement set of around 144 gigabytes ($= 164,808 \text{ pixels} \times 219,140 \text{ } m/z \text{ bins/pixel} \times 4 \text{ byte/bin}$). If now the TIMS mode of the instrument is turned on, effectively adding a spectral dimension separating molecular content in terms of drift time (dt) and turning the IMS experiment into an IM-IMS measurement, we will collect a m/z -by-dt matrix of intensity values per pixel rather than a m/z vector of intensity values. This will effectively multiply the data size of the standard IMS experiment times the number of ion mobility bins collected, which in this case would be 4001 dt bins. In short, the data collected during the IM-IMS measurement of the mouse pup, if it were stored exhaustively using standard approaches used for IMS data, would amount to around 578 terabytes of raw intensity values ($= 164,808 \text{ pixels} \times 4001 \text{ dt bins/pixel} \times 219,140 \text{ } m/z \text{ bins/dt bin} \times 4 \text{ byte/bin}$). Although this is but one example, if all modes of the timsTOF-instrument are used to full effect, which is physically possible with the instruments currently available in labs, measurement sets growing into the hundreds of terabytes of raw data will not be exceptional. The need to effectively deal with such dimensionality is a nontrivial task. Furthermore, with instrumental separation capabilities growing, there is an increasing abundance of low signal-to-noise ratio (SNR) measurements in such high-dimensional data sets. In the past,

this computational challenge has been mitigated by peak-picking spectra rather than retaining full profile descriptions of mass spectral and ion mobility signals. However, this incurs an inherent and sometimes pronounced loss of information, which in certain cases renders the effort of collecting the physical measurement in the first place mute. To make full ion mobility (IM) IMS, as a powerful generalization of MALDI-TOF IMS, a practical reality for molecular imaging in the various application domains where IMS is currently gaining a foothold, it will be essential to have mathematical methods available that allow us to store, analyze, and interact with the full profile descriptions and signals collected in such experiments.

3.2.5.2 Noise

The types of noise commonly associated with MALDI (tims)TOF measurements includes detector noise, i.e. Poisson/Gaussian distributed⁶ instrumental noise, experimental acquisition noise, e.g. drifts and biases forming intensity gradients over the image, and sample preparation and chemical matrix (multiplicative) noise^{7,8}. Some of these noise types can be attributed to genuine non-linear processes, while others are the result of complex, interactive, or insufficiently well-understood processes that e.g. take place at the atomic/molecular scale or that are difficult to model. In addition to these undesirable noise effects, high-dimensionality exacerbates cumulative problems such as noise accumulation⁹, further hampering subsequent statistical analysis. This can occur for example when the effects of sample preparation, of using different matrix solutions, of varying process temperatures and durations in the different steps of tissue preparation start accumulating along the different phases and levels of an IMS experiment. Overall, one can state that noise is a complex challenge in MALDI (tims)TOF acquisitions, which needs to be taken into account, particularly when combined with inherently complex biological samples such as human or mammalian tissue environments.

3.3 Integral field Spectroscopy: MUSE

Since 2014 the Multi-Unit Spectroscopic Explorer (MUSE) instrument performs integral field spectroscopy at the Very Large Telescope of the European Southern Observatory at Paranal, Chile¹⁰. First, we detail its importance in a global setting. Second, we discuss the instrumental principles, and the processing of the raw data. Finally, we discuss some of its challenges.

3.3.1 Introduction

MUSE is built as an alternative to classical approaches by both performing imaging and spectroscopy at the same time^{10,11}: *The idea is to merge into one instrument the best of the two capabilities: from the imaging world its field of view and high spatial resolution; and from the spectrograph's world its high resolving power and large spectral range.* One of its goals is to enable detection of bright emission lines, even when the continuum of those objects is very faint. Concurrently, this can be performed in a large field of view, making it ideal for surveys^{12–15}. The last few years, it has evolved to an instrument for general purpose exploration, ranging from exoplanet detection to studying nearby galaxies¹⁶.

3.3.2 Instrumental Principles

The MUSE instrument, depicted in **Fig. 3.4**, performs integral field spectroscopy experiments at the VLT. It has two modes: a wide field and narrow field mode, where the first mode is characterized by its large field of view and rough spatial sampling, and the second mode has a smaller field of view concurrent with finer spatial sampling. The wide field mode is ideal for large-scale studies, e.g. of the progenitors of normal nearby galaxies out to redshifts $z > 6$ and the identification of very faint sources detected in other bands¹⁰. The narrow field mode, on the other hand, enables detailed study of single objects at very high spatial resolution, e.g. the atmospheric signature detailing of Neptune¹⁷.

In its wide-field mode MUSE is capable of recording a $1' \times 1'$ patch of sky. The incoming light, captured by the VLT, is therefore distributed over 24 slicer integral field units (IFU). In **Fig. 3.4**, these can be seen as the 24 tube-like units (yellow & turquoise). The 24 so-called sub-fields of $60'' \times 2.5''$ are then fed into the integral field unit, where it is further sliced into 48 slices of $15'' \times 0.2''$. These slices are fed to the spectrograph and projected onto a CCD. Each IFU has a $4k^2$ CCD and is capable of detecting a wavelength range roughly between 480 to 900 nm. A schematic of this process is given in **Fig. 3.5**. As such, MUSE is able to concurrently deliver spectral information on 90,000 spaxels (spatial pixel elements) per exposure.




Figure 3.4 | Phase-A design of MUSE on the Nasmyth platform of a VLT unit telescope. Original picture and caption from **Figure 1**¹⁸.

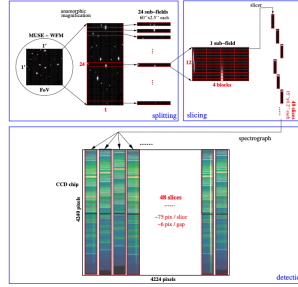


Figure 3.5 | A schematic view of the lightpath in the MUSE instrument, visualizing the splitting and slicing in the optics before the data is recorded by the CCD. The MUSE data reduction software has to reverse this process to remove the instrumental signature. Original picture and caption from **Figure 1**¹⁹.

3.3.3 Preprocessing

The processing of the raw CCD measurements is performed in a data pipeline^{16,20}, schematically depicted in a **Fig. 3.6**. It serves as a data reduction system able to remove the instrumental signature and provide fully calibrated spectra, images, and data cubes as end product. The data pipeline itself has multiple so-called recipes that can be applied consecutively on the raw data, e.g. for removing bias, dark current, flat-fielding, wavelength corrections, and handling twilight.

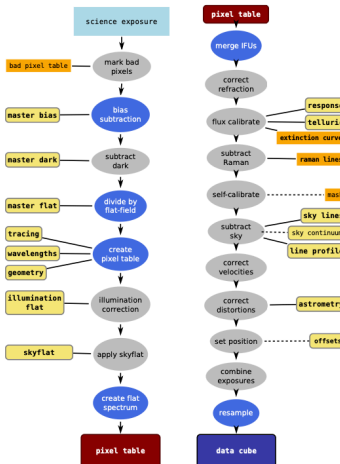


Figure 3.6 | **Left:** Basic processing from raw science data to the intermediate pixel table. **Right:** Post-processing from pixel table to the final data cube. Optional steps are marked grey, mandatory ones in blue. Manually created input files have an orange background, calibrations are highlighted. Inputs that are needed are connected with a solid line, dotted lines signify inputs that are not required. Original picture and caption from **Figure 1**¹⁹.

3.3.4 Challenges

Within MUSE two main challenges come to the forefront: noise and dimensionality. The first challenge is generally of higher urgency and importance, while the dimensionality of these data sets is not yet at the critical level that can be observed in MALDI-IMS and that could serve as a bottleneck in MUSE experiments.

3.3.4.1 Noise

Due to the counting statistics associated with the arrival of photons on the CCD, Poisson/Gaussian distributed noise is a first noise source. To circumvent its related problems, long integration times are performed, often by co-adding multiple exposures of a predefined integration time. The decision of the number of such exposures and concurrent integration times is therefore an important aspect to consider with regards to reaching the scientific goals for a selected target²¹. This is often referred to as the signal-to-noise ratio (SNR) and, when assuming only Poisson noise, it can be expressed as

$$\text{SNR} \propto \sqrt{n}, \quad (3.1)$$

where n is the number of measured photons. This shows that performing more and/or longer measurements, ideally, improves the SNR proportional to the square root of its integration time. This gets hampered, however, by other (systematic) noise sources such as night sky emissions, e.g. the OH forest signature, instrumental artifacts, wavelength misalignment, or data pipeline-introduced artifacts. The difficulty of correcting for these issues, lies in the definition of the (desired) signal, the (undesired) noise and background, their wavelength-dependency and their subsequent quantification. Some of these sources such as sky background residuals have been tackled in the past by specially developed tools, e.g. ZAP²² for the MUSE instrument, but remain sparsely applied and (most importantly) instrument-specific.

“We will differentiate between three components that add to give the photoelectron counts measured on detector pixels, namely signal, background, and noise. Signal is the count number coming from the object under study and background from any other light flux component, while noise is the statistical fluctuation of the count number. For example, for the study of the central nucleus of a galaxy, the signal comes from the nucleus itself and the background both from night-sky airglow and the other components of the galaxy (disk and bulge) integrated along the line of sight. Note that the signal versus background distinction is somewhat arbitrary, as it generally depends on the exact science goal of the observation.” - Bacon and Monnet²¹

It is clear that improving our ability to disassemble empirically acquired measurements into different types of subcomponents (by means of the methods explored in this thesis), will help advance the filtering out of undesirable (noise) components and the ability to focus on and isolate the desirable (signal) components of these measurements. This, in turn, will increase the SNR for signals of interest and lower the noise floor for reliable and robust signal detection in a noisy environment.

3.4 Summary

In this chapter, we identified the importance and goals, introduced the basic principles, and reviewed some of the processing steps of imaging mass spectrometry and integral field spectroscopy. Furthermore, we have identified some key challenges: for MALDI timsTOF this primarily consists of dimensionality reduction, where for the MUSE instrument noise reduction and increasing the SNR would have the larger impact. With those challenges in the back of our mind, we review the available literature in the next chapter, focus on both these aspects, and identify which tools are already available for tackling these problems.

Bibliography

- [1] Cornett, D. S., Reyzer, M. L., Chaurand, P. & Caprioli, R. M. Maldi imaging mass spectrometry: molecular snapshots of biochemical systems. *Nature methods* **4**, 828–833 (2007).
- [2] Wang, J., Yang, E., Chaurand, P. & Raghavan, V. Visualizing the distribution of strawberry plant metabolites at different maturity stages by maldi-tof imaging mass spectrometry. *Food Chemistry* **345**, 128838 (2021).
- [3] Anderson, D. M. *et al.* High-resolution matrix-assisted laser desorption ionization–imaging mass spectrometry of lipids in rodent optic nerve tissue. *Molecular vision* **19**, 581 (2013).
- [4] Neumann, E. K. *et al.* Spatial metabolomics of the human kidney using maldi trapped ion mobility imaging mass spectrometry. *Analytical Chemistry* **92**, 13084–13091 (2020).
- [5] Spraggins, J. M. *et al.* High-performance molecular imaging with maldi trapped ion-mobility time-of-flight (timstof) mass spectrometry. *Analytical chemistry* **91**, 14552–14560 (2019).
- [6] Keenan, M. R. & Kotula, P. G. Accounting for poisson noise in the multivariate analysis of tof-sims spectrum images. *Surface and Interface Analysis: An International Journal devoted to the development and application of techniques for the analysis of surfaces, interfaces and thin films* **36**, 203–212 (2004).
- [7] Hassell, K. M., LeBlanc, Y. C. & McLuckey, S. A. Chemical noise reduction via mass spectrometry and ion/ion charge inversion: amino acids. *Analytical chemistry* **83**, 3252–3255 (2011).
- [8] Van de Plas, R., Yang, J., Spraggins, J. & Caprioli, R. M. Image fusion of mass spectrometry and microscopy: a multimodality paradigm for molecular tissue mapping. *Nature methods* **12**, 366–372 (2015).
- [9] Fan, J., Han, F. & Liu, H. Challenges of big data analysis. *National science review* **1**, 293–314 (2014).
- [10] Bacon, R. *et al.* Probing unexplored territories with muse: a second generation instrument for the vlt. In *Ground-based and Airborne Instrumentation for Astronomy*, vol. 6269, 6269J (International Society for Optics and Photonics, 2006).
- [11] Bacon, R. *et al.* The muse second-generation vlt instrument. In *Ground-based and Airborne Instrumentation for Astronomy III*, vol. 7735, 773508 (International Society for Optics and Photonics, 2010).

-
- [12] Borisova, E. *et al.* Ubiquitous giant $\text{Ly}\alpha$ nebulae around the brightest quasars at $z \sim 3.5$ revealed with muse. *The Astrophysical Journal* **831**, 39 (2016).
- [13] Bacon, R. *et al.* The muse hubble ultra deep field survey-i. survey description, data reduction, and source detection. *Astronomy & Astrophysics* **608**, A1 (2017).
- [14] Herenz, E. C. *et al.* The muse-wide survey: A first catalogue of 831 emission line galaxies. *Astronomy & Astrophysics* **606**, A12 (2017).
- [15] Wisotzki, L. *et al.* Nearly all the sky is covered by $\text{Lyman-}\alpha$ emission around high-redshift galaxies. *Nature* **562**, 229–232 (2018).
- [16] Weilbacher, P. M. *et al.* The data processing pipeline for the muse instrument. *arXiv preprint arXiv:2006.08638* (2020).
- [17] Irwin, P. G. *et al.* Latitudinal variation in the abundance of methane (CH_4) above the clouds in neptune’s atmosphere from vlt/muse narrow field mode observations. *Icarus* **331**, 69–82 (2019).
- [18] Weilbacher, P. M., Roth, M. M., Pécontal-Rousset, A. & Bacon, R. The muse data reduction pipeline. *New Astronomy Reviews* **50**, 405–407 (2006).
- [19] Weilbacher, P. M. *et al.* Design and capabilities of the muse data reduction software and pipeline. In *Software and Cyberinfrastructure for Astronomy II*, vol. 8451, 84510B (International Society for Optics and Photonics, 2012).
- [20] Weilbacher, P. M. *et al.* The muse data reduction pipeline: status after preliminary acceptance europe. *arXiv preprint arXiv:1507.00034* (2015).
- [21] Bacon, R. & Monnet, G. *Optical 3D-spectroscopy for Astronomy* (John Wiley & Sons, 2017).
- [22] Soto, K. T., Lilly, S. J., Bacon, R., Richard, J. & Conseil, S. Zap-enhanced pca sky subtraction for integral field spectroscopy. *Monthly Notices of the Royal Astronomical Society* **458**, 3210–3220 (2016).

Chapter 4

Literature Review

This chapter introduces you to

- the objective and corresponding goals for this literature review;
- the basics of matrix decomposition;
- the concept of ℓ_1 -robustness;
- the extended linear mixture model and its validity for our purposes;
- the framework of analysis based on “atoms”;
- different classical and robust matrix decomposition methods.

4.1 Introduction

We present here a compact review on matrix decomposition literature. A selection of classical and robust linear-mixture-model-based methods is introduced, along with their assumptions, computational complexity and distinct applications in imaging.

4.1.1 Goals and objectives

The goal of this literature review is to provide a compact review of classical and robust matrix decomposition methodology for imaging applications by means of a unifying framework of notation and analysis. This will be pursued by summarizing the method’s underlying extended linear mixture model in comparable notation form, introducing you to our framework, and using the latter for reviewing the different methods from a couple of aforementioned viewpoints. By no means will the presented methodology be exhaustive. Instead, our method

selection is mainly based on their legacy use in applications for both imaging mass spectrometry, especially focusing on MALDI¹, and astronomical imaging. Finally, we note that only the decomposition of matrices is considered. Thus, we will not consider tensor methodology or tensor-centric extensions of the methods in this review. The latter is considered out-of-scope for this specific thesis, but we do see these extensions as a next stage in our research.

4.1.2 What's in a name?

Collaborative filtering, factor analysis, signal separation, system identification, dimensionality reduction, matrix decomposition, and matrix factorization, all share the same basic backbone of decomposing measurements. For clarity, we will primarily maintain the name matrix decomposition as it comes closest to the essence of this thesis is about: decomposing measurements collected in a matrix. The broad naming suggests however that this type of methodology covers a relatively wide spectrum of functions. Some important use cases are:

- **Data exploration/visualization:** increasing data set sizes challenge the exploration and visualization of complete data sets. Hence, reducing data sets to a small number of descriptive components, e.g. latent variables, enables exploration and visualization at least of the “most important” underlying relations in the data;
- **Denoising:** reducing noise, i.e. the part of the measurement that is not of interest, by decomposing measurements, on the grounds of discerning mathematical properties, into noise and signal, i.e. what is of interest, components;
- **Prediction/estimation:** using an underlying low-dimensional description of the measurements for predicting unknown values, e.g. Netflix Prize², or to estimate certain parameters that are inherent to the measured phenomena, e.g. system identification;
- **Dimensionality reduction:** increasing data set sizes challenges efficient and robust application of statistical analysis to such data sets. Phenomena like heterogeneity, noise accumulation, and spurious correlation can negatively influence their performance or even lead to false discovery³. Low-dimensional representations, obtained via matrix decomposition methods, can alleviate some of those problems such as the curse of dimensionality;
- **Feature selection:** given the ever-increasing number of measurements and modalities, manual feature selection has become tedious, often impractical, and sometimes impossible. At the same time, automated but traditional peak-picking, i.e. selecting highest intensity features in measurement sets, can lead to information loss due to an overly aggressive features

recognition approach. Hence, a more structured selection of features, as components of a matrix decomposition approach, could alleviate that loss via a low-dimensional representation, without necessarily sacrificing feature selection advantages and while still avoiding the need for human supervision;

- **Clustering:** identifying groups of similar measurements in data sets is often a desired operation to reduce the complexity of data sets. This operation can be incorporated directly into a matrix decomposition approach as it inherently decomposes measurements into linear mixtures of components with terms and factors.

Besides these more analytical use cases, there are also practical reasons for matrix decompositions to be applied. For example, in a traditional use of matrix decomposition, e.g. QR decomposition, see section 2.4.4, one exploits underlying structure to simplify subsequent matrix operations, e.g. preventing the creation of matrices that do not fit in computer memory⁴. Furthermore, matrix decomposition has been used to increase numerical stability and solution accuracy of certain methods, e.g. in solving linear systems⁴.

4.1.2.1 Terminology

We distinguish the following concepts:

- **Model:** We consider a model to represent the underlying relations between different matrix terms, and internal structure of those terms, e.g. a factorization model. These relations can be based on general knowledge, e.g. first principles, as well as assumptions. For example, a simple linear model could be $A = B + C$. In this case, A could be a measurement matrix, B a matrix containing signal, i.e. the part of the measurement that is of interest, C a matrix containing noise, i.e. the part of the measurement that is not of interest;
- **Method:** We consider a method to extend a particular model by imposing specific mathematical constraints on factors and terms to enable the retrieval of (unique) solutions suitable for a particular application. This often results in proposing an optimization problem with its corresponding parameters. Building on our previous example, we could add an assumption on the distribution of C , the noise matrix, when projected onto a one-dimensional space, such as that it should follow a Gaussian or Laplacian distribution;
- **Algorithm:** We consider an algorithm as a sequence of code lines being ran on a computer, that solves a method on a particular set of measurements. At its heart can lie a range of solvers from a simple iterative solver to a convex optimization solver, but it can also consist of a closed-form solution, e.g. an orthogonal projection.

4.1.3 Which matrix is suitable?

We need a matrix containing measurements: a measurement matrix (a.k.a. a data set, measurement set or data matrix). *We will assume that a column contains measurements of the same property and a row contains measurements of the same origin.* Examples of such matrices are given in Example 4.1.

Example 4.1: The three cases explain each one manner for constructing a measurement matrix. The used symbols represent: E an experiment, S a sample of a particular experiment, and P a sample's property. Finally, a “.” represents a measured property value.

Case A: single experiment of multiple samples

$$E_I \begin{array}{c} S_a \\ S_b \\ S_c \\ S_d \end{array} \begin{array}{|c|c|c|} \hline P_1 & P_2 & P_3 \\ \hline \cdot & \cdot & \cdot \\ \hline \cdot & \cdot & \cdot \\ \hline \cdot & \cdot & \cdot \\ \hline \cdot & \cdot & \cdot \\ \hline \end{array}$$

The matrix rows represent the different samples and the columns represent the measured properties.

Case B: multiple experiments of a single sample

$$S_a \begin{array}{c} E_I \\ E_{II} \\ E_{III} \\ E_{IV} \end{array} \begin{array}{|c|c|c|} \hline P_1 & P_2 & P_3 \\ \hline \cdot & \cdot & \cdot \\ \hline \cdot & \cdot & \cdot \\ \hline \cdot & \cdot & \cdot \\ \hline \cdot & \cdot & \cdot \\ \hline \end{array}$$

The measurement matrix encompasses multiple experiments of a single sample with multiple measured properties.

Case C: multiple experiments of multiple samples

$$\begin{array}{c} E_I \\ \\ E_{II} \end{array} \begin{array}{c} S_a \\ S_b \\ \\ S_a \\ S_b \end{array} \begin{array}{|c|c|c|} \hline P_1 & P_2 & P_3 \\ \hline \cdot & \cdot & \cdot \\ \hline \cdot & \cdot & \cdot \\ \hline \cdot & \cdot & \cdot \\ \hline \cdot & \cdot & \cdot \\ \hline \end{array}$$

The measurement matrix encompasses multiple experiments of multiple samples, for which multiple properties are measured.

However, note that many variations on these cases exist, each affecting the matrix decomposition in their own way. It is therefore essential to have a measurement matrix that is appropriate for the underlying model. This can be ob-

tained by selecting a model based on a particular measurement matrix, or by constructing a measurement matrix such that it fits the model at hand. The extended linear mixture model is introduced in section 4.2.

4.1.4 Why, what, when robust?

Classical matrix decomposition methods often perform poor in the presence of outliers. This has led to robust versions of such methods to be formulated. A simple analogy can be drawn between classical and robust matrix decomposition, and the arithmetic mean and median of a series. This analogy is elucidated in Example 4.2.

Example 4.2: Imagine we have a series of measurements x in the form of 5 real numbers. We now want to reduce that series to a single summarizing number, expressing what the series “generally” looks like. Therefore, we define the series’ mean and median:

$$\begin{aligned}x &= \{1, 4, 6, 9, 1000\}; \\ \mathbf{mean}(x) &= 204; \\ \mathbf{median}(x) &= 6.\end{aligned}\tag{4.1}$$

We could argue that the obtained median is closer to what most elements of x look like, while the mean seems to be skewed towards the largest value (1000), here considered an outlier. This is a direct result of the ability of the median to exclude that particular outlier from consideration. Also, observe that a distance measure arises: how do we define closeness?

In Example 4.2, three commonly attributed properties of an outlier are manifesting themselves: (1) its value is not in the line of the expected, e.g. a very high or low value; (2) it occurs on a sparse basis throughout the measurements, e.g. only 1 out of 5 measurements is very different from the others; (3) its occurrence can severely deteriorate the performance of statistical or other methods, e.g. the mean is skewed towards the outlier. Now, one could compare the classical matrix decomposition methods to a “mean”-based approach, while robust matrix decomposition methods can be seen as “median”-based viewpoints. This is a primarily the result of the use of ℓ_2 or Frobenius norm minimization in many classical methods as the distance measure on the approximation error. In such minimization formulations, large measurements, or in the case of matrix decomposition, outlier deviations between the original measurement matrix and the modeled matrix, i.e. error term, have a relatively larger influence on the minimization than smaller deviations.




Figure 4.1 | Comparison of ℓ_1 - and ℓ_2 -norm (respectively in red and blue) in a 1-dimensional representation, where $\|x\|_1 = |x|$ and $\|x\|_2^2 = x^2$ for $x \in \mathbb{R}$.

The latter is visualized in Figure 4.1. We observe that function values $f(x)$ corresponding to the ℓ_2 -norm (blue) get increasingly larger for larger input values of x . When minimizing an error modeled as $f(x) = \|x\|_2^2$, the decrease in function value will thus be higher when large errors are reduced^a, compared to small errors that only result in small changes in the objective function value. We can use this characteristic quite well if we want all deviations between our measurements and our model to be small. Similar to the outlier in Example 4.2, the presence of outliers, however, deteriorates the performance of the ℓ_2 -norm based methods. The minimization will then primarily try to eliminate those large error values. As the large but singular errors do not necessarily reflect what most measurement values look like, they can distort the outcome by leading the model fitting optimization astray. In classical matrix decomposition methods it is therefore often assumed that no outliers are present or that they were (manually) removed before processing. It is clear that whether a measurement matrix meets this assumption is often hard to determine beforehand, leading to a somewhat of a chicken-and-egg situation.

The trend towards big data³, furthermore hinders assessment of this assumption. First, the addition of new measuring modalities and properties such as to existing imaging instruments as well as the concurrent improvements in resolution and sensitivity have tended to also result in an increased likelihood of outliers to occur. Second, outliers might not always appear as clear cut as presented in Example 4.2, meaning that outlier removal mechanisms are generally not guaranteed to perform well under broad conditions. Nevertheless, given the large impact of outliers on the outcome of methods, we still want to be able to deal with them. Therefore, robust methodology, less sensitive to outliers, has been gaining in popularity.

In imaging applications, outliers can appear as the result of instrumental errors or due to defect sensor elements. They can also originate from natural processes. For example: in astronomical imaging, highly energized gamma rays of

^aexcluding values close to 0

extragalactic outbursts hitting telescopes may leave stripped patterns of saturated CCD^a pixels. Or, highly energetic line-emitting galaxies, e.g. Lyman-alpha emitters⁵, can manifest themselves in a few spatial and frequency bins. In imaging mass spectrometry, the nature of outliers can be multiplex: regions that contain more matrix solution^b or regions of spatially thicker tissue can locally lead to larger ion intensity responses, resulting in localized higher intensities measured.

4.1.4.1 ℓ_1 -robustness

The ℓ_0 -norm (a.k.a. cardinality, see section 2.4.1) is a measure that naturally arises in the characterization of sparsity, one of the properties of outliers, and as such can be exploited for outlier recognition and subsequent removal. Unfortunately, using this measure in an optimization problem's objective function results in an NP-hard problem formulation^c. As a result, a particular convex relaxation of the measure is commonly used instead: the ℓ_1 -norm (see section 2.3.5). In Figure 4.1, where the ℓ_1 -norm is given in red, one observes that all input values are weighted proportional to their value, in contrast to the earlier mentioned ℓ_2 -norm which weighted inputs not equal but proportional to the square of this value. This implies that a decrease in error value for a large error that is lowered, will be equal to the decrease in objective function value of a small error lowered by that same amount. This provides visually an intuition for the robustness of an ℓ_1 error norm against outliers and its decreased influence on the error minimization. Furthermore, under certain conditions^{6,7} its use in certain scenarios is equivalent to the ℓ_0 -norm. In addition, the ℓ_1 -norm is convex and has a gradient, in this case a subgradient, often enabling the recovery of solutions in polynomial time, and the guarantee for the retrieval of optimal and unique solutions under broad conditions⁸. *Therefore, we will mainly focus in this review on the ℓ_1 -norm based robust methodology, with an emphasis on sparse outlier robustness.* Note that many other robust techniques exist such as via multivariate trimming⁹, divergence measures¹⁰, or other sparsity measures¹¹. However, they often fail to guarantee performance under broad conditions, to be generalized, or cannot provide solutions within polynomial time⁷.

4.2 Extended Linear Mixture Model

We review methods that are based on a family \mathcal{F} of extended linear mixture models¹². They can be summarized as

$$\mathcal{F} : M = A + B + C, \text{ with } A = YZ^T, \quad (4.2)$$

^acharge-coupled device

^bin this context, matrix is not considered in a mathematical sense, but a naming for crystalline layer applied on a tissue section before analysis

^ca simple interpretation: it is impossible to solve this problem within reasonable time

where $M \in \mathbb{R}^{m \times n}$ represents a measurement matrix, and $A \in \mathbb{R}^{m \times n}$, $B \in \mathbb{R}^{m \times n}$, and $C \in \mathbb{R}^{m \times n}$ are matrix terms differing in mathematical properties. Note that \mathcal{F} also contains the family of linear mixture models, i.e. $C = 0$. We will refer to the linear mixture model as “LMM” and the extended linear mixture model as “XLMM”. The matrix term A is composed of a product of $Y \in \mathbb{R}^{m \times k}$ and $Z \in \mathbb{R}^{n \times k}$. This bilinear formulation of A as the product YZ^T , can also be interpreted as Y containing base vectors^a for the measurement representation, and Z containing the weights for the corresponding base vectors such that

$$YZ^T = \sum_{i=0}^k \mathbf{y}_i \mathbf{z}_i^T = \sum_{j=0}^n Y \mathbf{z}_j \mathbf{e}_j^T, \quad (4.3)$$

where \mathbf{e}_j is a Cartesian unit vector with the j -th entry equal to 1, and \mathbf{y} and \mathbf{z} with subscript i represent the columns of the corresponding matrices Y and Z , and subscript j indicates their rows. This type of factorization can be connected to the idea that perception of the whole is based on the perception of its parts¹³. The parts therefore define the whole, such that finding the parts, i.e. base vectors, is desired.

Finally, this model and its variations will underlie a method that further imposes structure on the different matrix terms and factors of the model. Structure can be thought of as some type of lower-triangular or banded matrix structure, the property of low-rank, a distribution in a one-dimensional space, sparsity, etc. Such requirements are translated into an optimization problem using matrix norms, measures, and constraints, such that the solution reflects the defined mathematical structures. Often, those structures will be traded-off against each other via adjustable model parameters.

4.2.1 Validity

In control theory, linear models are generally of interest if a system behaves as an LTI^b system and operates near a stable point. The difference between control and our imaging application is that no feedback mechanism is available. As such, we can only validate our model via the available measurements, expert knowledge, or through new measurements. It can be compared to system identification without being able to close the loop to evaluate its performance, e.g. via inspecting control performance or checking whether it is operating near that stable point.

First of all, we need to make sure that the chosen family \mathcal{F} of linear models is appropriate for the application and its measurement matrix. A first principles approach can for example be used as a starting point. This often results in complex models, e.g. including non-linearities, which leads to follow-up questions, e.g. whether all these (non-linear) effects are captured in the measurements and

^anot be confused with a basis, which is one particular set of base vectors.

^bLinear time-invariant

as such can contribute in the identification of the model, that should be taken into account. Often, terms are less or actually more decisive than initially modeled, resulting in model deviations from the actual measurements and ultimately leading to poor performance. In the worst case, a first principles method is not practical at all, as is the case for complex chemical reactions appearing in analytical chemistry and imaging mass spectrometry.

This problem can be alleviated by working the other way around, starting from a simple linear model and making it more complex if need be. Besides being parsimonious, i.e. Occam's razor¹⁴, it relieves overfitting of complex models to noisy measurements and has strong generalization properties¹⁵. The latter leads to a model that is often suited for a larger range of applications and measurements sets. This is in contrast to non-linear models that typically possess relatively poorer generalization properties than linear modeling¹⁶. After linear model fitting, we evaluate the outcome and improve the linear model through the addition of new or higher order terms based on a combination of empirical findings and theoretical insights. This process repeats itself cyclically, until an appropriate model for the application is constructed. A parallel can be drawn with Taylor expansions where the first components capture the linear behavior of a series or function. By adding more terms, the finer non-linear effects are captured. Hence, a linear model can often be seen as a first step in discovery. Eventually, we evaluate the model by inspecting departures from the chosen family of models \mathcal{F} to the available measurements or compare its performance on new measurement data.

4.3 Methods

In this section, we review different classical and robust methods. First, we introduce the framework of atoms in order to review the selection of methods.

4.3.1 Framework of Atoms

Following the idea that perception of the whole is based on perception of its parts¹³, we will analyze the parts that are the result of each particular method. Those parts are in our case the vectors that are stacked in the mixture model $A = YZ^T$, as it is often assumed to contain the results of interest. This bilinear formulation lends itself to be interpreted as a collection of base vectors and weights as was shown earlier (see equation 4.3). *For clarity across the different method descriptions, we will name those base vectors, "atoms", and a collection of base vectors, a dictionary.* The latter is in our case represented by the matrix Y , its columns on the other hand are its atoms. In some cases, Z can have a similar interpretation concurrently with Y . This terminology of atoms and dictionaries was first introduced in the context of wavelets¹⁷ and was later adopted for its use in basis pursuit¹⁸. The latter define an atom as discrete time signals of length

n . A collection of parametrized waveforms $\mathcal{D} = (\phi_\gamma : \gamma \in \Gamma)$ is called a dictionary, where atoms represent the waveforms ϕ_γ . Furthermore, note that this description generalizes the concepts of a design matrix, e.g. used for ordinary least squares, and bases as it assumes no inherent properties of atoms nor dictionaries.

For some methods a dictionary is required *a priori* such that only weights are to be calculated, e.g. Fourier transform, while for others both are sought, e.g. principal component analysis. Furthermore, methods can define a space where atoms and weights can or cannot reside in, e.g. in a positive orthant for non-negative matrix factorization. Methods can also impose conditions to guarantee optimality or to considering robustness against outliers. Our framework therefore focuses on understanding how those basic building blocks, i.e. atoms are formed and what they correspond to in a mathematical sense. This will not only enable comparison of methods, but also enable the translation to application domain specific information. Finally, the solutions of these methods are not to be considered exact (and the same is true for the underlying model). The crux, however, is to define the atoms and dictionaries such that they produce appropriate results for the application. As a first step, we will therefore explore different classical and robust methods and review them on their *atoms of robustness*.

4.3.2 Overview

Each method is introduced along with reference reading, its principles, and the nature of its atoms. Thereafter, we introduce the assumptions along with its time and space complexity (see section 2.4.4.2), if available. Finally, we mention a selection of papers in the context of both imaging mass spectrometry, especially focusing on MALDI¹, and astronomical imaging. For all methods we assume that a measurement matrix $M \in \mathbb{R}^{m \times n}$ is given and that for those specific cases a dictionary $Y \in \mathbb{R}^{m \times k}$ is available. We also suppose that $m \gg n$ and that the “true” matrix rank r of A is not known *a priori*. An overview of the reviewed methods is given in Table 4.1. Finally, we abbreviate the eigenvalue and singular value decompositions respectively as EVD and SVD.

4.3.2.1 Search Methodology

The methods were mainly retrieved from the review *Unsupervised machine learning for exploratory data analysis in imaging mass spectrometry*¹⁹ and the *Handbook of robust low-rank and sparse matrix decomposition: Applications in image and video processing*²⁰. The other papers and reference books in this review were mostly found via <http://scholar.google.com>. The application papers were sought via the same source by the following search terms:

- “*method-name*”,
- “*method-name*+“*application-name*”.

The method names are given in Table 4.1 and the used application names were *imaging*, *spectral imaging*, *imaging mass spectrometry*, *mass spectrometry imaging*, *malDI*, *astronomy*, *astronomical imaging*, *MUSE*, and *spectroscopy*. The selection of application papers is by no means exhaustive, and only reflects the reliable results obtained via above mentioned search terms.

Table 4.1 | A selection of matrix decomposition methods, divided into classical and robust methodology. Besides the methods, their abbreviation is given along with the year of initial mention, whether a dictionary is required a priori, the used error measure, e.g. norm, the underlying model, and a selection of application papers. Note that “.” implies that no application papers were found.

#	Method	Abbr.	Ref.	Year	Dict. Req.	Error Meas.	Model	Applications	Section
Classical Methodology									
1	Dictionary Learning	DL	[6,18]	1994	yes	ℓ_2	LMM	[21–26]	4.3.3.1
2	Common Spatial Pattern	CSP	[27]	1990	no	ℓ_2	LMM	[28–30]	4.3.3.2
3	CX/CUR Decomposition	CX/CUR	[31]	2006-2008	no	Frobenius	LMM	[32,33]	4.3.3.3
4	Independent Component Analysis	ICA	[34,35]	1980-1990	no	Negentropy	LMM	[36–39]	4.3.3.4
5	Maximum Autocorrelation Factorization	MAF	[40]	1984-1985	no	ℓ_2	LMM	[41,42]	4.3.3.5
6	Non-Negative Matrix Factorization	NMF	[13,43]	1994/1999	no	Frobenius	LMM	[38,44–47]	4.3.3.6
7	Ordinary Least Squares	OLS	[48]	~1795	yes	ℓ_2	LMM	[44,49–51]	4.3.3.7
8	Principal Component Analysis	PCA	[52–54]	1901/1933	no	ℓ_2	LMM	[19,55–58]	4.3.3.8
9	Ridge Regression	RR	[59]	1960s	yes	ℓ_2	LMM	[60]	4.3.3.9
10	Sparse Principal Component Analysis	SPCA	[61]	2006	no	ℓ_2	LMM	[62]	4.3.3.10
Robust Methodology									
11	Block-Sparse RPCA	RPCA-LBD	[63,64]	2011/2012	no	$\ell_{1,2}$	LMM	.	4.3.4.1
12	Bounded Principal Component Pursuit	BPCP	[65]	2019	no	ℓ_1	LMM	.	4.3.4.2
13	Manhattan NMF	MAHNMF	[66,67]	2008/2012	no	ℓ_1	LMM	.	4.3.4.3
14	Principal Component Pursuit	PCP	[7,68,69]	2009	no	ℓ_1	LMM	[70,71]	4.3.4.4
15	Robust Dictionary Learning	RDL	[72]	2013	yes	ℓ_1	LMM	.	4.3.4.5
16	Robust NMF	RNMF	[73]	2011	no	ℓ_1 , Frobenius	XLMM	.	4.3.4.6
17	Side-Information PCP	SIPCP	[74]	2017	Side Info.	ℓ_1	LMM	.	4.3.4.7
18	Stable Principal Component Pursuit	SPCP	[75]	2010	no	ℓ_1 , Frobenius	XLMM	.	4.3.4.8

4.3.3 Classical Methodology

In this section we present the classical methods in alphabetic order.

4.3.3.1 Dictionary Learning

Dictionary learning (DL) encompasses methods that seek for sparse measurement representations. The idea is that a measurement might not be sparsely represented in only one basis, such as in a Fourier transform, but becomes sparse when represented by a dictionary composed of different bases^{6,17}. This overdetermined dictionary of atoms allows for adaption by giving up the uniqueness of its solution. The literature mainly consists of two branches: (1) defining/finding an appropriate dictionary; (2) defining a sparse representation for a measurement given a dictionary (a.k.a. synthesis and analysis¹⁸). Some well-known methods are matching pursuit¹⁷, basis pursuit⁶, and morphological component analysis⁷⁶. Besides, other methods like K-SVD⁷⁷ try to combine both synthesis and analysis by concurrently updating a predefined dictionary. In this section, we introduce you to the second branch.

A sparse representation is often retrieved by imposing an ℓ_0 -norm on the weights. However, given the associated problems (see section 4.1.4.1), different solutions have been proposed. A first type of solution relaxes the cardinality measures by an ℓ_1 approximation⁶. A second type of solution uses a greedy approach¹⁷, where sparse representations are sought in an iterative way, one element at a time. A third solution class consists of using the so-called threshold technique, where an inner product, to find the corresponding weights, is thresholded. For a more thorough review on both branches and methods, we recommend the overview article *Dictionary Learning*⁷⁸ and a book chapter on *Dictionary learning for medical image denoising, reconstruction, and segmentation*⁷⁹ for a medical application-driven review.

“Linear expansions in a single basis, whether it is a Fourier, wavelet, or any other basis, are not flexible enough. A Fourier basis provided a poor representation of functions well localized in time, and wavelet bases are not well adapted to represent functions whose Fourier transforms have a narrow high frequency support. In both cases, it is difficult to detect and identify the signal patterns from their expansion coefficients, because the information is diluted across the whole basis. Similar examples can be found for any type of basis. Such decompositions are similar to a text written with a small vocabulary. Although this vocabulary might be sufficient to express all ideas, it requires to use circumvolutions that replace unavailable words by full sentences.” - Mallat, Zhang¹⁷

Methodology We will consider basis pursuit⁶ with its goal of obtaining sparse weights when a predefined dictionary is given. This method connects to robust formulations such as principal component pursuit (see 4.3.4.4) via its ℓ_1 -norm relaxation. The assumed underlying model is the LMM in vector formulation

$$\mathbf{m} = \mathbf{Y}\mathbf{z} + \mathbf{b}, \quad (4.4)$$

where $\mathbf{Y} \in \mathbb{R}^{m \times n}$ is the dictionary. For correct recovery we will assume that $\|\mathbf{b}\|_2^2 < \epsilon$, with ϵ being small⁶. In this formulation, the vector $\mathbf{m} \in \mathbb{R}^m$ is our measurement vector, $\mathbf{z} \in \mathbb{R}^n$ the sparse representation, and $\mathbf{b} \in \mathbb{R}^m$ represents an additive error vector. The latter takes into account that in practice often no exact recovery, i.e. $\mathbf{b} = \mathbf{0}$, can be retrieved. The problem is therefore stated as

$$\begin{aligned} \min_{\mathbf{z}} \quad & \|\mathbf{z}\|_0 \\ \text{subj.} \quad & \mathbf{m} = \mathbf{Y}\mathbf{z} + \mathbf{b}, \\ & \|\mathbf{b}\|_2^2 < \epsilon. \end{aligned} \quad (4.5)$$

The authors suggest the following convex relaxation⁶ to be solved

$$\begin{aligned} \min_{\mathbf{z}} \quad & \|\mathbf{z}\|_1 \\ \text{subj.} \quad & \|\mathbf{m} - \mathbf{Y}\mathbf{z}\|_2^2 \leq \epsilon. \end{aligned} \quad (4.6)$$

Finally, note that there are no restrictions to adapt it for the purpose of matrices, except for the additional computational burden.

Atoms The atoms of this method are defined by the given dictionary, which will remain unchanged over the optimization. Constructing an appropriate dictionary is hence important for the retrieval and sparsity of its solution. The latter is therefore application, if not measurement set, dependent. Furthermore, in contrast to methods such as ordinary least squares, the dictionary is allowed to be overdetermined. Examples of such dictionaries are the composition of different known bases such as orthogonal unit-length vectors, wavelets, sines, cosines, etc. For some dictionary learning methods⁷⁷, the dictionary is considered as an adaptive element to obtain even sparser solutions or to allow for adjustments with respect to the measurement set. In both cases, the obtained sparse weights in \mathbf{z} will not be unique.

Assumptions It has been proven⁸⁰ that for a noisy signal $\mathbf{m} = \mathbf{Y}\mathbf{z} + \mathbf{b}$, where $\|\mathbf{b}\|_2^2 < \epsilon$ and \mathbf{z} is sparse, i.e. $\|\mathbf{z}\|_0 < \frac{1}{4}(1 + \frac{1}{\mu})$, basis pursuit as proposed in the program given in equation 4.6 leads to a stable result such that $\|\hat{\mathbf{z}} - \mathbf{z}\|_2^2 \leq \frac{4\epsilon^2}{1 - \mu(4\|\mathbf{z}\|_0 - 1)}$, where μ represents the mutual coherence^a of atoms and $\hat{\mathbf{z}}$ is the estimated sparse representation.

^aa correlation measure

Complexity Different optimization algorithms have been proposed for solving basis pursuit in a primal as well as a dual formulation^{6,18,24,81}. Besides, *different dictionaries can impose drastically different computational burdens*¹⁸. For a thorough complexity review and comparison between the different solvers, we therefore refer the reader to the article named *Atomic Decomposition by Basis Pursuit*¹⁸.

Applications

- **Astronomical Imaging:** In a first paper²⁴, basis pursuit is applied for spectrum estimation on several synthetic data sets on radial velocity data from an extrasolar planetary system. In a second paper²¹, the concept of centered dictionary learning is introduced, which is based on basis pursuit. Centered dictionary learning updates the fixed dictionary, and centers the training patches to yield approximate translation invariance. Finally, its performance is compared to basis pursuit for astronomical image denoising. A third paper²², proposes a denoising technique that exploits the mathematical frameworks of sparse representations and coupled dictionary learning, and tests it on simulated noisy spectroscopic templates. Finally, a fourth paper²⁵ explores the use of dictionary learning techniques in the context of radio interferometry for astrophysical signals.
- **Imaging Mass Spectrometry:** In a first paper²³, compressed sensing principles are used to reduce the number of measurements needed, concurrent with peak picking the spectra via sparse representation principles, as proposed by basis pursuit. Reconstruction results are shown for a rat brain measurement set in the spectral and spatial domain. In a second paper²⁶, the authors show the use of dictionary learning for deconvolving IMS spectra. They adapt the method for Poisson distributed measurements, instead of the normal Gaussian distribution, and incorporate universally applicable information on common ion types of molecules in MALDI mass spectrometry.

4.3.3.2 Common Spatial Pattern

Common spatial pattern (CSP) was first introduced in 1990 in the context of extracting common spatial patterns underlying the EEGs from two human populations²⁷. The idea is to maximize the captured variance in one population, while concurrently minimizing it in the other to obtain the discriminating patterns with respect to both populations. The latter is optimal in an linear, least-square sense. This method is closely related to PCA and MAF, and can be used for classification and feature selection purposes.

Methodology The methodology consists of concurrently maximizing the variance in one population, while minimizing the variance of the other population.

Therefore, we need two measurement matrices, i.e. M_1 and M_2 , each consisting of a different set of measurements from the same or another population, i.e. origin. The following Rayleigh quotient is proposed²⁷

$$\begin{aligned} \max_z \quad & \frac{\|M_1 z\|_2^2}{\|M_2 z\|_2^2} \\ \text{subj.} \quad & \|z\|_2 = 1. \end{aligned} \quad (4.7)$$

Remark that this method reduces to a PCA approach when M_2 reduces to the identity matrix. The solution of this method is found by constructing a composite covariance matrix $C \in \mathbb{R}^{n \times n}$ as

$$\begin{aligned} P_{\perp}^i &= I - (\mathbf{1}\mathbf{1}^T/m)M_i, \\ X_i &= P_{\perp}^i M_i, \\ C_{X_i} &= X_i^T X_i, \\ C &= \frac{1}{2} \sum_{i=1}^2 C_{X_i}. \end{aligned} \quad (4.8)$$

In section 4.3.3.8, a thorough explanation of the different steps is given. Next, an EVD is performed along with the proposition of a composite transformation $P_C \in \mathbb{R}^{n \times n}$ as

$$\begin{aligned} C &= V \Sigma V^T, \\ P_C &= \Sigma^{-1/2} V^T. \end{aligned} \quad (4.9)$$

Finally, the composite transformation is applied on the two population covariance matrices and a second EVD is performed

$$\begin{aligned} \tilde{C}_{X_i} &= P_C C_{X_i} P_C^T \\ &= V_i \Sigma_i V_i^T. \end{aligned} \quad (4.10)$$

The composite transformation ensures that the sum of eigenvalues of both transformed covariance matrices is equal to $\mathbf{1}$ ^{27,82}. The matrix V_i now contains the atoms of interest and as such represents the dictionary. These eigenvectors constitute a rotation matrix and can thus be used to project our original measurement matrices on to obtain the weights, similar to PCA (see section 4.3.3.8, equation 4.26).

Atoms Conversely, the atoms obtained from CSP are aligned to the directions of maximal variance difference between our two measurement matrices M_1 and M_2 . The atoms in this case can be used to discriminate two populations from each other as the first few atoms of V_1 will maximize the captured variance of M_1 , while corresponding to the directions of minimal variance in M_2 . The inverse is true for V_2 and \tilde{C}_{X_2} . No claims on the uniqueness of the atoms and dictionaries are made.

Assumptions Although no explicit assumptions are made in the original paper²⁷, it is noted that homogeneity of the measurement matrices probably plays a large role in misclassification. The only implicit requirement for this method is that both measurement matrices contain measurements from a same origin. Remark that the modality could be different as long as the same property is measured. Finally, we suggest that an analysis based on the existing body of eigen-analysis is performed in order to formulate assumptions that guarantee performance under different conditions.

Complexity Similar to PCA, the complexity of CSP can be alleviated memory and time wise by performing SVDs instead of the EVD. Its time complexity will be mainly driven by this decomposition and as such reduces to $\mathcal{O}(mn^2)$ flops. Its spatial complexity $\mathcal{O}(mn)$ is of the same order of magnitude as the measurement matrix.

Applications

- **Astronomical Imaging:** The method is used by the same authors in three different papers²⁸⁻³⁰. In the respectively first, CSP is used in the context of planet signal extraction from direct imaging. In a second paper, CSP is utilized for planet detecting in direct-imaging observations. Finally, the authors propose a process based on CSP for circumstellar discs detection.
- **Imaging Mass Spectrometry:** No applications were found.

4.3.3.3 CX/CUR Decomposition

The CX and CUR decomposition were proposed in 2008 in the paper *Relative-error CUR matrix decompositions*³¹. The goal of this type of decompositions is dual: (1) obtaining low-rank decompositions in polynomial time; (2) making a decomposition insightful. The latter suggests that the atoms of the dictionary should not purely be linear combinations of properties, e.g. this is the case for SVD and NMF. Both goals are achieved by approximating the measurement matrix by targeted selection of its columns and/or rows, e.g. via subspace sampling, and using those as atoms. This approximation is achieved by randomized methods that are provided with a relative-error guarantee.

Methodology For both methods, we first need to define the matrix rank r of our low-rank approximation. The CX decomposition will approximate our input matrix as $M \approx CX$, while the CUR decomposition will approximate it as $M \approx CUR$. The quality of approximation for the CX decomposition³¹ follows with probability at least $1 - \delta$ such that

$$\left\| M - CC^\dagger M \right\|_F \leq (1 + \epsilon) \|M - M_k\|_F, \quad (4.11)$$

where $C \in \mathbb{R}^{m \times c}$ is a matrix consisting of the selected columns of M . The number of sampled columns c is directly related to the constant ϵ , δ and the desired rank k . In equation 4.11, M_k represents the optimal k -rank approximation of M , i.e. actual truncated SVD, this follows from the Eckhart-Young Theorem⁴. The CX decomposition is achieved in three steps: (1) defining sample probabilities; (2) finding the corresponding sampling and scaling matrices, respectively S_C and D_C , such that

$$C = MS_C D_C; \quad (4.12)$$

(3) finally calculating $X = C^\dagger M$.

The quality of approximation for the CUR decomposition³¹ follows with probability at least $1 - \delta$ such that

$$\|M - CUR\|_F \leq (1 + \epsilon) \|M - M_k\|_F, \quad (4.13)$$

where $C \in \mathbb{R}^{m \times c}$ is a matrix consisting of the selected columns of M , $R \in \mathbb{R}^{r \times n}$ is a matrix consisting of the selected rows of M , and $U \in \mathbb{R}^{c \times r}$ is a generalized inverse of the intersection between C and R . A CUR decomposition is also achieved in three steps, where a measurement matrix M and its corresponding column matrix C are assumed as input: (1) defining sample probabilities; (2) constructing row sampling and scaling matrices, respectively S_R and D_R , such that

$$R = D_R S_R^\top M; \quad (4.14)$$

(3) constructing U such that

$$U = (D_R S_R^\top C)^\dagger. \quad (4.15)$$

Finally, note that CX and CUR decompositions can be acquired through many different implementations using different sampling rules, e.g. uniform sampling, leverage score sampling, Gaussian projection, and different reconstruction methods.

Atoms The idea of CX/CUR decompositions is to use actual columns and rows, or at least scaled versions of those, to provide insight in the measurement matrix. As such, the atoms are the actual selected columns and rows in this case. They are contained in the dictionaries C and R . Hence, the selection procedure defines the dictionaries which are consequently not unique. Furthermore, no orthogonality or sparsity in the atoms is assumed, nor is there a non-negativity constraint.

Assumptions The only implicit assumption is that the measurement matrix allows for a low-rank approximation.

Complexity The time complexity of CX is in the order of $\mathcal{O}(mnc)$ flops and a space complexity of $\mathcal{O}(mc + cn)$. The time complexity of the linear time CUR is $\mathcal{O}(mn + nc + nr + mc^2)$ flops, the space complexity on the other hand is $\mathcal{O}(mc + nr + cr)$. Finally, note that the different implementations exist that mostly offer a trade-off between complexity and quality of approximation.

Applications

- **Astronomical Imaging:** No applications were found.
- **Imaging Mass Spectrometry:** In a first paper³², the CX and CUR decomposition are used for peak picking purposes on mouse brain and lung data. They show that *20 ions or 40 locations can be used to reconstruct the original matrix to a tolerance of 17% error for a widely studied image of brain lipids*. In a second paper³³, the authors showcase the use of a CX decomposition on a 1.1TB IMS measurement matrix of a plant sample from the *Peltatum* species.

4.3.3.4 Independent Component Analysis

The search for an independent component analysis (ICA) initiated in the 1980s, and resulted in a whole spectrum of methods^{34,35,83}. The main assumption is that different sources, i.e. atoms, are mixed together during the acquisition of the measurement matrix. As such, its goal is to find a transformation matrix to “unmix” and retrieve those atoms. In contrast to PCA with its uncorrelated atoms, ICA strives to find statistically independent atoms. Its applications consist of signal separation, e.g. cocktail party problem, dimensionality reduction, but also for detection problems. We advise a compact and up-to-date review⁸⁴ for more information on the identifiability issue that arises in ICA, the testing of the independent components, and extensions to tensors and non-negative models.

Methodology Most ICA methods, using an LMM under the hood, consist of a two-step approach. First, PCA or data centering and scaling is applied to obtain column-wise unit variance. Secondly, the actual ICA estimation is performed to find uncorrelated components that maximize non-Gaussianity⁸⁴. This non-Gaussianity is also required for the uniqueness and identifiability of its solution. FASTica⁸³ is a popular iterative algorithm for solving the ICA problem. After a whitening step, the following optimization is solved

$$\begin{aligned} \max_{\mathbf{Y}} \quad & \sum_{i=1}^n J_G(\mathbf{y}_i), \\ \text{subj.} \quad & \mathbb{E}[(\mathbf{y}_k^T \mathbf{z})(\mathbf{y}_j^T \mathbf{z})] = \delta_{jk}, \end{aligned} \tag{4.16}$$

where \mathbf{y} are the vectors that constitute the mixing matrix \mathbf{Y} , and $J_G(\cdot)$ is an approximated formulation of negentropy. Maximizing the later corresponds to

minimizing the mutual information. Both a fixed-point algorithm and a gradient method, e.g. Newton method, are possible solvers for the program given in equation 4.16⁸⁴. For the retrieval of the sources, i.e. Z , a pseudo-inverse of Y can be employed.

Atoms In this class of methods, Y is representing a transformation matrix (a.k.a. mixing matrix). The atoms are in this case captured by the dictionary Z and are approximately statistical independent from each other. Furthermore, through the optimization conditions and the statistical property of independence, the atoms are uncorrelated. Finally, the individual elements of each atom are assumed to follow a non-Gaussian distribution.

Assumptions The main assumption consists of the representation of measurements as statistically independent sources, following an underlying LMM. Besides, the following assumptions are made for identifiability purposes^{35,84}:

- the atoms of the dictionary Z are mutually statistically independent;
- the atoms of the dictionary Z have non-normal distributions;
- the mixing matrix Y is square and invertible.

Other assumptions may depend on the noise term properties. Finally, homoscedasticity of the measurements and sources is also required³⁵.

Complexity The time complexity for FASTica⁸⁵ is $\mathcal{O}(2(m+1)ri)$ flops per column, where r is the length of the atom and i corresponds to the number of iterations. Note that the convergence depends on the used cost function, which depends on the method and the underlying statistics of the data. Finally, for the retrieval of the dictionary Z , the extra cost of a pseudo-inverse is required.

Applications

- **Astronomical Imaging:** In a first paper³⁶, ICA is used in the context of cosmic microwave background data analysis. The method is shown to separate underlying spectral emissions from different astrophysical sources and to allow the estimation of a number of components.
- **Imaging Mass Spectrometry:** In a first paper³⁷, the authors compare the use PCA, NMF and ICA on imaging mass spectrometry data. In a second paper³⁸, ICA and NMF are explained and explored for identifying underlying factors in a mouse cerebellum IMS measurement sets. In a final paper³⁹, a similar set of methods is used for unmixing tablet compounds of MALDI IMS measurements.

4.3.3.5 Maximum Autocorrelation Factorization

Maximum autocorrelation factorization (MAF) was first proposed in 1984/1985 as alternative to principal component analysis for the transformation of multivariate spatial imagery⁴⁰. MAF maximizes the autocorrelation between neighboring measurements instead of maximizing the variance, such as for PCA. For example in the case of imaging, the neighboring measurements will consist of neighboring pixels in the individual images. The idea is that signal correlates on small scales, while noise does not. Hence, based on this discerning property a decomposition is proposed. However, it was noticed that this property might be erroneous in regions containing small spatial structures⁸⁶.

Methodology The idea of this method is the creation of an artificial noise covariance matrix without utilizing new or other measurements. To recover the noise covariance matrix shifted covariance matrices are subtracted from the original covariance matrix. The composite covariance matrix is defined as

$$C = 2C_X - C_X^\Delta - C_X^{-\Delta}, \quad (4.17)$$

where C_X is the covariance matrix as defined for PCA and

$$C_X^\Delta = \text{Cov}\{M, M_\Delta\}, \quad (4.18)$$

where Δ represents a spatially shifted matrix⁸⁷. Now, a Rayleigh quotient similar to the one used for CSP (see equation 4.7) is constructed, such that

$$\begin{aligned} \max_z \quad & \frac{z^T C z}{z^T C_X z} \\ \text{subj.} \quad & \|z\|_2 = 1 \end{aligned} \quad (4.19)$$

is optimized. When using centered measurements, this corresponds to maximizing the autocorrelation. This program can be solved via an EVD, similar to CSP and PCA. However, one could also directly multiply the measurement set with $C^{-1/2}$, calculate its covariance matrix, and subtract the error covariance matrix, i.e. a scaled identity matrix, and perform an SVD to retrieve the dictionary and corresponding atoms. Both methods result in the underlying LMM.

Atoms The atoms correspond to the directions of largest autocorrelation. As such, they are similar to the atoms retrieved for PCA with the difference that the influence of noise components is drastically reduced by spatial correlation. Similar to PCA, the atoms might be composed of positive as well as negative elements. Finally, both Y and Z can be regarded as dictionaries.

Assumptions The main assumption, besides the underlying LMM, is that noise and signal components are uncorrelated with respect to each other. Furthermore, it is assumed that noise does not correlate on small spatial scales, but signal does. Depending on its application similar assumptions hold as for PCA.

Complexity No complexity analysis is found in the literature. However, a back-on-the-envelope calculation results in the need for 3 matrix multiplications, 2 matrix summations and an EVD or SVD. Hence, the complexity will be in the order of complexity of the EVD or SVD.

Applications

- **Astronomical Imaging:** No applications were found.
- **Imaging Mass Spectrometry:** In a first paper⁴¹, the authors demonstrate the use of MAF in the context of identifying tumorous regions via spatially correlating components. In a second paper⁴², a similar context is used to prove that MAF is able to obtain spatially resolved tumor-specific mass spectra.

4.3.3.6 Non-Negative Matrix Factorization

Non-negative matrix factorization (NMF) was first proposed in 1994 as positive matrix factorization⁴³, and was later adapted and popularized in 1999 as NMF¹³. However, the main idea remained the non-negativity of retrieved dictionaries and weights. Through the years, numerous variations on the originally proposed method were researched. For an elaborate review, we refer the interested reader to the book^a *Nonnegative matrix and tensor factorizations: applications to exploratory multi-way data analysis and blind source separation*¹² and to a short overview paper⁸⁸.

Multivariate Curve Resolution Originally proposed in 1971 as self modeling curve resolution⁸⁹, multivariate curve resolution (MCR) is closely related to NMF. The obtained atoms and dictionaries of MCR are also non-negative and the bilinear formulation is similar to NMF. However, the prior knowledge on the error term is considered different. The latter translate itself into a different optimization formulation and solver. Its optimization formulation is given as

$$\begin{aligned} \min_{Y,Z} \quad & \left\| M - YZ^T \right\|_2 \\ \text{subj.} \quad & Y \geq 0, Z \geq 0, \end{aligned} \tag{4.20}$$

and solved for via an alternating least squares or a non-negative least squares⁹⁰. This is in contrast to NMF that is suggested⁹⁰ to be solved via a total least squares as a maximum likelihood estimator. In practice, one could consider MCR as a type of NMF. Note that the alternating least squares is also suggested for NMF¹².

^aIn section 1.2, different NMF versions are presented and in section 1.3 some basic solving approaches are described. Secondly, appendices 1.A-1.B offer a connection to SVD and PCA and some useful background information. Especially interesting is chapter 2, where the authors describe in detail different similarity measures and generalized divergences, going beyond the standard matrix norm definition (see section 2.4.1). For practical implementations chapter 3 to 6, can be reviewed.

Methodology First of all, it should be noted that NMF comes in a variety of colors and flavors. Each of the formulations uses available or assumed information on the matrix terms such as a particular noise distribution or the orthogonality of certain factors. Furthermore, similarity measures and divergences are used to translate distributional information into an optimization problem. Two popular methods are based on the Frobenius norm minimization⁴³ and the Kullback-Leibler divergence¹³. Also, note that different measures can be used for the different factors.

For the Frobenius-based formulation, with its underlying LMM, the following minimization is proposed

$$\begin{aligned} \min_{Y,Z} \quad & \frac{1}{2} \left\| M - YZ^T \right\|_F^2 \\ \text{subj.} \quad & Y \geq 0, Z \geq 0, \end{aligned} \quad (4.21)$$

where “ \geq ” is considered element-wise⁴³. This formulation can be seen as a natural extension to the non-negative least squares¹². Through its bilinear formulation it is convex with respect to Y and Z , although, not convex in both. As such, an alternating minimization is proposed in a two step scheme:

1. $Z^T \leftarrow [Y^\dagger M]_+$
2. $Y \leftarrow [MZ^\dagger]_+$

where \dagger represents the Moore-Penrose inverse, and $[\cdot]_+$ represents an element-wise operation setting negative values to 0 or a small parameter ϵ . After initialization, the iterative solver as stated above is performed until convergence criteria are met. However, due to its non-convexity convergence to the global minimum is not guaranteed, in contrast to for example PCA. Finally, note that the initialization of Y and Z has a large influence on the convergence behavior and its solution, such that it should be chosen appropriately¹².

Atoms Both dictionaries Y and Z contain non-negative atoms, which is useful for the interpretation of imaging measurement matrices. And although, originally proposed as learning by parts¹³ (see section 4.2 and 4.3.1) the uniqueness and correctness of those parts still remain open questions. A family of measurement matrices, i.e. the separable factorial articulation family, was proven to guarantee unique and “correct” solutions⁹¹. However, its family properties are difficult, if not impossible, to verify in a practical setting. Finally, additional constraints on the atoms can be added in the optimization, such as orthogonality or unit length requirements, to narrow down the solution space and improve the probability of its uniqueness¹².

Assumptions The main assumption is that an LMM along with its non-negative description is appropriate. Besides, one has to define the matrix rank a priori,

similar to most methods. Further assumptions are mostly application dependent and can be incorporated into the optimization such as the distribution of signal and noise terms and factors. Hence, the assumptions are also divergence or

Complexity Its time complexity largely depends on the initialization procedure, the used solver and the number of iterations until convergence. In an alternating least squares scheme, a least-square regression is applied twice along with a threshold step. Hence, it is of the order of $\mathcal{O}((2nm + r(n + m)/3)r^2i)$ flops, where i is the number of iterations and r is the assumed matrix rank. In case that $m \gg n$ and $r \ll n$, it reduces to $\mathcal{O}(2nmr^2i)$. The spatial complexity can be considered of $\mathcal{O}(mn)$.

Applications

- **Astronomical Imaging:** In a first paper⁴⁵, NMF is used for the analysis of molecular line spectra. The main idea is to exploit the ability of NMF to deal with the non-Gaussian nature of the profiles. The authors show that NMF is preferred over PCA or factor analysis and that its results are more meaningful. In a second paper⁴⁶, NMF is used in the context of direct imaging data of exoplanetary systems. Besides comparing its use for post-processing circumstellar disks, the authors show its use on Hubble Space Telescope STIS coronagraphic observations to extract circumstellar materials inside the primary ring. In a third paper⁴⁷, the authors use NMF to unmix sources in the context of MUSE, including prior information.
- **Imaging Mass Spectrometry:** In a first paper³⁸, NMF is compared and contrasted to ICA and PCA. A mouse cerebellum measurement matrix is used for comparison. NMF is mainly utilized to explore and identify underlying factors in the data. In a second paper⁴⁴, a non-negative least squares approach is pursued for the spatial querying of spatial/biochemical data coming from imaging mass spectrometry experiments.

4.3.3.7 Ordinary least squares

Originally attributed to Gauss⁹², ordinary least squares (OLS) is a type of least squares method primarily used in regression problems. Besides the ordinary type, with no weighting of features, weighted and regularized versions exist to cope with different assumptions, conditions or to preserve optimality properties. For example: the addition of ridge (see section 4.3.3.9) and LASSO⁹³ regularizers, but also the elastic net⁹⁴ to cope with collinearity and sparsity issues. Furthermore, its use is incorporated in more complex methods, e.g. NMF.

Methodology The essence of ordinary least squares, with its underlying LMM, is defining the error term in a least squares formulation as

$$\min_z \frac{1}{2} \|\mathbf{m} - Yz\|_2^2 \quad (4.22)$$

for vectors and

$$\min_Z \frac{1}{2} \left\| \left\| \mathbf{M} - YZ^T \right\|_F \right\|_F^2 \quad (4.23)$$

for matrices. In both cases, and under the assumption that the solution is bounded, i.e. $Y^T Y \succcurlyeq \mathbf{o}$, the solution can be formulated closed-form as

$$\begin{aligned} z &= (Y^T Y)^{-1} Y^T \mathbf{m}, \\ Z^T &= (Y^T Y)^{-1} Y^T \mathbf{M}. \end{aligned} \quad (4.24)$$

Atoms The atoms are captured by the predefined matrix Y , the dictionary (a.k.a. design matrix), while its weights, i.e. z and Z , are solved for in the above mentioned optimization. The latter can be seen as an orthogonal projection of the measurements onto the “solution” space, captured by Y . They are optimal in an ℓ_2 -sense.

Assumptions Most assumptions for the OLS are application dependent. However, some are recurring such as

- the underlying linear model is appropriate for the observed process;
- weak or strong exogeneity, e.g. $\mathbb{E}[\epsilon|Y] = \mathbf{o}$;
- no linear dependence of the regressors in Y , as such Y is full rank and its second-order moment is finite and positive semi-definite ($Y^T Y \succcurlyeq \mathbf{o}$);
- errors are homoscedastic and not autocorrelated (errors should be uncorrelated between different observations);
- normality of errors: $\epsilon|Y \sim \mathcal{N}(\mathbf{o}, \sigma^2 \mathbf{I})$;
- measurements are independent and identically distributed.

Complexity The time complexity and accuracy for this algorithm depends on the underlying method of normal equations. The method in the vector case requires $\mathcal{O}((m + k/3)k^2)$ flops⁴. For the matrix case, this scales linearly with the number of vectors in matrix Z , such that $\mathcal{O}(n(m + k/3)k^2)$ flops are required. The space complexity is no larger than $\mathcal{O}(mn)$.

Applications OLS and its variants have found their way in various application books ranging from system identification⁹⁵ to astronomy⁹⁶.

- **Astronomical Imaging:** In a first paper⁵⁰, the use of a sequential least squares fitting procedure is proposed for deconvolving radio astronomy images. In each iteration the response of the brightest source is removed, until the noisy residual image is left. In a second paper⁵¹, also on the subject of radio astronomy, a constrained least squares problem is solved for image formation using active set techniques. Furthermore, the authors show that its implementation via Krylov-subspace-based techniques allows for efficient recovery of solutions of the least squares problem.
- **Imaging Mass Spectrometry:** In a first paper⁴⁴, spatial querying of spatial/biochemical data of MALDI IMS nature is proposed via a non-negative least squares approach. In a second paper⁴⁹, a partial least squares approach is proposed for the detection and quantification of milk adulteration in MALDI experiments.

4.3.3.8 Principal Component Analysis

Principal component analysis (PCA) was independently discovered by Pearson⁵³ and Hotelling⁵², respectively in 1901 and 1933. In essence, PCA is a statistical method attempting to explain as much variance as possible in as few uncorrelated components as possible, in a minimum mean-squared sense. PCA minimizes the perpendicular distance between a data point and the principal component⁹⁷. This is in contrast to a least squares approach where the distance between measurement point and predicted point is minimized. In practice, it is often associated with the SVD and Karhunen-Loève Transform (KLT). Of the widely available literature we recommend an extensive review⁵⁴ of its history and use, a compact updated review⁹⁸ and a short refresher⁹⁹ for the interested reader.

The three Musketeers: SVD, KLT and PCA Different interpretations are given for the SVD, KLT and PCA. Gerbrands considers in his paper¹⁰⁰ the SVD as a purely deterministic method, that can be used for performing a KLT and PCA. Secondly, the SVDs truncation error is minimum in a least squares sense, while KLT and PCA, both statistical methods, are minimal in a mean-squared error sense. Following his interpretation, the KLT is originally described in digital signal processing literature and serves as an optimal orthogonal transform for signal representations. It therefore makes use of a basis transform that is obtained by an EVD of the covariance matrix of the measurements. This is similar to the transform performed by PCA, with the only difference that the latter is performed on mean-centered vectors, while for KLT this is not necessary. The mean-centering of the input matrix for PCA can be seen as a projection of the measurement matrix onto a space where the common offset is removed¹⁰¹.

“PCA is just SVD with some statistical window dressing.” - Schmidt¹⁰²

Methodology Classical PCA is performed on the covariance or correlation matrix of the measurement matrix and assumes the LMM to be underlying. Often, those population statistics are not known a priori, such that one falls back on the sample statistics with the underlying assumption of the sample representing the population in an adequate manner. As such, for an measurement matrix M , we define the transform on the covariance matrix as

$$\begin{aligned} P_{\perp} &= I - (\mathbf{1}\mathbf{1}^T/m)M, \\ X &= P_{\perp}M, \\ C_X &= X^T X, \\ C_X &= V\Sigma V^T, \end{aligned} \tag{4.25}$$

where $P_{\perp} \in \mathbb{R}^{m \times m}$ is an orthogonal projection, $\mathbf{1} \in \mathbb{R}^m$ is a vector with all entries equal to $\mathbf{1}$, $V \in \mathbb{R}^{n \times n}$ a rotation matrix containing the orthogonal eigenvectors of the EVD of $C_X \in \mathbb{R}^{n \times n}$. For the covariance matrix finally, $\Sigma \in \mathbb{R}^{n \times n}$ is a diagonal matrix containing the eigenvalues corresponding to the variances of the individual components. Note that the scaling in the covariance matrix was omitted. This does not influence the analysis however. Finally, we can truncate Σ to rank r , to account for a certain predefined variance and transform our original matrix by retrieving its corresponding column space

$$U = M V_r \Sigma_r^{-1/2}, \tag{4.26}$$

such that

$$M = U \Sigma_r^{1/2} V_r^T + B + \mathbf{1}\mu^T, \tag{4.27}$$

where B is a residual error matrix that consists of the residual terms of the truncation ($\sum_{i=r+1}^n \sqrt{\sigma_i} \mathbf{u}_i \mathbf{v}_i^T$) and $\mu = (\mathbf{1}/m)\mathbf{1}^T M$. A computationally less expensive⁴ and more accurate¹⁰³ route is preferred via a direct SVD or truncated SVD on the mean-centered measurement matrix X itself. PCA can then be formulated as

$$\begin{aligned} \min_{\Lambda} \quad & \|X - A\|_2 \\ \text{subj.} \quad & \mathbf{rank}(A) \leq k, \end{aligned} \tag{4.28}$$

or in a more classical formulation as

$$\begin{aligned} \max_z \quad & \|Xz\|_2^2 \\ \text{subj.} \quad & \|z\|_2 = \mathbf{1}. \end{aligned} \tag{4.29}$$

Finally, note that the EVD and SVD result in a focus on the diagonal elements of the covariance or correlation matrix, while the off-diagonal elements are not considered.

Atoms The atoms of this method correspond to the basis of the covariance or correlation matrix as retrieved by the EVD. The dictionary is hence represented by the r eigenvectors corresponding to the orthogonal directions of highest variance. As a result, the atoms are orthonormal with respect to each other. Through this procedure the individual atoms also provide an insight in what varies or correlates spatially as well as spectrally. This enables the exploration and analysis of underlying relations present in the measurement matrix. However, note that due to the nature of PCA negative correlations and thus negative entries are possible. The latter is not intuitive for imaging applications, which are inherently non-negative. Those negative-valued atoms can also result in negative values in the reconstruction. This is especially the case when the variance of signal is lower than the variance of noise, which most of the time result in the spread of information over several noisy atoms. Finally, note that the atoms are unique up to their sign⁴.

Assumptions The main assumption is that the underlying LMM is appropriate. Further assumptions depends on its particular use, such as dimensionality reduction, feature selection, i.e. the selection of the most useful variables, visualization of multi-dimensional data, identification of underlying variables, identification of groups of objects or of outliers¹². Depending on these applications the following assumptions can be related:

- sample statistics are mimicking population statistics;
- no outliers are present in the measurement matrix;
- noise behaves independent and identically distributed (i.i.d.);
- noise is uncorrelated with respect to signal;
- noise is uncorrelated with respect to itself;
- variance of atoms related to noise are lower than variance of atoms related to signal;
- noise distribution is symmetrical around zero, i.e. no systematic errors and error is homoscedastic¹⁰⁴.

Complexity The time-complexity is mainly driven by the underlying SVD, as such it is $\mathcal{O}(mn^2)$. The space complexity can be regarded this $\mathcal{O}(mn)$ when $r \ll n$ and $m \gg n$.

Applications

- **Astronomical Imaging:** PCA has a long legacy in astronomy and astronomical imaging. In a first paper⁵⁶, principal component analysis was used to estimate redshift of galaxies. It has also been used in the context of stellar classification in a second paper¹⁰⁵. In a third paper¹⁰⁶, PCA was utilized on 5869 galaxy spectra to classify the spectra in five classes that were further analyzed. Finally in an imaging related application, PCA has been proposed⁵⁷ for the removal of sky emission residuals.
- **Imaging Mass Spectrometry:** One of the first applications of PCA on multidimensional measurement matrices was through its application in genome-wide expression data¹⁰⁷ to provide so called eigengenes and eigenarrays that can be associated with observed genome-wide effects of regulators. In a second paper¹⁰⁸, PCA is used in the context of MALDI IMS in a combination with discriminant analysis *to identify changes in the mass spectra between regions*. In a third paper³⁸, PCA is compared and contrasted to ICA and NMF. A mouse cerebellum measurement matrix is used for comparison. In a final paper¹⁰⁹, PCA is used in the context of a MALDI IMS rat brain section data set of a Parkinson disease model and applied to reduce dimensionality.

4.3.3.9 Ridge Regression

Ridge regression (RR) dates back to the 1960s when it was proposed as an extension to OLS⁵⁹. Where the OLS can be regarded as an unbiased estimator, RR gives up this unbiasedness and trades it off with the estimator variance. This is especially of interest when the atoms in the dictionary are not orthogonal, such that the covariance matrix becomes ill-conditioned. For an extensive review of ridge regression with its bias and variance, we refer the reader to the following lecture notes¹¹⁰ and a more application-driven review¹¹¹.

Methodology Ridge regression, with underlying LMM, can be formulated as the following optimization

$$\min_z \frac{1}{2} \| \mathbf{m} - \mathbf{Y}z \|_2^2 + \lambda \| z \|_2^2 \quad (4.30)$$

for vectors or

$$\min_Z \frac{1}{2} \| \mathbf{M} - \mathbf{Y}Z^T \|_F^2 + \lambda \| \mathbf{Z} \|_F^2 \quad (4.31)$$

for matrices. One sees that a regularization term is added to the OLS minimization, such that the least squares solution is traded-off against a minimal z in ℓ_2 -sense. The regularization term $\lambda \| z \|_2^2$ is also called the ridge penalty. In both

cases, and under the assumption that the solution is bounded, i.e. $Y^T Y \succcurlyeq 0$, the solution is similar to OLS and can be formulated in closed-form as

$$\begin{aligned} z &= (Y^T Y + \lambda I)^{-1} Y^T m, \\ Z^T &= (Y^T Y + \Lambda)^{-1} Y^T M. \end{aligned} \quad (4.32)$$

The regularization parameters, i.e. λ and Λ are adjustable in such a way that the bias of the solution is enlarged, concurrent to a reduction in solution variance.

Atoms The atoms are captured by the predefined matrix Y , the dictionary (a.k.a. design matrix), while its weights, i.e. z and Z , are solved for in the above mentioned optimization. The orthogonal projection of the OLS is in this case regularized by a value λ or diagonal matrix Λ to prevent it from being ill-conditioned. From a heuristic point of view, this type of optimization will push the weights to become spread over most vector or matrix entries, such that large weights are punished (as explained in section 4.1.4).

Assumptions Most assumptions for the OLS hold for RR. However, the assumption of the atoms of the dictionary not to be collinear are relaxed and compensated by the regularization term.

Complexity The time complexity and accuracy for this algorithm depend on the underlying method of normal equations, similar to OLS. The method in the vector case requires approximately $\mathcal{O}((m + k/3)k^2)$ flops⁴. For the matrix case, this scales linearly with the number of vectors in matrix Z , such that around $\mathcal{O}(n(m + k/3)k^2)$ flops are required. The space complexity is slightly larger than $\Omega(mn + mk + kn)$.

Applications

- **Astronomical Imaging:** In this paper⁶⁰, ridge regression is used in the context of wavefront reconstruction, e.g. to remove deblurring and other atmospheric turbulence phenomena, to obtain high-resolution images of space objects from ground based telescopes.
- **Imaging Mass Spectrometry:** No applications were found.

4.3.3.10 Sparse Principal Component Analysis

Sparse principal component analysis (SPCA) was proposed in 2006 to overcome some of the problems of standard PCA⁶¹. The main goal of SPCA is to overcome atoms being formed as dense linear combinations of measurement properties, and as such prevent interpretability. It achieves this goal in a computationally less expensive manner as the earlier proposed SCoTLASS¹¹². Therefore, SPCA uses

a two-step alternating method, where first PCA is performed and secondly an elastic net problem is solved. These two steps continue iteratively until convergence is achieved. Finally, two recent papers^{113,114} propose a thorough review on the topic of SPCA, and as such are advised to the interested reader.

Methodology The methodology uses an underlying LMM and can be summarized in two steps. To initialize, PCA is applied on the measurement matrix M to group the first k principal components in Y . The first step is performed by solving the following elastic net problem

$$\begin{aligned} \min_{Y,Z} \quad & \left\| M - YZ^T M \right\|_F^2 + \lambda \|Z\|_F^2 + \|Z\theta\|_1 \\ \text{subj.} \quad & Y Y^T = I, \end{aligned} \quad (4.33)$$

where $\|\cdot\|_1$ is considered element-wise and $\theta \in \mathbb{R}^n$ are adjustable column weights. In practice, λ only serves for collinearity prevention⁶¹. On the other hand, θ can be adjusted and trades off the captured variance of the result and the sparsity of the atoms. Secondly, one fixes Z and solves for Y by performing an SVD on

$$(M^T M)Z = U \Sigma V^T \quad (4.34)$$

where $\tilde{A} = UV^T$ is reintroduced in the elastic net as Y . Those two steps are repeated until convergence is reached. Finally, the dictionary is normalized.

Atoms Sparsity can be traded off with the variance captured by the atoms and the sparsity of atoms themselves through the vector θ . In this method, the properties are hence sparsely represented in the atoms of Z . Note that their description is still dense in Y . As pointed out, the variance captured by this method is not straightforward as for PCA due to cross terms¹¹³. Hence, the atoms will not align with the directions of maximal variance as for PCA. Also note that atoms are not completely uncorrelated to each other anymore as the equality constraint in equation 4.33 is often relaxed or approximated. Finally, note that the uniqueness of the atoms is not guaranteed.

Assumptions No explicit assumptions are stated in the original paper⁷⁵. Implicit assumptions consist of the linear mixture-model to be valid and the sparse representation to be viable for the underlying measurement matrix. The latter also implies that a sparse representation is assumed to capture most of the variance of the original measurement matrix.

Complexity Its time complexity is extensively analyzed⁶¹. The two main drivers of the iterative complexity are the SVD that needs to be performed and the elastic net problem that needs to be solved. This results in the time complexity, when

- $n > m$: $\mathcal{O}(nm^2 + im^3)$, where i is the number of iterations.
- $m \gg n$: $\mathcal{O}(irmnJ + J^3)$, where i is the number of iterations, r the pre-defined rank, and J the number of nonzero coefficients for the elastic net.

Finally, its spatial complexity is $\mathcal{O}(mn)$, as it needs to store several variants in the size order of the original measurement matrix.

Applications

- **Astronomical Imaging:** In a single encountered paper⁶², a particular set of face-on, central galaxies is analyzed via a clustering method, incorporating SPCA, to identify galaxy groups of similar mass and morphology that can be linked through a particular life cycle and level of activity.
- **Imaging Mass Spectrometry:** No applications were found.

4.3.4 Robust Methodology

In this section we present the robust methodology in alphabetic order.

4.3.4.1 Block-Sparse and Low-Rank Matrix Decomposition

Robust principal component analysis based on block-sparse and low-rank matrix decomposition (RPCA-LBD) was first proposed in 2011 as an extension to PCP⁶³. In addition, in 2012 the outlier pursuit (OP) method⁶⁴ was proposed, based on the same principle: instead of considering all columns of M to contain outliers, it assumes that most columns of the error term are zero columns. The idea behind this rational is that when *few columns of the measurement matrix are generated by mechanisms different from the rest of the columns, the existence of these outlying columns tends to destroy to low-rank structure of the measurement matrix*⁶³. For example: corrupt or defect sensors, but also contaminated experiments or in collaborative filtering applications the existence of malicious agents⁶⁴.

Methodology RPCA-LBD is fairly similar to PCP and uses an underlying linear mixture model. The proposed convex program is

$$\begin{aligned} \min_{A, B} \quad & \|A\|_* + \kappa(1 - \lambda) \|A\|_{2,1} + \kappa\lambda \|B\|_{2,1} \\ \text{subj.} \quad & M = A + B. \end{aligned} \quad (4.35)$$

The term $\kappa(1 - \lambda) \|A\|_{2,1}$ ensures that the recovered A has exact zero columns corresponding to outliers. The augmented Lagrange multiplier methods or its inexact version is proposed to solve for this convex program. The main difference

to PCP is that the error term, in this case B , may have a low-rank structure, where PCP breaks under such conditions.

OP proposes two convex programs: (1) ideal case where no other noise components are required, i.e. LMM; (2) noisy case, when an extra noise component is assumed, i.e. XLMM. The respectively first convex program, is similar to RPCA-LBD and formulated as

$$\begin{aligned} \min_{A,B} \quad & \|A\|_* + \lambda \|B\|_{1,2} \\ \text{subj.} \quad & M = A + B. \end{aligned} \quad (4.36)$$

The respectively second convex program, is similar to the proposition of SPCP and formulated as

$$\begin{aligned} \min_{A,B} \quad & \|A\|_* + \lambda \|B\|_{1,2} \\ \text{subj.} \quad & \|M - A - B\|_F \leq \epsilon. \end{aligned} \quad (4.37)$$

In both cases λ is an optimization parameter. Finally, ϵ is adjustable to the noise level, similar to SPCP. Both OP methods are solved via a proximal gradient algorithm. Note the discrepancy in notation for RPCA-LBD and OP between the $\|\cdot\|_{1,2}$ and $\|\cdot\|_{2,1}$. In practice, the same optimization is pursued: the sum of ℓ_2 -column norms is calculated. See section 2.4.1 for a review on these type of norms.

Atoms The atoms of all methods described above will contain zeros in those entries corresponding to columns that are considered outliers. The atoms are retrieved by an SVD on low-rank matrix A . Given the incoherence principles (see assumptions), the atoms are highly unlikely to align with the Cartesian unit vectors. As such, the information in atoms is spread over the whole range of their entries. Furthermore, the low-rank description results in only few atoms to be retrieved. Finally, only OP guarantees that the atoms and dictionaries are unique.

Assumptions For both methods it is assumed that the matrix A is low-rank and that B contains only few non-zero columns that correspond to outliers. OP has an additional incoherence condition on its columns, relaxing the conditions of PCP to

$$\max_j \left\| V^T e_j \right\|_2^2 \leq \frac{\mu r}{n}, \quad (4.38)$$

where e_j is a Cartesian unit vector with the j -th entry equal to 1, μ is the incoherence, and r is the matrix rank. This assures the column support to be spread out, necessary for the exact retrieval. Furthermore, it is stated that OP *recovers the column space exactly and identifies exactly the indices of columns corresponding to outliers not lying in the recovered column space, as long as the fraction of corrupted*

points γ satisfies

$$\frac{\gamma}{1 - \gamma} \leq \frac{c}{\mu r}, \quad (4.39)$$

where c is a predefined constant⁶⁴. Finally, remark that the method analysis and proof of existence for OP is more rigorous than the proof given for RPCA-LBD.

Complexity No complexity study was performed for both methods, however, the implementation of RPCA-LBD was shown to require between 10 to 25 times more iterations than the standard PCP⁶³. Hence, its complexity might be problematic when scaling to larger measurement matrices.

Applications Although no applications papers were found, the application of these methods might be helpful for imaging applications such as imaging mass spectrometry for the detection of outlying ion images. In astronomical imaging, it could be used for the recovery from sky subtracting and residual errors.

4.3.4.2 Bounded Principal Component Pursuit

Robust principal component analysis with non-sparse errors or bounded principal component pursuit (BPCP) was proposed in 2019 from an economics point of view⁶⁵. The idea is to extend the existing body of theory around PCP to cope with different error models. Therefore, a theoretical argument is developed without sparsity assumption or the existence of moments of the error matrix. As a result it is able to deal with fat-tailed random error distributions, such as a Cauchy distribution, while maintaining the low-rank approximation.

Methodology BPCP, using an LMM as underlying model, consists of the following convex program

$$\begin{aligned} \min_{A, B} \quad & \|A\|_* + \lambda \|B\|_1 \\ \text{subj.} \quad & M = A + B, \\ & \|A\|_\infty \leq \alpha, \end{aligned} \quad (4.40)$$

where both $\|\cdot\|_1$ and $\|\cdot\|_\infty$ are considered element-wise. Remark that an additional constraint is proposed in comparison to PCP for technical reasons. It bounds the max entry of A by α and is allowed to be dependent on m and n . The augmented Lagrangian multiplier method is used for solving a relaxed problem where α is set to be a large number. In a second step it is verified whether the solution satisfies the inequality. Note, that the consistency between both the relaxed and non-relaxed programs is considered future work by the authors⁶⁵.

Atoms The atoms are retrieved by an SVD on low-rank matrix A , similar to PCP. Given the incoherence principles (see assumptions), the atoms are highly unlikely to align with the Cartesian unit vectors. As such, the information in atoms is spread over the whole range of their entries. Furthermore, the low-rank description results in only few atoms to be retrieved. Finally, the method guarantees that the atoms and dictionaries are unique.

Assumptions Two sets of assumptions are proposed. The first set of assumptions states that

1. the matrix A and B should be statistically independent;
2. the elements of B are independent and zero median;
3. the set of densities of the elements of B are equicontinuous and uniformly bounded away from 0 at 0.

The respectively second statement can be compared to the LAD estimator. For the latter to work the errors essentially only need to have zero median. The second set of assumptions consists of the incoherence principles, being

$$\begin{aligned} \max_i \left\| \mathbf{u}^T \mathbf{e}_i \right\|_2^2 &\leq \frac{\mu r}{m}, \\ \max_j \left\| \mathbf{v}^T \mathbf{e}_j \right\|_2^2 &\leq \frac{\mu r}{n}, \\ \left\| \mathbf{u} \mathbf{v}^T \right\|_\infty &\leq \sqrt{\frac{\mu r}{mn}}. \end{aligned} \quad (4.41)$$

with the addition of

$$\|A\|_\infty \leq \alpha. \quad (4.42)$$

Finally, α is allowed to increase to ∞ together with m and n .

Complexity No complexity analysis is available. However, for the complexity analysis of the augmented Lagrangian multiplier method for the same program see section 4.3.4.8.

Applications

- **Astronomical Imaging:** No applications were found.
- **Imaging Mass Spectrometry:** No applications were found.

4.3.4.3 Manhattan Non-Negative Matrix Factorization

Manhattan non-negative matrix factorization (MAHNMF) was first proposed⁶⁶ in 2008 as a semi-definite program. Its idea is to model the heavy tailed Laplace noise that is often related to sparse outliers via a Manhattan norm, i.e. ℓ_1 -norm. The latter replaces the Kullback-Leibler divergence or Euclidean distance that are typically used in NMF for the optimal handling of Poisson and Gaussian noise. This is achieved concurrently to the usual non-negativity constraint. However, due to its slow convergence and scalability issues it was not until 2012 that a practical optimization method was available⁶⁷. Therefore, a combined alternating minimization scheme was utilized. Besides, the authors also proposed extensions to MAHNMF, such as through the integration of box-constraints, manifold regularization, group sparsity, and the use of an elastic net induced MAHNMF and symmetric MAHNMF. Conclusively, a short performance review¹¹⁵ shows its ability to robustly deal with different noise types with respect to classical NMF methods. The latter also proposes errors bounds for MAHNMF.

Methodology The method, with underlying LMM, consists of minimizing the Manhattan distance or element-wise ℓ_1 -norm of

$$\begin{aligned} \min_{Y,Z} \quad & \left\| M - YZ^T \right\|_1 \\ \text{subj.} \quad & Y \geq 0, Z \geq 0. \end{aligned} \tag{4.43}$$

This optimization is solved via Nesterov's smoothing method combined with an alternating minimization strategy. The main idea is that the program in equation 4.43 can be minimized column-wise, where a non-negative least absolute deviation regression problem is solved. Besides, a rank-one residual iteration method is also proposed for small scale problems⁶⁷.

Atoms The individual entries of atoms are all non-negative. The latter makes them ideal for imaging applications, which are inherently non-negative. Furthermore, the dictionaries are retrieved in a robust manner, such that they are not prone to sparse outliers. It can furthermore be shown that they are maximum likelihood optimal with respect to Laplacian noise¹¹⁵. However, note that through the non-convexity of the bilinear optimization, the dictionaries and atoms are not unique. The latter is in correspondence to the classical NMF methods.

Assumptions It is assumed for MAHNMF that a low-rank LMM is able to describe the underlying relations of the measurement matrix. It can be viewed as the maximum likelihood estimator with respect to Laplacian noise, such that it performs optimal with respect to Laplacian noise¹¹⁵. Furthermore, it has been shown to perform well under a wide range of noise types: occlusions, salt and pepper and Gaussian noise¹¹⁵.

Complexity The time complexity of the i -th iteration of the alternating minimization scheme with Nesterov's smoothing is given to be $\mathcal{O}(2mn(r+1)\sqrt{i})$ flops, where r is the defined matrix rank. Therefore, the total time complexity is estimated to be $\mathcal{O}(2mn(r+1)\sum_{i=1}^T\sqrt{i})$ where T is the total number of iterations. It has been shown empirically, that $T \leq 100$ suffices for convergence⁶⁷. The spatial complexity is in the order of $\Omega(mn)$.

Applications

- **Astronomical Imaging:** No applications were found.
- **Imaging Mass Spectrometry:** No applications were found.

4.3.4.4 Principal Component Pursuit

Principal component pursuit (PCP) was proposed in 2009 as an idealized version of robust PCA^{7,68,69}. Its main idea consists of exactly decomposing a measurement matrix into a low-rank plus sparse error representation. This is achieved through a relaxed convex program, avoiding the use of the usual covariance or correlation matrix such as in PCA. The latter also differentiates it from previously defined robust PCA methods, e.g. ROBPCA¹¹⁶. This convex relaxation is achieved through both the relaxation of rank in nuclear norm¹¹⁷ and of the ℓ_0 -norm into the ℓ_1 -norm⁹¹. Both are discussed in a survey paper¹¹⁸, reviewing the mathematics of sparsity which lies at the basis of this research field. Since its discovery, many variations and adaptations have been proposed, such as SPCP and BPCP. We recommend an overview²⁰ encompassing a large number of methods and solvers based on PCP. Through the years, the theory behind PCP itself has been refined amongst others for exact retrieval with dense noise¹¹⁹, allowing for an arbitrary large but fixed fraction of corrupted entries, and with respect to the quality of approximations, when exact recovery is not guaranteed^{120,121}.

Methodology Suppose that a measurement matrix consists of the addition of a low-rank and sparse term where the rank of the low-rank component is not known a priori. Then, the exact terms can be recovered, through an LMM, using the optimization defined as

$$\begin{aligned} \min_{A,B} \quad & \text{rank}(A) + \gamma \|B\|_0 \\ \text{subj.} \quad & M = A + B, \end{aligned} \tag{4.44}$$

where γ is a regularization parameter. Given its non-convex nature, no exact nor efficient solution is known. Therefore, the following tractable and relaxed optimization^{7,68,69} is proposed

$$\begin{aligned} \min_{A,B} \quad & \|A\|_* + \lambda \|B\|_1 \\ \text{subj.} \quad & M = A + B, \end{aligned} \tag{4.45}$$

where $\lambda = 1/\sqrt{\max(m, n)}$ and which is shown to exactly recover the solution with high probability under certain conditions⁷. This program can be solved e.g. via an interior point method, semi-definite program⁶⁹ or augmented Lagrangian multiplier method¹²². Finally, note that λ can be tuned to achieve the best possible result⁷.

Atoms In contrast to classical PCA, principal component pursuit does not align the atoms towards the direction of largest variance. When the measurement matrix is mean-centered and/or whitened, PCP can be seen as a deterministic method on the covariance/correlation matrix. However, note that even in this situation the atoms will not align to the directions of maximal variance as proposed by PCA. However, PCP does offer the advantage of being robust against sparse outliers. The atoms are retrieved by an SVD on low-rank matrix A . Given the incoherence principles (see assumptions), the atoms are highly unlikely to align with the Cartesian unit vectors. As such, the information in atoms is spread over the whole range of their entries. Furthermore, the low-rank description results in only few atoms to be retrieved. Finally, the method guarantees that the atoms and dictionaries are unique.

Assumptions To make the problem meaningful, it is vital to impose that the low-rank term A is not sparse and vice versa that the sparse matrix is not low-rank. Therefore, two conditions are proposed⁷. The respectively first set of conditions (a.k.a. incoherence conditions) states that if $A = U\Sigma V^T = \sum_{i=1}^r \sigma_i u_i v_i^T$ where $A \in \mathbb{R}^{m \times n}$, we can define the incoherence condition which imposes that the low-rank component A is not sparse as

$$\begin{aligned} \max_i \left\| u^T e_i \right\|_2^2 &\leq \frac{\mu r}{m}, \\ \max_j \left\| v^T e_j \right\|_2^2 &\leq \frac{\mu r}{n}, \\ \left\| uv^T \right\|_\infty &\leq \sqrt{\frac{\mu r}{mn}}. \end{aligned} \tag{4.46}$$

The concept of incoherence was previously introduced in the context of matrix completion¹²³. Also note that these conditions are stated to be sufficient to ensure the exact solution recovery and have been refined in later papers^{120,121}.

“The role of the coherence in this theory is also very natural, and can be understood when thinking about the prediction of movie ratings. Here, we can imagine that the complete matrix of ratings has (approximately) low rank because users’ preferences are correlated. Now the reason why matrix completion is possible under incoherence is that we can exploit correlations and infer how a specific user is going to like a movie she has not yet seen, by examining her ratings and learning about her general preferences, and inferring how other users with such preferences have rated this particular item. Whenever we have users or small groups of users that are very singular in the sense that their ratings are orthogonal to those of all other users, it is not possible to correctly predict their missing entries. Such matrices have large coherence.” - Candes¹¹⁸

Secondly, one should also make sure that the sparse matrix is not low-rank. This is prevented through the second condition which assumes that the support of the sparsity pattern of the sparse term is selected uniformly at random. Note that in practice validating these assumptions often fails as the ground truth is required a priori. If both conditions are met, then exact recovery is guaranteed with probability at least $1 - c_1 m^{-10}$ provided that

$$\begin{aligned} \mathbf{rank}(A) &\leq \rho_r n \mu^{-1} (\log m)^{-2}, \\ c &\leq \rho_s m n \end{aligned} \tag{4.47}$$

where ρ_r and ρ_s are some positive constants, μ is the largest incoherence that fulfills the incoherence conditions, and c is the cardinality of the sparsity term.

Complexity The convex program can be solved by an interior point method, however its time complexity is too high for high-dimensional measurement matrices $\mathcal{O}(n^6)$. Other algorithms have been proposed such as an accelerated proximal gradient⁷, iterative thresholding and augmented Lagrange multiplier methods¹²². Furthermore, it was also proposed to reduce the complexity of the SVD to a thresholded SVD¹²². The inexact augmented Lagrangian multipliers method (IALM) has a complexity of $\mathcal{O}(mnr)$ flops per iteration where r is the estimated matrix rank. Heuristic findings (see chapter 6) show that the algorithm has often converged within 30 iterations. The spatial complexity is $\Omega(mn)$.

Applications

- **Astronomical Imaging:** In a first paper⁷⁰, principal component pursuit is used for exoplanet detection in high-contrast imaging. The authors compared different implementations for detection on synthetic and real data. In a second paper⁷¹, principal component pursuit is used for the retrieval

of the redshifted 21 cm signal from neutral hydrogen. As it is weak in comparison to foreground radiation from the Milky Way as well as extragalactic radio sources, they applied the method on the frequency covariance matrix to separate both foreground and the 21 cm signal.

- **Imaging Mass Spectrometry:** No applications were found.

4.3.4.5 Robust Dictionary Learning

Robust dictionary learning (RDL) was proposed in 2013 as a robust dictionary learning method⁷². The idea comes close to K-SVD⁷⁷ (see section 4.3.3.1), where beside the weights also the dictionary is updated, with the difference of modeling sparse outliers. Both an online and batch version are proposed⁷².

Methodology Similar to BP an underlying LMM is assumed. However, instead of using an ℓ_2 -norm, the ℓ_1 -norm is used on the error term. The following adjustment results in the following optimization formulation

$$\min_{Y, z} \|m - Yz\|_1 + \lambda \|z\|_1, \quad (4.48)$$

where λ is a regularization parameter. The suggested algorithm for solution retrieval is an alternating optimization where first the sparse representation, i.e. z , is sought, where-after, the dictionary is updated. Finally, note that the optimization formulation is closely related to the programs proposed in *Dense Error Correction via ℓ_1 -Minimization*¹²⁴.

Atoms The atoms of this method are preset by the dictionary, which will change over the optimization. Constructing an initial dictionary is hence important for the retrieval and sparsity of the solution, however, not as decisive as for BP. We refer to section 4.3.3.1 for a short review of possible dictionaries. Conclusively, the obtained sparse weights in z will not be unique.

Assumptions No assumptions are made. The only implicit assumption is that the dictionary is appropriate and that the measurement matrix can be described through an LMM.

Complexity No theoretical complexity analysis is performed.

Applications Although this method has no rigorous body of theory, the convex formulation might be of interest due to its connections to OLS, RR, BP, and LASSO.

- **Astronomical Imaging:** No applications were found.
- **Imaging Mass Spectrometry:** No applications were found.

4.3.4.6 Robust Non-Negative Matrix Factorization

Similar to MAHNMF, robust non-negative matrix factorization (RNMF) formulates a method that accounts for sparse outliers⁷³. However, in contrast to MAHNMF it is not maximum likelihood optimal for Laplacian noise. It can be considered as a combination of NMF with Euclidean distance measure and SPCP, including an additional noise term to account for dense noise.

Methodology RNMF, with its underlying XLMM, is formulated as

$$\begin{aligned} \min_{Y, Z, B} \quad & \left\| M - YZ^T - B \right\|_F^2 + \lambda \|B\|_1 \\ \text{subj.} \quad & Y \geq 0, Z \geq 0, \end{aligned} \quad (4.49)$$

where the $\|\cdot\|_1$ is an element-wise ℓ_1 -norm, and λ is a regularization parameter, similar to PCP. However, no particular setting is advised for λ . In the original paper⁷³, the authors have performed studies with a λ value between 0 and 2. An alternating non-negative quadratic program is proposed with soft-thresholding step to solve the program presented in equation 4.49.

Atoms The atoms and its entries will be non-negative, but not uniquely defined due to the bilinear formulation. The dictionary of Y will consist of unitary atoms. Furthermore, both the atoms of Y and Z are considered robust against sparse outliers and Gaussian noise. For an elaborate review of NMF atoms we refer to section 4.3.3.6.

Assumptions The sole implicit assumption is that the underlying XLMM is appropriate and that the non-negativity is appropriate for the application. No explicit assumptions are made in the original paper⁷³.

Complexity No complexity analysis is provided. Furthermore, no convergence rate is known such that we can only provide the iterative time complexity and spatial complexity. As previously mentioned, an alternating non-negative quadratic program is proposed with soft-thresholding step, where first Y and Z are solved for with fixed S followed by solving for S through a soft-thresholding operator. The alternating non-negative quadratic program can be best compared to the order of time complexity of NMF, hence reducing to $\mathcal{O}(nmr^2i)$, where r corresponds to the predefined matrix rank, and i accounts for the number of iterations required. The spatial complexity can be considered $\mathcal{O}(mn)$.

Applications

- **Astronomical Imaging:** No applications were found.
- **Imaging Mass Spectrometry:** No applications were found.

4.3.4.7 Side-Information Principal Component Pursuit

In many practical applications the assumptions of PCP are not met, resulting in degenerate or suboptimal solutions⁷⁴. Therefore, domain-dependent prior knowledge is introduced in the formulation to alleviate those assumptions^{125,126}. The original idea was to incorporate prior information in the form of noiseless column and/or row spaces, i.e. assumed to be exact. The latter can be compared to basis pursuit where the dictionary stays unchanged. It was not until 2017 that this idea was elaborated on to allow for noisy approximation⁷⁴.

Methodology Two convex programs within the framework of PCP are proposed. The first one assumes that the prior knowledge is given as an approximation matrix $W \in \mathbb{R}^{m \times n}$, such that

$$\begin{aligned} \min_{A,B} \quad & \|A\|_* + \kappa \|A - W\|_* + \lambda \|B\|_1 \\ \text{subj.} \quad & M = A + B. \end{aligned} \quad (4.50)$$

A second program adds row and column information in the form of $X \in \mathbb{R}^{m \times k}$ and $Y \in \mathbb{R}^{n \times k}$, such that

$$\begin{aligned} \min_{A,B} \quad & \|A\|_* + \kappa \|A - D\|_* + \lambda \|B\|_1 \\ \text{subj.} \quad & M = XAY^T + B, \\ & XWY^T = D. \end{aligned} \quad (4.51)$$

The latter incorporates the assumption that A and D are bilinear mappings of the recovered low-rank matrix $L = XAY^T$. In both notations κ and λ are hyperparameters. κ is considered as a quality measure of the available side information: when set to a large value the information is proper, while for a small value it is considered improper. The value of $\lambda = 1/\sqrt{\max(m, n)}$ is given as empirical guidance based on extensive experiments⁷⁴. Remark that when only the approximation matrix W is available, the second convex program can be recast into the respectively first program. For both programs an alternating direction method of multipliers is proposed as solver.

Atoms For the respectively first program, prior knowledge nudges the dictionaries, obtained via an SVD on the low-rank term A , into the direction of that actual prior knowledge matrix W . In the respectively second program, X and Y , the bilinear mappings, exploit prior knowledge in the form of column and row spaces. The latter also reduces the space where the low-rank component resides to $\mathbb{R}^{d \times d}$. For this program both dictionaries can be viewed as rotated and scaled versions of X and Y . Finally, it is suggested to learn those matrices by a dictionary learning approach¹²⁷.

Assumptions The LMM is assumed to be appropriate for this application, along with its low-rank and sparse interpretation. No further explicit assumptions are made, however, from a heuristic point of view it is assumed that the given side-information is to some extent proper for the application.

Complexity For both programs the time complexity is dominated by the SVD. This operation requires $\mathcal{O}(mn^2)$ flops and needs to be performed each iteration. Although no convergence analysis is performed, an indication of convergence is given for the KKT^a conditions which are fulfilled within a maximum of 1000 iterations. Note that this is far more than the complexity of PCP, i.e. a factor 30. Finally, assuming that $d \ll m$, the spatial complexity is $\mathcal{O}(mn)$.

Applications

- **Astronomical Imaging:** No applications were found.
- **Imaging Mass Spectrometry:** No applications were found.

4.3.4.8 Stable Principal Component Pursuit

Stable principal component pursuit (SPCP) is proposed as an extension to PCP by concurrently modeling small entry wise-noise⁷⁵. The latter is achieved by relaxing the original convex program. The idea is that classical PCA is able to deal with small entry-wise noise, where PCP is not. Hence, the authors coin it as *the first result that shows the classical PCA, optimal for small i.i.d. noise, can be made robust to gross sparse errors*.

Methodology The method, with underlying XLMM, is proposed as

$$\begin{aligned} \min_{\mathbf{A}, \mathbf{B}} \quad & \|\mathbf{A}\|_* + \lambda \|\mathbf{B}\|_1 \\ \text{subj.} \quad & \|\mathbf{M} - \mathbf{A} - \mathbf{B}\|_F \leq \delta, \end{aligned} \quad (4.52)$$

where λ is a regularization parameter, similar to PCP, and δ is a parameter that allows for the relaxation of exact recovery of $\mathbf{M} = \mathbf{A} + \mathbf{B}$. The authors show that when $\lambda = 1/\sqrt{\max(m, n)}$ a stable estimate of \mathbf{A} and \mathbf{B} is retrieved. Therefore, the dual is often solved, which is given by

$$\min_{\mathbf{A}, \mathbf{B}} \quad \|\mathbf{A}\|_* + \lambda \|\mathbf{B}\|_1 + \frac{1}{2\mu} \|\mathbf{M} - \mathbf{A} - \mathbf{B}\|_F^2, \quad (4.53)$$

instead of the primal, formulated in equation 4.52. For the latter, $\mu = \sqrt{2n}\sigma$ where each entry of the noise term \mathbf{C} is assumed to be i.i.d. $\sim \mathcal{N}(0, \sigma^2)$ ⁷⁵.

^aKarush-Kuhn-Tucker

Atoms The atoms of this method are, similarly to PCP, retrieved by an SVD on the low-rank matrix A . Again, given the incoherence principles (see assumptions), the atoms are highly unlikely to align with the Cartesian unit vectors. As such, the information in atoms is spread over the whole range of their entries. Furthermore, the low-rank description results in only few atoms to be retrieved and guarantees the atoms and dictionaries to be unique. Beside the robustness against outliers, the atoms and dictionaries are shown to be stable against small entry-wise noise in the measurement matrix. Although the method is coined as a PCA robust against outliers, the atoms are not aligned along the directions of highest variance as for PCA. Hence, we consider it as an extension to PCP rather than PCA.

Assumptions Since its results are based on PCP, its assumptions are similar. The same incoherence conditions are required along with conditions on the support of the entries of the sparse term. In addition, two conditions are proposed. The first assumption is on the entry-wise noise term. It states that C should i.i.d. on each entry as small for some value $\delta \leq \epsilon$. If all these conditions are met, then SPCP satisfies

$$\|A_{\text{exact}} - A\|_F^2 + \|B_{\text{exact}} - B\|_F^2 \leq Cmn\delta^2 \quad (4.54)$$

for some constant C with high probability.

Complexity The original paper⁷⁵ proposed an accelerated proximal gradient method, but is outperformed by an alternating direction method of multipliers^{128,129}. For the latter no complexity analysis is provided, however, the authors show through computational tests that the iterative time complexity is of the order of an SVD. Hence, the estimated time complexity reduces to $\mathcal{O}(mn^2)$ per iteration. However, through the use of a truncated SVD that can be lowered. From a practical point of view, we note that a vector sorting is performed each iteration on the k largest values of a vector of length mn . Its complexity corresponds to $\mathcal{O}(k \log k)$, where k corresponds to the cardinality of the sparse term. In case k becomes large this might overtake the complexity of the SVD.

Applications

- **Astronomical Imaging:** No applications were found.
- **Imaging Mass Spectrometry:** No applications were found.

4.4 Summary

In this literature review, we introduced the basics of matrix decompositions from the perspective of atoms. We explained the concept of ℓ_1 -robustness with respect

to sparse outliers and introduced the family of extended linear mixture models together with its validity for imaging applications. Furthermore, 18 different classical and robust matrix decomposition methods were analyzed in the framework of atoms, along with the presentation of some of their applications in imaging context, in particular in the fields of imaging mass spectrometry and astronomical imaging.

The novelty of this literature review lies in the inclusion of the framework of atoms into the broader context of matrix decompositions for analysis and comparison of methods. A second accomplishment was the discovery of promising classical methods that have an application legacy in only one of the considered imaging modalities, but not in both. This knowledge can be transferred. Thirdly, the review of ℓ_1 -robust matrix decomposition methodology has indicated a substantial gap between method literature itself and application-driven literature from both fields. This knowledge gap can be closed. Hence, we conclude this review with the acquisition of two possible research directions.

Bibliography

- [1] Cornett, D. S., Reyzer, M. L., Chaurand, P. & Caprioli, R. M. Maldi imaging mass spectrometry: molecular snapshots of biochemical systems. *Nature methods* **4**, 828–833 (2007).
- [2] Bennett, J., Lanning, S. *et al.* The netflix prize. In *Proceedings of KDD cup and workshop*, vol. 2007, 35 (Citeseer, 2007).
- [3] Fan, J., Han, F. & Liu, H. Challenges of big data analysis. *National science review* **1**, 293–314 (2014).
- [4] Golub, G. H. & Van Loan, C. F. *Matrix computations*, vol. 3 (JHU press, 2012).
- [5] Wisotzki, L. *et al.* Nearly all the sky is covered by lyman- α emission around high-redshift galaxies. *Nature* **562**, 229–232 (2018).
- [6] Chen, S. & Donoho, D. Basis pursuit. In *Proceedings of 1994 28th Asilomar Conference on Signals, Systems and Computers*, vol. 1, 41–44 (IEEE, 1994).
- [7] Candès, E. J., Li, X., Ma, Y. & Wright, J. Robust principal component analysis? *Journal of the ACM (JACM)* **58**, 1–37 (2011).
- [8] Donoho, D. L. For most large underdetermined systems of linear equations the minimal ℓ_1 -norm solution is also the sparsest solution. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences* **59**, 797–829 (2006).
- [9] Gnanadesikan, R. & Kettenring, J. R. Robust estimates, residuals, and outlier detection with multiresponse data. *Biometrics* 81–124 (1972).
- [10] Cichocki, A. & Amari, S.-i. Families of alpha-beta-and gamma-divergences: Flexible and robust measures of similarities. *Entropy* **12**, 1532–1568 (2010).
- [11] Hurley, N. & Rickard, S. Comparing measures of sparsity. *IEEE Transactions on Information Theory* **55**, 4723–4741 (2009).
- [12] Cichocki, A., Zdunek, R., Phan, A. H. & Amari, S.-i. *Nonnegative matrix and tensor factorizations: applications to exploratory multi-way data analysis and blind source separation* (John Wiley & Sons, 2009).
- [13] Lee, D. D. & Seung, H. S. Learning the parts of objects by non-negative matrix factorization. *Nature* **401**, 788–791 (1999).
- [14] Blumer, A., Ehrenfeucht, A., Haussler, D. & Warmuth, M. K. Occam's razor. *Information processing letters* **24**, 377–380 (1987).

- [15] Cox, D. R. & Hinkley, D. V. *Theoretical statistics* (CRC Press, 1979).
- [16] Van Der Maaten, L., Postma, E. & Van den Herik, J. Dimensionality reduction: a comparative. *J Mach Learn Res* **10**, 13 (2009).
- [17] Mallat, S. G. & Zhang, Z. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on signal processing* **41**, 3397–3415 (1993).
- [18] Chen, S. S., Donoho, D. L. & Saunders, M. A. Atomic decomposition by basis pursuit. *SIAM review* **43**, 129–159 (2001).
- [19] Verbeek, N., Caprioli, R. M. & Van de Plas, R. Unsupervised machine learning for exploratory data analysis in imaging mass spectrometry. *Mass spectrometry reviews* (2019).
- [20] Bouwmans, T., Aybat, N. S. & Zahzah, E.-h. *Handbook of robust low-rank and sparse matrix decomposition: Applications in image and video processing* (CRC Press, 2016).
- [21] Beckouche, S., Starck, J.-L. & Fadili, J. Astronomical image denoising using dictionary learning. *Astronomy & Astrophysics* **556**, A132 (2013).
- [22] Fotiadou, K., Tsagkatakis, G., Moraes, B., Abdalla, F. B. & Tsakalides, P. Denoising galaxy spectra with coupled dictionary learning. In *2017 25th European Signal Processing Conference (EUSIPCO)*, 498–502 (IEEE, 2017).
- [23] Bartels, A., Trede, D., Alexandrov, T. & Maass, P. Hybrid regularization and sparse reconstruction of imaging mass spectrometry data. In *10th international conference on Sampling Theory and Applications (SampTA 2013)*, 189–192 (2013).
- [24] Chen, S. S. & Donoho, D. L. Application of basis pursuit in spectrum estimation. In *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP'98 (Cat. No. 98CH36181)*, vol. 3, 1865–1868 (IEEE, 1998).
- [25] Wiaux, Y., Jacques, L., Puy, G., Scaife, A. M. & Vandergheynst, P. Compressed sensing imaging techniques for radio interferometry. *Monthly Notices of the Royal Astronomical Society* **395**, 1733–1742 (2009).
- [26] Harn, Y.-C., Powers, M., Shank, E. A. & Jovic, V. Deconvolving molecular signatures of interactions between microbial colonies. *Bioinformatics* **31**, i142–i150 (2015).
- [27] Koles, Z. J., Lazar, M. S. & Zhou, S. Z. Spatial patterns underlying population differences in the background eeg. *Brain topography* **2**, 275–284 (1990).

- [28] Shapiro, J., Savransky, D., Ruffio, J.-B., Ranganathan, N. & Macintosh, B. Detecting planets from direct-imaging observations using common spatial pattern filtering. *The Astronomical Journal* **158**, 125 (2019).
- [29] Shapiro, J., Ranganathan, N., Savransky, D., Ruffio, J.-B. & Macintosh, B. Planet signal extraction from direct imaging using common spatial pattern filtering. In *Techniques and Instrumentation for Detection of Exoplanets VIII*, vol. 10400, 104001O (International Society for Optics and Photonics, 2017).
- [30] Shapiro, J., Savransky, D., Ruffio, J.-B. & Macintosh, B. Common spatial pattern filtering for detection of circumstellar discs. In *Optical and Infrared Interferometry and Imaging VI*, vol. 10701, 107012G (International Society for Optics and Photonics, 2018).
- [31] Drineas, P., Mahoney, M. W. & Muthukrishnan, S. Relative-error cur matrix decompositions. *SIAM Journal on Matrix Analysis and Applications* **30**, 844–881 (2008).
- [32] Yang, J., Rubel, O., Mahoney, M. W. & Bowen, B. P. Identifying important ions and positions in mass spectrometry imaging data using cur matrix decompositions. *Analytical chemistry* **87**, 4658–4666 (2015).
- [33] Gittens, A. *et al.* Matrix factorizations at scale: A comparison of scientific data analytics in spark and c+ mpi using three case studies. In *2016 IEEE International Conference on Big Data (Big Data)*, 204–213 (IEEE, 2016).
- [34] Jutten, C. & Herault, J. Blind separation of sources, part i: An adaptive algorithm based on neuromimetic architecture. *Signal processing* **24**, 1–10 (1991).
- [35] Comon, P. Independent component analysis, a new concept? *Signal processing* **36**, 287–314 (1994).
- [36] Delabrouille, J., Cardoso, J.-F. & Patanchon, G. Multidetector multicomponent spectral matching and applications for cosmic microwave background data analysis. *Monthly Notices of the Royal Astronomical Society* **346**, 1089–1102 (2003).
- [37] Hanselmann, M. *et al.* Concise representation of mass spectrometry images by probabilistic latent semantic analysis. *Analytical chemistry* **80**, 9649–9658 (2008).
- [38] Siy, P. W. *et al.* Matrix factorization techniques for analysis of imaging mass spectrometry data. In *2008 8th IEEE International Conference on BioInformatics and BioEngineering*, 1–6 (IEEE, 2008).

- [39] Gut, Y. *et al.* Application of chemometric algorithms to maldi mass spectrometry imaging of pharmaceutical tablets. *Journal of pharmaceutical and biomedical analysis* **105**, 91–100 (2015).
- [40] Switzer, P. Min/max autocorrelation factors for multivariate spatial imagery. *Computer science and statistics* (1985).
- [41] Jones, E. A. *et al.* Multiple statistical analysis techniques corroborate intratumor heterogeneity in imaging mass spectrometry datasets of myxofibrosarcoma. *PLoS one* **6**, e24913 (2011).
- [42] Balluff, B. *et al.* De novo discovery of phenotypic intratumour heterogeneity using imaging mass spectrometry. *The Journal of pathology* **235**, 3–13 (2015).
- [43] Paatero, P. & Tapper, U. Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics* **5**, 111–126 (1994).
- [44] Van de Plas, R., Pelckmans, K., De Moor, B. & Waelkens, E. Spatial querying of imaging mass spectrometry data: a nonnegative least squares approach. In *Neural Information Processing Systems Workshop on Machine Learning in Computational Biology* (2007).
- [45] Juvela, M., Lehtinen, K. & Paatero, P. The use of positive matrix factorization in the analysis of molecular line spectra. *Monthly Notices of the Royal Astronomical Society* **280**, 616–626 (1996).
- [46] Ren, B., Pueyo, L., Zhu, G. B., Debes, J. & Duchêne, G. Non-negative matrix factorization: robust extraction of extended structures. *The Astrophysical Journal* **852**, 104 (2018).
- [47] Meganem, I., Deville, Y., Hosseini, S., Carfantan, H. & Karoui, M. S. Extraction of stellar spectra from dense fields in hyperspectral data cubes using non-negative matrix factorization. In *2011 3rd Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, 1–4 (IEEE, 2011).
- [48] Gauss, F. Theory of the motion of heavenly bodies moving about the sun in conic sections (english transl. by ch davis), reprinted 1963 (1809).
- [49] Nicolaou, N., Xu, Y. & Goodacre, R. Maldi-ms and multivariate analysis for the detection and quantification of different milk species. *Analytical and bioanalytical chemistry* **399**, 3491–3502 (2011).
- [50] Högbom, J. Aperture synthesis with a non-regular distribution of interferometer baselines. *Astronomy and Astrophysics Supplement Series* **15**, 417 (1974).

- [51] Sardarabadi, A. M., Leshem, A. & van der Veen, A.-J. Radio astronomical image formation using constrained least squares and krylov subspaces. *Astronomy & Astrophysics* **588**, A95 (2016).
- [52] Hotelling, H. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology* **24**, 417 (1933).
- [53] Pearson, K. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **2**, 559–572 (1901).
- [54] Jolliffe, I. T. Principal components in regression analysis. In *Principal component analysis*, 129–155 (Springer, 1986).
- [55] Marchetti, A. *et al.* The vimos public extragalactic redshift survey (vipers): spectral classification through principal component analysis. *Monthly Notices of the Royal Astronomical Society* **428**, 1424–1437 (2013).
- [56] Glazebrook, K., Offer, A. R. & Deeley, K. Automatic redshift determination by use of principal component analysis. i. fundamentals. *The Astrophysical Journal* **492**, 98 (1998).
- [57] Soto, K. T., Lilly, S. J., Bacon, R., Richard, J. & Conseil, S. Zap-enhanced pca sky subtraction for integral field spectroscopy. *Monthly Notices of the Royal Astronomical Society* **458**, 3210–3220 (2016).
- [58] Van de Plas, R. *et al.* Prospective exploration of biochemical tissue composition via imaging mass spectrometry guided by principal component analysis. In *Biocomputing 2007*, 458–469 (World Scientific, 2007).
- [59] Hoerl, A. E. & Kennard, R. W. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* **12**, 55–67 (1970).
- [60] Chu, Q., Jefferies, S. & Nagy, J. G. Iterative wavefront reconstruction for astronomical imaging. *SIAM Journal on Scientific Computing* **35**, S84–S103 (2013).
- [61] Zou, H., Hastie, T. & Tibshirani, R. Sparse principal component analysis. *Journal of computational and graphical statistics* **15**, 265–286 (2006).
- [62] Yesuf, H. M. & Ho, L. C. Gas content regulates the life cycle of star formation and black hole accretion in galaxies. *The Astrophysical Journal* **901**, 42 (2020).
- [63] Tang, G. & Nehorai, A. Robust principal component analysis based on low-rank and block-sparse matrix decomposition. In *2011 45th Annual Conference on Information Sciences and Systems*, 1–5 (IEEE, 2011).

- [64] Xu, H., Caramanis, C. & Sanghavi, S. Robust pca via outlier pursuit. *IEEE transactions on information theory* **58**, 3047–3064 (2012).
- [65] Bai, J. & Feng, J. Robust principal component analysis with non-sparse errors. *arXiv preprint arXiv:1902.08735* (2019).
- [66] Lam, E. Y. Non-negative matrix factorization for images with laplacian noise. In *APCCAS 2008-2008 IEEE Asia Pacific Conference on Circuits and Systems*, 798–801 (IEEE, 2008).
- [67] Guan, N., Tao, D., Luo, Z. & Shawe-Taylor, J. Mahnmf: Manhattan non-negative matrix factorization. *arXiv preprint arXiv:1207.3438* (2012).
- [68] Wright, J., Ganesh, A., Rao, S., Peng, Y. & Ma, Y. Robust principal component analysis: Exact recovery of corrupted low-rank matrices via convex optimization. In *Advances in neural information processing systems*, 2080–2088 (2009).
- [69] Chandrasekaran, V., Sanghavi, S., Parrilo, P. A. & Willsky, A. S. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization* **21**, 572–596 (2011).
- [70] Gonzalez, C. G. *et al.* Low-rank plus sparse decomposition for exoplanet detection in direct-imaging adi sequences-the llsg algorithm. *Astronomy & Astrophysics* **589**, A54 (2016).
- [71] Zuo, S., Chen, X., Ansari, R. & Lu, Y. 21 cm signal recovery via robust principal component analysis. *The Astronomical Journal* **157**, 4 (2018).
- [72] Lu, C., Shi, J. & Jia, J. Online robust dictionary learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 415–422 (2013).
- [73] Zhang, L., Chen, Z., Zheng, M. & He, X. Robust non-negative matrix factorization. *Frontiers of Electrical and Electronic Engineering in China* **6**, 192–200 (2011).
- [74] Xue, N., Panagakis, Y. & Zafeiriou, S. Side information in robust principal component analysis: Algorithms and applications. In *Proceedings of the IEEE International Conference on Computer Vision*, 4317–4325 (2017).
- [75] Zhou, Z., Li, X., Wright, J., Candes, E. & Ma, Y. Stable principal component pursuit. In *2010 IEEE international symposium on information theory*, 1518–1522 (IEEE, 2010).
- [76] Starck, J.-L., Moudden, Y., Bobin, J., Elad, M. & Donoho, D. Morphological component analysis. In *Wavelets XI*, vol. 5914, 59140Q (International Society for Optics and Photonics, 2005).

- [77] Aharon, M., Elad, M. & Bruckstein, A. K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on signal processing* **54**, 4311–4322 (2006).
- [78] Tomic, I. & Frossard, P. Dictionary learning. *IEEE Signal Processing Magazine* **28**, 27–38 (2011).
- [79] Tong, T., Caballero, J., Bhatia, K. & Rueckert, D. Dictionary learning for medical image denoising, reconstruction, and segmentation. In *Machine learning and medical imaging*, 153–181 (Elsevier, 2016).
- [80] Donoho, D. L., Elad, M. & Temlyakov, V. N. Stable recovery of sparse overcomplete representations in the presence of noise. *IEEE Transactions on information theory* **52**, 6–18 (2005).
- [81] Huang, J., Berger, C. R., Zhou, S. & Huang, J. Comparison of basis pursuit algorithms for sparse channel estimation in underwater acoustic ofdm. In *OCEANS'10 IEEE SYDNEY*, 1–6 (IEEE, 2010).
- [82] Fukunaga, K. *Introduction to statistical pattern recognition* (Elsevier, 2013).
- [83] Hyvarinen, A. Fast and robust fixed-point algorithms for independent component analysis. *IEEE transactions on Neural Networks* **10**, 626–634 (1999).
- [84] Hyvärinen, A. Independent component analysis: recent advances. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **371**, 20110534 (2013).
- [85] Zarzoso, V., Comon, P. & Kallel, M. How fast is fastica? In *2006 14th European Signal Processing Conference*, 1–5 (IEEE, 2006).
- [86] Keenan, M. R. & Smentkowski, V. S. Simple statistically based alternatives to maf for tof-sims spectral image analysis. *Surface and interface analysis* **43**, 1616–1626 (2011).
- [87] Larsen, R. Decomposition using maximum autocorrelation factors. *Journal of Chemometrics: A Journal of the Chemometrics Society* **16**, 427–435 (2002).
- [88] Gillis, N. The why and how of nonnegative matrix factorization. *Regularization, optimization, kernels, and support vector machines* **12**, 257–291 (2014).
- [89] Lawton, W. H. & Sylvestre, E. A. Self modeling curve resolution. *Technometrics* **13**, 617–633 (1971).

- [90] Albuquerque, C. D. & Poppi, R. J. Detection of malathion in food peels by surface-enhanced raman imaging spectroscopy and multivariate curve resolution. *Analytica Chimica Acta* **879**, 24–33 (2015).
- [91] Donoho, D. & Stodden, V. When does non-negative matrix factorization give a correct decomposition into parts? *Advances in neural information processing systems* **16**, 1141–1148 (2003).
- [92] Bühler, W. K. *Gauss: a biographical study* (Springer Science & Business Media, 2012).
- [93] Tibshirani, R. The lasso method for variable selection in the cox model. *Statistics in medicine* **16**, 385–395 (1997).
- [94] Zou, H. & Hastie, T. Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)* **67**, 301–320 (2005).
- [95] Bruls, J., Chou, C. T., Haverkamp, B. & Verhaegen, M. Linear and non-linear system identification using separable least-squares. *European Journal of Control* **5**, 116–128 (1999).
- [96] Isobe, T., Feigelson, E. D., Akritas, M. G. & Babu, G. J. Linear regression in astronomy. *The astrophysical journal* **364**, 104–113 (1990).
- [97] Lever, J., Krzywinski, M. & Altman, N. Points of significance: Principal component analysis (2017).
- [98] Jolliffe, I. T. & Cadima, J. Principal component analysis: a review and recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **374**, 20150202 (2016).
- [99] Lever, J., Krzywinski, M. & Altman, N. Points of significance: model selection and overfitting (2016).
- [100] Gerbrands, J. J. On the relationships between svd, klt and pca. *Pattern recognition* **14**, 375–381 (1981).
- [101] Bro, R. & Smilde, A. K. Centering and scaling in component analysis. *Journal of Chemometrics* **17**, 16–33 (2003).
- [102] Schmidt, D. Introducing float: Single precision floats for r (2017).
- [103] Cardoso, J.-F., Delabrouille, J. & Patanchon, G. Independent component analysis of the cosmic microwave background. In *4th International Symposium on Independent Component Analysis and Blind Signal Separation (ICA03)* (2003).

- [104] Bro, R. & Smilde, A. K. Principal component analysis. *Analytical Methods* **6**, 2812–2831 (2014).
- [105] Singh, H. P., Gulati, R. K. & Gupta, R. Stellar spectral classification using principal component analysis and artificial neural networks. *Monthly Notices of the Royal Astronomical Society* **295**, 312–318 (1998).
- [106] Folkes, S. *et al.* The 2df galaxy redshift survey: spectral types and luminosity functions. *Monthly Notices of the Royal Astronomical Society* **308**, 459–472 (1999).
- [107] Alter, O., Brown, P. O. & Botstein, D. Singular value decomposition for genome-wide expression data processing and modeling. *Proceedings of the National Academy of Sciences* **97**, 10101–10106 (2000).
- [108] McCombie, G., Staab, D., Stoeckli, M. & Knochenmuss, R. Spatial and spectral correlations in maldi mass spectrometry images by clustering and multivariate analysis. *Analytical chemistry* **77**, 6118–6124 (2005).
- [109] Verbeeck, N. *et al.* Connecting imaging mass spectrometry and magnetic resonance imaging-based anatomical atlases for automated anatomical interpretation and differential analysis. *Biochimica et Biophysica Acta (BBA)-Proteins and Proteomics* **1865**, 967–977 (2017).
- [110] van Wieringen, W. N. Lecture notes on ridge regression. *arXiv preprint arXiv:1509.09169* (2015).
- [111] Draper, N. R. & Smith, H. *Applied regression analysis*, vol. 326 (John Wiley & Sons, 1998).
- [112] Jolliffe, I. T., Trendafilov, N. T. & Uddin, M. A modified principal component technique based on the lasso. *Journal of computational and Graphical Statistics* **12**, 531–547 (2003).
- [113] Camacho, J., Smilde, A., Saccenti, E. & Westerhuis, J. All sparse pca models are wrong, but some are useful. part i: Computation of scores, residuals and explained variance. *Chemometrics and Intelligent Laboratory Systems* **196**, 103907 (2020).
- [114] Camacho, J., Smilde, A., Saccenti, E., Westerhuis, J. & Bro, R. All sparse pca models are wrong, but some are useful. part ii: Limitations and problems of deflation. *Chemometrics and Intelligent Laboratory Systems* 104212 (2020).
- [115] Liu, T. & Tao, D. On the performance of manhattan nonnegative matrix factorization. *IEEE transactions on neural networks and learning systems* **27**, 1851–1863 (2015).

- [116] Hubert, M., Rousseeuw, P. J. & Vanden Branden, K. Robpca: a new approach to robust principal component analysis. *Technometrics* **47**, 64–79 (2005).
- [117] Fazel, M. Matrix rank minimization with applications (2002).
- [118] Candès, E. J. Mathematics of sparsity (and a few other things). In *Proceedings of the International Congress of Mathematicians, Seoul, South Korea*, vol. 123 (Citeseer, 2014).
- [119] Ganesh, A., Wright, J., Li, X., Candès, E. J. & Ma, Y. Dense error correction for low-rank matrices via principal component pursuit. In *2010 IEEE international symposium on information theory*, 1513–1517 (IEEE, 2010).
- [120] Hsu, D., Kakade, S. M. & Zhang, T. Robust matrix decomposition with sparse corruptions. *IEEE Transactions on Information Theory* **57**, 7221–7234 (2011).
- [121] Agarwal, A., Negahban, S., Wainwright, M. J. *et al.* Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions. *The Annals of Statistics* **40**, 1171–1197 (2012).
- [122] Lin, Z., Chen, M. & Ma, Y. The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. *arXiv preprint arXiv:1009.5055* (2010).
- [123] Candès, E. J., Romberg, J. & Tao, T. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on information theory* **52**, 489–509 (2006).
- [124] Wright, J. & Ma, Y. Dense error correction via ℓ^1 -minimization. *IEEE Transactions on Information Theory* **56**, 3540–3560 (2010).
- [125] Sagonas, C., Panagakis, Y., Zafeiriou, S. & Pantic, M. Raps: Robust and efficient automatic construction of person-specific deformable models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1789–1796 (2014).
- [126] Chiang, K.-Y., Hsieh, C.-J. & Dhillon, I. Robust principal component analysis with side information. In *International Conference on Machine Learning*, 2291–2299 (2016).
- [127] Patel, V. M., Wu, T., Biswas, S., Phillips, P. J. & Chellappa, R. Dictionary-based face recognition under variable lighting and pose. *IEEE Transactions on Information Forensics and Security* **7**, 954–965 (2012).
- [128] Aybat, N. S., Goldfarb, D. & Iyengar, G. Fast first-order methods for stable principal component pursuit. *arXiv preprint arXiv:1105.2126* (2011).

- [129] Aybat, N. S. & Iyengar, G. An alternating direction method with increasing penalty for stable principal component pursuit. *Computational Optimization and Applications* **61**, 635–668 (2015).

Part II

Research

Chapter 5

A Comparative Study under Inexact Recovery Conditions

After this chapter you will be able to

- apply ℓ_1 -robustness theory in a practical set-up
- understand several (in)exact recovery conditions
- recognize different quantitative and qualitative measures
- recognize shortcomings and breakdown phenomena
- reproduce this study for other measurement matrices

5.1 Introduction

This chapter translates the theory, presented in the literature review, to a practical level. Therefore, we review a selection of methods in terms of their theoretical conditions through the use of a synthetic measurement matrix.

5.1.1 Goals and Objectives

The aim of this study is to get a better understanding of robust methods' principles through quantitative and qualitative inspection under different inexact recovery conditions on a synthetic measurement matrix. Therefore, we will use classical methodology as yardstick, not to claim the robust methods' superiority, but rather to give perspective and a way to relate its performance to more broadly studied methods. We consider direct comparison of classical and robust methods to be often sensitive to bias, with the synthetic measurement matrix often

assumed to be close to reality, and with claiming no superiority in general but only in a few limited cases. Rather, our objective is to focus on inexact recovery conditions and the existence of their corresponding breakdown phenomena. We inspect the numerical and visual manifestation of such discrepancies between what the methods expect and what they are confronted with in a real-world data set by means of a synthetic measurement matrix that can be adjusted as needed, and relate those observations to a widely used and well studied yardstick. By no means do we claim to list such phenomena exhaustively, nor to specify their exact boundaries of occurrence. In the end, the obtained numerical and qualitative phenomena remain measurement matrix-specific. The idea, however, is to make the reader aware of some of the mechanisms underlying phenomena that arise under inexact recovery conditions, such as those conditions that can arise in case for real measurement matrices. These findings subsequently answer research subquestion A.1.

A. Can matrix decomposition methods based on the ℓ_1 -robustness principles efficiently reduce massively multivariate and high-dimensional spectral imaging measurement matrices to a lower-dimensional representation while incurring minimal information loss?

A.1. How do these methods perform with respect to classical dimensionality reduction methods under inexact recovery conditions: (a) quantitatively, i.e. in terms of efficiency, compression ratio and information loss; (b) qualitatively, i.e. what break down phenomena appear in the atoms, dictionaries, and matrix terms?

5.1.2 Overview

First, the experimental design is discussed in the following section. This includes the selection of methods, the generation of a synthetic low-rank matrix and noise term(s), the introduction to multiple measures for quantitative and qualitative performance analysis, and the provision of additional settings and parameters of this study. In section 5.3, the quantitative results are presented, followed by a qualitative performance analysis in section 5.4. Finally, we draw our conclusions and provide our contributions and recommendations for further investigation.

5.2 Experimental Design

In this section, we state the experimental design which is comprised of the selection of methods, the generation of a synthetic measurement matrix via the construction of a synthetic low-rank and noise term, and the definition of several measures. Finally, we provide the additional settings for this study.

5.2.1 Method Selection

In the past, PCA and NMF have been studied extensively in the context of imaging applications (see chapter 4) such that they form logical yardstick methods. Furthermore, they are related to multiple robust methods: RPCA-LBD, BPCP, MAHNMF, PCP, RNMF, SIPCP, and SPCP. From the latter selection we can exclude RPCA-LBD, BPCP, and SIPCP, respectively because it is computationally too expensive, because it consists of a theoretical extension to PCP, and finally because it uses side information or predefined dictionaries aiding the process of retrieval. To further limit the scope of this study and alleviate uniqueness related problems, we exclude the non-negative constrained methods^a. Besides, a large body of literature is presented on PCP and SPCP regarding parameter settings and exact recovery conditions. The latter further simplifies the method assumptions considered for this study. An overview of the selected methods along with their specifications is given in **Tab. 5.1**.

Table 5.1 | Selection of matrix decomposition methods. Besides method abbreviation, the corresponding assumptions are given along with noise robustness, time and spatial complexity, and a reference to our literature review. We assume that $M \in \mathbb{R}^{m \times n}$, i is the number of iterations required for convergence and r is the “true” matrix rank. The time complexity is given for a full SVD, i.e. LAPACK GESVD or GESDD (left), as well as randomized SVD (right).

Method	Asm.	Robust.	$\mathcal{O}(\cdot)$	$\Omega(\cdot)$	Lit. Ref.	
PCA	Sample statistics, no outliers, noise is i.i.d. and uncorrelated to signal, noise variance is low	small noise	i.i.d.	mn^2/rmn	mn	4.3.3.8
PCP	Incoherence conditions and noise conditions	outliers / dense noise		$imn^2/irmn$	mn	4.3.4.4
SPCP	Incoherence conditions and noise conditions	outliers, small i.i.d. noise		$imn^2/irmn$	mn	4.3.4.8

These methods make use of an underlying (extended) linear mixture model consisting of

$$\begin{aligned} M &= A + B + C, \\ &= XY^T + B + C, \end{aligned} \quad (5.1)$$

where $M \in \mathbb{R}^{m \times n}$ is the synthetic measurement matrix, $A \in \mathbb{R}^{m \times n}$ is assumed to represent a low-rank matrix, $B \in \mathbb{R}^{m \times n}$ a dense noise matrix, and $C \in \mathbb{R}^{m \times n}$ a sparse outlier matrix. Note that for PCA, $C = 0$ and for PCP, $B = 0$. To obtain a synthetic measurement matrix we therefore need to construct a low-rank term and one or two noise terms. We will assume that signal, i.e. what is of interest, is captured by A , while noise is both modeled by B and C .

^atheir implementations are, however, available from Appendix B

5.2.2 Synthetic Low-Rank Matrix

The method selection allows us to focus on the recovery conditions as specified for PCP and SPCP, both sharing the incoherence conditions on the low-rank term. Their exact recovery guarantees with respect to noise types are different, however. Therefore, we fix the incoherence conditions on the synthetic low-rank matrix, assuring that no signal will be considered sparse and consecutively ends up in the sparse outlier term. This also allows us to solely focus on the effect of different noise types on method performance. For this, we make use of Hadamard matrices.

5.2.2.1 Hadamard Matrix

The Hadamard matrix is a type of square matrix that is often used in experimental design¹. Solely containing -1 and 1 entries, all its rows and columns are orthogonal. A Hadamard matrix is generated as

$$H_0 = [1], H_i = \begin{bmatrix} H_{i-1} & -H_{i-1} \\ H_{i-1} & H_{i-1} \end{bmatrix}, \quad (5.2)$$

where $i \in \mathbb{N}^+$ and it only comes in $\mathbb{R}^{2^i \times 2^i}$ flavors.

5.2.2.2 Low-Rank Matrix Construction

For our low-rank matrix construction, we start by defining the left and right singular vectors of an SVD model (i.e. $A = U\Sigma V^T$). To do so, we first scale the columns of a Hadamard matrix to obtain orthonormality. Then, by selecting the r first columns we reconstruct

$$\begin{aligned} A &= U\Sigma V^T, \\ &= \bar{H}_{i,r} \Sigma (\bar{H}_{j,r})^T, \end{aligned} \quad (5.3)$$

where $\bar{H}_{i,r} \in \mathbb{R}^{2^i \times r}$ represent the matrix consisting of the first r columns of the column-normalized Hadamard matrix $\bar{H}_i \in \mathbb{R}^{2^i \times 2^i}$ and $\bar{H}_{j,r} \in \mathbb{R}^{2^j \times r}$ represent the matrix consisting of the first r columns of the column-normalized Hadamard matrix $\bar{H}_j \in \mathbb{R}^{2^j \times 2^j}$. Note that $m = 2^i$ and $n = 2^j$. Finally, we fix the uniqueness of the low-rank matrix up to signs by setting its singular values to r unique, evenly spaced samples in a range between 1 and 0.9.

5.2.2.3 Incoherence

The reason for the Hadamard matrix lies in its optimality with respect to the incoherence conditions. Recall the incoherence conditions that assert that the

singular vectors are not too singular, and that the joint of singular vectors (i.e. UV^T) does not produce “spiky” results, stated as

$$\begin{aligned} \max_i \left\| U^T e_i \right\|_2^2 &\leq \frac{\mu r}{m}, \\ \max_j \left\| V^T e_j \right\|_2^2 &\leq \frac{\mu r}{n}, \\ \left\| UV^T \right\|_\infty &\leq \sqrt{\frac{\mu r}{mn}}. \end{aligned} \quad (5.4)$$

We reformulate these to their lower-bound equality to obtain

$$\begin{aligned} \mu_u &= \max_i \left\| U^T e_i \right\|_2^2 \frac{m}{r}, \\ \mu_v &= \max_j \left\| V^T e_j \right\|_2^2 \frac{n}{r}, \\ \mu_\infty &= \left\| UV^T \right\|_\infty^2 \frac{mn}{r}. \end{aligned} \quad (5.5)$$

By substituting the information on \bar{H}_i and \bar{H}_j given to be

$$\begin{aligned} \bar{H}_i &= \begin{bmatrix} \frac{H_{i-1}}{\sqrt{2^i}} & -\frac{H_{i-1}}{\sqrt{2^i}} \\ \frac{H_{i-1}}{\sqrt{2^i}} & \frac{H_{i-1}}{\sqrt{2^i}} \end{bmatrix}, \\ \bar{H}_j &= \begin{bmatrix} \frac{H_{j-1}}{\sqrt{2^j}} & -\frac{H_{j-1}}{\sqrt{2^j}} \\ \frac{H_{j-1}}{\sqrt{2^j}} & \frac{H_{j-1}}{\sqrt{2^j}} \end{bmatrix}, \end{aligned} \quad (5.6)$$

in the lower-bound equality conditions, and using the fact that $m = 2^i$, $n = 2^j$, and r columns are selected, we obtain

$$\begin{aligned} \mu_u &= r \left(\frac{\pm 1}{\sqrt{2^i}} \right)^2 \frac{2^i}{r} = 1, \\ \mu_v &= r \left(\frac{\pm 1}{\sqrt{2^j}} \right)^2 \frac{2^j}{r} = 1, \\ \mu_\infty &= \left(r \frac{\pm 1}{\sqrt{2^i}} \frac{\pm 1}{\sqrt{2^j}} \right)^2 \frac{2^{i+j}}{r} = r. \end{aligned} \quad (5.7)$$

Hence, we see that both μ_u and μ_v are fixed to 1 (i.e. fully incoherent) and independent of the chosen rank. The joint incoherence μ_∞ is, however, rank bound so choosing a small rank r assures that the incoherence conditions are met.

5.2.2.4 Parameters

The matrix parameters for the experiment are specified in **Tab. 5.2**. The constructed low-rank matrix $A \in \mathbb{R}^{256 \times 128}$ is of rank 10.

Table 5.2 | Parameters related to the synthetic low-rank matrix. i and j are related to the Hadamard matrix, r is the matrix rank.

Parameter	Value
i	8
j	7
r	10

5.2.3 Synthetic Noise Term

We are interested in those cases where the methods break down, i.e. where the conditions set on the noise model are not met. The latter is namely what we expect to be true in practical applications. For example, for PCP it is stated that the support set of the sparse noise matrix is uniformly distributed among all sets of cardinality c . This implies that the unknown corruption pattern should be randomly distributed, with no restrictions on its signs or magnitudes². Alternatively, the same program is also able to exactly recover from dense noise with the stronger assumption that signs are randomly sampled from an i.i.d. distribution³. Note that the latter is not in addition to sparse noise. That situation, however, is accounted for by SPCP, dealing with small i.i.d. noise in addition to the sparsely distributed noise⁴.

5.2.3.1 Dense Noise Term

In imaging applications, additive Gaussian noise is often assumed to be present in measurement matrices. It is inherently connected to the Poisson distribution, which in its turn is closely related to the principles of most imagers, e.g. appearing in the form of shot noise. Furthermore, several effects of Gaussian noise have been researched in the past for PCA⁵. Secondly, speckle noise is also a recurring noise type in the context of imaging applications⁶⁻⁸. Hence, we considered Gaussian and speckle noise separately as dense noise terms. For our experiment, we varied the intensity of both noise types and inspected their influence on the retrieved terms and their components. In **Tab. 5.3** an overview of the dense noise scenarios is given. We only considered zero mean noise for speckle as well as Gaussian noise. Finally, note that the noise variances were chosen heuristically and that for speckle noise, correlation between signal and noise is possible.

Gaussian Noise For a matrix $N \in \mathbb{R}^{m \times n}$ generated following an independent Gaussian distribution with mean 0 and variance σ^2 , its variance is related to the Frobenius norm of its matrix as

$$\sigma^2 = \frac{1}{mn} \|N\|_F^2. \quad (5.8)$$

Hence, we scaled the variance of the Gaussian noise matrix as given in **Tab. 5.3** with the variance of the synthetic low-rank matrix A to obtain a reasonable level of noise.

Speckle Noise The speckle noise used in this study is obtained by entry-wise multiplying a Gaussian noise matrix with corresponding mean and variance with the synthetic low-rank matrix A as

$$B = N \circ A, \quad (5.9)$$

where \circ is the Hadamard product, the element-wise product of two matrices, N is a Gaussian noise matrix and A is the synthetic low-rank matrix. Note that for the speckle noise term we do not scale the variance due to its inherent multiplicative nature. The variances used for this study are also given in **Tab. 5.3**.

Table 5.3 | Parameters related to the noise matrices and the different dense noise cases under study. Besides the noise type descriptions, the different cases with their mean and variance are given. Note that the given variance of the Gaussian noise type has to be scaled.

Noise Type	Intensity	Case	Mean	Variance
Gaussian	High	1	0	2×10^0
		2	0	1×10^0
	↓	3	0	1×10^{-2}
		4	0	1×10^{-3}
		Low	5	0
Speckle	High	1	0	1×10^{-1}
		2	0	5×10^{-2}
	↓	3	0	1×10^{-2}
		4	0	5×10^{-2}
		Low	5	0

5.2.3.2 Sparse Noise Term

Besides the dense noise term, we studied cases with additional fixed sparse outliers. As discussed in section 4.1.4, these outliers are inherently present in most imaging applications. The sparse noise term is modeled as 'salt and pepper' noise with a fixed amount of 5% of the total number of entries $m \times n$ afflicted, and consisting of -1 and 1 entries (equally weighted). Hence, we obtained a total of 20 cases to examine: 5 Gaussian, 5 Gaussian with outliers, 5 speckle, and 5 speckle with outliers.

5.2.4 Measure Definition

The first set of measures focuses on quantitative aspects such as information loss, compression ratio, and efficiency. The second set of measures focuses on the

quality of the retrieved low-rank matrix via the visual inspection of atoms and low-rank reconstructions.

5.2.4.1 Quantitative Measures

We considered compression ratio, information loss, and efficiency, and translate each aspect into one or more measures.

Compression Ratio The compression ratio is represented by the retrieved matrix rank of the low-rank term. Given our fixed matrix rank of 10 (see **Tab. 5.2**) and the orthogonality of the columns and rows of Hadamard matrices, the compression ratio should be interpreted as the method's ability to retrieve this true rank or at least consistently approach it. Besides, matrix rank is informative about the obtainable compression. When the rank $r \ll m, n^a$, the compression ratio (CR) is given by

$$\text{CR} = \frac{\text{old size}}{\text{new size}} = \frac{mn}{(m+n)r}. \quad (5.10)$$

In practice, this one “true” rank is considered application, if not measurement matrix, dependent, however.

Information Loss For measuring the information loss, we use five measures: a low-rank relative offset, a noise relative offset, and three incoherence measures. The first two are defined as

$$\epsilon_{\text{low-rank}} = \frac{\|\mathbf{A} - \hat{\mathbf{A}}\|_{\text{F}}}{\|\mathbf{A}\|_{\text{F}}} \times 100\%, \quad (5.11)$$

and

$$\epsilon_{\text{noise}} = \frac{\|(\mathbf{B} + \mathbf{C}) - (\hat{\mathbf{B}} + \hat{\mathbf{C}})\|_{\text{F}}}{\|\mathbf{B} + \mathbf{C}\|_{\text{F}}} \times 100\%, \quad (5.12)$$

where $\hat{\cdot}$ represents the retrieved matrices, and \mathbf{A} , \mathbf{B} and \mathbf{C} are the generated synthetic matrix terms. The Frobenius norm is similar to the ℓ_2 -norm when the matrix is stretched into a vector. This insight alerts that its use might become problematic when scales become large or when the error under investigation is more complex than solely small in all samples (see section 4.1.4). An alternative could be the use of a Chebyshev norm (i.e. ℓ_∞) in combination with other norms to measure different aspects. However, given the scale of our matrix and the interest in sufficiently small deviations, we will stick to the Frobenius norm.

Secondly, we use the three incoherence measures as defined in Eq. 5.4. The incoherence provides measures of sparsity of the singular vectors and enables comparison to the calculated incoherences in Eq. 5.7. The latter enables a review of the retrieved dictionaries.

^athis is mostly true for low-rank matrix approximations

Efficiency We reduce efficiency to the time needed for the method to converge without taking the quality of the outcome into account. This simplification is made due to a possible and probable mismatch in optimal parameters and stopping criteria, falsifying any conclusions. Hence, we will stick to running time, which we consider a ball park measure.

5.2.4.2 Qualitative Measures

We review two qualitative measures: the reconstructed images of the low-rank term and the absolute cosine similarity of the retrieved left singular vectors.

Reconstructed Images The reconstructed images enable a visual inspection and give the ability to link visual to numerical results. It consists of picking a single column $\hat{\mathbf{a}}_i \in \mathbb{R}^{256}$ of our retrieved low-rank matrix $\hat{\mathbf{A}} \in \mathbb{R}^{256 \times 128}$, and reshaping it to an image as $\hat{\mathbf{A}}_{\text{image}} \in \mathbb{R}^{16 \times 16}$. This is analogous to the reconstruction of images for spectral imaging modalities when processed via matrix decompositions.

Absolute Cosine Similarity This measure returns the correlation between the atoms of the retrieved dictionary with respect to the atoms of the true dictionary. Given the orthonormal nature of the SVD, the absolute cosine similarity is defined as

$$|\cos(\widehat{\mathbf{U}}\mathbf{U})| = |\hat{\mathbf{U}}^T\mathbf{U}|, \quad (5.13)$$

where \mathbf{U} are the true left singular vectors of the low-rank term \mathbf{A} and $\widehat{\mathbf{U}}$ are the retrieved left singular vectors of the retrieved low-rank term $\hat{\mathbf{A}}$. Note that $|\cdot|$ is considered element-wise. Lying in a range between 0 and 1, a value close to 0 reflects a low correlation, i.e. close to orthogonality, while values close to 1 reflect vectors aligning in the same direction. The orthonormality property of the left singular vectors implies that $\mathbf{U}^T\mathbf{U} = \mathbf{I}$.

5.2.5 Additional Settings

Method Implementation For PCA the covariance flavor is chosen, for PCP we use the inexact augmented Lagrangian multiplier method (IALM)⁹, and for SPCP the alternating direction method with increasing penalty (ADMIP)¹⁰ is opted.

Parameter Setting For both PCP and SPCP the λ parameter was fixed to the usual $1/\sqrt{m}$, while for PCA the rank was set equal to the rank retrieved by PCP, to make comparison meaningful. For SPCP, δ is set heuristically to $\|\mathbf{M}\|_F/n$ to estimate the dense noise level (see Eq. 5.8).

Stopping Criteria The stopping criteria for PCP are set to

$$\begin{aligned} \frac{\|M - A - B\|_F}{\|M\|_F} &\leq 1 \times 10^{-7}, \\ \mu \frac{\|\hat{B} - B\|_F}{\|M\|_F} &\leq 1 \times 10^{-5}, \end{aligned} \quad (5.14)$$

where μ is an algorithmic variable. For SPCP the stopping criterion is defined as

$$\frac{\|(\hat{A}_{i+1}, \hat{B}_{i+1}) - (\hat{A}_i, \hat{B}_i)\|_F}{\|(\hat{A}_i, \hat{B}_i)\|_F + 1} \leq 5 \times 10^{-6} \delta, \quad (5.15)$$

where i is the iteration step. For both methods a maximum of 100 iterations is set.

Noise Term We used the scikit-image package¹¹ to generate the discussed noise patterns.

Computational Specification Floats with a precision of 64 bits were used and the underlying SVD for all methods is the OpenBLAS-LAPACK GESVD accessed via SciPy¹². All experiments were performed in parallel in Python 3.8.3¹³ in solely and separate CPU based processes. Finally, the computer system specifics are given in **Tab. 5.4**.

Table 5.4 | Computer system specifications. The operating system, number of cores and their clock speed, and the size of available memory is stated.

OS	Cores	Clock Speed	Memory Size
Windows 10 Pro	52 (104 logical)	2.7 GHz (3.3 GHz overclock)	1.5 TB

Quantitative Performance Analysis Each case (20 in total) for the quantitative performance analysis is simulated a 100 times. Finally, the mean, maximum, and minimum values of each measure are retrieved. Given the fixed low-rank measurement matrix, the sole variable in each simulation is the noise matrix, random by its nature.

5.3 Quantitative Performance Analysis

In this section we compare the quantitative performances of the aforementioned methods and shortly analyze some of their behavior. We propose links between these numeric phenomena and particular parameter settings. Note, that we do not claim these links to be causal, rather we want to point out that such settings might contribute to these phenomena.

5.3.1 Recovered Rank

The recovered rank for the three methods and four noise scenarios is depicted in Fig. 5.1. A first remark is that in most Gaussian noise cases, with as well as without outliers, the rank is overestimated by PCP and SPCP. This is especially true for the Gaussian noise cases and the high Gaussian noise with outlier cases. We observe a reduction in this behavior when the intensity is lowered. To a lesser extend, the same trend is visible for all speckle noise cases.

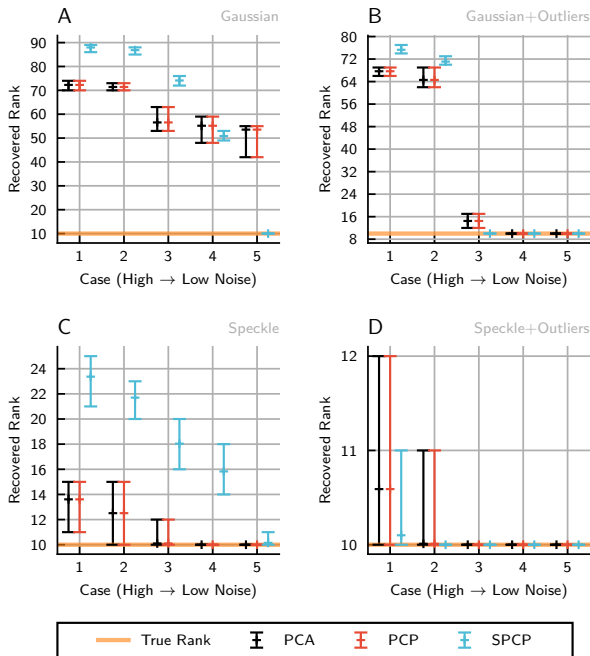


Figure 5.1 | Recovered rank for the cases: (A) Gaussian noise, (B) Gaussian noise with additional outliers, (C) Speckle noise, (D) Speckle noise with additional outliers. The + sign represents the mean of 100 simulations, and is given along with maximum and minimum values. The true rank is given as reference point.

This might stem from the fact that for our synthetic measurement matrix the speckle noise term is nearly Gaussian distributed with a lower intensity in comparison to the Gaussian noise term. Furthermore, comparing PCP to SPCP, we observe this effect to be more severe for the latter than the former in high intensity noise cases. In the outlier cases we see that this overestimation of the rank is reduced, probably due to a better fit of the sparse noise term, leading to less “leakage” to the low-rank term. These two observations lead to the proposal of adjustments to the λ and δ parameters. In case of PCP, a higher λ value will probably lead to a better fit, while for SPCP, a lower λ value and adjusted δ setting will

allow for better results. Note that the rank of PCA is coupled to the rank retrieved by PCP, such that no comments on it are possible here.

5.3.2 Low-Rank Relative Offset

While the rank reflects on the compression, it does not comment on the the quality of the retrieved low-rank matrix term. Hence, the low-rank relative offset for the three methods and four noise scenarios is depicted in **Fig. 5.2**.

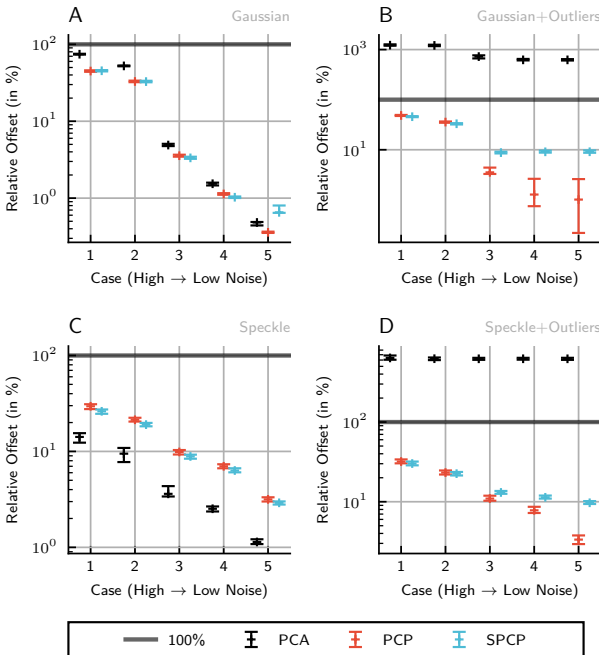


Figure 5.2 | Low-rank relative offset for the cases: (A) Gaussian noise, (B) Gaussian noise with additional outliers, (C) Speckle noise, (D) Speckle noise with additional outliers. The + sign represents the mean of 100 simulations, and is given along with maximum and minimum values. The 100% offset line references to the case of $\hat{\Lambda} = 0$.

We observe that PCA performs remarkably well for the speckle noise cases confirming its approximate Gaussian distribution, although that the speckle noise can be correlated to signal. Secondly, we see that for both non-outlier scenarios PCA, PCP, and SPCP closely follow the same downward trend which seems to be a logical observation: a reduced amount of noise makes the retrieved low-rank term to be better retrievable. This probably also follows from the reduced overestimation in rank. The latter is also true for PCP and SPCP in the outlier scenarios. However, SPCP's performance seems to flatten towards low intensity noise cases. This is probably due to the δ setting, driven by $\|M\|_F$, being overly dependent on

the outliers. Note that SPCP and PCP perform equally well for most cases.

5.3.3 Noise Relative Offset

Besides the fit of the low-rank term, we can also inspect the fit of the noise term. The noise relative offset for the three methods and four noise scenarios is therefore depicted in **Fig. 5.3**.

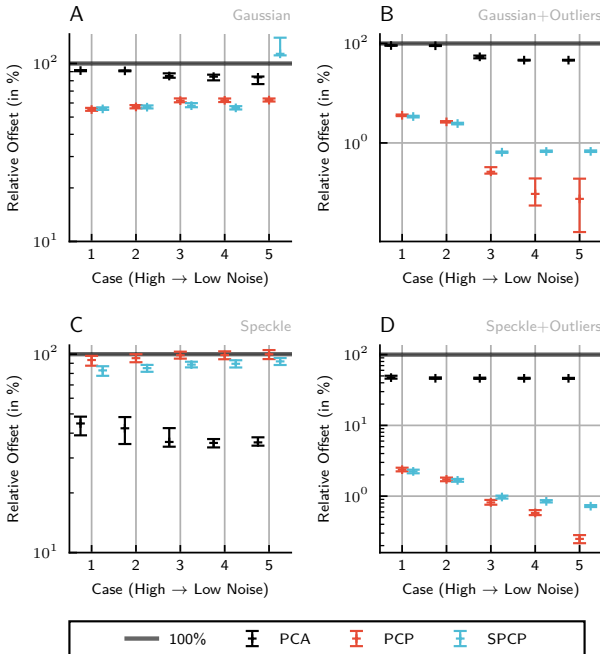


Figure 5.3 | Noise relative offset for the cases: (A) Gaussian noise, (B) Gaussian noise with additional outliers, (C) Speckle noise, (D) Speckle noise with additional outliers. The + sign represents the mean of 100 simulations, and is given along with maximum and minimum values. The 100% offset line references to the case of $\hat{B} + \hat{C} = 0$.

A first remark is the bad fit in most non-outlier cases for all methods. This is probably a direct consequence of the rank overestimation. Secondly, we see the same flattening behavior for SPCP. The large offsets for the Gaussian noise are probably due to the relatively small Frobenius norm of the noise term itself, i.e. B , with respect to the captured noise, i.e. $\hat{B} + \hat{C}$. This effect concurrently leads to large offset values, as previously discussed, and will probably lead to false conclusions. The slight increase in trend for the low intensity noise cases might confirm this presumption. Finally, in the outlier scenarios we see a relatively good fit for PCP and SPCP. This is explained by the fact that the Frobenius norm of the noise term is mainly filled up by the sparse outlier term, i.e. C . This confirms that the

sparse outliers are indeed captured by the sparse outlier term, i.e. \hat{C} .

5.3.4 Incoherence

The mutual incoherence of the retrieved low-rank term for the three methods and four noise scenarios is depicted in **Fig. 5.4**.

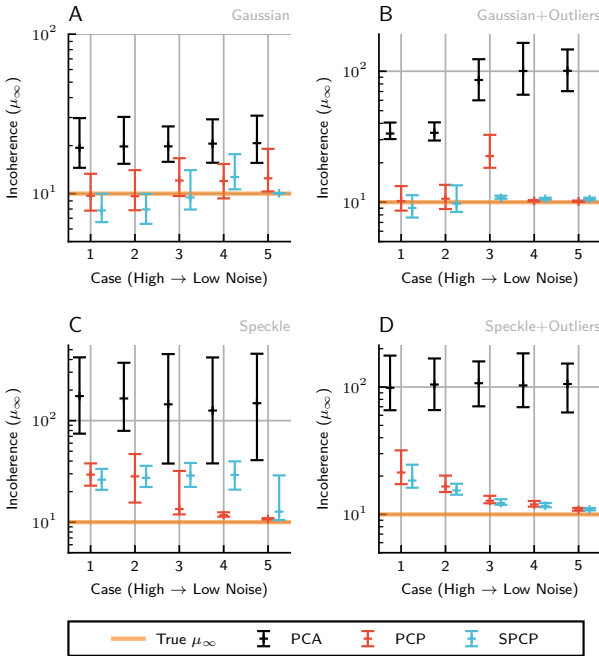


Figure 5.4 | Joint incoherence for the cases: (A) Gaussian noise, (B) Gaussian noise with additional outliers, (C) Speckle noise, (D) Speckle noise with additional outliers. The + sign represents the mean of 100 simulations, and is given along with maximum and minimum values. The “True μ_∞ ” reflects the calculated incoherence as given in Eq. 5.7

The same conclusions can be drawn for the μ_u and μ_v incoherences, such that we refer the interested reader to Appendix A.1 for those. Firstly and most remarkably, PCP and SPCP, although not always correctly recovering rank, seem to approximate the true incoherence better than PCA. Furthermore, their spread is lower than PCA (note the logarithmic scale). The latter is especially visible in the speckle noise scenario. Whether this phenomenon is due to a discrepancy in the conditions of PCA being met (e.g. covariance and underlying statistics) or a true result has to be further investigated. Note that a low incoherence might also affect the compression positively, as “spiky” singular vectors tend to only carry local information, while incoherent vectors carry global information. Secondly, we remark that the incoherence of PCP for case 3 in the scenario of Gaussian

noise with outliers seems to be consistently outlying the trend, although only by around a factor 2. Whether this is due to the slight overestimation in rank or a noise related phenomenon is not clear.

5.3.5 Running Time

Finally, the running time for the three methods and four noise scenarios is depicted in **Fig. 5.5**. Blindly looking at the different scenarios could lead to the conclusion that SPCP outperforms PCP. However, PCP performs more iterations due to the set stopping conditions. Although their similar time complexity, SPCP converges faster. Whether this is the case when similar stopping criteria are set, has to be investigated. A second remark is that SPCP has a slightly higher cost per iteration which does not manifest itself in this small scale experiment.

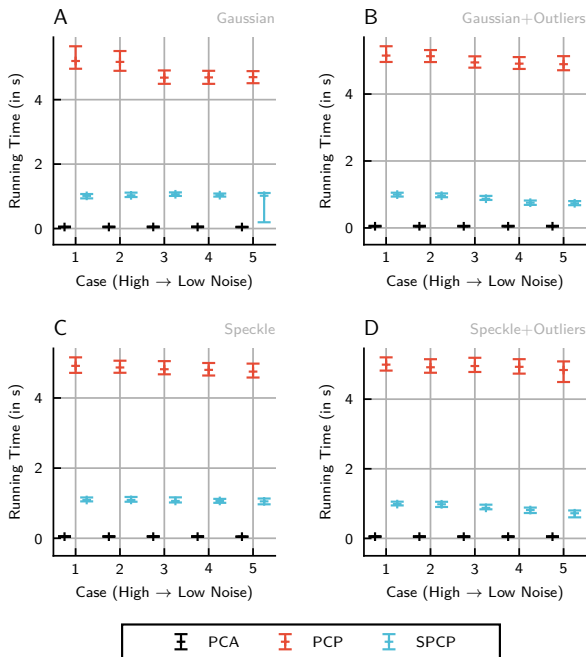


Figure 5.5 | Running time for the cases: (A) Gaussian noise, (B) Gaussian noise with additional outliers, (C) Speckle noise, (D) Speckle noise with additional outliers. The + sign represents the mean of 100 simulations, and is given along with maximum and minimum values.

An extra sorting of matrix values is namely required with a cost of $\mathcal{O}(k \log(k))$, where k is the number of non-zero matrix entries of C . On larger scales, this increases the time per iteration drastically. The only conclusion we can draw is that PCP and SPCP are both several factors slower than PCA. This might be

alleviated by setting less restrictive stopping criteria, by using a warm start (extra information required), or using a randomized SVD.

5.4 Qualitative Performance Analysis

In this section, we compare the qualitative performances of the aforementioned methods and visually analyze some of their behavior. We propose links between these visual phenomena and particular parameter settings. Again, we do not claim that these links are causal.

5.4.1 Reconstructed Images

In **Fig. 5.6**, the reconstruction of the 100th column of the five cases of Gaussian noise with outliers is depicted for a single simulation. A first remark is that PCA is completely distorted due to the outliers and exhibits ghosting effects, i.e. due to truncation artifacts. Secondly and similar to **Fig. 5.2**, SPCP does not seem to improve towards lower intensity noise cases (e.g. 4 and 5), while PCP does.

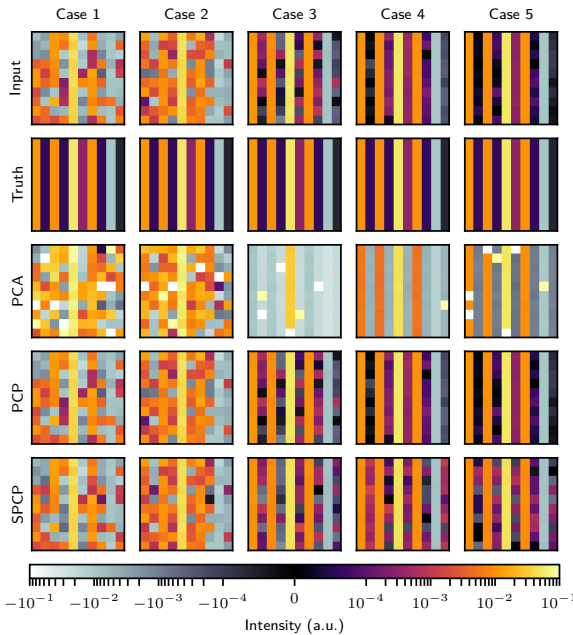


Figure 5.6 | Reconstructed images. From left to right the different cases of Gaussian noise with outliers (from high to low) and from top to bottom respectively the 100th column as image of the originally generated matrix (i.e. M), the ground truth (i.e. A), PCA, PCP and SPCP (i.e. \hat{A}). This reflects the results of one single simulation.

We also remark that both PCP and SPCP seem to be able to grasp the high intensity “stripe” patterns quite well for all noise intensities and also perform reasonable on the lower intensity stripes, in contrast to PCA where only the high energetic stripes are captured. The latter directly follows from its minimization formulation. However, we remark that in the high intensity noise cases, i.e. 1 and 2, PCP and SPCP don’t seem to visually improve the input to the ground truth. And finally, note that the small percentage difference in **Fig. 5.2**, results in reasonable visual difference for the low SNR regions. The latter shows that the Frobenius norm should be appreciated in its ability to report large deviations, not the small.

5.4.2 Absolute Cosine Similarity

In **Fig. 5.7**, the absolute cosine similarity of the left singular vectors of the five cases of speckle noise without outliers is depicted for a single simulation. Although the low-rank relative offset is small, as is shown in **Fig. 5.2**, and the recovered rank is close to the true rank, the left singular vectors space of the retrieved components do not align for all cases to the ground truth. For a perfect retrieval, we expect the vectors to align to form the identity matrix such as visualized in the first row. However, we observe that for the high intensity noise cases, i.e. 1 and 2, the vectors are poorly aligned. Secondly, PCP and SPCP seem to struggle more than PCA to align in the high noise cases, but improve towards the low-noise cases. Remark that the first singular vector is not grasped by PCA, i.e. it is off-diagonal and not diagonal. It is only captured by later vectors (not visualized).




Figure 5.7 | Absolute cosine similarity. From left to right the different cases of speckle noise without outliers (from high to low) and from top to bottom respectively the absolute cosine similarity between first 10 columns of the left singular values of the ground truth and (1) the ground truth; (2) PCA; (3) PCP and (4) SPCP. This reflects the results of one single simulation. Values of 0 correspond to no correlation, while values of 1 indicate perfect alignment between vectors. Optimally, a diagonal with 1 values is seen, such as in the first row.

5.5 Conclusions

In this chapter we reviewed a selection of robust methods under inexact recovery conditions, i.e. with respect to different noise types, and compared them to our yardstick, PCA. We compared PCP and SPCP to PCA on a total of 20 noise cases: 5 Gaussian, 5 Gaussian with outliers, 5 speckle, and 5 speckle with outliers.

The recovered rank suggests that both PCP and SPCP might benefit from an improved λ setting, as the rank is consistently overestimated in most non-outlier noise scenarios. Secondly, the low-rank relative offset suggests that SPCP could further benefit from an adjusted δ setting. The latter is manifesting itself through a flattening trend in performance towards smaller intensity noise cases. On the other hand, the incoherence results suggest that the singular vectors of PCP and SPCP are consistently incoherent. The latter probably positively affects the ability to compress measurement matrices efficiently. The same incoherence results revealed a consistent outlying value for PCP in an intermediate intensity case of the Gaussian noise with outliers. Whether this is due to a slight overestimation in rank or noise related phenomenon, is not clear. Nevertheless, in contrast to PCA the spread on the incoherence values is small for both PCP and SPCP. Finally, the running time shows that SPCP and PCP are several factors slower than

PCA. This can be partially alleviated by setting less restrictive stopping criteria, by using a warm start, or using a randomized SVD.

The qualitative performance analysis visually confirms the flattening trend of SPCP for the cases of low intensity Gaussian noise with outliers. It also suggests that PCP and SPCP improve the retrieval of low SNR patterns, but do fail under high intensity noise. Finally, the absolute cosine similarity suggests that the left singular vectors retrieved by PCP and SPCP are highly affected by high intensity speckle noise, lowering this intensity increases their ability to align their left singular vectors with the ground truth.

Contributions We contributed the following:

- Object-oriented Python implementations of several methods
- Shown the use of robust methods under inexact recovery conditions
- Suggested the use of the Hadamard matrices for fixing incoherence
- Proposed different quantitative and qualitative measures
- Recognized shortcomings and identified breakdown phenomena
- Made this study reproducible

Recommendations Further research could include:

- Reviewing varying incoherence conditions under fixed noise cases, e.g. through a Hadamard matrix, by increasing rank.
- Optimizing λ and δ parameter setting to improve inexact recovery performance
- Researching the effects of non-negative synthetic measurement matrices
- Non-negative matrix decomposition methods
- Review the link of PCA conditions to incoherence

Bibliography

- [1] Gentle, J. E. Matrix algebra. *Springer texts in statistics*, Springer, New York, NY, doi **10**, 648–0 (2017).
- [2] Candès, E. J. & Recht, B. Exact matrix completion via convex optimization. *Foundations of Computational mathematics* **9**, 717 (2009).
- [3] Ganesh, A., Wright, J., Li, X., Candès, E. J. & Ma, Y. Dense error correction for low-rank matrices via principal component pursuit. In *2010 IEEE international symposium on information theory*, 1513–1517 (IEEE, 2010).
- [4] Zhou, Z., Li, X., Wright, J., Candès, E. & Ma, Y. Stable principal component pursuit. In *2010 IEEE international symposium on information theory*, 1518–1522 (IEEE, 2010).
- [5] Lichtert, S. & Verbeek, J. Statistical consequences of applying a pca noise filter on eels spectrum images. *Ultramicroscopy* **125**, 35–42 (2013).
- [6] Tur, M., Chin, K.-C. & Goodman, J. W. When is speckle noise multiplicative? *Applied optics* **21**, 1157–1159 (1982).
- [7] Racine, R., Walker, G. A., Nadeau, D., Doyon, R. & Marois, C. Speckle noise and the detection of faint companions. *Publications of the Astronomical Society of the Pacific* **111**, 587 (1999).
- [8] Van de Plas, R., Yang, J., Spraggins, J. & Caprioli, R. M. Image fusion of mass spectrometry and microscopy: a multimodality paradigm for molecular tissue mapping. *Nature methods* **12**, 366–372 (2015).
- [9] Lin, Z., Chen, M. & Ma, Y. The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. *arXiv preprint arXiv:1009.5055* (2010).
- [10] Aybat, N. S. & Iyengar, G. An alternating direction method with increasing penalty for stable principal component pursuit. *Computational Optimization and Applications* **61**, 635–668 (2015).
- [11] van der Walt, S. *et al.* scikit-image: image processing in Python. *PeerJ* **2**, e453 (2014). URL <https://doi.org/10.7717/peerj.453>.
- [12] Virtanen, P. *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17**, 261–272 (2020).
- [13] Van Rossum, G. & Drake, F. L. *Python 3 Reference Manual* (CreateSpace, Scotts Valley, CA, 2009).

Chapter 6

An Application-Driven Framework for Parameter Optimization

After this chapter you will be able to

- understand nuances in ground truth information loss measures
- follow the underlying idea of heuristic information loss measures
- translate domain-specific phenomena into quantitative measures
- reproduce our application-driven framework
- formulate and tune your own framework

6.1 Introduction

This chapter introduces an application-driven framework for parameter optimization for matrix decomposition methods. It consists of obtaining an optimal set of parameters through the optimization of a composite measure via the combination of heuristic information loss measures. Finally, we showcase the use of such framework on a synthetic measurement matrix, similar to chapter 5.

6.1.1 Goals and Objectives

The aim of this chapter is to lay out an application-driven framework for parameter optimization. In second instance, we showcase its use for matrix decompo-

sition methodology. In literature, we find different solutions for such parameter optimizations, e.g. the rank determination problem in PCA¹⁻³. These solutions, however, mostly consider optimality with respect to one generic measure, not taking into account the origin or application of the measurement matrix. Furthermore, the optimality of such solutions is likely not to be guaranteed due to the corresponding assumptions failing in practical applications. Therefore, we consider optimality as application-dependent, i.e. different applications on the same measurement matrix require a different set of parameters. For example with PCA, when one is interested in only a few first components for data exploration or visualization, a small rank is sufficient. In contrast, when the main task is denoising, one might want to preserve signal patterns as much as possible, leading to higher ranks.

Our objective hence consists of exploiting domain- and application-specific knowledge and translating that into heuristic information loss measures that can be traded off by the user depending on its application. The optimum of this trade-off is then projected onto a set of method parameters. We showcase such framework, propose three different applications and analyze their use via different ground truth information loss measures. The latter is possible as the experiments will be performed on a similar synthetic measurement matrix as proposed in chapter 5. These findings subsequently formulate the answer to research sub-question A.2. Finally, note that the validity of the outcomes will be considered in chapters 8 and 9.

A. Can matrix decomposition methods based on the ℓ_1 -robustness principles efficiently reduce massively multivariate and high-dimensional spectral imaging measurement matrices to a lower-dimensional representation while incurring minimal information loss?

A.2. Can a framework be built to efficiently tune the regularization parameters of the above-mentioned methods, but customized for a specific application?

6.1.2 Overview

Section 6.2 provides the experimental design, including method selection, the selection of the synthetic low-rank matrix and synthetic noise terms. In that same section, several information loss measures are defined, both including and excluding the availability of the ground truth. Then our framework is proposed along with its solver. Finally, the additional settings are provided for reproducibility purposes. In section 6.3, we review through our framework three case studies each representing different applications on a synthetic measurement matrix. Finally, we draw our conclusions and provide our contributions and recommendations for further investigation.

6.2 Experimental Design

In this section, we provide the experimental design which is comprised of the selection of methods, the selection of a synthetic measurement matrix via the construction of a synthetic low-rank and noise term, the definition of several information loss measures, and the proposition of a framework with solver. Finally, we provide the additional settings for this study.

6.2.1 Method Selection

For simplicity we restrict ourselves to PCP due to its sole λ parameter. Nevertheless, the framework can be scaled to $n_p \in \mathbb{N}_0$ parameters. Our intuition teaches us that in PCP, λ can be used to trade off the sparsity of the recovered noise term \hat{B} and the sparsity of singular values of \hat{A} . For some cases, it has been mathematically proven that λ is dependent on noise conditions^{4,5}. We numerically show that it can also improve recovery when the incoherence conditions are not met. To complete the method selection, the underlying model is the LMM which is described as

$$\begin{aligned} M &= A + B, \\ &= XY^T + B, \end{aligned} \tag{6.1}$$

where $M \in \mathbb{R}^{m \times n}$ is the synthetic measurement matrix, $A \in \mathbb{R}^{m \times n}$ is assumed to represent a low-rank matrix, $B \in \mathbb{R}^{m \times n}$ a noise matrix. Secondly, note that A can also be decomposed following the SVD as $A = U\Sigma V^T$. We will assume that signal, i.e. what is of interest, is captured by A , while noise is modeled by B . We denote the recovered matrices with a hat: $\hat{\cdot}$. Finally, note that B is not restricted to sparse noise only, as was previously discussed.

6.2.2 Synthetic Low-Rank Term

The same synthetic low-rank term is chosen as in chapter 5. However, in practical measurement matrices, it is likely that not all signal components are fully incoherent, as was assumed earlier. Therefore, we reduce the incoherence component-wise while maintaining the orthogonality property of our low-rank term. This is achieved by gradually setting a larger fraction of column entries of $U \in \mathbb{R}^{256 \times 10}$ and $V \in \mathbb{R}^{10 \times 128}$ to 0. The range of entries set to 0 is respectively given in **Tab. 6.1** for U and in **Tab. 6.2** for V .

Table 6.1 | Range of column entries of U set to 0.

Column	1	2	3	4	5	6	7	8	9	10
Range	\emptyset	1 : 25	1 : 57	1 : 81	1 : 113	1 : 137	1 : 169	1 : 193	1 : 225	1 : 249

Table 6.2 | Range of row entries of V set to o .

Row	1	2	3	4	5	6	7	8	9	10
Range	\emptyset	1 : 16	1 : 25	1 : 40	1 : 49	1 : 64	1 : 73	1 : 88	1 : 97	1 : 112

6.2.3 Synthetic Noise Term

The same synthetic noise terms are generated as in chapter 5. Hence, our study includes 20 different noise cases. They are grouped as Gaussian noise (5), Gaussian noise with outliers (5), speckle noise (5) and speckle noise with outliers (5), and have different intensities ranging from high to low (see **Tab. 5.3**).

6.2.4 Measure Selection

Central to our measure selection is the availability of the ground truth. In practical applications this ground truth is not available, such that prior information on the solution space is required to evaluate the method's outcome. To narrow down the scope of this study, we solely focus on the use of information loss measures. However, nothing limits the use of other measures. Finally, we consider sticking to one loss measure sensitive to bias as only a single aspect of the recovered term is reviewed. Therefore, we review multiple ground truth and heuristic information loss measures.

6.2.4.1 Ground Truth Information Loss

We define a total of five ground truth information loss measures, i.e. when the ground truth is available: rank, low-rank relative offset, maximum absolute offset, median ℓ_1 absolute offset and the mean absolute cosine similarity. For calculating the rank, the ground truth is not needed. However, having information on the true rank we included it as a ground truth information loss measure. Besides, many more measures can be formulated depending on the domain of application and the aspect of the recovered term to be captured. The measure abbreviations, definitions and aspect description are given in **Tab. 6.3**. For an in-detail description of the ranges of those measures and the corresponding interpretation of values are given in **Tab. A.1** (see Appendix A.2). Note that there is certain redundancy in the choice of these measures. Even if they grasp different aspects of the recovered terms, we did not make a choice on the criteria of statistical dependency or correlation. This affects the framework, and might be a topic for further research. Secondly, note that there is always a loss of information when using summarizing measures, whether it is averaging, summing or taking the median. Nevertheless, one could specify a measure that acts on a local scale, e.g. when only a few entries are of interest, or when some a priori information is known on some parts of the given measurement matrix. We did not consider the latter here.

Table 6.3 | Selection of ground truth information loss measures. Besides the measure name, the abbreviations, definition and a short aspect description are given.

#	Measure	Abbr.	Definition	Aspect
1	Rank	rank	$\text{rank}(\hat{A})$	Measures the dimension of low-rank term, linked to information content dimension and variability.
2	Low-Rank Relative Offset	lr	$\frac{\ A - \hat{A}\ _F}{\ A\ _F} \times 100\%$	Measures the global closeness of fit between the ground truth and recovered low-rank term.
3	Maximum Absolute Offset	mao	$\ A - \hat{A}\ _\infty$	Measures the maximal offset in absolute value of the low-rank entries. Reflects the worst-case entry-wise offset.
4	Median ℓ_1 Absolute Offset	ml1	$\text{median}(\ A - \hat{A}\ _{\ell_1})$	Measures the median absolute column offset. Reflects the average absolute column-wise offset.
5	Mean Absolute Cosine Similarity	macs	$1 - \text{mean}(\text{diag}(\hat{U}^T U))$	Measures the median absolute cosine similarity of the first 10 components. Reflects the average quality of the first 10 recovered atoms in the left singular values.

6.2.4.2 Heuristic Information Loss

In practical applications the use of measures using ground truth information is an additional problem, as it is not available. Hence, we fall back on heuristic information loss measures. We select four measures: rank, cardinality of B , the Frobenius residual and the structure in B . The measure definitions and aspect description are given in **Tab. 6.4**. For an in-detail description of the ranges of those measures and the corresponding interpretation of values are given in **Tab. A.1** (see Appendix A.2).

Table 6.4 | Selection of heuristic information loss measures. Besides the measure name, the abbreviations, definition and a short aspect description are given.

#	Measure	Definition	Aspect
1	Rank	$\text{rank}(\hat{A})$	Measures the dimension of the low-rank term, linked to information content dimension and variability. Naturally arises in the non-convex formulation of PCP.
2	Cardinality of B	$\ \hat{B}\ _0$	Measures the sparsity of the noise term. Naturally arises in the non-convex formulation of PCP.
3	Structure in B	$\text{mean}(I_{\hat{B}})$	Measures “structure” in the columns of \hat{B} . The structure can be adjusted to the measurement matrix. See <i>Moran’s I</i> below for definition of $I_{\hat{B}}$.
4	Frobenius Residual	$\text{tr}(\hat{A}^T \hat{B})$	Measures the uncorrected correlation between \hat{A} and \hat{B} . See <i>Frobenius Residual</i> below.

Suggestions for other heuristic measures are signal-to-noise ratio estimations, all kinds of correlations, e.g. Fourier ring correlation⁶, but also captured variance measures, e.g. the mean absolute correlation coefficient of the loadings and the mean absolute correlation coefficient of the scores⁷, and image quality measures, e.g. structural similarity index measure⁸, mutual information measures, etc.

Frobenius Residual The Frobenius residual can be best understood from the following expansion, using $M = A + B$,

$$\begin{aligned}
 \|M\|_F^2 &= \text{tr}(M^T M), \\
 &= \text{tr}((A + B)^T (A + B)), \\
 &= \text{tr}(A^T A) + \text{tr}(A^T B) + \text{tr}(B^T A) + \text{tr}(B^T B), \\
 &= \text{tr}(A^T A) + 2 \text{tr}(A^T B) + \text{tr}(B^T B), \\
 &= \|A\|_F^2 + \|B\|_F^2 + 2 \text{tr}(A^T B).
 \end{aligned} \tag{6.2}$$

We can further expand the latter term by rewriting it as

$$\begin{aligned}
 \text{tr}(A^T B) &= \sum_{i=1}^n \mathbf{a}_i^T \mathbf{b}_i, \\
 &= \sum_{i=1}^n \cos(\widehat{\mathbf{a}_i \mathbf{b}_i}) \|\mathbf{a}_i\|_2 \|\mathbf{b}_i\|_2.
 \end{aligned} \tag{6.3}$$

Hence, this Frobenius residual can be seen as a mean uncorrected correlation measure between the matching columns of \hat{A} and \hat{B} . This column-wise evaluation is natural for imaging applications as columns represent images. Furthermore, for PCA this term is 0 due the fact that A and B reside in orthogonal spaces, resulting in a zero inner product. This term will be approximately 0 when the noise in \hat{B} is coherent. For other noise types this is likely not to be the case. Therefore, this property can be exploited, e.g. to detect signal creation which would manifest itself as a negative value.

Moran's I Moran's I can be used as a measure of spatial autocorrelation. It has been exploited in the past in combination with PCA to include spatial information⁹. The measure is defined as

$$I = \frac{c_n \sum_i \sum_j w_{ij} (x_i - \bar{x})(x_j - \bar{x})}{c_w \sum_i (x_i - \bar{x})^2} \tag{6.4}$$

where w_{ij} is a weighting term, $c_w = \sum_i \sum_j w_{ij}$, x_i represent the column entries of \mathbf{x} , \bar{x} is the column mean, and c_n corresponds to the length of \mathbf{x} . This measure is thus calculated for each column of \hat{B} , where $\mathbf{x} = \hat{B}_i$. The weighting term can be set to any pattern that suits the application. Common patterns are 4- or 8-connected regions or 2D-Gaussian functions.

6.2.5 Framework

Our framework consists of linearly combining heuristic information loss measures into a composite measure which is consecutively minimized. This composite measure is proposed as approximation to ground truth measures, which are not available. Therefore, we propose the unconstrained optimization

$$\min_{\lambda} \alpha^T \bar{\mathbf{f}}(\lambda) \quad (6.5)$$

where α are the user-set, application-dependent weights for the corresponding selection of normalized heuristic measures $\bar{\mathbf{f}}(\cdot)$, which might also be application-dependent. For our selection, Eq. 6.5 is equivalent to

$$\min_{\lambda} \alpha_0 \overline{\mathbf{rank}(\hat{\mathbf{A}}(\lambda))} + \alpha_1 \overline{\|\hat{\mathbf{B}}(\lambda)\|_0} + \alpha_2 \overline{\mathbf{tr}(\hat{\mathbf{A}}^T(\lambda)\hat{\mathbf{B}}(\lambda))} + \alpha_3 \overline{\mathbf{mean}(I_{\hat{\mathbf{B}}}(\lambda))} \quad (6.6)$$

where α_i are the entries of α , and $\hat{\mathbf{A}}$ and $\hat{\mathbf{B}}$ are λ dependent, and the overline $\bar{\cdot}$ represents the filtering and min-max normalization. Note that although our proposed framework is linear in \mathbf{f} , it is non-linear and non-convex with respect to λ . An example of such composite over a range of λ values is given in **Fig. 6.1** along with the heuristic measures. Note that the measures are min-max normalized and filtered through a Savitzky-Golay filter prior to the weighting and combination. This is respectively set in case to keep the weights simple and to smooth out steep jumps due to computational differences in convergence.

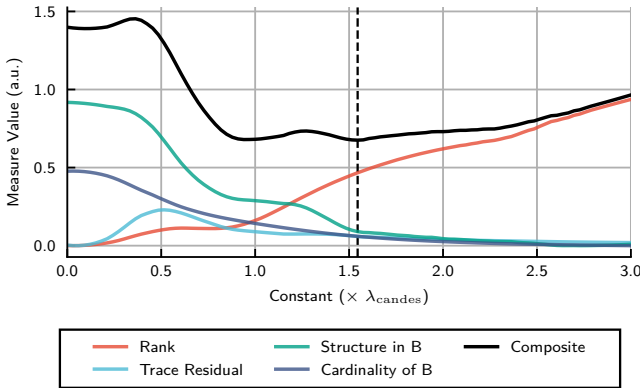


Figure 6.1 | Exemplary composite measure constructed from a linear combination of min-max normalized and filtered heuristic information loss measures. The dashed line marks the minimum of the composite measure. The noise term consists of the 4th case of speckle noise.

6.2.5.1 Grid Search

We propose a simple grid search to solve for Eq. 6.5. Note that less computational expensive solutions exists for such optimizations, e.g. Bayesian approaches¹⁰, or

solutions exploiting structure¹¹. However, the latter is considered for future research. The main idea of the grid search consists of extensively calculating the measures for a range of λ values and as such reconstructing the composite function $f(\lambda)$. After defining the weights α , the minimal value of the inner product is found via an extensive list search. Finally, this minimal value is projected onto the parameter space to obtain the corresponding optimal λ . Therefore, an equidistant sampling of a range of λ values is chosen. As earlier mentioned, a $\lambda_{\text{candes}} = 1/\sqrt{m}$ was originally proposed⁵. Heuristically, this point works excellent as a starting point, such that we define our grid around this point as

$$\lambda = c\lambda_{\text{candes}}, \quad (6.7)$$

where c is in a range between $[0, 3]$. This experimentally proves to be sufficient for good solutions. For our purposes, we sampled 100 equidistant points in the $[0, 3]$ range. Again, this experimentally proves to be sufficient for sampling the dynamics of the underlying measures. However, note that this can be coarsened or refined if needed. As such, our framework requires roughly $\mathcal{O}(qm^2/qm)$ flops, where q is the number of samples.

6.2.5.2 Defining Case Studies

We define three case studies: maximal compression, maximal signal retention and a trade-off. The respectively first can be linked to imaging mass spectrometry where often measurement matrices are too large to be stored altogether such that achieving maximal compression is desired. The second is related to applications in astronomical imaging where denoising with minimal signal loss is often required. Thirdly, we define a trade-off between those cases to show the consistency of our framework. The weights of the first two case studies are found by trial-and-error, while the trade-off is obtained by summing the ℓ_2 -normalized first two case studies. The weights of all three case studies are given in **Tab. 6.5**.

Table 6.5 | Weights for the three case studies. The case studies consists of maximal compression, maximal signal retention and a trade-off between both. The weights are defined as for Eq. 6.6.

Weight	Maximal Compression	Maximal Signal Retention	Trade-Off
α_0	1.00	1.00	0.94
α_1	0.00	1.00	0.23
α_2	0.00	4.00	0.92
α_3	1.00	-1.00	0.48

A graphical reconstruction of the measures for the different case studies is depicted in **Fig. 6.2**. From this reconstruction, we observe that the maximal compression has the lowest rank of all three cases. Secondly, we observe the

maximal signal retention corresponds to a small low-rank relative offset, which is desired. Finally, the trade-off seems to find the point that trades-off both good compression and good retrieval accuracy. The latter is remarkably close to the point where the maximal absolute offset and low-rank relative offset are close to minimal.

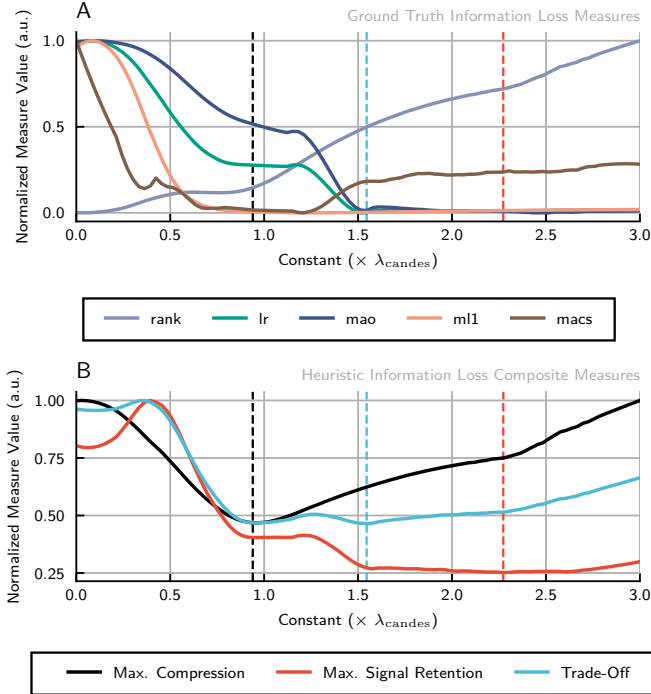


Figure 6.2 | Measures over the grid search: (A) normalized ground truth information loss measures; (B) normalized composite measures. The dashed line marks the different minima of the corresponding composite measures. The noise term for both (A) and (B) consists of the 4th case of speckle noise.

6.2.6 Additional Settings

Method Implementation For PCP we use the inexact augmented Lagrangian multiplier method (IALM)¹² was utilized.

Stopping Criteria The stopping criteria for PCP are set to

$$\begin{aligned} \frac{\|M - A - B\|_F}{\|M\|_F} &\leq 1 \times 10^{-7}, \\ \mu \frac{\|\hat{B} - B\|_F}{\|M\|_F} &\leq 1 \times 10^{-5}, \end{aligned} \quad (6.8)$$

where μ is an algorithmic variable. A maximum of 100 iterations is set.

Noise Term We used the scikit-image package¹³ to generate the discussed noise patterns.

Moran's I We use an equal weighted 4-connected pattern.

Savitzky-Golay Filter The Savitzky-Golay filter¹⁴ was utilized from SciPy package¹⁵. The window length is set to 15 and the order of the polynomial set to 2. Its parameters were set through trial-and-error.

Computational Specification Floats with a precision of 64 bits were used and the underlying SVD for all methods is the OpenBLAS-LAPACK GESVD accessed via SciPy¹⁵. All experiments were performed in parallel in Python 3.8.3¹⁶ in separate, purely CPU-based processes. Finally, the computer system specifics are given in **Tab. 6.6**.

Table 6.6 | Computer system specifications. The operating system, number of cores and their clock speed, and the size of available memory is stated.

OS	Cores	Clock Speed	Memory Size
Windows 10 Pro	52 (104 logical)	2.7 GHz (3.3 GHz overclock)	1.5 TB

Case Studies Each noise case (20 in total) for the different case studies is simulated 10 times. Finally, the mean, maximum and minimum values of each measure are retrieved. Given the fixed low-rank measurement matrix the sole variable in each simulation is the noise matrix, random by its nature.

6.3 Case Studies

The λ settings for the different noise cases and case studies, along with the optimal λ setting with respect to the low-rank relative offset are given in **Fig. 6.3**. Firstly, we remark that for all noise cases the spread for the maximal signal retention (light blue) is significant. The spread for the maximal compression (red) is small in comparison to the other case studies, and the trade-off case (orange) takes the middle grounds. Given the higher impact on the maximal signal retention and in-depth analysis attribute this phenomenon to the structure in B measure. It appears to be caused by the small image size, i.e. only 16×16 , being sensitive to the 4-connected pattern used in the measure.

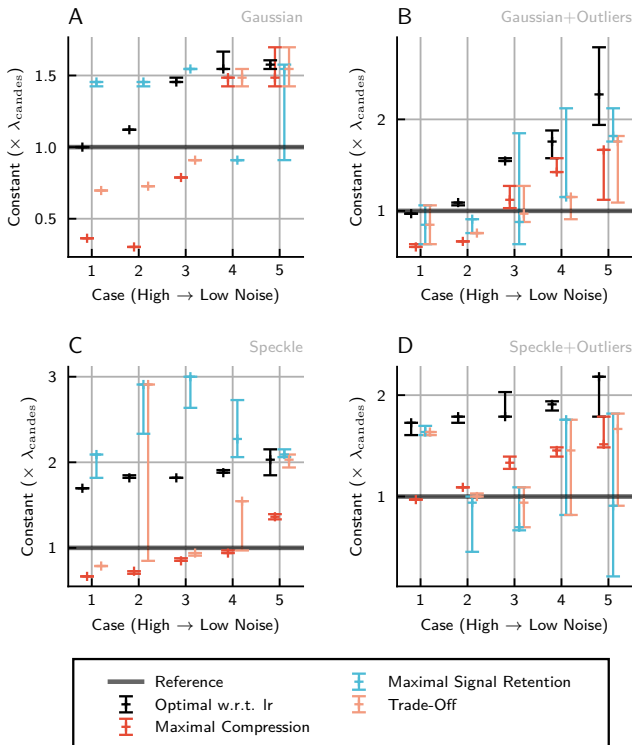


Figure 6.3 | λ setting for the cases: (A) Gaussian noise, (B) Gaussian noise with additional outliers, (C) Speckle noise, (D) Speckle noise with additional outliers. The + sign represents the mean of 10 simulations, and is given along with maximum and minimum values. The optimal λ setting with respect to the low-rank relative offset is also depicted.

We also remark that the spread for the optimal λ with respect to the low-rank relative offset (black) is small for large intensity noise and becomes larger for low intensity noise. This is caused by computational artifacts: the low noise level results in the “best fit” to become heavily dependent on the actual noise pattern. A second observation is that the spread for the speckle noise cases is larger than for the Gaussian noise cases. Reviewing some of the composite measures in **Fig. 6.4**, shows us that this spread is actually caused by a similar inconsistent behavior. For the Gaussian case (A in **Fig. 6.4**), the consistency over different simulations seems to be decent, while for the speckle with outlier noise case (B in **Fig. 6.4**), it seems to be “noisy”. We speculate that it is caused by one of measures being sensitive to this type. The latter could be explained through speckle noise being multiplicative and hence showing similar patterns as the underlying low-rank term which are consecutively identified as structure in B.

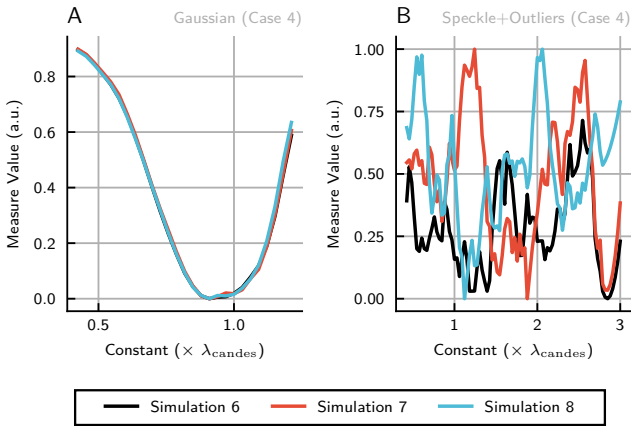


Figure 6.4 | Composite measures for maximal signal retention over 3 different simulations (of the 10 in total) for the cases: (A) the 4th Gaussian noise, (B) 4th Speckle noise with outliers.

The trade-off does not show consistent behavior. However, we observe that the maximal compression consistently sits below the low-rank relative offset optimal λ . Furthermore, for Gaussian and speckle noise it seems that the maximal signal retention case sits consistently above that same λ . We conclude with the fact that the framework should not only be considered application-dependent, but also noise type dependent. Note that this spread makes interpretation of the case study results as presented below more difficult as outliers might have large effects on the presented measures.

6.3.1 Maximal Compression

For the maximal compression case study, comparative results for the five ground truth information loss measures are depicted in **Fig. 6.5**. It shows the relative

difference with respect to the λ_{candes} setting. For simplicity, we will name the latter in this section the trivial setting. Firstly, we observe that in most cases the rank of the retrieved low-rank matrix is lower than the trivial setting. This is most of the time traded off against worsening the other measures, except for the *macs*. The latter could be logically explained due to a more compact notation leading to a more accurate description of those first components, and is especially true in the high intensity Gaussian noise cases. Besides, this property is advantageous for data visualization or exploration of the first few components.

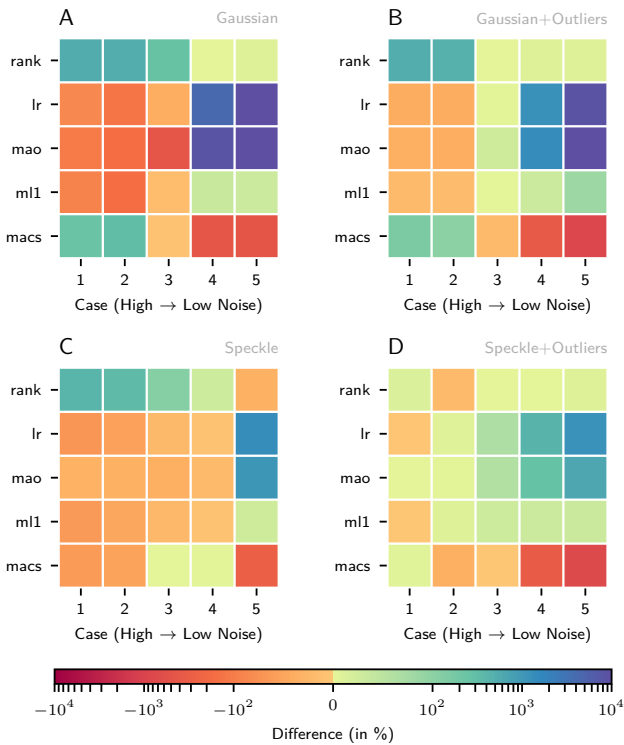


Figure 6.5 | Relative difference for the maximal compression case study with respect to the trivial setting for different ground truth information loss measures for the noise cases: (A) Gaussian noise, (B) Gaussian noise with additional outliers, (C) Speckle noise, (D) Speckle noise with additional outliers.

We also note that a lower rank seems to be advantageous for the low intensity noise cases of most noise types. This is remarkable as the rank and low-rank relative offset fit in those cases is already decent (see section 5.3.1, **Fig. 5.1**). Rather, these large relative differences are caused by the relative aspect. In absolute terms not much improvement is achieved.

6.3.2 Maximal Signal Retention

For the maximal signal retention case study, comparative results for the five ground truth information loss measures are depicted in **Fig. 6.6**. As expected, we see that for most cases the rank increases with respect to the trivial setting. Remarkably, this increase in rank does not always result in a better *lr* and *mao* fits, with exception for the speckle noise cases where a consistent improvement takes place. However, for the same speckle noise the *mI1* measure worsens. Whether the latter is a true phenomenon or a computational artifacts needs to be further researched.

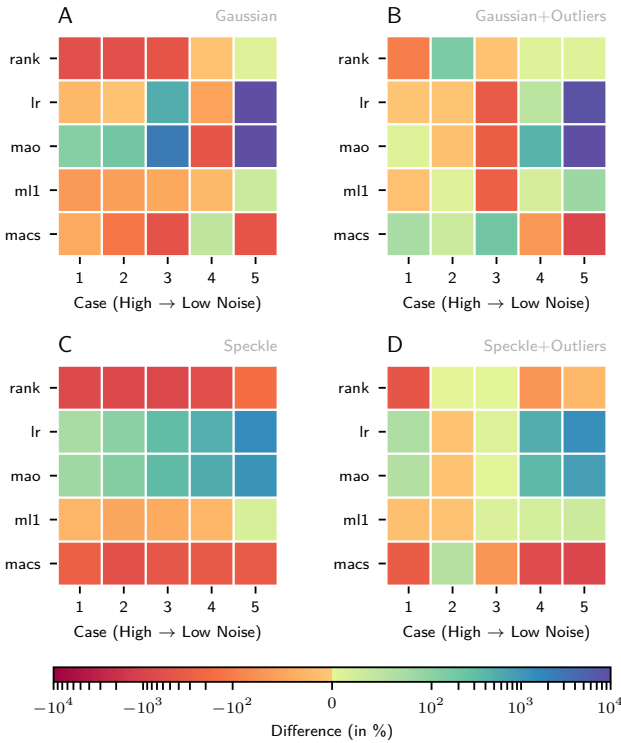


Figure 6.6 | Relative difference for the maximal signal retention case study with respect to the trivial setting for different ground truth information loss measures for the noise cases: (A) Gaussian noise, (B) Gaussian noise with additional outliers, (C) Speckle noise, (D) Speckle noise with additional outliers.

6.3.3 Trade-Off

Finally, for the trade-off case study, comparative results for the five ground truth information loss measures are depicted in **Fig. 6.7**. Here, we observe mixed results: some parts worsen and some ameliorate. This seems logical given its

combined nature. However, note that the variance on the λ setting for this case study is large and thus results can be skewed. Nevertheless, we have included this section for completeness.

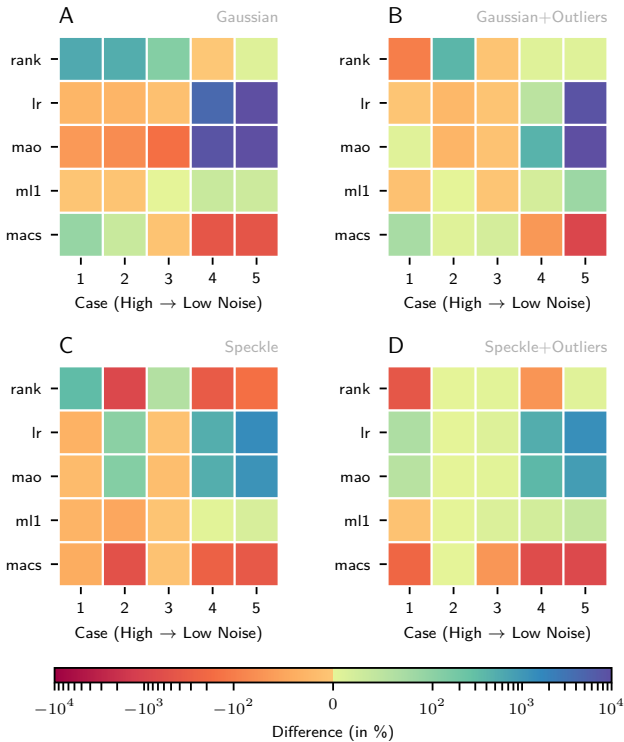


Figure 6.7 | Relative difference for the trade-off case study with respect to the trivial setting for different ground truth information loss measures for the noise cases: (A) Gaussian noise, (B) Gaussian noise with additional outliers, (C) Speckle noise, (D) Speckle noise with additional outliers.

6.4 Conclusions

We reviewed the use of several ground truth and heuristic information loss measures and suggested other measures that can be used concomitantly. In second instance, we proposed a framework for parameter optimization along with a grid search as solution to our framework's non-linear optimization problem. The associated computational cost is estimated to be $\mathcal{O}(qmn^2/qm\eta)$ flops, where q is the number of samples. We have shown its use through different case studies and analyzed both the parameter setting for those and their performance with respect to five ground truth information loss measures. The latter was performed

on a total of 20 noise cases: 5 Gaussian, 5 Gaussian with outliers, 5 speckle and 5 speckle with outliers. From this analysis we conclude that the measures are not only application, but also noise type dependent. Furthermore, we have shown that certain measures can be sensitive to the noise. This has to be taken into account when using the framework. We conclude with the fact that this framework is a first step towards a truly data-driven parameter setting.

Contributions We contributed the following:

- Showcased the use of a framework to define optimality
- Suggested the use of different ground truth and heuristic information loss measures
- Implemented several of those measures and defined different case studies
- Shown some of the vulnerabilities of this approach
- Made this study reproducible

Recommendations Further research could include:

- Taking measure variance into account / Include a sensitivity analysis for the measures
- Further exploit structure to increase the speed of the framework, e.g. similar to the Pareto front optimization¹¹
- Make the grid search adaptive to improve accuracy
- Review statistical dependency and correlation of measures
- Review the use of more or different measures
- Research the influence of the weighting setting
- Research the influence of filter on measures
- Research the influence of the normalization step: min-max normalization might influence the process
- Research the possibilities for making the framework data-driven
- Review the λ -setting for a synthetic application domain mimicking measurement matrix

Bibliography

- [1] Wold, S. Cross-validatory estimation of the number of components in factor and principal components models. *Technometrics* **20**, 397–405 (1978).
- [2] Valle, S., Li, W. & Qin, S. J. Selection of the number of principal components: the variance of the reconstruction error criterion with a comparison to other methods. *Industrial & Engineering Chemistry Research* **38**, 4389–4401 (1999).
- [3] Ulfarsson, M. O. & Solo, V. Selecting the number of principal components with sure. *IEEE Signal Processing Letters* **22**, 239–243 (2014).
- [4] Ganesh, A., Wright, J., Li, X., Candes, E. J. & Ma, Y. Dense error correction for low-rank matrices via principal component pursuit. In *2010 IEEE international symposium on information theory*, 1513–1517 (IEEE, 2010).
- [5] Candès, E. J., Li, X., Ma, Y. & Wright, J. Robust principal component analysis? *Journal of the ACM (JACM)* **58**, 1–37 (2011).
- [6] Banterle, N., Bui, K. H., Lemke, E. A. & Beck, M. Fourier ring correlation as a resolution criterion for super-resolution microscopy. *Journal of structural biology* **183**, 363–367 (2013).
- [7] Camacho, J., Smilde, A., Saccenti, E. & Westerhuis, J. All sparse pca models are wrong, but some are useful. part i: Computation of scores, residuals and explained variance. *Chemometrics and Intelligent Laboratory Systems* **196**, 103907 (2020).
- [8] Hore, A. & Ziou, D. Image quality metrics: Psnr vs. ssim. In *2010 20th international conference on pattern recognition*, 2366–2369 (IEEE, 2010).
- [9] Jombart, T., Devillard, S., Dufour, A.-B. & Pontier, D. Revealing cryptic spatial patterns in genetic variability by a new multivariate method. *Heredity* **101**, 92–103 (2008).
- [10] Snoek, J., Larochelle, H. & Adams, R. P. Practical bayesian optimization of machine learning algorithms. *arXiv preprint arXiv:1206.2944* (2012).
- [11] Van Den Berg, E. & Friedlander, M. P. Probing the pareto frontier for basis pursuit solutions. *SIAM Journal on Scientific Computing* **31**, 890–912 (2009).
- [12] Lin, Z., Chen, M. & Ma, Y. The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. *arXiv preprint arXiv:1009.5055* (2010).
- [13] van der Walt, S. *et al.* scikit-image: image processing in Python. *PeerJ* **2**, e453 (2014). URL <https://doi.org/10.7717/peerj.453>.

- [14] Savitzky, A. & Golay, M. J. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry* **36**, 1627–1639 (1964).
- [15] Virtanen, P. *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17**, 261–272 (2020).
- [16] Van Rossum, G. & Drake, F. L. *Python 3 Reference Manual* (CreateSpace, Scotts Valley, CA, 2009).

Chapter 7

Acceleration and Scaling of PCP

After this chapter you will be able to

- perform a computational cost analysis for PCP
- understand the challenges related to accelerating PCP
- understand the challenges related to scaling PCP
- review the performance of different SVDs with respect to PCP

7.1 Introduction

This chapter essentially focuses on obtaining PCP's decomposition terms more efficiently with an accuracy comparable to the originally proposed algorithm¹. This is necessary for handling larger measurement matrices, but also for our framework, presented in the previous chapter, to be practically applicable. One can think of performing these calculations in memory and compute resource restricted environments such as your personal laptop or a high-performance computing cluster. In the second part of this chapter, we propose an out-of-memory implementation of PCP able to scale to measurement matrices that run into the terabyte range. The latter enables us to step up our game from small toy problems, as reviewed in the previous chapters, to measurement matrices commonly found in practical applications such as for imaging mass spectrometry and astronomical imaging. A case study thereof is presented in chapter 8.

7.1.1 Goals and Objectives

The aim of this chapter is to improve IALM, a commonly used implementation of PCP, to enable analysis on real-life matrices. Initially, we review the acceleration of PCP through its time cost. This is achieved via a computational complexity analysis, including a review of the underlying SVD. This ensures that the available resources are used as efficiently as possible, but also that the right resources are allocated. To achieve the latter, we focus on Python as programming language by using readily available packages and libraries. In a second instance, we review the process of scaling the algorithm to allow for measurement matrices that are running into the terabyte range. This will be achieved through the use of the `dask` package², enabling native scaling in Python. These findings subsequently formulate the answer to research subquestion A.3.

A. Can matrix decomposition methods based on the ℓ_1 -robustness principles efficiently reduce massively multivariate and high-dimensional spectral imaging measurement matrices to a lower-dimensional representation while incurring minimal information loss?

A.3. Can a memory and compute resources restricted version of such approaches be developed, and how does it perform with respect to the original batch version that assumes the entire data set can be held in memory?

7.1.2 Overview

In section 7.2, the experimental design is elaborated on, including algorithm selection, the analysis of the computational cost and the techniques used for accelerating and scaling PCP. Finally, it also provides the specifications of the synthetic measurement matrix that is used throughout this chapter, along with the measures for comparison and the additional settings. In section 7.3 and 7.4, we respectively present our accelerating performance analysis and scaling performance analysis. Finally, our conclusions are presented in 7.5, along with the contributions and recommendations.

7.2 Experimental Design

The experimental design is comprised of the algorithm selection and its computational cost analysis, accelerating and scaling techniques, the presentation of the synthetic measurement matrix and measure selection. Finally, additional settings are provided in the last subsection.

7.2.1 Algorithm Selection

Firstly, we select PCP as the method to be accelerated and scaled. This is a logical choice as it forms the basis for many ℓ_1 -robustness methods. More specifically, its algorithmic build-up is similar to e.g. SPCP and SIPCP, and as such the proposed scaling and acceleration techniques can be used for those methods too. As a side note, a second challenge for SPCP is that for larger measurement matrices its cost is dominated by a sorting algorithm^a.

Different algorithms are proposed to solve the convex formulation of PCP (see Eq. 4.45), such as an interior point method, accelerated proximal gradient method (APG) and augmented Lagrangian multiplier method (ALM). However, the inexact augmented Lagrangian multiplier method (IALM), also known as alternating directions, is preferred. It alleviates some of the time and memory costs related to PCP, mainly due to its reduced number of SVDs, while improving accuracy of the recovered terms with respect to APG and ALM³ (its validity and optimality are guaranteed by theorem 2 and 3³). This implementation has furthermore been shown to be fastest converging among a whole range of PCP algorithms on a selection of practical examples^{4,5}. Therefore, we considered IALM as starting point for acceleration and scaling.

MATLAB implementations are available for different PCP algorithms, such as IALM, under open-source license^b. However, our experience learns us that those algorithms should be considered as proof-of-concepts.

7.2.1.1 A Theoretical Computational Cost Analysis

In this section, we propose a theoretical computational cost analysis. We want to emphasize that the theoretical cost and the actual cost of the implementation might still differ^c. This can be attributed e.g. to the BLAS level of the algorithmic operations, linear algebra package and math kernel, data locality, etc. In case of an out-of-memory implementation, i.e. when the memory cost is higher than actual memory size, or implementations for distributed systems this time cost might actually be dwarfed by the disk write/read time or inter-node communication bandwidth. Therefore, this cost analysis should be considered as a ballpark exercise.

“However, as we have repeatedly seen, flops are now a poor metric for performance. It matters whether flops are in compute-intensive Level 3 BLAS or not.” - Dongarra et al.⁶

^aour implementation of SPCP via ADMIP is already an improvement with respect to the available code through a cleaner object-oriented Python implementation with different optimized parts.

^bfrom <https://github.com/andrewssobral/lrslibrary>

^cfor practical considerations a cpu or memory profiler might be useful: some example profiles of PCP are given in Fig. A.3 and Fig. A.4

For most implementations of PCP the memory cost is $\Omega(mn)$. Therefore, we consider the time cost, i.e. $\mathcal{O}(\cdot)$, as it is most restraining for acceleration and to some extent the scaling. In **Tab. 7.1**, the cost per iteration is given along with the cost for each individual calculation step within that iteration and its corresponding LAPACK⁷ routine. In the last column of **Tab. 7.1**, the relative cost is reflected for a measurement matrix $M \in \mathbb{R}^{m \times n}$ of size $m = 10000$, $n = 100$ and rank $r = 10$. The latter reflects the weight of the individual calculation steps on the iteration.

Table 7.1 | Computational cost per iteration of IALM. A description of the different steps within the iteration are given along with their corresponding LAPACK routine, related time cost in absolute terms and time cost in relative terms. For both costs, we assume that $m \gg n$. S is a soft-threshold operator. [†] for values of $m = 10000$, $n = 100$ and $r = 10$.

Description	LAPACK Routine	Cost $\mathcal{O}(\cdot)$	Relative Cost [†]
<i>Solve $A_{k+1} = \arg \min_A L(A, B_k, Y_k, \mu_k)$</i>			
$Q = M - B_k + \mu_k^{-1} Y_k$		$3mn$	0.005
$[U, S, V] = \text{svd}(Q)$	dgesvd	$6mn^2 + 16n^3$	0.956
$A_{k+1} = US_{\mu_k^{-1}}(S)V^T$	dgemm	mnr	0.016
<i>Solve $B_{k+1} = \arg \min_B L(A_{k+1}, B, Y_k, \mu_k)$</i>			
$B_{k+1} = S_{\lambda \mu_k^{-1}}(M - A_{k+1} + \mu_k^{-1} Y_k)$		$4mn$	0.006
<i>Update Lagrangian multipliers</i>			
$Y_{k+1} = Y_k + \mu_k(M - A_{k+1} - B_{k+1})$		$4mn$	0.006
<i>Calculate convergence criteria</i>			
$\ M - A - B_{k+1}\ _F$	dlange	$4mn$	0.006
$\ B_k - B_{k+1}\ _F$	dlange	$3mn$	0.005
<i>Total cost per iteration</i>		$6mn^2 + 16n^3 + (18 + r)mn$	1

It is observed from the absolute and relative cost columns in **Tab. 7.1** that the SVD is the bottleneck for the implementation. As already suggested^{1,3}, this cost can be alleviated by using a so-called truncated SVD as available from PROPACK⁸. Unfortunately, no stable wrappers exist for PROPACK in Python. Nevertheless, SciPy⁹ offers a similar SVD from ARPACK¹⁰ in Python. For an overview of Arnoldi/Lanczos SVD methods and packages, we advise the reader to *A survey of software for sparse eigenvalue problems*¹¹. Our practical experience teaches us that IALM with such SVDs has convergence issues due to triplets: eigenvalues/singular values that are very close to each other. These especially occur for large measurement matrices and lead to a break down of the convergence and accuracy of IALM.

Therefore, we identify four possible solutions to address the discussed bot-

tleneck: (a) change the inherent algorithm; (b) address the SVD cost; (c) use better/faster linear algebra routines e.g. via GPU; (d) alleviate costs by using a lower precision. Though our Python library HannibalSpecter has a divide-factor-conquer implementation¹² available (first solution type), we only review the last three solutions.

7.2.2 Accelerating Techniques

As indicated, the SVD is the bottleneck for acceleration of the algorithm. Hence, we propose to solve this issue on a practical level by:

1. **Addressing and reducing the cost of the SVD:** identify and use different SVD methods readily available in Python packages, and comparing their outcome in a practical setting. Recently, a similar type of research has been conducted for the influence of the SVD on PCA¹³.
2. **Use better/faster routines:** the use of an optimized kernel¹³, e.g. Intel Math Kernel, but also the use of linear algebra packages such as PLASMA and MAGMA¹⁴ exploiting parallelism and GPUs. The latter does not guarantee advantages over a CPU based approach due to the bidiagonal reductions of the GESVD (inherently iterative). Other SVD methods, however, might take advantage of the GPU. Hence, we provide a fully GPU-based IALM version via the CuPy package¹⁵.
3. **Alleviate costs by going to lower precision:** switching from the default 64 bits precision to a 32 bits notation, both alleviates memory and CPU use. Of course, the convergence of IALM should be guaranteed in 32 bits^b.

7.2.2.1 SVD implementations

We identify SVD implementations that are readily available from Python libraries and implemented some ourselves in HannibalSpecter. Note that this selection is not exhaustive and other implementations might still be available within Python^c. An overview is given in **Tab. 7.2**. Each implementation is attributed to a method type corresponding to its underlying methodology. Note that within method types different implementations exist, e.g. for the randomized SVD type different iterators exist. Hence, those implementations might still result in distinct outcomes.

^ae.g. see blog post <http://markus-beuckelmann.de/blog/boosting-numpy-blas.html>

^ban interesting blog post on precision: <https://librestats.com/2015/04/17/floating-point-arithmetic-is-hilarious/>

^ce.g. a recently found, but not included SVD method is <http://www.cs.wm.edu/~andreas/software/>

Table 7.2 | Selection of SVD implementations. Besides the name, the processing unit type, the underlying linear algebra package and python library are given. Finally, a time cost estimation is given and the method type of the SVD. For the time cost estimation, we assumed that $m \ll n$. [†] we could only find the computational cost for a square $n \times n$ matrix, where S is the number of sweeps for the Jacobi method.

Name	Proc.	Lin. Alg. Package	Python Library	Cost $\mathcal{O}(\cdot)$	Method Type
arpack	CPU	ARPACK	SciPy	$r m n$	Arnoldi/Lanczos iterations
cupy_gesvd	GPU	cuSOLVER	CuPy	$m n^2$	Bidiagonal reduction with QR iteration
cupy_randomized	GPU	cuSOLVER	Hannibal Specter	$r m n$	Randomized SVD
fbpca	CPU	LAPACK	fbpca	$r m n$	Randomized SVD
lobpcg	CPU	LAPACK	SciPy	$r m n$	Block Preconditioned Conjugate Gradient Method
numpy_gesdd	CPU	LAPACK	NumPy	$m n^2$	Divide and Conquer
pytorch	CPU	LAPACK	PyTorch	$m n^2$	Divide and Conquer
pytorch_cuda	GPU	cuSOLVER	PyTorch	$S n^{3\dagger}$	One-sided Jacobi method
pytorch_randomized	CPU	LAPACK	PyTorch	$r m n$	Randomized SVD
randomized	CPU	LAPACK	Scikit-learn	$r m n$	Randomized SVD
recycling_randomized	CPU	LAPACK	Hannibal Specter	$r m n$	Randomized SVD
scipy_gesdd	CPU	LAPACK	SciPy	$m n^2$	Divide and Conquer
scipy_gesvd	CPU	LAPACK	SciPy	$m n^2$	Bidiagonal reduction with QR iteration
svdecon	CPU	LAPACK	Hannibal Specter	$m n^2$	Eigenvalue Decomposition

7.2.3 Scaling Techniques

For scaling purposes, we identify two cases: (1) when the measurement matrix is large but sits comfortably in memory, a.k.a. the in-memory case; (2) when memory is too small to house all matrices needed for IALM, a.k.a. the out-of-memory case. For the respectively first case, the main bottleneck remains the computational cost of the SVD, while for the second case an extra complexity raises in the form of the memory cost. For both cases, we advise the interested reader a review on *The singular value decomposition: Anatomy of optimizing an algorithm for extreme scale*⁶.

7.2.3.1 In-Memory Implementation

For in-memory implementations we can fall back on the same solutions as earlier presented due the bottleneck being the same as for acceleration. However,

while scaling in-memory some additional problems appear. For example, the divide-and-conquer approach as used in the LAPACK routine GESDD becomes out-of-memory as it requires large submatrices for its matrix partitioning step⁶. Experiments point out that this happens in Python at around 30 GB of raw data with a column-row ratio of around 10 (see error message **Fig. A.5**). The convergence with Arnoldi/Lanczos methods also breaks down^a. Even when the tolerance is decreased, number of Lanczos vectors are increased, a starting vector is provided, the number of singular values are increased, and the maximum number of iterations is increased, this problem is not alleviated. We believe this is due to triplets of singular values becoming more apparent in larger measurement matrices. Hence, for the in-memory scaling we mainly focus on randomized methods to alleviate the costs. Furthermore, randomized methods enable leveraging between accuracy and speed through their iterators and can be used for out-of-memory scaling too.

7.2.3.2 Out-Of-Memory Implementation

An out-of-memory implementation can be performed both on a single machine or a distributed system. For the latter, many linear algebra packages are available, such as ScaLAPACK¹⁶, DPLASMA¹⁷, and P-ARPACK¹⁸. However, we did not consider to walk that route as dedicated distributed packages are not appropriate for the single machine that we have at our disposal^b. Therefore, we review the use of the `dask` package², which enables native scaling in Python. Furthermore, it is also deployable on distributed systems and even supercomputers^c.

Dask Dask is a Python package that allows users to produce out-of-memory implementations through lazy calculations. By chunking matrices into smaller blocks, it performs actions on blocks. Hence, it only allows for the number of blocks to be in memory that memory is capable of and the calculations require. Additionally, it has a tall-and-skinny QR implementation¹⁹ available, such that a randomized SVD can easily be performed for tall-and-skinny matrices of the terabyte size. Note that it is of the MapReduce kind, such that it is not expected to be more efficient than implementations available from distributed linear algebra packages. Furthermore, their original randomized SVD implementation had several smaller issues which we helped resolve^d. Finally, note that `dask` does not scale the SVD endlessly. It will therefore be mainly bound by the sampling size.

^aan example of such event is given in **Fig. A.6**

^be.g. in-memory SVD implementations outperform their distributed counterparts in efficiency (e.g. see ScaLAPACK, LAPACK comparison⁶)

^c<https://www.olcf.ornl.gov/2020/04/29/new-data-analysis-software-tools-aim-at-addressing-the-convergence-of-hpc-and-ai/>

^dsee <https://github.com/dask/dask/issues/6757>

Alleviating Memory The original algorithm requires around 5 to 10 times the matrix size in peak memory. To reduce some of this cost we considered the following solutions:

- Going from a 64 bits to a 32 bits notation
- Removing the second convergence criterion (on the noise matrix), releases B_k to be saved each iteration
 - In case this criterion is necessary, save A_k in decomposed form and replace B_k by $M - A_k$
- Keep A_{k+1} in its decomposed form, i.e. not multiplying it out

Note that the latter two solutions increase computational costs, but when disk read/write is dominant, this saves time.

7.2.4 Synthetic Measurement Matrix

For the experiments we used a slightly different synthetic low-rank term, and further simplified the noise term to solely cover fixed salt and pepper noise. Both guarantee that the exact recovery conditions are met, and that we can utilize the low-rank relative offset to inspect convergence behavior.

7.2.4.1 Synthetic Low-Rank Term

We use a similar low-rank as proposed in the original paper¹: $X \in \mathbb{R}^{n \times r}$ and $Y \in \mathbb{R}^{m \times r}$ are matrices with entries independently sampled respectively from an $\mathcal{N}(0, 1/n)$ and $\mathcal{N}(0, 1/m)$ distribution.

7.2.4.2 Synthetic Noise Term

We solely use sparse noise term, as mentioned in the previous chapters. It consists of salt and pepper noise with an incidence of 5% of the total number of entries, with a random basis, and a minimum value of -1 and maximum value of 1 . This is also similar to the approach as originally proposed¹.

7.2.5 Measure Selection

We compare the implementations on mainly two measures: low-rank relative offset and running time.

Low-Rank Relative Offset We use the low-rank relative offset, similar to previous chapters, to inspect whether the implementation convergences to the ground truth. It is defined as

$$\epsilon_{\text{low-rank}} = \frac{\|A - \hat{A}\|_F}{\|A\|_F} \times 100\%. \quad (7.1)$$

Running Time We also use the heuristic of timing of the IALM implementation to measure the computational complexity. Of course, many factors have already been brought up that might influence this running time. Furthermore, it also depends on the machine load, concurrently available cores, and other underlying settings. Hence, we should again consider this measurement as a ballpark estimate.

7.2.6 Additional Settings

Parameter Setting For IALM the λ parameter was fixed to the usual $1/\sqrt{m}$, ρ was set to 1.6 for the in-memory experiments and 1.3 for the out-of-memory ones.

Singular Values Estimation The estimation of the number of singular values, required for truncated and randomized SVD methods, is estimated as

$$sv_{i+1} = \begin{cases} \min(svp + dn, n) & \text{if } svp = sv_i, \\ svp + 1 & \text{if } svp < sv_i, \end{cases} \quad (7.2)$$

where svp equals the rank of A_k and sv is the estimated number of singular values. We use $sv_0 = 10$, close to what was originally proposed³. For all matrices in this chapter we use $d = 0.05$. However, for larger matrices a smaller value is advised, e.g. $d = 0.005$. The number of oversamples was set to 10, if not stated otherwise.

Stopping Criteria The stopping criteria for PCP are set to

$$\begin{aligned} \frac{\|M - A - B\|_F}{\|M\|_F} &\leq 1 \times 10^{-7}, \\ \mu \frac{\|\hat{B} - B\|_F}{\|M\|_F} &\leq 1 \times 10^{-5}, \end{aligned} \quad (7.3)$$

concurrently to a fixed number of 100 iterations. These stopping criteria are named respectively convergence criteria 1 and 2.

Noise Term We used the scikit-image package²⁰ and the dask package² to generate the discussed noise patterns.

Computational Specification A precision of both 32 and 64 bits were used along with an OpenBLAS kernel. Most experiments were performed parallel in Python 3.8.3²¹ in solely and separate CPU based processes. We fixed the number of CPU threads, to prevent over-allocation and preserve comparability, to

OPENBLAS_NUM_THREADS = 1. For GPU-based calculations, we performed experiments in sequential order. Finally, Python package version of additional packages are given in **Tab. 7.3**.

Table 7.3 | Used Python package versions.

Package	Version
CUDA	11.2.r11.2
CuPy	8.4.0
fbpca	1.0
HannibalSpecter	0.0.1
Numpy	1.20.0
PyTorch	1.7.1 + cu110
SciKit-learn	0.23.1
SciPy	1.6.0

We noticed that in NumPy 1.18.5 there are problems with respect to the SVD, leading to a break down in convergence. We are not sure why this appears. We have extensively tested all other SVD implementation and our algorithm, such that we suspect this is a NumPy/SciPy related bug. In earlier versions and later versions such as 1.20.0 this problem does not appear. Finally, the computer system specifics are given in **Tab. 7.4**. In addition, the GPU used is an NVIDIA Quadro RTX 5000 with 16 GB of dedicated GPU memory.

Table 7.4 | Computer system specifications. The operating system, number of cores and their clock speed, and the size of available memory is stated.

OS	Cores	Clock Speed	Memory Size
Windows 10 Pro	52 (104 logical)	2.7 GHz (3.3 GHz overclock)	1.5 TB

7.3 Acceleration Performance Analysis

In this section, we analyze the performance of IALM with the different SVD methods, as listed in **Tab. 7.2**. Note that these results remain measurement matrix and computer bound due to e.g. the influence of m/n ratio, true matrix rank, but also CPU/GPU specifications and machine settings. In **Fig. 7.1**, we provide the results of our performance experiments in the form of violin plots for 32 (black) and 64 (red) bits. Each IALM with SVD method is simulated 100 times, resulting in distributional information, a minimum, a median, and a maximum. For this experiment, the number of iterations is fixed to 100 to better reflect the cost per iteration, reducing the relative overhead cost. We expect the latter not to influence the timing results drastically as a comparable convergence accuracy is reached in most cases. Although, some truncated and randomized methods might obtain a better convergence rate. (see Appendix B, **Fig. A.7**). Finally, for all experiments we consider IALM with the LAPACK GESVD as standard, and all implementations with prefix `cupy` are fully GPU-based. Note that for the latter only `cupy_randomized` and `cupy_gesvd` are fully GPU-based SVD methods.

We made the following observations:

- convergence of IALM is guaranteed with both 32 and 64 bits;
- all methods converge to ground truth, except for the randomized method of PyTorch, this is probably due to its inaccurate power iterator;
- most 32 bit versions are around a factor 2 faster than the 64 bit versions;
- NumPy and SciPy GESDD implementation seem to be differently accessed;
- large difference in timing for the GESDD between 32 and 64 bits, are probably be due to the overhead of the creation of a large submatrix;
- `arpack` and `randomized` (with QR iterator) seem to be fastest and most accurate;
 - convergence break down for `arpack` only occurs through triplets, here not the case;
- GPU-based versions are generally slower than CPU-based versions, except for the `pytorch_randomized`. We believe this phenomenon occurs due to initiation, communication and/or overhead costs, and the small matrix size leading to less gains with respect to CPU-based versions;
 - for larger matrices, GPU-based implementations are generally several factors faster;
 - CuPy has no clean transition to PyTorch requiring the matrix to be routed via the CPU;

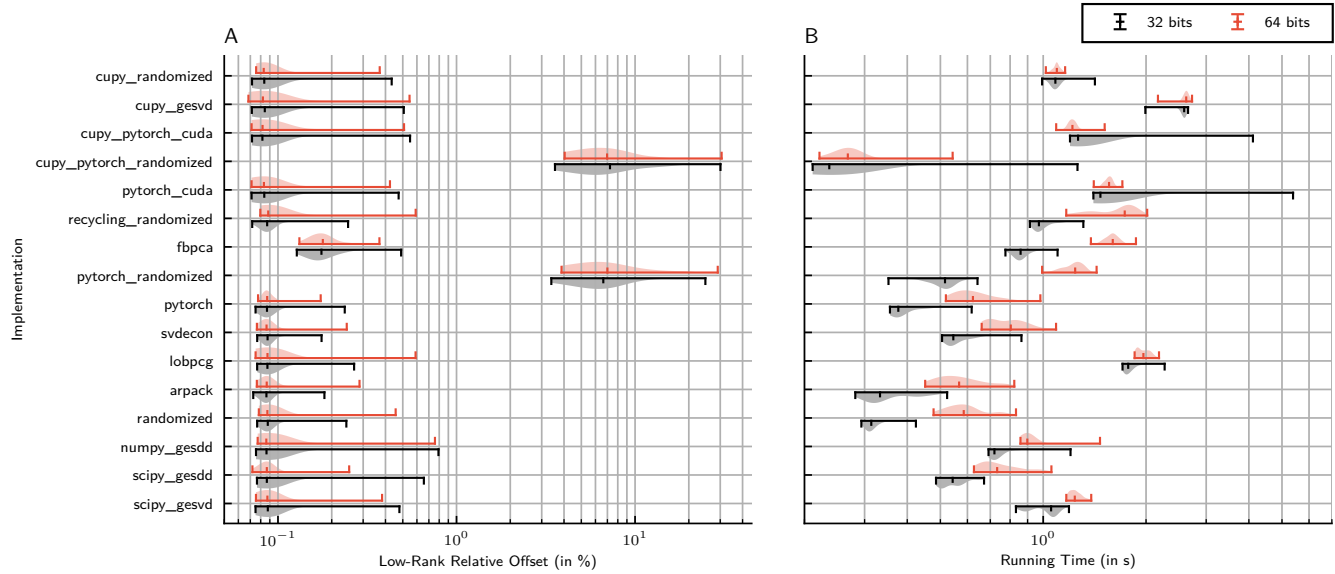


Figure 7.1 | Results for IALM with different SVD implementations. In (A) the low-rank relative offset is represented and in (B) the running time is given. All experiments are repeated 100 times, and their corresponding violin plots including densities, minimum, median and maximum are given for 32 and 64 bit implementations. Note that both (A) and (B) are plotted on a logarithmic scale. A fixed number of 100 iterations was used as stopping criterion. The input matrix $A \in \mathbb{R}^{1000 \times 1000}$ is generated as explained in the experimental design and is of rank 10.

7.3.1 A Randomized Approach

As earlier remarked, the randomized SVD approaches offer a similar cost and speed as Arnoldi/Lanczos based methods, but have the advantage that accuracy can be traded off through a so-called iterator. Nevertheless, during practical experiments on larger matrices we remarked a similar detrimental behavior as with `arpack`. We believe that similar to ARPACK's triplets, the accuracy of some randomized SVDs is not adequate to safeguard the convergence. The latter is dependent on the spectrum of singular values and can be improved through different iterators^{22,23}. More specifically, this breakdown behavior for randomized methods was observed to be caused by the sampling step sometimes missing a substantial part of the range of the measurement matrix. Therefore, a remedy was sought: recycling singular vectors.

“The intuition behind these techniques (ed. iterators) is that the singular vectors associated with small singular values interfere with the calculation, so we reduce their weight relative to the dominant singular vectors by taking powers of the matrix to be analyzed.” - Halko et al.²²

7.3.1.1 Recycling

Recycling can be seen as a warm start of the SVD method²⁴. Its main assumption is that the range of the low-rank term changes slowly over iterations²⁵, such that the range space of a previous iteration can be used as a starting point for the next iteration. This can be achieved through a partial orthogonalization of the new range with respect to the left singular vectors of the previous iteration²⁶ or through an orthogonal projection of the new sampled range space with respect to the left singular vectors of the previous iteration²⁵. Although both methods are similar, the respectively first is provided with an error bound, extending the originally proposed idea to build up the range space in an iterative manner²². The error associated to the randomized method can be obtained from

$$\|A - QQ^T A\|_2 \leq \epsilon, \quad (7.4)$$

where Q is the estimated range space associated to the dominant singular values and ϵ is the error bound. Although, the latter is suggested to be implemented in level 3 BLAS, neglecting the small overhead associated to it, no implementations were readily available. To simplify the above mentioned approaches, and simplifying the later out-of-memory implementation, we recycled the right singular values of the previous iteration in sampling step as

$$\Omega = [V \quad \Omega_n], \quad (7.5)$$

where $V \in \mathbb{R}^{n \times r}$ contains the right singular vectors of the previous iteration, and $\Omega_n \in \mathbb{R}^{n \times (s_{v_{i+1}} - r)}$ is a random sampling matrix. The estimated range space can then be obtained from $Q = A\Omega$, with or without iterator. We argue that this approach is comparable to the other methods as long as Ω_n sufficiently samples the space orthogonal to V , which is guaranteed by its random nature. The mathematical proof is considered for future work. For now, we show its use heuristically.

7.3.1.2 What iterator is best?

Besides recycling singular vectors, an iterator can be used to further reduce the error bound. Especially, when the singular value spectrum is slowly decaying²². Therefore, different iterators have been proposed, such as the power iterator, QR and LU iterators. However, while iterators increase the accuracy of the obtained approximation, they also increase the computational cost.

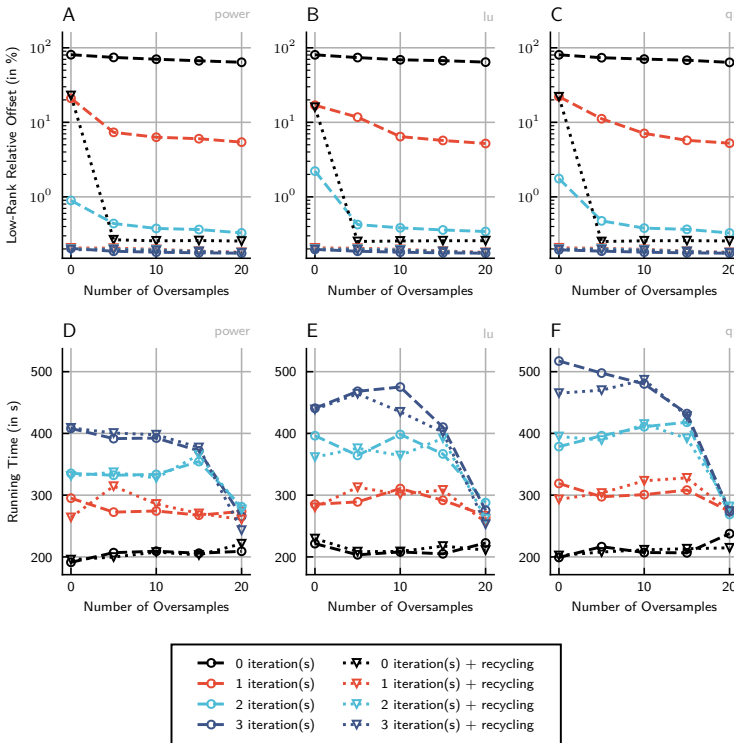


Figure 7.2 | Results for IALM with randomized SVD with different iterators and settings. IALM is set with a fixed number of 50 iterations on measurement matrix $M \in \mathbb{R}^{5000 \times 500}$ of rank 50. The first row of plots represent the low-rank relative offset for (A) the power iterator; (B) the LU iterator; (C) the QR iterator. The second row of plots represent the running time for (A) the power iterator; (B) the LU iterator; (C) the QR iterator.

Hence, it is key to use an appropriate iterator, and define an appropriate number of iterations (not to be confused with the iterations of IALM). The latter, however, is measurement matrix bound^a.

In Fig. 7.2, we provide the low-rank relative offset and running time of IALM with randomized SVDs with different settings, i.e. iterators, number of iterations, and number of oversamples. Furthermore, we compare those settings on both a simple iterator and the iterator including recycling. Furthermore, observe that when the number of iterations is set to 0, no iterator is in place and similar results should be across the different iterators. We observe the following:

- when the number of oversamples is set to 0, there is a bias. This is logical as Ω_n will be empty for those cases such that the estimated range is locked by the previous estimation;
- LU and QR iterators behave similarly. However, it has been proven that LU is slightly better in time cost while obtaining similar error bounds²³;
- recycling safeguards the convergence for cases;
- recycling does not affect the time cost;
- for all cases no iterator is fastest, but results in the largest error;

Finally, the attentive reader remarks that the running time decreases for larger numbers of oversamples. We believe this is not a true discovery, but rather a computational issue. Those samples are namely obtained near the end of the parallel calculations, such that more cores are available. Hence, we believe that somehow computations are not completely performed on only one core. However, further investigation is needed. For earlier presented results in Fig. 7.1, this phenomenon is likely to be minimal due to the repeated number of calculations.

7.4 Scaling Performance Analysis

In this section, we compare the influence of the row-column ratio, m/n , and size of the measurement matrix via row size m on the convergence of IALM with different SVD methods. In the first part, we review in-memory scaling, while in the second part we compare out-of-memory implementations.

7.4.1 In-Memory Implementations

As already mentioned, when it is possible to stay in-memory, stay in-memory. This can be achieved through partly alleviating costs, as suggested in section 7.2.3.2. Nevertheless, when scaling the measurement matrix, the standard LAPACK GESVD tends to become too expensive to perform. Therefore, truncated

^asee Fig. A.8 for such an example

and randomized SVD methods can come to the rescue. In Fig. 7.3, we present results for 5 methods: `scipy_gesvd`, `scipy_gesdd`, `arpack`, and two randomized implementations with QR iterator, respectively performing one and two iterations and both including recycling.




Figure 7.3 | Results for IALM with different SVD methods. The following SVD methods are presented: (A) `scipy_gesvd`; (B) `scipy_gesdd`; (C) `arpack`; (D) randomized with 1 QR iteration; (E) randomized with 2 QR iterations. These are compared on low-rank relative offset, final rank, rank flips, running time, number of iterations, and both convergence criteria as earlier specified. All experiments are performed in 32 bits precision and the matrix rank is set to 10% of $\min(m, n)$, with a minimum of 1.

Besides the low-rank relative offset, we also provide the rank at final iteration, the number of rank flips that occur^a, the running time, the number of iterations

^athis reflects the number of times the rank over the IALM iterations has decreased, instead of monotonically increased

required for convergence, if achieved, and both convergence criteria at the final iteration. We observe the following for:

- **Low-Rank Relative Offset (1):** the value of around 10% on small, skinny matrices of all methods with respect to ground truth (red left top corner), is due to methodical nature. Due to the small matrix size, the true low-rank term inherently becomes more sparse, hence violating incoherence principles;
- **Low-Rank Relative Offset (2):** truncated and randomized methods obtain better results for fat matrices, but fail when matrices are fat and large (bottom right corner). The latter is believed to be due to theorem 2 and 3^3 not being met with respect to the μ parameter. This can be resolved by adjusting the ρ -parameter accordingly, influencing the convergence rate;
- **Final Rank:** all IALM versions converge to the right final rank;
- **Rank Flip** for truncated and randomized rank flips occur. However, given the right final rank and the small low-rank relative offset in those areas, it is believed to be convergence behavior, similar to a gradient descent dancing around a minimum, hence not detrimental. It might indicate that for those settings the algorithm converged after fewer steps;
- **Running Time (1):** GESVD is by far the slowest method especially when matrices become large and fat;
- **Running Time (2):** GESDD performs better time-wise. However, it has some problems with memory as indicated earlier, making it inadequate for larger matrix sizes;
- **Running Time (3):** One QR iteration is slightly faster than 2, but overall both are in the same ballpark range;
- **Number of Iterations and Convergence (1):** for most matrices all versions run to the 100 iteration mark. This is probably due to the other two convergence criteria being too rigid. This can be alleviated by defining them matrix dependent;
- **Number of Iterations and Convergence (2):** randomized and truncated versions generally achieve better convergence, this has also been remarked previously for other randomized methods²⁷;
- **Number of Iterations and Convergence (3):** as earlier mentioned, the truncated and randomized methods do not converge for tall fat matrices;

7.4.2 Out-Of-Memory Implementations

We have two out-of-memory implementations available, one with a `numpy_gesdd` at its heart and a randomized method with recycling step. The `numpy_gesdd`-version is roughly bound by the memory need for a single SVD. In comparison, for the in-memory implementations memory is concurrently required for the other matrices, e.g. Y and B . The randomized approaches can scale further, and are bound by the tall-skinny QR¹⁹ at the heart of the QR iterator. Therefore, the sampling step has a large influence on the eventual memory bound and the speed of the overall SVD and algorithm. Note, that we only implemented IALM with randomized methods based on the QR iterator, as no out-of-core LU decomposition is readily available in the `dask` package. Nevertheless, this already allows us to reach the Terabyte scale on our machine.

We provide a similar comparison as for the in-memory implementations, but simplify the number of results. In **Fig. 7.4**, we present the four implementations: (A) in-memory `scipy_gesvd`; (B) out-of-memory `numpy_gesdd`; (C) out-of-memory randomized with 1 QR iteration; (D) out-of-memory randomized with 2 QR iterations. Note that all calculations are performed in memory to be able to fairly assess and compare all methods without becoming data locality dependent.




Figure 7.4 | Results for IALM with different SVD methods. The following SVD methods are presented: (A) in-memory `scipy_gesvd`; (B) out-of-memory `numpy_gesdd`; (C) out-of-memory randomized with 1 QR iteration; (D) out-of-memory randomized with 2 QR iterations. These are compared on low-rank relative offset, final rank, rank flips and running time. All experiments are performed in 32 bits precision and the matrix rank is set to 10% of $\min(m, n)$, with a minimum of 1. We fixed the number of iterations to 50.

We observe the following:

- **Low-Rank Relative Offset (1):** all versions converge, except for small, skinny matrices (red left top corner) as earlier discussed due to violating incoherence conditions;
- **Low-Rank Relative Offset (2):** out-of-memory implementations hit a lower-bound, this is believed to be due to a different ρ setting, 1.3 instead of 1.6. Nevertheless, it resolves the convergence for fat matrices. Further research on the influence of ρ -parameter is required;
- **Final Rank:** all versions converge to the right final rank;
- **Rank Flip:** remarkably no rank flips occur. This confirms that in **Fig. 7.3** the rank flips occur at final iterations, and are not detrimental.
- **Running Time (1):** out-of-memory implementations have a small overhead with respect to the in-memory implementation. The overhead becomes especially apparent for smaller matrices (left upper triangle) for all out-of-memory methods;
- **Running Time (2):** out-of-memory implementations prove to be competitive with the in-memory version for larger and fatter matrices (right lower triangle);

7.5 Conclusions

We have reviewed the theoretical computational cost for IALM and provided different improvements with respect to the standard implementation. Additionally, we have addressed these improvements, time- as well as memory-related. In a second phase, we have practically implemented these solutions and compared them over different settings and criteria.

For acceleration purposes, we provided results for different SVD methods in both 32 and 64 bits precision. We shortly reviewed a randomized approach and analyzed the different iterators, along with proposing a simplified recycling step. The recycling step is heuristically found to be contributing to the iterative stability while leaving the time cost unchanged and remaining simple to implement. To conclude the acceleration section, we advise a 32 bits precision randomized implementation with recycling step and appropriate iterator, depending on the singular value spectrum. For smaller matrices a truncated SVD in the form of ARPACK or PROPACK is also possible, but one has to closely monitor the convergence process for those methods. For larger matrices, GPU-based approaches based on randomized SVD methods show to be promising.

For scaling purposes, we have assessed different truncated and randomized methods and provided a comparison with respect to the standard in-memory

GESVD. We conclude that scaling is possible with both ARPACK and randomized methods, but that the latter are preferred as they easily allow to trade off speed and accuracy. It was found that the ρ and μ parameter are of importance to safeguard convergence for those cases. For out-of-memory implementations, enabling scaling to the terabyte range, it was shown that randomized methods are a viable solution. Our analysis indicates that such implementations have reasonable overhead when used in-memory with respect to the in-memory LAPACK GESVD and thus can be considered optimally implemented.

To conclude, this chapter contributed to the practical use of ℓ_1 -robustness methodology in real-life applications by proposing PCP implementations enabling analysis on measurement matrices in reasonable time achieving equal or even better low-rank relative offsets. This ultimately leads to overall reduced time and computational costs for any project using PCP in its pipeline. Furthermore, we proved that out-of-memory implementations are viable and perform equally well with respect to low-rank relative offset and running time in comparison to in-memory versions, but are able to perform in compute resources restricted environments.

Contributions We contributed the following:

- Theoretically and practically analyzed IALM and its SVD methods
- Suggested a novel, simplified type of recycling
- Provided an overview of available SVD methods in Python
- Implemented different SVD methods in Hannibal Specter
- Fixed and improved out-of-memory SVD implementation of dask
- Implemented different out-of-memory implementations able to scale to the terabyte range

Recommendations Further research could include:

- Reviewing the influence of ρ and μ -parameter on convergence behavior
- Comparing out-of-memory implementations with divide-and-factor implementations
- Exploring fully randomized methods, such as the GoDec algorithm²⁷
- Reviewing the influence of the ϵ -parameter on the convergence of randomized methods
- Comparing performance with distributed SVD algorithms²⁸
- Implementing a fully GPU-based dask out-of-core implementation

- Investigating the scaling performances when going out-of-core
- Reviewing different sampling techniques for randomized SVD methods
- Comparing out-of-memory implementations with streaming versions²⁹
- Investigating lowering the precision level further, e.g. to 16 bits
- Reviewing the evolution of GPU use on randomized SVD implementations³⁰

Bibliography

- [1] Candès, E. J., Li, X., Ma, Y. & Wright, J. Robust principal component analysis? *Journal of the ACM (JACM)* **58**, 1–37 (2011).
- [2] Dask Development Team. *Dask: Library for dynamic task scheduling* (2016). URL <https://dask.org>.
- [3] Lin, Z., Chen, M. & Ma, Y. The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. *arXiv preprint arXiv:1009.5055* (2010).
- [4] Bouwmans, T. & Zahzah, E. H. Robust pca via principal component pursuit: A review for a comparative evaluation in video surveillance. *Computer Vision and Image Understanding* **122**, 22–34 (2014).
- [5] Bouwmans, T., Javed, S., Zhang, H., Lin, Z. & Otazo, R. On the applications of robust pca in image and video processing. *Proceedings of the IEEE* **106**, 1427–1457 (2018).
- [6] Dongarra, J. *et al.* The singular value decomposition: Anatomy of optimizing an algorithm for extreme scale. *SIAM review* **60**, 808–865 (2018).
- [7] Anderson, E. *et al.* *LAPACK users' guide* (SIAM, 1999).
- [8] Larsen, R. M. Lanczos bidiagonalization with partial reorthogonalization. *DAIMI Report Series* (1998).
- [9] Virtanen, P. *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17**, 261–272 (2020).
- [10] Lehoucq, R. B., Sorensen, D. C. & Yang, C. *ARPACK users' guide: solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods* (SIAM, 1998).
- [11] Hernandez, V., Roman, J., Tomas, A. & Vidal, V. A survey of software for sparse eigenvalue problems. *Universitat Politècnica De València, SLEPs technical report STR-6* (2009).
- [12] Mackey, L., Talwalkar, A. & Jordan, M. I. Divide-and-conquer matrix factorization. *Advances in neural information processing systems* **24** (2011).
- [13] Tsuyuzaki, K., Sato, H., Sato, K. & Nikaido, I. Benchmarking principal component analysis for large-scale single-cell rna-sequencing. *Genome biology* **21**, 1–17 (2020).
- [14] Agullo, E. *et al.* Numerical linear algebra on emerging architectures: The plasma and magma projects. In *Journal of Physics: Conference Series*, vol. 180, 012037 (IOP Publishing, 2009).

- [15] Nishino, R. & Loomis, S. H. C. Cupy: A numpy-compatible library for nvidia gpu calculations. *31st conference on neural information processing systems* 151 (2017).
- [16] Blackford, L. S. *et al. ScaLAPACK users' guide* (SIAM, 1997).
- [17] Bosilca, G. *et al.* Flexible development of dense linear algebra algorithms on massively parallel architectures with dplasma. In *2011 IEEE International Symposium on Parallel and Distributed Processing Workshops and Phd Forum*, 1432–1441 (IEEE, 2011).
- [18] Maschho, K. J. & Sorensen, D. A portable implementation of arpack for distributed memory parallel architectures. In *Proceedings of the Copper Mountain Conference on Iterative Methods*, vol. 1 (1996).
- [19] Benson, A. R., Gleich, D. F. & Demmel, J. Direct qr factorizations for tall-and-skinny matrices in mapreduce architectures. In *2013 IEEE international conference on big data*, 264–272 (IEEE, 2013).
- [20] van der Walt, S. *et al.* scikit-image: image processing in Python. *PeerJ* **2**, e453 (2014). URL <https://doi.org/10.7717/peerj.453>.
- [21] Van Rossum, G. & Drake, F. L. *Python 3 Reference Manual* (CreateSpace, Scotts Valley, CA, 2009).
- [22] Halko, N., Martinsson, P.-G. & Tropp, J. A. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review* **53**, 217–288 (2011).
- [23] Szlam, A., Kluger, Y. & Tygert, M. An implementation of a randomized algorithm for principal component analysis. *arXiv preprint arXiv:1412.3510* (2014).
- [24] Lin, Z. & Wei, S. A block lanczos with warm start technique for accelerating nuclear norm minimization algorithms. *arXiv preprint arXiv:1012.0365* (2010).
- [25] Li, Y. & Yu, W. A fast implementation of singular value thresholding algorithm using recycling rank revealing randomized singular value decomposition. *arXiv preprint arXiv:1704.05528* (2017).
- [26] Oh, T.-H., Matsushita, Y., Tai, Y.-W. & Kweon, I. S. Fast randomized singular value thresholding for low-rank optimization. *IEEE transactions on pattern analysis and machine intelligence* **40**, 376–391 (2017).
- [27] Zhou, T. & Tao, D. Godec: Randomized low-rank & sparse matrix decomposition in noisy case. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011* (2011).

-
- [28] Schmidt, D. A survey of singular value decomposition methods for distributed tall/skinny data. In *2020 IEEE/ACM 11th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems (ScalA)*, 27–34 (IEEE, 2020).
- [29] Tropp, J. A., Yurtsever, A., Udell, M. & Cevher, V. Streaming low-rank matrix approximation with an application to scientific simulation. *SIAM Journal on Scientific Computing* **41**, A2430–A2463 (2019).
- [30] Lu, Y. *et al.* Reducing the amount of out-of-core data access for gpu-accelerated randomized svd. *Concurrency and Computation: Practice and Experience* **32**, e5754 (2020).

Chapter 8

Case Studies I: Imaging Mass Spectrometry

The key result of this chapter is that our methodology can be used for

- the efficient setting of an application-driven λ -parameter, outperforming the predefined λ_{candes}
- compression purposes, with the ability to improve on PCA's compression by a factor 1.73
- increasingly large in-memory MALDI measurement matrices due to its optimal implementation with compression ratios up to and beyond 10^3
- MALDI timsTOF measurement matrices requiring > 1 TB memory for processing (a first)

8.1 Introduction

In this chapter, we explore several use cases of robust matrix decomposition methods for imaging mass spectrometry, specifically the MALDI (tims)TOF-type of molecular imaging, and focus especially on what these methods can bring in terms of addressing the dimensionality reduction challenges inherent to this application domain.

8.1.1 Goals and Objectives

The aim of this chapter lies in the practical application of the developed methodology for analytical dimensionality reduction and data compression purposes for MALDI timsTOF imaging mass spectrometry. We explore this by means of different case studies, each highlighting particular aspects or delivering a proof-of-concept. First, we review the use of the application-driven framework for hyperparameter setting on a relative small molecular imaging measurement matrix. Secondly, we compare robust methods' ability to compress with traditional PCA on similar measurement matrix. Third, we gradually scale up the size of our real-world example matrices to a range requiring > 100 GB of memory, but still in memory, and compare the performance of our developed strategies. In a final case study, we present the results of a MALDI timsTOF measurement matrix requiring > 1 TB of memory via our out-of-memory implementation. These developments will be essential in making full ion mobility (IM) IMS, a powerful generalization of MALDI-TOF IMS that collects two spectral dimensions per pixel (or 4-mode data tensors), and that can yield petabyte measurements using standard techniques, a practical reality. The findings of this chapter subsequently formulate the answer to research subquestion B.1.

B. How can the application-driven methodology we have developed, be used to address application domain-specific challenges?

B.1. In molecular imaging, and specifically imaging mass spectrometry high-dimensionality and *massive data set footprints* are the most pressing bottleneck. How can we apply the developed methodology to measurement matrices of MALDI TOF and MALDI timsTOF nature, respectively of mouse kidney and mouse pup origin, specifically for *dimensionality reduction and compression purposes*?

8.1.2 Overview

In the next section, we introduce the experimental design for all four case studies, including the measurement matrices, and the used models and methods. Finally, we shortly describe the nomenclature and provide additional settings. Thereafter, the four case studies are presented:

Case Study A: investigates the use of our application-driven framework for hyperparameter setting in a molecular imaging context and consecutively compares multiple cases quantitatively and qualitatively;

Case Study B: compares the achievable compression of our robust decomposition method to that of PCA, when an application-optimal λ is set as

found in previous case study;

Case Study C: researches the in-memory scalability of our developed strategy and compares processed measurement matrices on different scales w.r.t. running time, singular values, and quality of retrieval;

Case Study D: demonstrates a first-ever robust matrix decomposition of a MALDI-timsTOF measurement set, an operation that would normally require > 1 TB memory.

Finally, we draw conclusions and provide our contributions and recommendations for further investigation.

8.2 Experimental Design

In this section, we elaborate on the experimental design. First, we introduce the different measurement sets. Second, we provide the methods and model that we use for each case study. Finally, we provide the additional settings regarding stopping criteria.

8.2.1 Acknowledgements

The mouse kidney dataset was acquired by William Perry PhD (Vanderbilt University) and the tissue sample was provided by Eric Skaar PhD (Vanderbilt University). The mouse pup dataset was acquired by Katerina Djambazova (Vanderbilt University). All data were acquired at Vanderbilt University under the supervision of Eric Skaar PhD, Jeffrey Spraggins PhD, and Richard Caprioli PhD.

8.2.2 Materials

Acetic acid, 1,5-diaminonaphthalene (DAN), ammonium formate, hematoxylin, and eosin were purchased from Sigma-Aldrich Chemical Co. (St. Louis, MO, USA). HPLC-grade ethanol was purchased from Fisher Scientific (Pittsburgh, PA, USA).

8.2.3 Sample Preparation

A one-week old C57BL/6 control mouse pup was snap-frozen at -80 °C, shaved over dry ice, and cryo-sectioned at 20 μm thickness, using a CM3050 S cryostat (Leica Biosystems, Wetzlar, Germany). The tissue was thaw-mounted onto a conductive indium tin oxide coated glass slide (Delta Technologies, Loveland, CO, USA). Approximately 500 mg of DAN was sublimed at 130 °C and 24 mTorr for 3.5 min onto the tissue surface for a final density of 1.0 mg/cm². The *S. aureus* infected mouse kidney was obtained from the Skaar Laboratory (Vanderbilt University) and snap-frozen at -80 °C and cryo-sectioned at 10 μm thickness, using

a CM3050 S cryostat. The tissue was thaw-mounted onto a conductive indium tin oxide coated glass slide. DAN in a solution of 90% acetonitrile was sprayed at a surface density of 3.6 $\mu\text{g}/\text{mm}^2$ at 85 °C; 1:1 ratio of carbonate buffer (pH 10.3) and 250 mM sodium acetate in a solution of 30% methanol sprayed at a surface density of 6.8 $\mu\text{g}/\text{mm}^2$ at 85 °C.

8.2.4 MALDI TIMS-IMS

The mouse pup data was acquired using the timsToF fleX mass spectrometer (Bruker Daltonik, Bremen, Germany) in trapped ion mobility (TIMS) mode of operation with an ion transfer time of 100 μs , pre-pulse storage time of 8 μs , and a collision RF of 2,000 Vpp, a TIMS funnel 1 (accumulation) RF of 450 Vpp, a TIMS funnel 2 (analysis) RF of 400 Vpp, a multipole RF of 400 Vpp, and a collision cell entrance voltage of 300 V. Tissue imaging data (164,808 pixels) were collected at 50 μm pixel size, using 200 shots per pixel and 48% laser power in positive ionization mode from m/z 300 to 1,200. The TIMS scan time was set to 400 ms, with reduced mobility ($1/K_0$) range of 0.4 to 1.9 Vs/cm².

The mouse kidney data was acquired using the timsToF fleX mass spectrometer (Bruker Daltonik, Bremen, Germany) in QTOF mode of operation. Tissue imaging data (161,547 pixels) were collected at 15 μm pixel size using 500 shots per pixel and 53% laser power in negative ionization mode from m/z 400 to 1,400.

8.2.5 Data Processing

The mouse pup data were exported into a custom binary file format optimized for storage and speed of analysis of the ion mobility-IMS data. Each frame/pixel contains between 10,000-100,000 centroid peaks that span the acquisition range of m/z 300-1,200 and $1/K_0$ 0.4-1.9 (Vs)/cm² with 221,888 and 4,001 bins in the mass spectrometry and ion mobility dimensions, respectively. The processing pipeline requires common m/z and $1/K_0$ axes, hence individual centroid peaks were inserted at their correct bin positions along the mass spectrometry and ion mobility dimensions; missing values were set to zero. Following the conversion process, a mean mass spectrum of the entire data set was generated, and peak picked. A total of 284 features were selected and extracted to generate ion mobility-rich ion images.

The mouse kidney data were exported into the custom binary file format mentioned earlier. Data were acquired in the QTOF mode, hence the ion mobility dimension is not present, in which case we introduce a secondary dimension by enforcing the data set to contain one ion mobility bin. This is carried out to ensure efficient storage and data processing without having any impact on the actual data. Following the conversion process, a mean mass spectrum of the entire data set was generated and peak picked resulting in several different data set dimensions (135, 152, 300, 573 and 656), as shown in **Tab. 8.1**. Also, larger

versions of the data matrix were generated by selecting peaks in the mean mass spectrum using different quantiles.

8.2.6 Measurement Matrix

An overview of the measurement matrices is given in **Tab. 8.1**. Note that the mouse pup measurement matrix is in 16-bit data format. However, the available svds are only capable of handling 32-bit values, necessitating that the actual data size should be doubled. Finally, the corresponding mean ion images are depicted in **Fig. 8.1**.

Table 8.1 | Table of measurement matrices. The measurements are given per case study, along with the origin, the mode of acquisition and whether the data set is binned, the corresponding matrix size and data size, and the datatype. The matrix sizes are obtained by vectorizing the spatial components.

Case	Origin	Mode	Binned	Matrix Size	Data Size	Datatype
A	Mouse Kidney	QTOF	Yes	161547×135	0.087 GB	32 bits integer
B	Mouse Kidney	QTOF	Yes	161547×573	0.37 GB	32 bits integer
C	Mouse Kidney	QTOF	Yes	161547×135	0.087 GB	32 bits integer
C	Mouse Kidney	QTOF	Yes	161547×152	0.098 GB	32 bits integer
C	Mouse Kidney	QTOF	Yes	161547×300	0.194 GB	32 bits integer
C	Mouse Kidney	QTOF	Yes	161547×573	0.37 GB	32 bits integer
C	Mouse Kidney	QTOF	Yes	161547×656	0.424 GB	32 bits integer
C	Mouse Kidney	QTOF	No	161547×6588	4.26 GB	32 bits integer
C	Mouse Kidney	QTOF	No	161547×21958	14.19 GB	32 bits integer
C	Mouse Kidney	QTOF	No	161547×43915	28.38 GB	32 bits integer
D	Mouse Pup	TIMS	No	$164808 \times 284 \times 4001$	374.54 GB	unsigned 16 bits integer

8.2.7 Methods and Model

In this chapter, we restrict ourselves to the use of PCP for simplicity. Nevertheless, in the future SPCP might help improve the results as it is inherently able to deal with small i.i.d. noise. The underlying model we employ is the LMM, which is

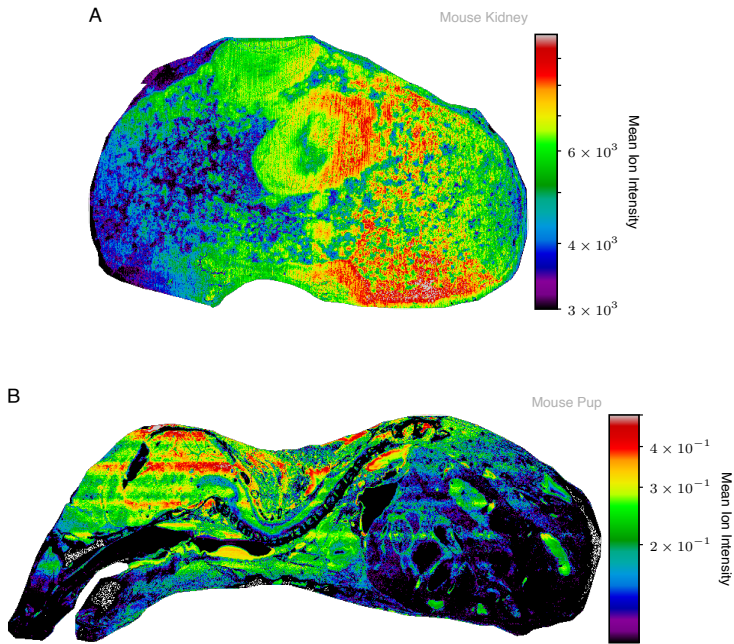


Figure 8.1 | Mean ion images of the different IMS measurement sets. The different panels are: (A) mouse kidney as used in case study A, B, and C; and (B) mouse pup as used in case study D. The ion intensity is in arbitrary (relative) units.

described as

$$\begin{aligned} M &= A + B, \\ &= XY^T + B, \end{aligned} \tag{8.1}$$

where $M \in \mathbb{R}^{m \times n}$ is the measurement matrix, $A \in \mathbb{R}^{m \times n}$ represents a low-rank matrix, and $B \in \mathbb{R}^{m \times n}$ is a noise matrix. Secondly, note that A can also be decomposed following the SVD as $A = U\Sigma V^T$. We will assume that signal, i.e. what is of interest, is captured by A , while noise and residuals are approximated by B .

8.2.8 Nomenclature

MALDI-TOF provides data cubes with two spatial dimensions and one spectral, m/z feature, dimension. Furthermore, MALDI timsTOF provides a fourth dimension, drift time. Each of these dimensions can be condensed, e.g. by taking mean or median values over all values along that particular dimension. We will use the following nomenclature:

Ion Image: a single image showing the spatial distribution and abundance of a single ion or m/z feature;

Mean Ion Image: the spectral dimension is condensed to a single image by taking the mean value of all values along the spectral dimension(s);

Vectorization: consists of linearly transforming a matrix into a single vector. For example, when each individual ion image “channel” is considered a matrix, i.e. a two-dimensional set of values, vectorization consists of stacking the columns of that matrix on top of each other so that the ion image is condensed to a single vector. Spatial information, e.g. spatial nearness, is lost in this process. Performing vectorization on each ion image or m/z feature channel individually results in a set of vectors, i.e. a matrix with one condensed spatial direction and one spectral dimension. This corresponds to going back the pixel table as originally obtained in the MUSE pipeline;

Reconstruction: when we talk about reconstruction we mean going from vector form to image, undoing the vectorization;

Score: corresponds to a particular reconstructed left singular vector (U), containing spatial information. It can be seen as a “synthetic” spectral channel image. At the same time, and together with its corresponding right singular vector, it represents a rank-1 decomposition of the original input matrix, i.e. a full set of spectra. Hence, it contains particular information from the original measurement set;

Loading: corresponds to a particular reconstructed right singular vector (V), containing spectral information. It can be seen as a “synthetic” spectrum. At the same time, and together with its corresponding left singular vector, it represents a rank-1 decomposition of the original input matrix, i.e. full set of spectra. Hence, it contains particular information from the original measurement set;

Components: correspond to a score, or a loading, or a pair of those. Together with its corresponding singular value, it represents a rank-1 matrix. The sum of all components, i.e. of all rank-1 matrices reconstructs the low-rank term. Each component carries a particular part of the information from the low-rank term.

8.2.9 Additional Settings

Method Implementation For PCP, the inexact augmented Lagrangian multiplier method (IALM)¹ was utilized.

Stopping Criteria The stopping criteria for PCP are set to

$$\begin{aligned} \frac{\|M - A - B\|_F}{\|M\|_F} &\leq 1 \times 10^{-5}, \\ \mu \frac{\|\hat{B} - B\|_F}{\|M\|_F} &\leq 1 \times 10^{-5}, \end{aligned} \quad (8.2)$$

where μ is an algorithmic variable. A maximum of 100 iterations is set.

Moran's I We use an equal weighted 4-connected pattern, similar to previous chapters, as our measure of spatial autocorrelation.

Spectral Correlation For the spectral correlation, we use a simple 2-connected pattern, where we only consider neighboring pixels. This corresponds roughly to assuming correlation on a 1.1 m/z value range (assuming an equidistant grid along the spectral dimension).

Savitzky-Golay Filter The Savitzky-Golay filter² was utilized from the SciPy package³. The window length is set to 15 and the order of the polynomial is set to 2. Its parameters were set through trial-and-error.

Computational Specification Floats with a precision of 32 bits were used and the underlying linear algebra is performed via OpenBLAS-LAPACK accessed via NumPy and SciPy³. All experiments were performed in Python 3.8.3⁴ in purely CPU-based processes. Finally, the computer system specifics are given in **Tab. 8.2**.

Table 8.2 | Computer system specifications. The operating system, number of cores and their clock speed, and the size of available memory is stated.

OS	Cores	Clock Speed	Memory Size
Windows 10 Pro	52 (104 logical)	2.7 GHz (3.3 GHz overclock)	1.5 TB

8.3 Case Study A: Application-Driven Parameter Setting

There have been a few early investigations of PCP for MALDI-TOF IMS⁵⁻⁸, however, no hyperparameter optimization was performed. Hence, in this first case study, we focus on the practical application of our application-driven framework on a MALDI-TOF IMS data set. We first define the measures for our framework and define three different applications, similar to the ones presented in chapter 6, along with the corresponding weights. In second instance, we analyze the quantitative and qualitative behavior of the different solutions obtained with these settings and show the superior performance of our optimized setting versus the commonly-used predefined $\lambda = \lambda_{\text{candes}}$.

8.3.1 Additional Experimental Design

In this section, we provide additional experimental design settings.

Method We use PCP with `scipy_gesdd`. The PCA version used, is covariance-based, and its rank was set to the rank value obtained by PCP. The λ parameter is linearly sampled 100 times between a range of 0 to $3\lambda_{\text{candes}}$.

Assumptions For this case study, we assume the variability of our framework to be small, i.e. the spread on the λ -factor for different realizations is minimal. As we only have one realization available, we are not able to check this assumption. However, one could think of synthetically obtaining multiple realizations by sampling and processing the measurement matrix multiple times, under the assumption that the sampled matrix represents the whole accordingly. Note that additional problems might appear with respect to spatial and spectral correlation measures as not all data points will be connected when sampled.

8.3.1.1 Quantitative Measures

To define quantitative measures, we constructed a “leakage table” to define those types of signal that are not effectively captured by PCP. These are mainly determined by its incoherence principles, which is why we subdivide our table of potentially uncaptured signal spectrally and spatially into two classes: (1) coherent, (2) weakly incoherent. The combinations of those define 4 different types of signal that might leak from the low-rank component into the residual term B. Hence, defining measures to detect those specifically is of importance. The leakage table is given in **Tab. 8.3**. In addition to the four types listed there, we define a 5th type of potentially leaked signal that represents other non-defined (at least by this table) types.

Table 8.3 | Leakage table for imaging mass spectrometry. Represents the property of signals, i.e. weakly incoherent or coherent, that are not captured by the low-rank term A and consecutively defines 4 types of such signals.

		Spectral	
		Coherent	Weakly Incoherent
Spatial	Coherent	Type 1	Type 2
	Weakly Incoherent	Type 3	Type 4

We identified measures that can be used to quantify these different leaked signal types:

- **Type 2:** Isotope correlation / Spectral Moran's I
- **Type 3:** Spatial Moran's I
- **Type 4:** Correlation Matrix

Type 1 and type 5 are hard to quantify and in most cases will require domain-specific expert knowledge to isolate and detect. This can, for example, be by manual inspection of the residual B term (e.g. by a trained pathologist in molecular imaging of disease).

Isotope correlation Besides, the previously introduced rank, and Moran's I for measuring "structure" in individual images, we introduce the use of a spectral autocorrelation to detect isotopes. This measure is defined similar to the spatial autocorrelation of Moran's I, with the difference that it acts on the spectra and uses a simple weighing scheme where the weights are set to 1 for m/z bin distances within 1.1 m/z value of the original. This uses prior knowledge on the isotopes of molecules. The measure is defined as

$$I^{\text{spec}} = \frac{c_n}{c_w} \frac{\sum_i \sum_j w_{ij} (x_i - \bar{x})(x_j - \bar{x})}{\sum_i (x_i - \bar{x})^2} \quad (8.3)$$

where w_{ij} is a weighting term, $c_w = \sum_i \sum_j w_{ij}$, x_i represent the column entries of \mathbf{x} , \bar{x} is the column mean, and c_n is the number of spectral bins. As such, we obtain the spectral autocorrelation of each spectrum, more specifically focussing on the molecular isotopes. A value of 1 will correspond to high correlation within a distance of 1.1 m/z , e.g. capturing a lot of isotopes, while values close to 0 report low correlating m/z features.

8.3.2 Framework

In Fig. 8.2, a selection of measures is visualized that can drive our hyperparameter setting framework. These measure values have been min-max normalized. From those, we selected six specific measures to optimize towards three distinct applications, similar to chapter 6:

- **Maximal Compression application:** We want a highly compressed low-rank term A , while still guaranteeing “good” reconstruction of the original measurement matrix;
- **Maximal Signal Retention application:** We want signal retention of as much as possible;
- **Trade-Off application:** Defines a trade-off between both cases above.

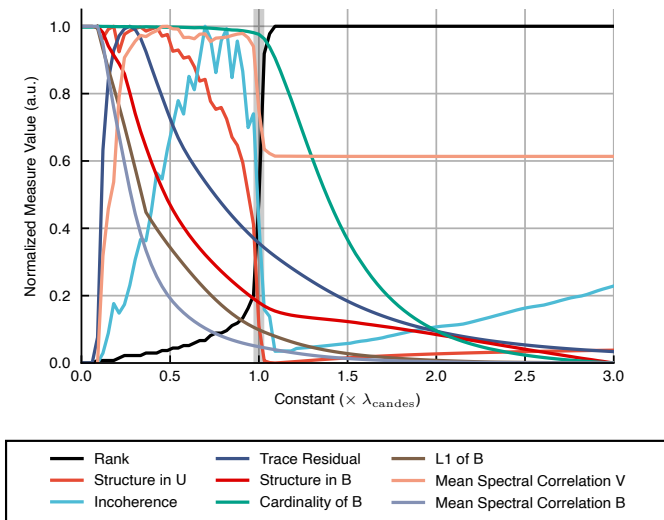


Figure 8.2 | Min-max normalized measures over the grid search. The λ parameter is linearly sampled 100 times between a range of 0 to $3\lambda_{\text{candes}}$.

Weight Definition The weights for the used measures, making the optimization application-specific, are given in Tab. 8.4. These were heuristically tuned to perform optimal for this measurement matrix. This was done by defining a large range of measures and selected those describing the underlying patterns of interest.

Table 8.4 | Measures and weights for the three case studies. The application case studies consist of maximal compression, maximal signal retention, and a trade-off between both.

Measure	Definition	Weight	Max. Compr.	Max. Sign. Ret.	Trade-Off
Rank	$\text{rank}(\hat{A})$	α_0	1	0	0.33
Frobenius Residual	$\text{tr}(\hat{A}^T \hat{B})$	α_1	0	1	0.13
Structure in B (Spatial)	$\text{mean}(\hat{I}_{\hat{B}})$	α_2	1	4	0.85
Cardinality of B	$\ \hat{B}\ _0$	α_3	0	-1	-0.13
Structure in V (Spectral)	$\text{mean}(I_V^{\text{spec}})$	α_4	1	0	0.33
Structure in B (Spectral)	$\text{mean}(I_B^{\text{spec}})$	α_5	0	1	0.13

8.3.3 Results

The results consists of the presentation of the different λ -settings, a short quantitative analysis of the different optimal settings, and a conclusive qualitative analysis.

8.3.3.1 λ -setting

In **Fig. 8.3**, the three composite measures are given along with their respective minima (dashed vertical line). We observe that all values are in the neighborhood of $\lambda = 1\lambda_{\text{candes}}$, the optimal setting for low-rank with perfectly sparse noise. Therefore, one could see our maximal compression setting and maximal signal retention setting as the extreme points of an interval of useful λ parameters. Secondly, we remark that our trade-off is slightly smaller than $\lambda = 1\lambda_{\text{candes}}$.

8.3.3.2 Quantitative Analysis

The quantitative analysis of the three defined cases is depicted in **Fig. 8.4**. We observe the following:

- the optimal “maximal compression” λ -value (vertical dashed black line) is situated at the first plateau for the mean spectral correlation of V, resides on the down slope of the structure in U measure, and sits before the steep increase in rank. This corresponds to what is expect for maximal compression;
- the optimal, “maximal signal retention” λ -value (vertical dashed red line) coincides with the first λ for which the normalized rank is equal to 1, and sits at the final plateau for mean spectral correlation in V and at the leg of the structure in B. The latter shows that a decrease in structure in B will be harder to obtain while the quality is guaranteed, as less energy will also be captured by the B term. This corresponds to what is expected for the maximal signal retention case;

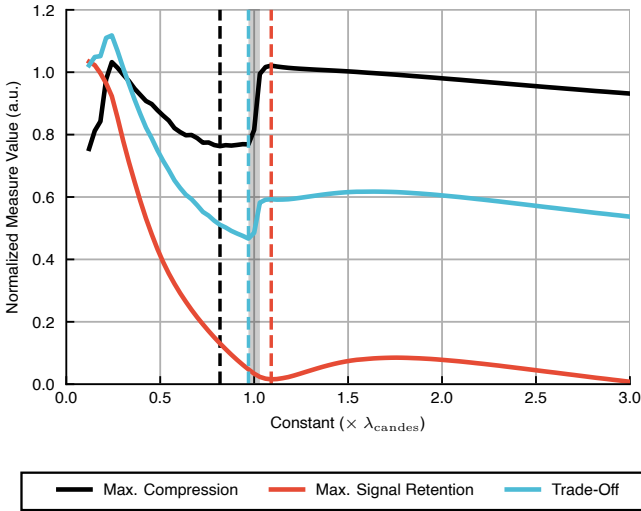


Figure 8.3 | Composite measures corresponding to the maximal compression, maximal signal retention and trade-off case studies over the grid search. The dashed line marks the different minima of the corresponding composite measures. The vertical grey strip corresponds to the predefined λ -setting⁹.

- the $\lambda = 1$, is fairly high in normalized rank (i.e. ≈ 0.5), where the structure in \mathbf{U} is low. The latter signifies that the mean spatial component is not very structured. Hence, this setting is not desirable. We show with our trade-off case that both can be improved by setting a slightly smaller λ -value.

Secondly, we review the normalized ℓ_2 -norm of the separate ion images in \mathbf{A} and \mathbf{B} , i.e. their columns, with respect to the columns of the measurement matrix \mathbf{M} in **Fig. 8.5**. This gives us an insight into where the changes take place while increasing or decreasing the λ -value. We observe that most ion images change only marginally for the different cases. Only for a few m/z features a difference is visible. This might indicate, however, that small intensity noise is added to most of those m/z features, for a better fit in the few m/z bins that were matched poorly for low λ values. It suggests that by increasing the λ -value, the structural components remain to some degree unchanged, while the components that are added are more coherent, i.e. only affecting certain m/z features, but do contain noise. The qualitative analysis gives further insight in these processes.

8.3.3.3 Qualitative Analysis

The observation that the application-driven λ mostly affects a few coherent components, is reviewed in this section by a qualitative analysis. Therefore, in **Fig.**

8.6, different reconstructions are given at a particular m/z value of 762.50397. From top to bottom we see respectively the ion image of raw data, the reconstruction image of the low-rank component, and the reconstruction image of the sparse term. Remember that for these reconstructions the same matrix logic holds such that $M = A + B$ is also true for the columns and thus ion images. From left to right, the maximal compression case, the trade-off, and the maximal signal retention case are given. We observe that all three images of the retrieved low-rank term A appear similar except for the shoulder on the right side that gains intensity, when the λ -value is increased. This confirms the phenomena as previously observed in the quantitative analysis. Nevertheless, the residual ion images, i.e. B , show that for all three cases the residuals in these areas remain high. A zoomed in version of the right shoulder for the same cases of the same m/z feature is given in **Fig. 8.7**. We assume that this “shoulder” in the residual term can be associated to a physical biological feature due to its intensity and spatial localization. However, we should be careful as some of these visually apparent spatial localizations might appear due to multiplicative noise. Hence, one should be careful with respect to spatial and spectral correlation measures as aforementioned. Finally, we review the statement of the increase of noise for increasing λ -values, for a m/z feature with little change over the three cases. Such case is depicted in **Fig. 8.8**. When we look at a zoomed in reconstructed image of m/z value 464.16748, we don't see large changes in the reconstructed low-rank term A ion images. However, in the residual term B , we see a slight decrease in intensities when increasing the λ -value. However, due to the large intensities in A the addition of noise is not visually apparent.




Figure 8.4 | Quantitative measures calculated on A , and its factors U and V , and B over the grid search. The dashed line marks the different minima of the corresponding composite measures as in **Fig. 8.3**. The vertical gray strip corresponds to the predefined λ -setting⁹.

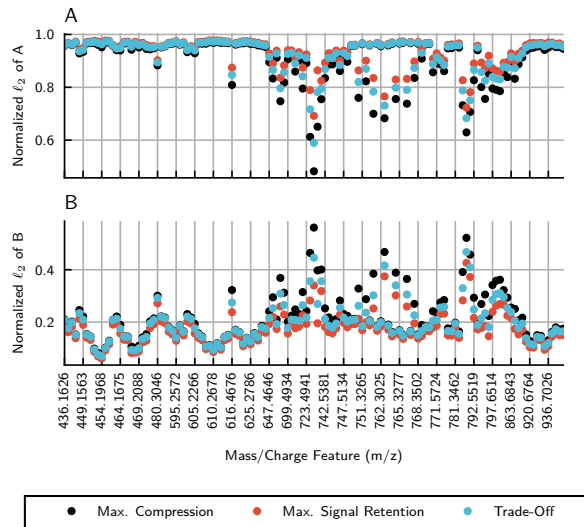


Figure 8.5 | Relative energy captured per ion image in A and residual B for different cases. Normalized ℓ_2 -norm is used on the separate ion images in A and B , i.e. their columns, with respect to the columns of the measurement matrix M . The subplots represent: (A) the normalized ℓ_2 of A , and (B) the normalized ℓ_2 of B .




Figure 8.6 | Reconstructed ion images for the three different cases at m/z value 762.50397. The different columns depict: (A) the maximal compression case; (B) the trade-off case, and (C) the maximal signal retention case. The units are arbitrary.




Figure 8.7 | Zoomed in reconstructed ion images of the right shoulder for the three different cases at m/z value 762.50397. The different columns depict: (A) the maximal compression case; (B) the trade-off case, and (C) the maximal signal retention case. The units are arbitrary.




Figure 8.8 | Zoomed in reconstructed ion images of the middle abscess for the three different cases at m/z value 464.16748. The different columns depict: (A) the maximal compression case; (B) the trade-off case, and (C) the maximal signal retention case. The units are arbitrary.

8.3.4 Discussion

In this case study, we reviewed the use of our developed framework for the application-driven hyperparameter setting of PCP for MALDI-TOF IMS data. We introduced three applications and demonstrated some of their quantitative as well as qualitative behavior.

Most importantly, we note that we are able to improve over the predefined and wide-spread λ_{candes} parameter, in compression as well as in retrieved quality of the corresponding spatial components in U . Secondly, we have shown that for all cases, signal leaking might still be taking place due to signal components behaving coherent, i.e. being only spatially or spectrally active in certain m/z values or pixels. The latter can not be avoided by altering the λ setting. Nevertheless, we have shown that the three defined application cases result in distinct solutions with specific favorable properties, and that there is value in optimizing towards a specific application goal.

To improve the quality of the fit and analogous to the idea as reviewed in section 4.2.1, the areas of signal that are still visible in the error term B can also be removed by other methods. The signal is sparse and mostly highly energetic such that a (robust) sparse PCA method could be more appropriate. Note that the latter connects to the central idea in basis pursuit, where a basis in a single fixed dictionary, e.g. the Fourier transform's dictionary, might not be appropriate to sparsely describe all phenomena. We have tried such methods and obtained reasonable results. However, its computational complexity is high and thus not useful for the high-dimensional data that we face. Furthermore, we tried several other techniques such as a "waterfall" strategy where we post-process the residual term B again by PCP, however, with a different λ setting. This was not successful. In second instance, we tried to peak pick those regions spatially and spectrally to create new incoherent bases and process this reduced set with PCP. Again, this gave no satisfying results as it appears that some parts still remain coherent.

8.4 Case Study B: A Comparative Study between PCP and PCA

In this case study, we compare PCP to the traditional PCA on a MALDI-TOF IMS data set with a slightly larger size than in the first case study. For this comparison we are especially interested in the achievable compression ratio. Therefore, we define several measures to defined the quality of the retrieved low-rank component. We show that the latter is necessary as the use of the Frobenius and ℓ_2 norms, e.g. as used in PCA for measuring the total captured variance, would distort the results obtained for PCP with its removal of outlying values.

8.4.1 Additional Experimental Design

In this section, we provide additional experimental design settings.

Method We use PCP with `scipy_gesdd`. The PCA version used, is covariance-based, and its rank was set to the rank value obtained by PCP. The λ parameter is linearly sampled 100 times between a range of 0 to $3\lambda_{\text{candes}}$.

Measures We compare PCP and PCA on the basis of traditional measures such as the Frobenius and nuclear norm. However, we will base our main comparison on the spatial and spectral correlation as previously defined. Furthermore, we introduce an additional quality check in the form of the correlation matrix to catch type 4 errors (see **Tab. 8.3**), i.e. the ion image correlation on the residual term B . It is defined as

$$\rho_{B,B} = \frac{\text{cov}(B, B)}{\sqrt{\sigma_B \sigma_B}}, \quad (8.4)$$

where cov is the covariance operator and σ represents the standard deviation matrix. Values close to -1 and 1 show that anti-correlation/correlation patterns are present in the residual matrix. This is expected when this residual term does not solely consist of uncorrelated ion images. The latter can be an indicator of signal leakage, especially when multiple bins are correlating. Again, note that this might also indicate that the “true” noise is correlated, which would be the case for multiplicative noise.

8.4.2 Results

The results consist of the presentation of a global comparison based on the usual Frobenius and nuclear norms, followed by the comparison of achievable compression ratio based on spatial and spectral correlation scores. Finally, we review the ion image correlation of the residual term B .

Frobenius and Nuclear Norm Comparison Traditionally, the Frobenius norm is used to express the variance captured by PCA. This is a logical measure as one operates on the covariance matrix of the measurement matrix M . However, this is only valid under the assumption that outliers are removed *a priori*. When outliers are not properly removed, PCA tries to fit them and starts compensating skewing away from the true underlying signal. This is detrimental for the low-rank description. We argue that the nature of these outlying values is not necessarily noise-related (see last paragraph of section 4.1.4). The Frobenius and nuclear norm for both PCP and PCA are given in **Fig. 8.9**. Note that the x-axis is over λ_{candes} , which is not a linear mapping of the underlying rank. Furthermore, remember that in the singular value space the nuclear norm corresponds to the

ℓ_1 -norm on the singular values where the ℓ_2 -norm is represented by the Frobenius norm (see chapter 2). As such, the Frobenius norm is considered to be more prone to high-valued singular values than the nuclear norm, similar to Fig. 4.1. We made the following observations:

- it is generally true that $\|\cdot\|_* \geq \|\cdot\|_F$ (can be proven mathematically);
- the Frobenius norm of PCA has a steep increase at low constant values, i.e. between 0 and 0.2, and almost directly peaks to its final value. This indicates that the first few singular values for PCA are defining most “energy” with respect to the whole measurement matrix;




Figure 8.9 | Frobenius and nuclear norm for both PCP and PCA over λ range from 0 to $3\lambda_{\text{candes}}$. The vertical gray strip corresponds to the predefined λ -setting⁹.

- PCP grows slowly towards the stable value for both Frobenius and nuclear norm. Hence, those first few singular vectors captured by PCA are probably coherent;
- the Frobenius and nuclear norm keeps increasing, even after having reached full rank, i.e. at around $1\lambda_{\text{candes}}$.

Compression Rate Secondly, we are interested in the achievable compression ratio and its quality of retrieval of the low-rank term A . In Fig. 8.10, we compare the mean structure in U (via Moran’s I) and the mean structure in V (via isotope

correlation, i.e. spectral correlation). Both are the corresponding left and right singular vectors of A ($A = U\Sigma V^T$). These measures give a general impression of the quality of the average retrieved components of A . We assume that these measures are additive and (to some extent) disjoint. We observe the following:

- PCP’s left and right singular values score on average better “structurally”;
- PCA scores worse on average but spatially overtakes PCP at our trade-off case (light blue dashed line);
- the mean structure both decrease for PCA and PCP, this can be attributed to the addition of less structured components to the low-rank term A .




Figure 8.10 | Mean spatial and spectral correlation, i.e. structure in U and spectral correlation V , for PCP and PCA over λ range from 0.5 to $1.5\lambda_{\text{candes}}$. The vertical gray strip corresponds to the predefined λ -setting⁹. The vertical dashed lines correspond to the application cases of the previous case study: (black) maximal compression; (red) maximal signal retention; (light blue) trade-off.

When we translate those average spatial and spectral scores to cumulative ones, we obtain a plot as viewed in **Fig. 8.11**. We can utilize these scores to compare PCP and PCA and define their ability to compress by comparing equal reconstruction scores spatially as well as spectrally. Based on the spatial reconstruction score, i.e. the cumulative structure in U , the largest measured difference in rank between PCA and PCP to achieve a similar cumulative structure in U is 22 for a rank of 30 for PCP and 52 for PCA at $\lambda = 0.818$. The latter coincides with the maximal compression case study (proving its efficiency). This roughly doubles,

i.e. a factor 1.73, the compression ratio from 10.98 for PCA to 19.03 PCP. Note that this result does remain specific to this measurement matrix and the used score. Nevertheless, it does illustrate the ability of PCP to compress beyond the abilities of PCA.




Figure 8.11 | Cumulative spatial and spectral correlation, i.e. cumulative structure in U and cumulative spectral correlation V , for PCP and PCA over λ range from 0.5 to $1.5\lambda_{\text{candes}}$. The vertical gray strip corresponds to the predefined λ -setting⁹. The vertical dashed lines correspond to the application cases of the previous case study: (black) maximal compression; (red) maximal signal retention; (light blue) trade-off.

Finally, to further accumulate evidence on the aforementioned observations we review the absolute cosine correlation of the components of PCP in comparison to PCA in **Fig. 8.12**. Values close to 0 (blue) correspond to low correlation and, as such, reports little common information in these columns. Values close to 1 (red) have a lot of common information. Values close to 0.5 (gray) share some information. In the depiction, we remark two structures:

- the first band-like structure runs for the first few PCA columns to the roughly 20th component of U_{PCP} . This indicates that the information of the roughly first 5 components of U_{PCA} are partially spread among those components in PCP;
- the second tail-like structure that runs from roughly 0 to the 40th component of U_{PCP} and the whole column side of U_{PCA} . This tail confirms our story on PCP and PCA. PCP is able to capture way more information

in its first few components then PCA is able to: we give up some sparse components to gain a higher compression.




Figure 8.12 | Absolute cosine similarity between PCA and PCP on the left singular vectors, i.e. spatial information. Values of 0 correspond to no correlation, while values of 1 indicate perfect alignment between vectors. The λ -setting is equal to the maximal signal-retention case.

Ion Image Correlation Finally, we compare the ion image correlation of both PCP and PCA, depicted in **Fig. 8.13**. This correlation and especially regions of correlation are not desired as they are indicators of type 4 signal leakage. We observe the following:

- both the correlations for PCA and PCP are relatively low, probably indicates a good retrieval;
- PCA has more negative correlation, which is believed to be a result of its compensation behavior as aforementioned;
- the correlation for PCA is mainly at the low side of the m/z features, while for PCP it is in mid range;
- comparing PCP's ion correlation to the results obtained in the previous case study in **Fig. 8.5**, we observe that most small m/z features and high features are well recovered while some outlying m/z features seem to pop-up in the mid range, see **Fig. 8.13**;

- the correlations for PCP are mainly positive. Whether this is a direct result of the absence preprocessing, e.g. mean-centering, is not clear.




Figure 8.13 | Ion image correlation of the residual term B . Given its symmetry, we plot both the ion image correlation for PCA and PCP as one square. The correlation for PCA is depicted in the left top triangle, while the correlation for PCP is given in the right bottom triangle. The λ -setting is equal to the maximal compression case.

8.4.3 Discussion

In this case study, we have compared PCP to PCA on its ability to compress MALDI-TOF IMS data. We have shown that at most a 1.73-fold improvement in compression rate can be obtained with PCP in comparison to PCA (for this specific data set), while still obtaining a similar spatial reconstruction score. This corresponds to a compression ratio of 19.03 for PCP compared to 10.98 for PCA. Concurrently, this observation coincides with our framework's optimal λ for a maximal compression application. This result is further elaborated by reviewing the absolute cosine correlation of the left singular vectors of the low-rank term A and ion image correlation of the residual term B . Both suggest that PCP is able to compress to a higher level, by giving up spatially as well as spectrally coherent, i.e. sparse, energetic components. Finally, other summarizing reconstruction scores can be used to compare the compression. We suggest e.g. the mean absolute correlation coefficient of the loadings and the mean absolute correlation coefficient of the scores¹⁰, suggested earlier in Chapter 6.

8.5 Case Study C: Scale-Dependent Comparison

In this case study, we review the in-memory scalability of our methodology in combination with our framework's hyperparameter optimization for the trade-off case. Therefore, we compare different MALDI-TOF measurement sets of increasing size on the basis of running time and quality of their low-rank term retrieval, and provide reconstruction scores to assess.

8.5.1 Additional Experimental Design

In this section, we provide further experimental design settings.

Method We use PCP with randomized svd, including recycling of singular vectors with 2 QR iterations. The λ parameter is set to the trade-off case of case study A with a $\lambda = 0.97\lambda_{\text{candes}}$.

Assumptions We assume that our λ setting remains stable for increasing measurement matrix sizes, i.e. it provides similar results.

8.5.2 Results

The results consist of the presentation of a global comparison based on running time, rank and cardinality, singular values and the first score and loading. Secondly, we provide spatial and spectral measures to assess the reconstruction.

8.5.2.1 General Comparison

In this subsection, we compare our retrieved low-rank matrix A from a global perspective.

Running Time The running time is depicted in **Fig. 8.14**. We remark the linear slope due to the use of the randomized SVD: a linear increase in time proportional to the measurement matrix.

Rank and Cardinality The relative rank and cardinality are given in **Fig. 8.15**. The rank is retrieved from A, while the cardinality of B is considered. We observe the following:

- the relative rank reduces drastically for increasing measurement matrix sizes. This is the result of a quasi-static absolute rank for all cases. This confirms our assumption of our λ setting being stable;
- the relative cardinality decreases to around 60%, this is mainly due to the fact that our largest three measurement matrices are not binned.




Figure 8.14 | Running time for different MALDI-TOF measurement matrices plotted on logarithmic scale. This plot corresponds to a linear slope in time, which corresponds to our theoretical $\mathcal{O}(m \cdot n \cdot r)$.




Figure 8.15 | The relative rank and cardinality are depicted for different MALDI-TOF measurement matrices. We remark that the relative rank, i.e. relative to n , decreases linearly with growing sizes. The relative cardinality, i.e. relative to $m \cdot n$ decreases slightly towards larger measurement matrices to around 60%.

Note that both evolutions towards sparsity, in terms of rank and cardinality, increase the validity of using PCP as they increasingly fulfill the conditions as originally defined⁹. The achieved compression ratios in order of increasing data set sizes are 7.1, 7.6, 13.0, 24.8, 27.2, 234.4, 623.6, and 1113.8. The latter suggests that the 30 GB matrix can be compressed effectively (and with minimal information loss) to around 30 MB.

Singular Values Spectrum In Fig. 8.16, the singular value spectrum for each retrieved low-rank matrix A is depicted. Note again that the largest three sized matrices are not binned. We observe the following:

- the binned sets result in higher singular values of the first few components. This is a logical result as they will contain more “energy” due to the binning;
- smaller singular values increase with data set sizes, i.e. from the 5th singular value and on, and influence the decay. This can indicate that faint phenomena are better captured by the larger data sets. These faint phenomena are thus probably only detectable when using unbinned versions.




Figure 8.16 | Singular values for the retrieved low-rank matrix A of different MALDI-TOF measurement matrices for $\lambda = 0.97\lambda_{\text{candes}}$.

First Score and Loading The first score and loading of the retrieved low-rank matrix A for each measurement matrix is respectively depicted in **Fig. 8.17** and **Fig. 8.18**. We observe the following:

- the first pair of singular vectors are expected positive. This is true for SVD where such components are called the Perron-Frobenius vectors (see section 2.4.4.2). Whether this is also true for PCP is not clear;
- data set A and D, and B, C and E are binned differently. This can be linked to the similarity in their spectra in **Fig. 8.16**;
- all scores are visually similar, except for B, C and E;
- in the scores, we see a gradient in vertical direction with high intensive areas (red) at the top of the image, going to lower values at the bottom. Whether this can be linked to instrumental artifacts is not clear;
- for the loadings, the intensity is generally lower for the unbinned versions.

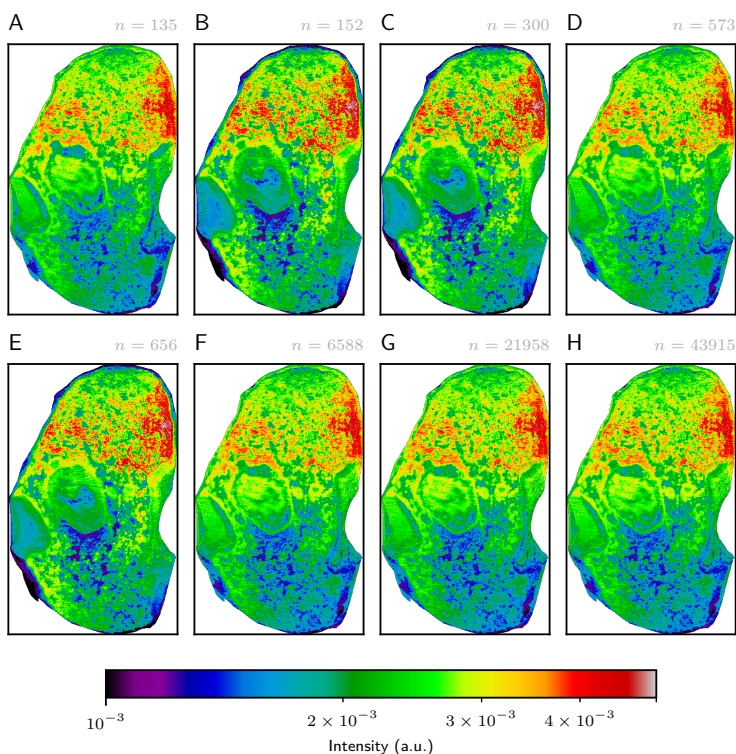


Figure 8.17 | First scores for the retrieved low-rank matrix A of different MALDI-TOF measurement matrices for $\lambda = 0.97\lambda_{\text{candes}}$. The subplots correspond to: (A) $n = 135$; (B) $n = 152$; (C) $n = 300$; (D) $n = 573$; (E) $n = 656$; (F) $n = 6588$; (G) $n = 21958$, and (H) $n = 43915$.




Figure 8.18 | First loadings for the retrieved low-rank matrix A of different MALDI-TOF measurement matrices for $\lambda = 0.97\lambda_{\text{candes}}$. The green lines under the m/z features indicate which spectra are included in the measurement matrix. The subplots correspond to: (A) $n = 135$; (B) $n = 152$; (C) $n = 300$; (D) $n = 573$; (E) $n = 656$; (F) $n = 6588$; (G) $n = 21958$, and (H) $n = 43915$.

8.5.2.2 Structural Errors

Finally, we compare the retrieved residuals B , similar to previous case studies, to inspect the leakage of signal from the low-rank term. Therefore, correlation measures are used. They provide a spatial and spectral visualization of regions that might be affected by signal leakage. From **Fig. 8.19**, we observe that the high-dimensional nature of the largest three data matrices influences the usability of the correlative measure. Growing size results in the measure to become singular. Secondly, we see that generally the leakages occurs in the same spatial regions regardless of the measurement matrix size.

Finally, we can also review the spatial autocorrelation per ion image. The latter is depicted in **Fig. 8.20**. We observe that this measure is not affected by the dimensionality. This is logical as it is not the spatial size that is increased, but rather the spectral size. Secondly, we see that the unbinned versions (F, G, H) possess a lower spatial autocorrelation. It is not clear whether this is the result of the binning or a “true” observation.




Figure 8.19 | Spectral correlation of the residual matrix B of different MALDI-TOF measurement matrices for $\lambda = 0.97\lambda_{\text{candes}}$. The subplots correspond to: (A) $n = 135$; (B) $n = 152$; (C) $n = 300$; (D) $n = 573$; (E) $n = 656$; (F) $n = 6588$; (G) $n = 21958$, and (H) $n = 43915$. The intensity corresponds to the spectral correlation as aforementioned in this chapter.




Figure 8.20 | Spatial correlation of the residual matrix B of different MALDI-TOF measurement matrices for $\lambda = 0.97\lambda_{\text{candes}}$. The green lines under the m/z features indicate which spectra are included in the measurement matrix. The subplots correspond to: (A) $n = 135$; (B) $n = 152$; (C) $n = 300$; (D) $n = 573$; (E) $n = 656$; (F) $n = 6588$; (G) $n = 21958$, and (H) $n = 43915$.

8.5.3 Discussion

In this case study, we reviewed the in-memory scalability of our methodology in combination with our framework's hyperparameter setting optimized for the trade-off case. We have shown that our implementations scale linearly with measurement matrix size as theoretically predicted by $\mathcal{O}(mnr)$. We have shown that a fixed λ -setting gives similar results when the measurement matrix is scaled. Except, its relative cardinality and rank decrease, the corresponding compression ratio grows, and its smaller singular values increase in value. Finally, we have shown that the structural errors remain similar, but that our spectral auto-correlation measure increasingly suffers from the high-dimensional nature. Finally, we have shown for our largest in-memory measurement matrix the ability to compress its corresponding 30 GB data set size into a version requiring less than 30 MB on disk.

8.6 Case Study D: Terabyte Scale

This case study delivers preliminary results on a MALDI timsTOF measurement matrix requiring > 1 TB of memory for processing via PCP. As far as we are aware, this is a first for MALDI timsTOF data. We gradually scaled up the measurement matrix to a size of $164,808 \times 568,142$ (which under normal circumstances would require around 1.38 TB of memory to house the different matrices involved in the algorithm), which we present in this study. At the time of this experiment, the available memory was 1 TB. This includes half of the available timsTOF measurement matrix, vectorizing the m/z and drift time dimension.

8.6.1 Additional Experimental Design

In this section, we provide additional experimental design settings.

Method We use PCP with randomized svd, not including recycling of singular vectors without iterator. The latter was opted for to relieve some of the computational costs associated with it, and especially the out-of-memory QR implementation that is available from `dask`¹¹ requires reasonable disk read/write time. The λ parameter is set to $\lambda = 1\lambda_{\text{candes}}$.

Assumptions We assume that the MALDI timsTOF measurement matrix is suitable for a low-rank plus sparse/dense noise description.

Dask Settings For `dask`¹², we set the chunk size to 8002×8002 , which corresponds to chunks of about 128 MB. Furthermore, we defined the number of threads per worker to 52 and the number of workers to 1. This is optimal as we work on a single machine with 52 cores. However, it might be advantageous to

reduce this to 50 threads to prevent oversubscribing. Furthermore, it is of utmost importance to set the number of threads for OpenBLAS to 1 to prevent oversubscribing of the threads via OpenBLAS. We excluded the use of processes. Finally, it is also advised to set the upper limit for fuse width to 5. This has been proven useful for the SVD^a.

8.6.2 Results

It took seven days to process this size of matrix on the available compute and memory resources. The largest factor contributing to this time was spent on reading and writing to and from the disk (already an SSD-disk in this case). We obtained a low-rank matrix of rank 13. The latter is considered low for this measurement matrix. We believe this might be caused by: (a) the exclusion of an iterator for the randomized svd; (b) the very sparse spectral nature of this measurement matrix, and/or (c) we only consider half of the m/z dimension. Nevertheless, we share here some of its scores and loadings and some retrieved denoised ion images.

Scores and Loadings In Fig. 8.21, 8.22, and 8.23, the score and loading respectively associated to the first, second, and eight component is depicted. From these, we observe that the scores align to physical regions within the mouse pup. However, the observed loadings are sparse. We believe this might be due to the occurrence of a few very high intensity spectral features caused by the peak picking that took place on this data set. Nevertheless, smaller singular values and associated scores and loadings seem to be less sparse along the spectral dimension (see e.g. Fig. 8.23). The observed sparsity might also be associated to performing PCP on the whole mouse pup, effectively making organ-specific molecules inherently coherent with respect to the massive size of the whole mouse pup. This could be partially corrected for by normalizing the data before processing or spatially restricting the analysis to certain organ subareas.




Figure 8.21 | Score and loading associated to the first component of PCP for a MALDI timsTOF measurement matrix to a size of $164,808 \times 568,142$ with $\lambda = 1\lambda_{\text{candes}}$.

^asee <https://blog.dask.org/2020/05/13/large-svds> note that these experiments were performed before our corrections to the dask implementation (see Chapter 7)




Figure 8.22 | Score and loading associated to the second component of PCP for a MALDI timsTOF measurement matrix to a size of $164,808 \times 568,142$ with $\lambda = 1\lambda_{\text{candes}}$.




Figure 8.23 | Score and loading associated to the eight component of PCP for a MALDI timsTOF measurement matrix to a size of $164,808 \times 568,142$ with $\lambda = 1\lambda_{\text{candes}}$.

Ion Images In **Fig. 8.24**, two sets of ion images (top and bottom) are presented associated to the different vectorized spectral features. From left to right, the measurement matrix (A), the retrieved denoised low-rank term (B), and the retrieved residual term (C) are given. In these images, the denoising capabilities of PCP are apparent, and present promising results for further investment into and enhancement of our techniques.




Figure 8.24 | Example ion images associated to (A) the measurement matrix, (B) the retrieved denoised low-rank term, and (C) the retrieved residual term. The intensity units are arbitrary.

Disk: Read - Write As aforementioned, the disk read and write speed becomes the main bottleneck for data sets of this size. This is depicted in **Fig. 8.25**. The horizontal axis represents the running time, while the vertical axis represents task channels, i.e. threads (52 in total). The plot consists of hundreds of thousands of colored blocks, each associated to a specific task performed on a chunk in a specific thread. White areas indicate no tasks being performed, which would ideally not be occurring, while orange blocks are associated to disk read and write tasks.




Figure 8.25 | Task stream of dask for the first 18 hours of our optimization. Horizontal axis consists the running time, while vertical axis represents task channels, i.e. threads. Each colored block is associated to a specific task performed on a chunk. White areas indicate no tasks (i.e. this is to be reduced) and small orange blocks are associated to disk read and write tasks.

The orange blocks are the cause of larger white areas within the different task channels, which directly relates to our set-up with one spill directory on a single, separate SSD hard disk. In the future, this bottleneck could be partially alleviated by an efficient RAID set-up to specifically increase those disk read and write speeds.

8.6.3 Discussion

In this case study, we presented the proof-of-concept results of a MALDI timsTOF measurement matrix normally requiring > 1 TB of memory for processing via PCP. As far as we are aware, this is a first for MALDI timsTOF data. We gradually scaled up the measurement matrix to a size of $164,808 \times 568,142$ (which under normal circumstances would require around 1.38 TB of memory to house the different matrices involved in the algorithm). At the time of this experiment, the available memory was 1 TB. This includes half of the available timsTOF measurement matrix, vectorizing the m/z and drift time dimension. Hence, we have shown that our developed methodology is capable of performing PCP out-of-

memory. From the scores and loadings we observe that physical spatial regions seem to be captured by the low-rank term, while the concurrent spectral dimension is observed to be sparse. We note that the “signal”, i.e. what is of biological interest, associated to organ-specific tissue can be retrieved by the residual term for this mouse pup MALDI timsTOF measurement matrix. Secondly, we have shown the denoising capabilities of our method for different ion images. Finally, we have discussed the current bottleneck of this methodology, which consists of the disk read/write time. To further scale this problem two promising directions are available: (a) the use of a proper RAID configuration on the spill disk, and (b) the use of GPUs for matrix-matrix heavy calculations, in combination with the out-of-memory capabilities of dask.

8.7 Conclusions

In this chapter, we presented four different case studies related to the use of our developed methodology in a molecular imaging context, for use on MALDI (tims)TOF-type IMS measurement matrices and with a focus on dimensionality reduction and compression.

In a first case study, we reviewed the use of our developed framework for the application-driven hyperparameter setting of PCP for MALDI-TOF IMS data. We introduced three applications and demonstrated some of their quantitative as well as qualitative behavior. We demonstrated its improvement over the common, predefined, and wide-spread λ_{candes} parameter, in compression as well as in retrieved quality of the corresponding spatial components in \mathbb{U} . Secondly, we have shown that the three defined application cases result in distinct solutions with specific favorable properties. In a second case study, we have compared PCP to PCA on its ability to compress MALDI-TOF IMS data. We have shown that a 1.73-fold improvement in compression rate can be obtained with PCP in comparison to PCA, while obtaining a similar spatial reconstruction score. This results in a compression ratio of 19.03 for PCP instead of 10.98 for PCA. Concurrently, this observation coincides with our framework’s optimal λ for a maximal compression application. In the third case study, we reviewed the in-memory scalability of our methodology in combination with our framework’s hyperparameter optimization for the trade-off case. We have shown that our implementations scale linearly with measurement matrix size as theoretically predicted by $\mathcal{O}(mnr)$. We have shown that a fixed λ -setting gives similar results when the measurement matrix is scaled. And, we have shown for our largest in-memory measurement matrix the ability to compress its corresponding 30 GB data set size into a version requiring less than 30 MB on disk, while retaining spatial and spectral reconstruction scores. Finally, in the fourth case study, we presented the proof-of-concept results of a MALDI timsTOF measurement matrix normally requiring > 1 TB of memory for processing via PCP, which as far as we know is a first. We gradually scaled up the measurement matrix

to a size of $164,808 \times 568,142$ (otherwise requiring around 1.38 TB of memory). Furthermore, we have shown its denoising capabilities. Finally, we have discussed the current bottleneck for further scaling of the methodology related to disk read/write time.

Hence, we have shown that for molecular imaging, and specifically imaging mass spectrometry with its high-dimensionality and massive data sets, our developed methodology can be applied for dimensionality reduction and compression purposes, while also providing measures to detect any leakage of signals in the process.

Bibliography

- [1] Lin, Z., Chen, M. & Ma, Y. The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. *arXiv preprint arXiv:1009.5055* (2010).
- [2] Savitzky, A. & Golay, M. J. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry* **36**, 1627–1639 (1964).
- [3] Virtanen, P. *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17**, 261–272 (2020).
- [4] Van Rossum, G. & Drake, F. L. *Python 3 Reference Manual* (CreateSpace, Scotts Valley, CA, 2009).
- [5] El Aalamat, Y. *et al.* Improved computational analysis of imaging mass spectrometry data through sparse intensity variation denoising. (Talk) *The 63rd Conference on Mass Spectrometry and Allied Topics (ASMS 2015)* (2015).
- [6] El Aalamat, Y. *et al.* Enhancement of imaging mass spectrometry data through removal of sparse intensity variations. (Poster) *NVMS/BSMS 2016 Rolduc Meeting* (2016).
- [7] El Aalamat, Y. *et al.* Enhancement of imaging mass spectrometry data through removal of sparse intensity variations. *Benelux Bioinformatics Conference (BBC 2015)* (2015).
- [8] El Aalamat, Y. *et al.* Efficient noise reduction in imaging mass spectrometry data using robust pca. *The 62nd Conference on Mass Spectrometry and Allied Topics (ASMS 2014)* (2014).
- [9] Candès, E. J., Li, X., Ma, Y. & Wright, J. Robust principal component analysis? *Journal of the ACM (JACM)* **58**, 1–37 (2011).
- [10] Camacho, J., Smilde, A., Saccenti, E. & Westerhuis, J. All sparse pca models are wrong, but some are useful. part i: Computation of scores, residuals and explained variance. *Chemometrics and Intelligent Laboratory Systems* **196**, 103907 (2020).
- [11] Benson, A. R., Gleich, D. F. & Demmel, J. Direct qr factorizations for tall-and-skinny matrices in mapreduce architectures. In *2013 IEEE international conference on big data*, 264–272 (IEEE, 2013).
- [12] Dask Development Team. *Dask: Library for dynamic task scheduling* (2016). URL <https://dask.org>.

Chapter 9

Case Studies II: Integral-Field Spectroscopy

The key result of this chapter is that our methodology can be used for

- co-adding purposes with increased signal-to-noise ratio
- reducing data artifacts to improve the detectability of faint objects
- exploring spiral galaxies for subsequent measurement of galaxy kinematics
- improving signal-to-noise ratios of line emissions

9.1 Introduction

In this chapter, we explore the usability of robust matrix decomposition methods for integral field spectroscopy and especially the MUSE instrument¹ at ESO's Very Large Telescope (VLT).

9.1.1 Goals and Objectives

The nature of this chapter is exploratory due to the recent introduction (2014) of integral field spectroscopy and its imaging capabilities, and due to the instrument-specific nature of methodology within astronomy. This underlines the novelty of the application of (robust) matrix decomposition methods for such purposes. Hence, the aim of this chapter is to explore the use of robust matrix decomposition methods for noise reduction purposes to integral field spectroscopy, with a focus on the MUSE instrument. We achieve this by presenting the versatility of

such methods through different case studies applied on different levels ranging from low-level reduction, to high-level data exploration and visualization, and domain-specific application. For all these cases the emphasis lies on the methodological aspect along with its noise reducing capabilities. These findings subsequently formulate the answer to research subquestion B.2.

B. How can the application-driven methodology we have developed, be used to address application domain-specific challenges?

B.2. In astronomical imaging, and specifically integral field spectroscopy *low signal-to-noise ratios and instrumental artifacts* are the most pressing bottleneck. How can we apply the developed methodology to measurement matrices of the MUSE instrument picturing a spiral galaxy (AM2210-262), the Hubble Ultra Deep-Field (UDF-06), and different exposures of the Hubble Deep Field South (HDFS), *for noise reduction purposes?*

9.1.2 Overview

In the next section, we introduce the experimental design for all four case studies, including the measurement matrices, and the used models and methods. Finally, we shortly describe the nomenclature and provide additional settings. Thereafter, the four case studies are presented:

Case Study A: investigates the use of robust matrix decomposition methods for co-adding 54 exposures of the Hubble Deep Field South to achieve higher SNR;

Case Study B: presents the identification of instrumental artifacts for subsequent removal to improve detectability of faint objects, e.g. faint line emitting galaxies;

Case Study C: explores the galaxy kinematics of a well-extended source through matrix decomposition;

Case Study D: presents a high-level use of robust matrix decomposition for improved line emission detection and spatial retrieval of those emissions.

Finally, we draw our conclusions and provide our contributions and recommendations for further investigation.

9.2 Experimental Design

In this section, we elaborate on the experimental design. First, we introduce the different measurement sets. Second, we provide the methods and model that

we use for each case study. Finally, we provide the additional settings regarding stopping criteria.

9.2.1 Measurement Matrix

The different measurement sets are provided by the Leibniz Institute (AIP), Potsdam (GER). They consist of reduced data processed by the MUSE pipeline². All of them went through the standard recipes, except for the exposures of Hubble Deep Field South (HDFS) we have non sky background subtracted versions to our disposal. An overview is given in **Tab. 9.1**.

Table 9.1 | Table containing basic information on the measurement matrices. The measurements are given per case study, along with the object of observation, the number of exposures available, the integration time per exposure, the corresponding matrix size and data size and the datatype. All measurement sets are recorded using the wide field mode. The matrix sizes are obtained by vectorizing the spatial components

Case	Object	Num. of Exp.	Int. Time [s]	Matrix Size	Data Size	Datatype
A	HDFS	54	1,800	$54 \times 3862 \times 1600$	1.27 GB	32 float bits
B	UDF-06	1	36,000	99745×3682	1.47 GB	32 float bits
C	AM2210-262	1	1680	147293×3681	2.17 GB	32 float bits
D	UDF-06	1	36,000	1560×3682	0.02 GB	32 float bits

Given their reduced nature, it is possible that negative flux values occur in the data sets. Those are a residual of some of the subtraction steps in the pipeline. Finally, also note that the object spectra inside the measurement sets have possible different redshifts, adding to the complexity of processing. The corresponding white light images are available from **Fig. 9.1**. For our experiments, we removed the 0^{th} order sky, by removing the median sky signal and found by taking the spatial median of each monochromatic layer, similar to ZAP³. Note that this is similar to mean subtracting, with the difference that median subtraction is a non-linear step, while mean subtraction corresponds to a simple projection⁴.

9.2.2 Methods and Model

We restrict ourselves to the use of PCP for simplicity. Nevertheless, SPCP might improve the results as it is able to deal with small i.i.d. noise. The underlying model is the LMM which is described as

$$\begin{aligned} M &= A + B, \\ &= XY^T + B, \end{aligned} \tag{9.1}$$




Figure 9.1 | White light images of the different measurement sets. The different subplots are: (A) subfield of HDF5, containing galaxies and line emitting galaxies, used in case study A; (B) UDF-06, Hubble Ultra-Deep Field, used in case study B; (C) AM2210-262, spiral galaxy in the Piscis Austrinus constellation, used in case study C, and (D) subfield of UDF-06 used in case study D, corresponds to the galaxy in middle of (B). Flux is given in units of 10^{-20} erg/s/cm²/Angstrom.

where $M \in \mathbb{R}^{m \times n}$ is the synthetic measurement matrix, $A \in \mathbb{R}^{m \times n}$ is assumed to represent a low-rank matrix, $B \in \mathbb{R}^{m \times n}$ a noise matrix. Secondly, note that A can also be decomposed following the SVD as $A = U\Sigma V^T$. We will assume that signal, i.e. what is of interest, is captured by A , while noise and residuals are approximated by B .

9.2.3 Nomenclature

MUSE provides data cubes with two spatial dimensions and one spectral dimension. Furthermore, when multiple exposures are available a fourth dimension of time is added to those three. Each of these dimensions can be condensed, e.g. by taking mean or median values over all values in that particular dimension. We will use the following nomenclature:

White light image: the spectral dimension is condensed to a single image by taking the mean value of all values in the spectral dimension;

Vectorization: consists of linearly transforming a matrix into a single vector. For example, when each monochromatic layer is considered a matrix,

i.e. a two-dimensional set of values, vectorization consists of stacking the columns of that matrix on top of each other so that the monochromatic layer is condensed to a single vector. Spatial information, e.g. spatial nearness, is lost in this process. Performing vectorization on each monochromatic layer results in a set of vectors, i.e. a matrix with one condensed spatial direction and one spectral dimension. This corresponds to going back the pixel table as originally obtained in the MUSE pipeline;

Reconstruction: when we talk about reconstruction we mean going from vector form to image, undoing the vectorization;

Spectral Channel Map: corresponds to a single monochromatic layer;

Score: corresponds to a particular reconstructed left singular vector (U), containing spatial information. It can be seen as a “synthetic” spectral channel map. At the same time, and together with its corresponding right singular vector, it represents a rank-1 decomposition of the original input matrix, i.e. full set of spectra, not to be confused with the matrix notation of a monochromatic layer. Hence, it contains particular information of the original measurement set;

Loading: corresponds to a particular reconstructed right singular vector (V), containing spectral information. It can be seen as a “synthetic” spectrum. At the same time, and together with its corresponding left singular vector, it represents a rank-1 decomposition of the original input matrix, i.e. full set of spectra, not to be confused with the matrix notation of a monochromatic layer. Hence, it contains particular information of the original measurement set;

Components: correspond to a score, or loading, or a pair of those. Together with its corresponding singular value, it represents a rank-1 matrix. The sum of all components, i.e. of all rank-1 matrices reconstructs the low-rank term. Each component carries a particular part of the information from the low-rank term.

9.2.4 Additional Settings

Method Implementation For PCP we use the inexact augmented Lagrangian multiplier method (IALM)⁵ was utilized.

Stopping Criteria The stopping criteria for PCP are set to

$$\begin{aligned} \frac{\|M - A - B\|_F}{\|M\|_F} &\leq 1 \times 10^{-5}, \\ \mu \frac{\|\hat{B} - B\|_F}{\|M\|_F} &\leq 1 \times 10^{-4}, \end{aligned} \tag{9.2}$$

where μ is an algorithmic variable. A maximum of 50 iterations is set.

Moran's I We use an equal weighted 4-connected pattern, similar to previous chapters, as measure of spatial auto-correlation.

Spectral Correlation For the spectral correlation we use a simple 2-connected pattern, where we only consider neighboring pixels. This corresponds roughly to assuming correlation on a 0.125 nm scale (assuming equidistant grid in spectral direction).

Savitzky-Golay Filter The Savitzky-Golay filter⁶ was utilized from SciPy package⁷. The window length is set to 15 and the order of the polynomial set to 2. Its parameters were set through trial-and-error.

Computational Specification Floats with a precision of 32 bits were used and the underlying linear algebra is performed via OpenBLAS-LAPACK accessed via NumPy and SciPy⁷. All experiments were performed in Python 3.8.3⁸ in purely CPU-based processes. Finally, the computer system specifics are given in **Tab. 9.2**.

Table 9.2 | Computer system specifications. The operating system, number of cores and their clock speed, and the size of available memory is stated.

OS	Cores	Clock Speed	Memory Size
Windows 10 Pro	52 (104 logical)	2.7 GHz (3.3 GHz overclock)	1.5 TB

9.3 Case Study A: SNR Improvement by Co-Adding through Matrix Decomposition

In this first case, we study the use of matrix decomposition methods for the co-adding of multiple exposures. The collaborative point of view of such methodology, where multiple pixels within monochromatic layers aid the process of retrieving the continuum spectrum, is considered advantageous with respect to time averaging. Secondly, the temporal differences between exposures are considered coherent within the robust methodology framework, such that those local effects are assumed to be end up in the residual term, further enhancing the SNR. This could be interesting with an eye towards the detection of faint continuum spectra of detected line emitting galaxies. The incoherence principles of robust methods also guarantee in such case that sparse outliers, i.e. strong line emissions, are captured in a residual term and thus prevent the distortion of the faint spectra themselves.

9.3.1 Additional Experimental Design

In this section, we provide additional experimental design settings.

Data Reduction We used a data cubes where no sky method is applied (HIERARCH ESO PRO REC 2 PARAM14).

Preprocessing We subtracted the 0^{th} order sky background, by taking the median of each monochromatic layer and subtracting it from all spectra.

Method We use PCP with `scipy_gesdd`. The λ was set to $0.6\lambda_{\text{candes}}$, through a “maximal compression” setting of our developed framework. The PCA version used, is covariance-based, and its rank was set to the rank value obtained by PCP.

Stacking For this methodological study, we use 54 exposures the Hubble Deep Field South (HDFS) of MUSE. Each exposure cube is first vectorized to obtain a matrix with one spatial and one spectral dimension. Next, the exposures are stacked in spectral direction. This ensures we maintain temporal data on the spectra. The exposures can also be stacked in spatial direction. The latter ensures that temporal data on the spatial evolution are maintained (see stacking in **Example 4.1**, section 4.1.3).

Assumptions We assume that spectral dimensions are aligned properly between exposures.

9.3.2 Results

In this section, we present our results. We compare PCA and PCP results along with time averaging, both with a mean and median approach. Secondly, we compare the first two scores and loadings of PCP and PCA, review the spectral improvement and finally show some retrieved spectra of a single galaxy when processing a different number of exposures. We achieved a compression ratio of 75.58 with respect to the reduced data, and a compression ratio of 1.40 with respect to time averaged results, while retaining temporal information.

PCP-PCA Comparison In **Fig. 9.2**, the white light images and corresponding mean spectra are depicted for the denoised data for PCA (top) and PCP (bottom). We observe that the white light image of PCP is visually cleaner than the retrieved PCA version, in which a grid-like structure is remarked. Furthermore, the galaxy is better delimited and its faint edges are visible for PCP. When looking at both corresponding mean spectra, this takes all depicted spaxels into account, we see that PCA is mainly affected by the OH forest, where PCP has smaller OH residuals. Probably, these grid lines correlate with those OH residuals, such that PCA fits on those instead of the galaxy.

Time Averaging The same data is also time averaged via a mean and median approach. For both, the results are given in **Fig. 9.3**. The mean approach (top) behaves similarly to PCA spectrally, although its lowered continuum. Its white light image on the other hand indicates that spatial patterns are better maintained. Secondly, we see a similar mean spectrum for time median (bottom) as for PCP with the exception of some outlying values at around 560 and 840 nm for the median approach. Comparing its white light image with that of PCP we observe that an arm-like feature is missing on the top-left of the largest galaxy. Whether this is a true feature or an artifact of the background grid is not clear.

Score and Loadings The aforementioned observations for PCA and PCP are further confirmed by reviewing their respective first two scores and loadings. These are synthetic components that are retrieved from the decomposition. When the data set is well retrieved, we expect the scores and loadings to (partially) align to galaxy kinematics or physical differences within the galaxy. The first score and loading often “set the scene” and can be interpreted as spatial and spectral components that best match to the whole spatial and spectral dimension. For MUSE, this often consists of the continuum spectrum. The first pair of scores and loadings are depicted in **Fig. 9.4** and the second pair in **Fig. 9.5**. For PCA (both top), we see that the data set is poorly correlating to the galaxy, where PCP (both bottom) makes use of the information. The second components of PCP, actually unravels galaxy kinematics, predominantly an emission line at around 570 nm.

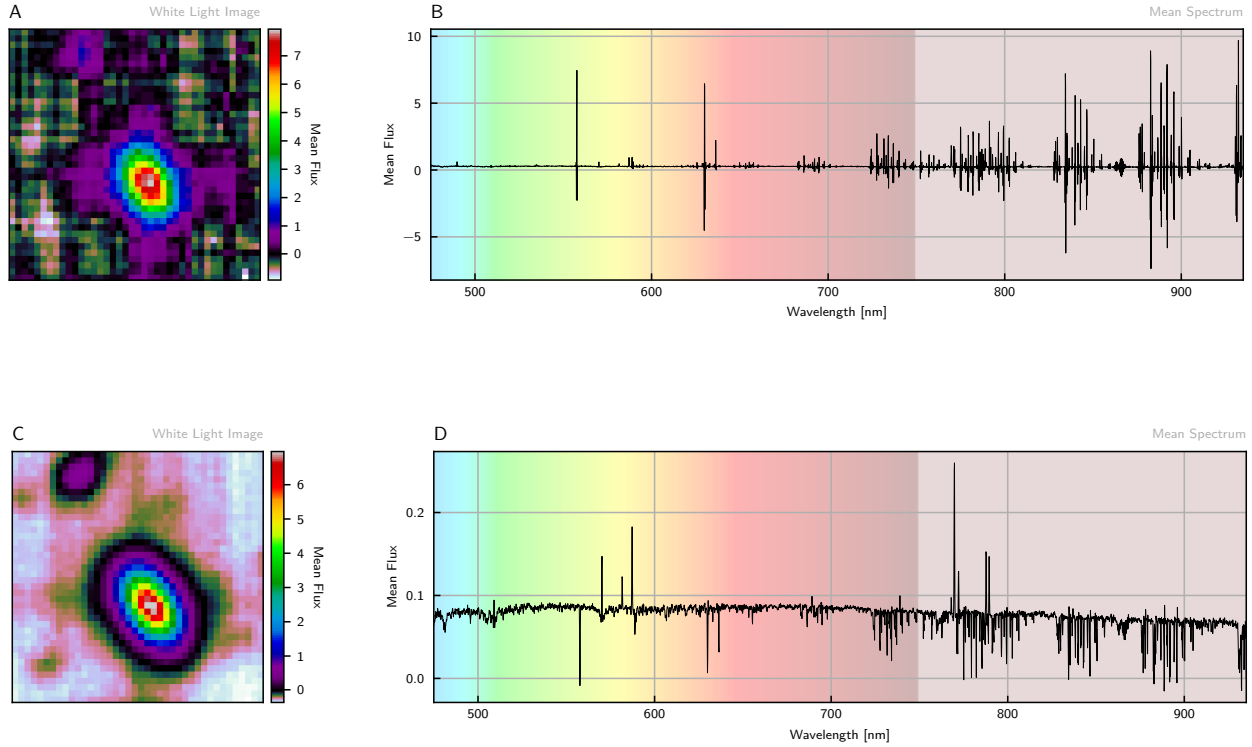


Figure 9.2 | White light images and corresponding mean spectra for denoised measurement sets for PCA (A, B) and PCP (C, D) for a subfield of the HDFs. Flux is given in units of 10^{-20} erg/s/cm²/Angstrom.

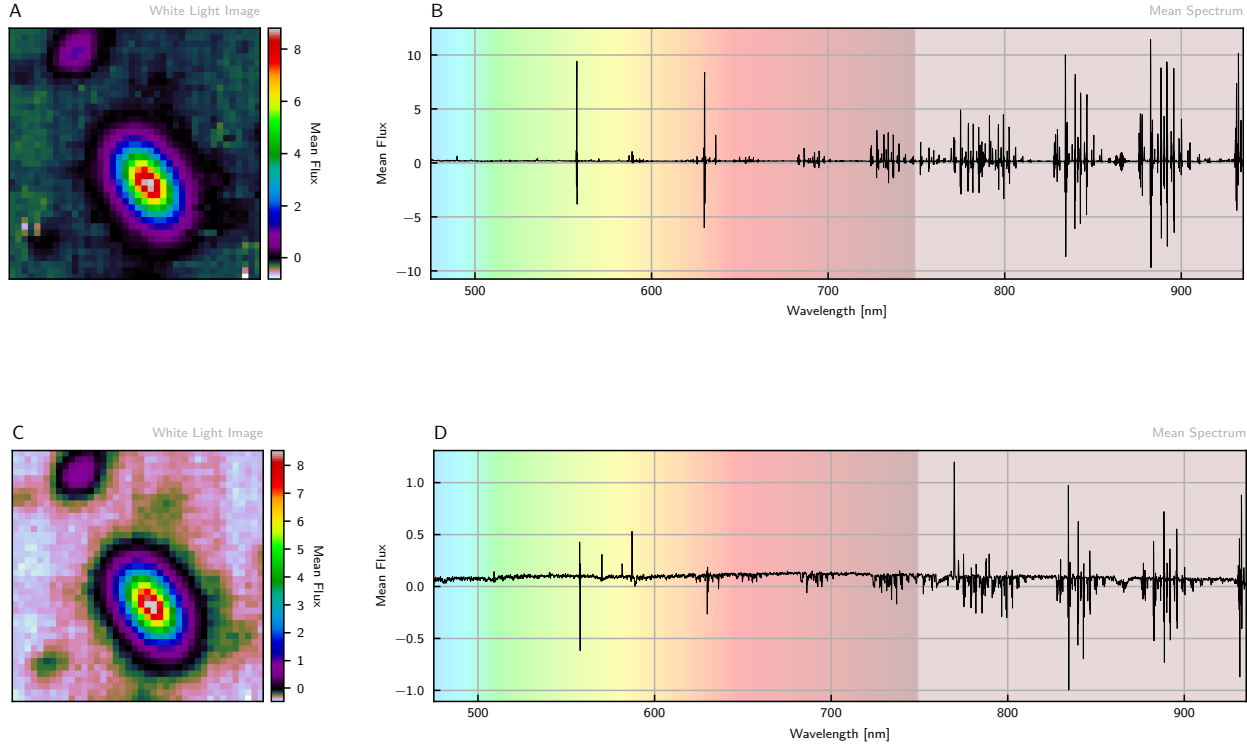


Figure 9.3 | White light images and corresponding mean spectra for denoised measurements sets for time mean (A, B) and time median (C, D) co-adding for a subfield of the HDFs. Flux is given in units of 10^{-20} erg/s/cm²/Angstrom.

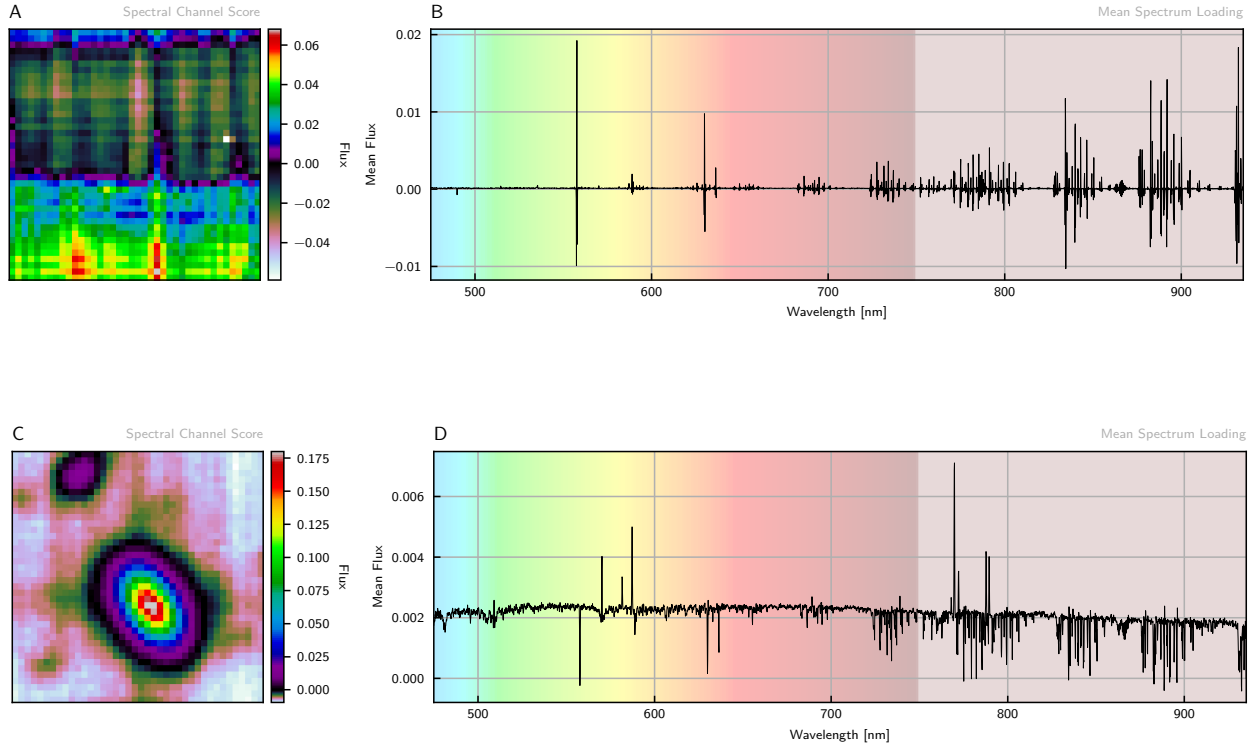


Figure 9.4 | Scores and loadings for the first component of PCA (A, B) and PCP (C, D) for a subfield of the HDFs. Flux is given in normalized units.

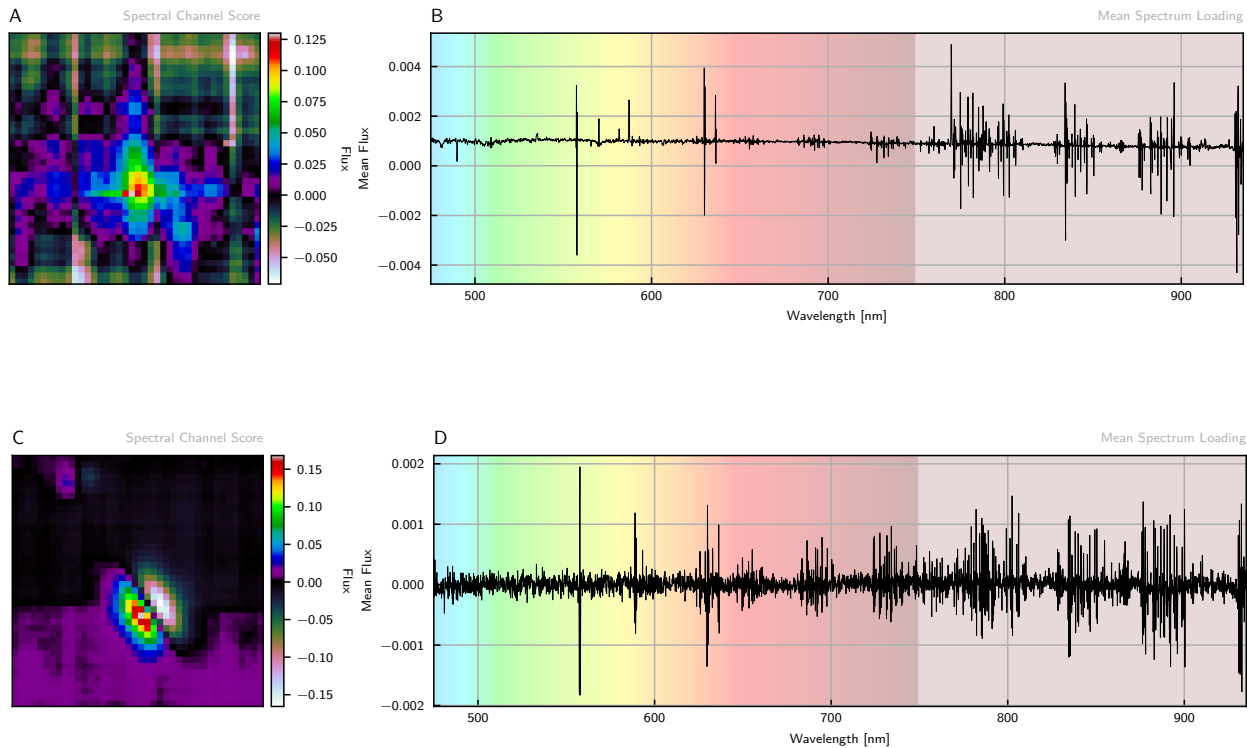


Figure 9.5 | Scores and loadings for the second component of PCA (A, B) and PCP (C, D) for a subfield of the HDFS. Flux is given in normalized units.

9.3.2.1 Spectral Improvement

In this experiment, we compared the relative difference between PCP, applied on a number of reduced exposures, and time averaging, both time median and mean on the full 54 cubes. Therefore, we restricted ourselves to the spectra related to the center of the largest galaxy in the data, as reviewed in the previous experiment. For PCP, we used the spatial mean, i.e. the mean of the monochromatic layer, of those spectra and for the time averaging, we utilized both mean and median approaches. We used the following equation to define the relative difference

$$\epsilon = \frac{\|\text{PCP} - \text{reference spectrum}\|_1}{\|\text{reference spectrum}\|_1} \times 100\%, \quad (9.3)$$

we used the ℓ_1 -norm to prevent overfitting on the OH forest residuals. The results are depicted in **Fig. 9.6**. First, we remark that after the co-addition of 5 to cubes a plateau is reached with a small upwards include for the time median and mean approaches with spatial averaging through its mean. Whether this is due to the λ setting preventing new information to be grasped or the fact that no new information is entered is not clear. The latter can also indicate that PCP reaches a lower noise level of removed temporal-unstructured artifacts.




Figure 9.6 | Relative difference in % for different exposure settings. The relative difference is calculated from Eq. 9.3. Low values correspond to a good match to the traditional co-adding method when all 54 cubes are used.

Variability The variability of the retrieved spectra via our method in a flat wavelength region between 662.466 and 681.216 nm is depicted in **Fig. 9.7**. We observe that PCP has the ability to achieve lower variability while using less exposures. When using all 54 exposures are used, a 1.2 – 2.4 fold reduction in variability is observed.




Figure 9.7 | Standard deviation of region between 662.466 and 681.216 nm. Low values correspond to a low variability, references for traditional co-adding methods are given when all 54 cubes are used.

9.3.2.2 Spectral Comparison

Finally, we compare the obtained spectra, again of the center region of the largest galaxy and compare the obtained spectra for both PCP and traditional time averaging approach, when only a limited number of exposures is used. The results are depicted in **Fig. 9.8**. We remark the following:

- OH forest residuals seem to be more reduced for the PCP approach, with respect to the traditional approach;
- The variability of PCP appears smaller with respect to time averaging;
- For low wavelengths < 600 nm the flux obtained by PCP is consistently below the time averaging flux;
- For wavelengths in the range 500 – 970 nm the flux obtained by PCP is consistently above the time averaging flux.

The low wavelength range mismatch can be linked to observations that are presented in case study B. However, whether the other effects can be related to the lower noise level of removed temporal-unstructured artifacts, i.e. structural biases, is not yet clear.




Figure 9.8 | Total flux for different number of exposures for co-adding purposes for both PCP and time averaging. The different subplots consist of: (A) the co-adding of 2 exposures; (B) the co-adding of 4 exposures; (C) the co-adding of 6 exposures; (D) the co-adding of 41 exposures. Black lines correspond to the PCP method, where red consists of the traditional time averaging method. Flux is given in units of 10^{-20} erg/s/cm²/Angstrom.

9.3.3 Discussion

In this case study, we showed the abilities of robust matrix methods for co-adding purposes. They show the promise for increased SNR, especially for the reduction of structural noise that differs between exposures. The latter can consist of instrumental signatures, but also sky background residuals. Both are valuable for the quantization of faint continuum spectra associated to detected line emissions of line-emitting galaxies. But also for the detection of line-emitting galaxies that appear to be overlapping with the OH forest and possible related residuals after subtraction. Finally, they show the ability to achieve lower variability while using less exposures. When using all 54 exposures a 1.2 – 2.4 fold reduction in variability is observed.

9.4 Case Study B: Improving Detectability by Reducing Level of Artifacts

This case investigates the identification of a faint residual that is present in the reduced MUSE data. A similar observation was made in MUSE variance cubes^{9,10}, where it was found that $\sigma_{\text{emp}} = 1.15 - 1.20\sigma_{\text{pipe}}$, where σ_{emp} is the empirically calculated standard deviation and σ_{pipe} is the standard deviation that is provided by the reduction pipeline. This implies that variance cubes provided by the pipeline underestimate the true variances, believed to be caused by hidden (non-Gaussian) systematics such as imperfect flat-fielding or imperfect sky-subtraction⁹. We present the identification of spatial and corresponding spectral signatures of different of those artifacts via PCP for subsequent subtraction.

9.4.1 Additional Experimental Design

In this section, we provide additional experimental design settings.

Masking We masked the brightest objects by manually annotating the white light image. However, automated threshold methods, sigma-clipping methods or catalog-based masking of the (brightest) objects is possible. Nevertheless, for PCP, with its incoherence principles, small outliers such as line emitting galaxies or spatially small objects, i.e. consisting of only a few spaxels, will not affect the process.

Preprocessing We subtracted the 0^{th} order sky background, by taking the median of each monochromatic layer of the masked cube and subtracting it from all spectra.

Method We use PCP with randomized SVD, with QR iteration and including recycling step. The number of iterations was set to 2 and the number of oversamples to 20. The λ was set to $1\lambda_{\text{candes}}$. The PCA version used, is covariance-based, and its rank was set to the rank value obtained by PCP.

Assumptions The faint background structure is low-rank and meets incoherence principles. Secondly, we assume that faint line-emitting galaxies and other faint objects in the “background” are not low-rank.

9.4.2 Results

In this section, we first present the identification of these artifacts spectrally as well as spatially and point out their presence in the data cube. Concurrently, we compare the results with PCA that is not able to recover the spectral and spatial components accurately and efficiently.

9.4.2.1 Artifact Identification

When processing the masked MUSE measurement set through PCP and PCA, we hope to obtain the artifacts in the low-rank term as they form a grid-like structure over the spatial plane. The obtained white light images and mean spectra of those are given in **Fig. 9.11**. We observe the following:

- the white light image of PCA suggest outliers are present in the data (the mean spectrum confirms this at around 570 nm);
- the white light image of PCP indicates that it grasps the grid-like structure as seen in **Fig. 9.1** (B);
- the mean spectrum of the low-rank term obtained by PCA is affected by the presence of this outlier;
- for PCP, and to a lesser extend in PCA, we observe a small artifact for low wavelengths, i.e. < 700 nm. A small offset with respect to the zero flux line;
- both PCP and PCA seem to fit the residual OH forest features.

Singular values The singular values represent the energy captured by each of its singular vectors and corresponding rank-1 representation. In short, it indicates the importance of the corresponding score and loading, i.e. left and right singular vectors, within the whole measurement set. The scores and loadings are decomposed factors of the low-rank term and form synthetic images and spectra that are expected to align, to some extend, to physical phenomena within the data. The singular values corresponding to PCA and PCP are given in **Fig. 9.9**. We observe the following:

- the first singular value is a full order higher for PCA than PCP. This might indicate PCA overfits on outlying values within the data;
- PCP's decay is fast, in contrast to PCA. This is advantageous for compression purposes;
- slow decay of PCA can be attributed to its compensating behavior for mismatch in the first rank-1 approximation. This behavior results in PCA to need a lot of components to fully describe the background, with the increased possibility of grasping faint objects in the process.

For the scores and loadings of the first three components of PCP and PCA we refer the reader to the addendum, and more specifically to **Fig. A.9, A.10, A.11**.

Impact The impact of these observations is confirmed when comparing the retrieved artifact to the unmasked MUSE data cube. This is obtained similarly to the variance calculation, by calculating the mean and median spectra, i.e. the mean or median taken for each monochromatic layer, of both retrieved artifacts and unmasked MUSE data cube.




Figure 9.9 | Singular values of the retrieved low-rank approximation for PCA (black, triangle) and PCP (red, circle). PCA's singular values show slow decay with respect to those of PCP. Furthermore, the large offset in the first singular value results in compensation behavior for PCA. This detrimental effect is not observed for PCP.

The relative impact (ϵ) is then calculated by dividing them by one another:

$$\begin{aligned}
 \epsilon_{\text{med}} &= \left| \frac{\text{median spectrum}_{\text{PCP}}}{\text{median spectrum}_{\text{unmasked}}} \right| \times 100\%, \\
 \epsilon_{\text{mean}} &= \left| \frac{\text{mean spectrum}_{\text{PCP}}}{\text{mean spectrum}_{\text{unmasked}}} \right| \times 100\%.
 \end{aligned} \tag{9.4}$$

From **Fig. 9.10**, depicting this comparison, we notice that 10 – 25% of total flux contained in the monochromatic layers can be attributed to this artifact. Furthermore, we observe a similar pattern as obtained in previous observations (see **Fig. 2⁹**). Subtracting this artifact, identified via PCP, could be a first step for improving detectability of faint objects in the non-masked areas due to its concurrent spatial identification of this grid-like structure. Secondly, one could argue these impact numbers are skewed by the large number of the spaxels considered in the masking. However, when reviewing a smaller subfield of the UDF-06 and a single galaxy within that subfield, as given in **Fig. 9.12**, we observe consistency in its mean spectra and the actual artifact, suggesting that the galaxies are also affected by a similar artifact.




Figure 9.10 | Relative impact of artifact, obtained by PCP, and compared to mean and median of monochromatic layers. For each wavelength the relative impact is given following Eq. 9.4.

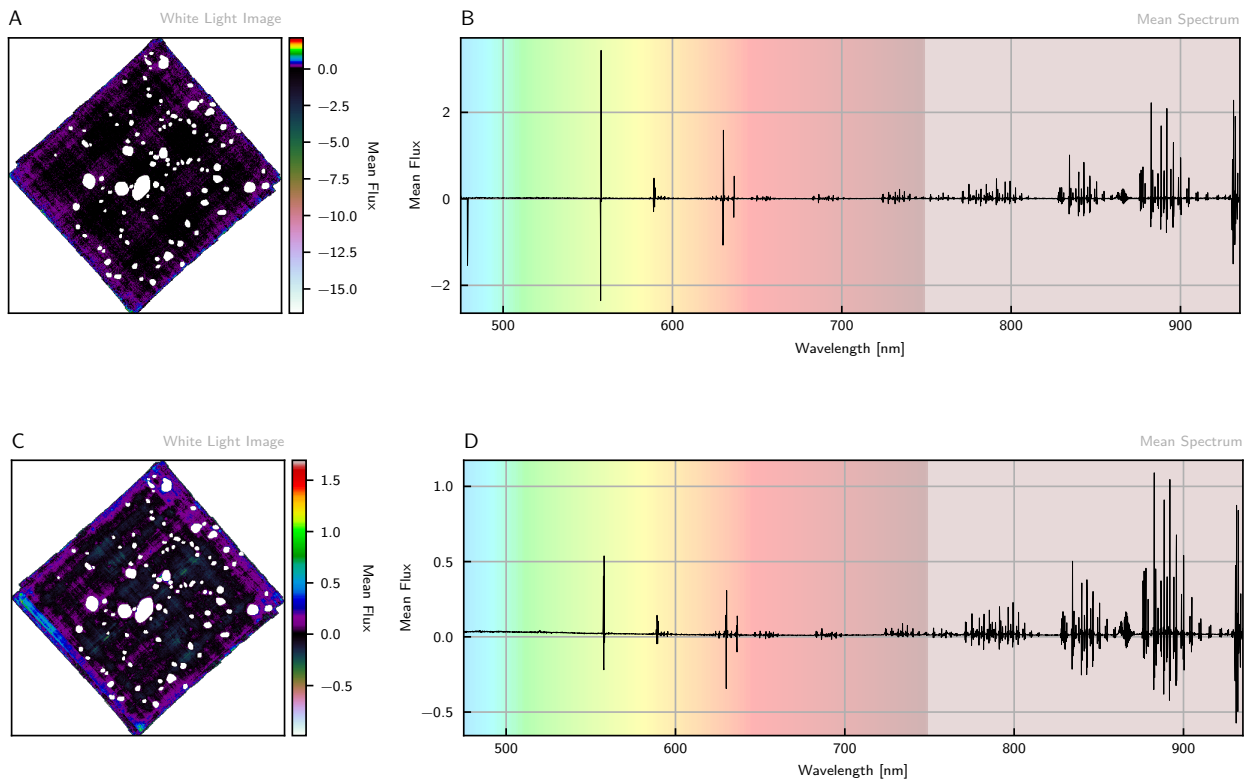


Figure 9.11 | White light images and corresponding mean spectra for denoised measurement sets for PCA (A, B) and PCP (C, D) for a subfield of the HDFs. Flux is given in units of 10^{-20} erg/s/cm²/Angstrom.

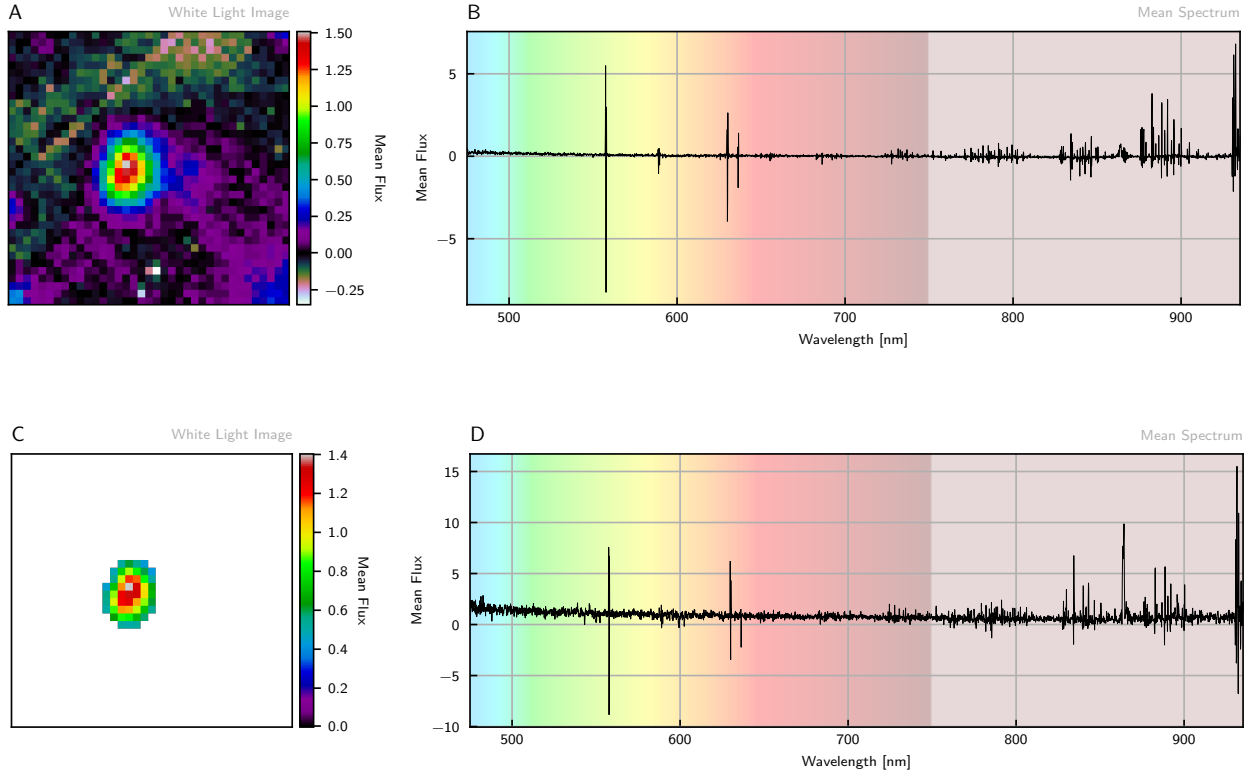


Figure 9.12 | White light images and corresponding mean spectra for reduced data of subfield of UDF-06 (A, B) and cut-out of galaxy in the same subfield of UDF-06 (C, D). Flux is given in units of $10^{-20}\text{erg/s/cm}^2/\text{Angstrom}$.

9.4.3 Discussion

We showed the use of PCP for identification of artifacts. Furthermore, we showed that PCA is not able to retrieve such artifacts due to overfitting on sparse outliers, believed to be line-emitting galaxies. Using more components to improve results of PCA is thought to be detrimental for faint background structures. Secondly, we showed that both PCA and PCP are able to identify OH forest residuals. The latter is, however, already addressed in a non-linear fashion by ZAP³, via a similar approach of generalized principal component analysis¹¹. The artifacts for low wavelengths, i.e. < 700 nm, were addressed in previous papers by a running median^{9,12} filter, known to be introducing new artifacts. Therefore, we propose this novel approach for improving detectability in masked areas. Our approach has the advantage of concurrently providing spatial and spectral signatures. Future direction could be the direct application of line detection tools, e.g. LSDcat¹³, on those masked area and comparison to this running mean approach.

Removing Artifacts from Masked Objects In second instance, we tried to reduce those artifacts from masked objects. This is an ongoing project. So far we remarked the following:

- Noise and signal space are not orthogonal, which mean that orthogonal projections are not useful for removal/reduction of these artifacts (ZAP uses this property by artificially correcting spaces to become nearly orthogonal);
- Noise and signal space are not disjoint, hence oblique projections are not appropriate;
- We explored the use of regression type of solutions, e.g. ridge regression, this was not successful;
- Tried to reduce them via SI-PCP: similar problems occurred as with orthogonal projections;
- Tried to match artifacts with object spectra via their OH forest peaks, on only the OH forest peaks, by removing the running median from the artifact spectra. The latter is observed to break the dimensionality of the artifacts, with the result that there is no unique solution and no good fits are obtained.

The main problem is that artifacts are correlated to object signal. One solution might be to identify the origin of artifacts in the pipeline and improve those sections. However, it is not clear whether the identified artifacts are a direct result of the pipeline or can be attributed to the instrument itself. We believe the translation of this problem into a custom optimization might be a possible direction.

9.5 Case Study C: Exploring Galaxy Kinematics through Component Evaluation

In this case study, we review matrix decomposition methods as exploratory tool for data visualization purposes. By use of the rank-1 descriptions of the low-rank term in our linear model, we explore data-specific relations. In some cases these relations align with galaxy kinematics that can be exploited in further studies. Inversely, the same is true for sparse phenomena, e.g. line emissions, captured by residual term. The latter will be further analyzed in the next case study.

9.5.1 Additional Experimental Design

In this section, we provide additional experimental design settings.

Preprocessing We subtracted the 0^{th} order sky background, by taking the median of each monochromatic layer of the masked cube and subtracting it from all spectra.

Method We use PCP with randomized SVD, with QR iteration and including recycling step. The number of iterations was set to 2 and the number of oversamples to 10. The λ was set to $0.9\lambda_{\text{candes}}$, through a “maximal compression” setting of our developed framework. The PCA version used, is covariance-based, and its rank was set to the rank value obtained by PCP.

Assumptions The main assumption in this study is that physical phenomena, e.g. related to galaxy kinematics, align or partially align to the components, i.e. scores and loadings, as retrieved from the low-rank term.

9.5.2 Results

In this section, we propose the different components obtained by PCP and PCA and compare some of the aspects of both methods. Again, we want to emphasize that these scores and loadings are not necessarily completely aligning to a real physical process. However, they show (synthetic) spectral components that correlate to (synthetic) spatial components. They form the factors by which the method decomposes the terms (see **Example 2.3** of section 2.4.4 for an analogy).

Component Exploration In this section, we compare different score and loadings components, that we manually picked based on their spatial and spectral patterns. We summarize our findings per pair of components.

For **Fig. 9.13**:

- The first score and loading are usually scene setting, i.e. one could see it as a rough mean/median of the data in spectral and spatial direction. This pair of score and loading corresponds to the largest singular value and thus represents the most “energetic” relation within the data found with this method;
- For PCA and PCP we observe similar results, except for an emission peak at around 670 nm that is higher for PCA than PCP.

For **Fig. 9.14**:

- From this second component and on the scores and loadings add or subtract specific information from the first scene setting pair. As such, the following components carry mostly spatial as well as spectral local information or faint structures not captured in the first pair;
- For both PCA and PCP we see a similar loading and score containing a spatial slope from top left to right bottom, and a spectral slope. Note that the positive parts of this spatial slope correlate to the positive parts of the spectral slope, and anti-correlate to the negative parts of the spectral slopes. The same is true for the negative spaxels correlating to the negative parts of the spectral slope. This means that in regions < 700 nm flux is added to the top part of the central galaxy and subtracted from the bottom, while the inverse is true for wavelengths > 700 nm;

For **Fig. 9.15**:

- For PCA, one can see a line emission at around 670 nm and 680 nm and some smaller signatures before and after;
- For PCP, one observes that the line emissions are not as clear. This can be considered to be caused by multiple true physical relations within the data being entangled due the extraction of sparse phenomena within PCP;
- PCP seems to grasp multiple line emissions, being noisy spatially as well as spectrally at 600 nm, 670 nm and 680 nm, but also at 880 nm and 890 nm, this is probably a result of its preference for incoherent structures;
- The line emissions observed by PCA are considered sparse for PCP, such that they end up in the residual term and not the low-rank term. When we apply spatial and spectral correlation measures on this residual term of PCP, we observe these features appearing. The latter is depicted in the bottom row of **Fig. 9.16**;

For **Fig. 9.17**:

- For PCA, multiple line emissions are observed, with a peak at around 570 nm,
- For PCP, again multiple physical aspects seems to get entangled within the data;




Figure 9.13 | Scores and loadings for the first component of PCA (A, B) and PCP (C, D) for the spiral galaxy AM2210-262. Flux is given in normalized units.




Figure 9.14 | Scores and loadings for the second component of PCA (A, B) and PCP (C, D) for the spiral galaxy AM2210-262. Flux is given in normalized units.

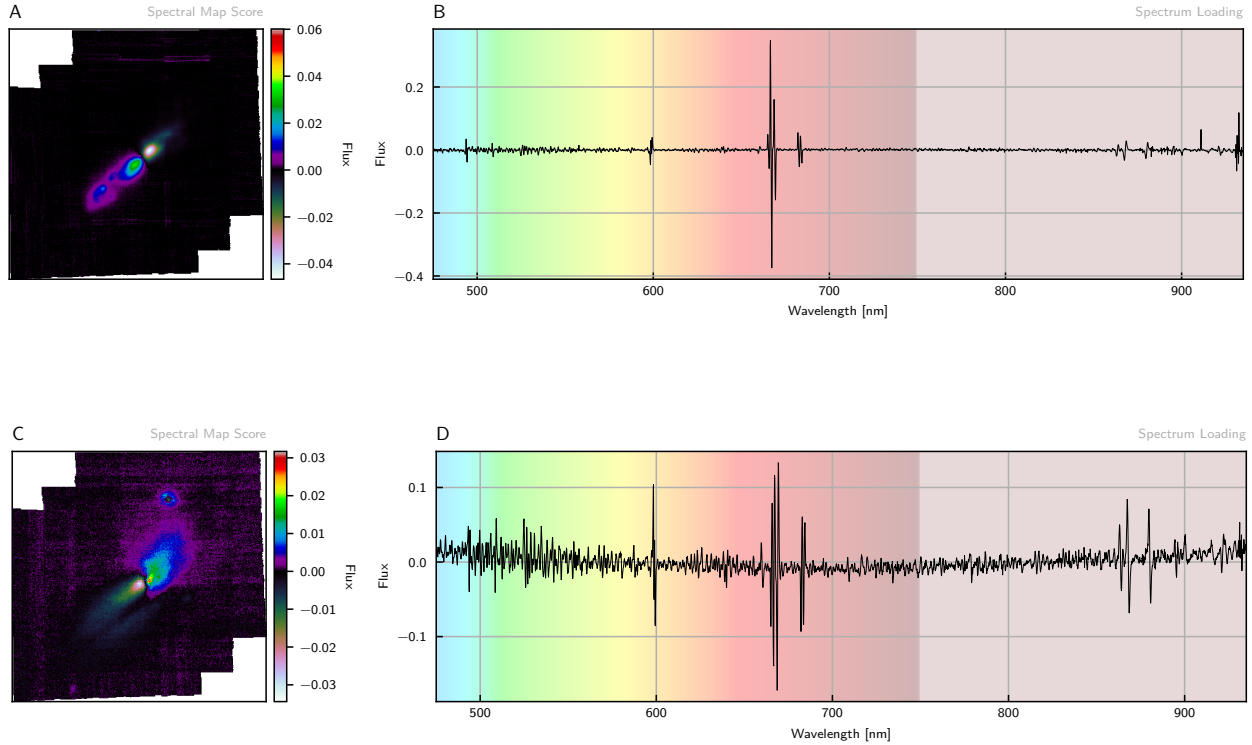


Figure 9.15 | Scores and loadings for the seventh component of PCA (A, B) and fifth component PCP (C, D) for the spiral galaxy AM2210-262. Flux is given in normalized units.

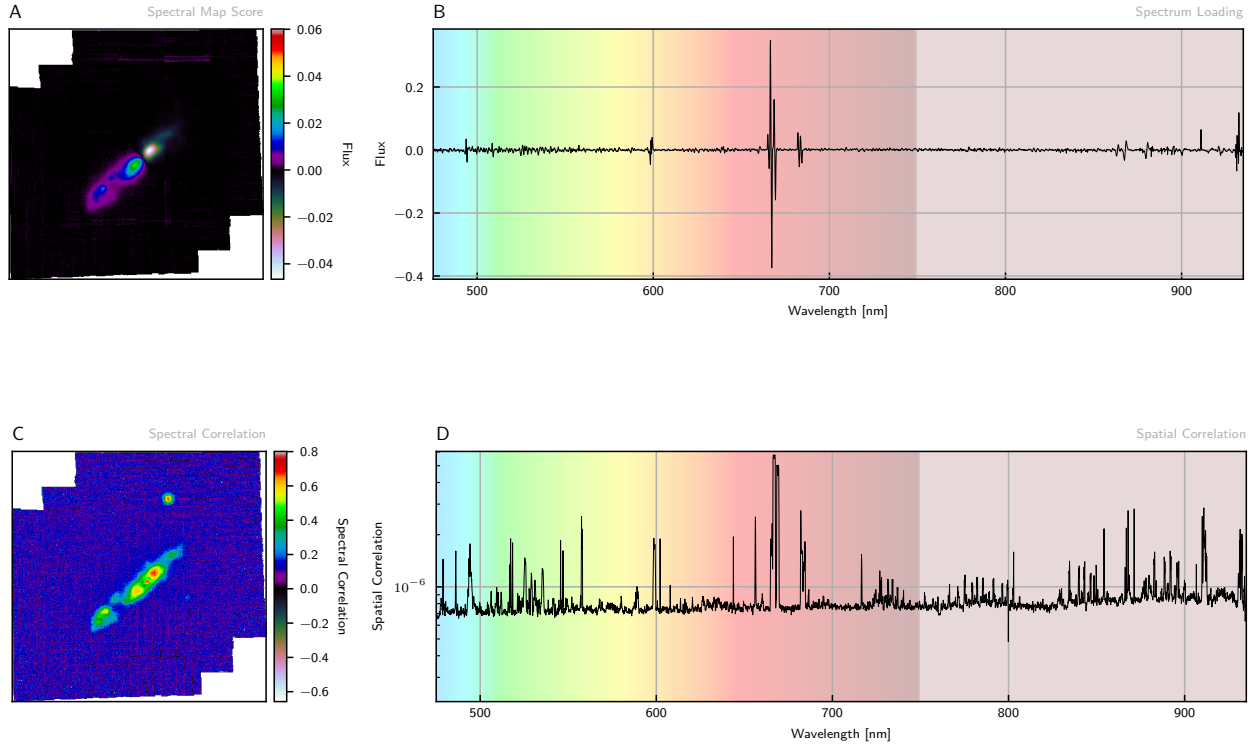


Figure 9.16 | Scores and loadings for the seventh component of PCA (A, B) and spectral and spatial correlation of residual term of PCP (C, D) for the spiral galaxy AM2210-262. Flux is given in normalized units.

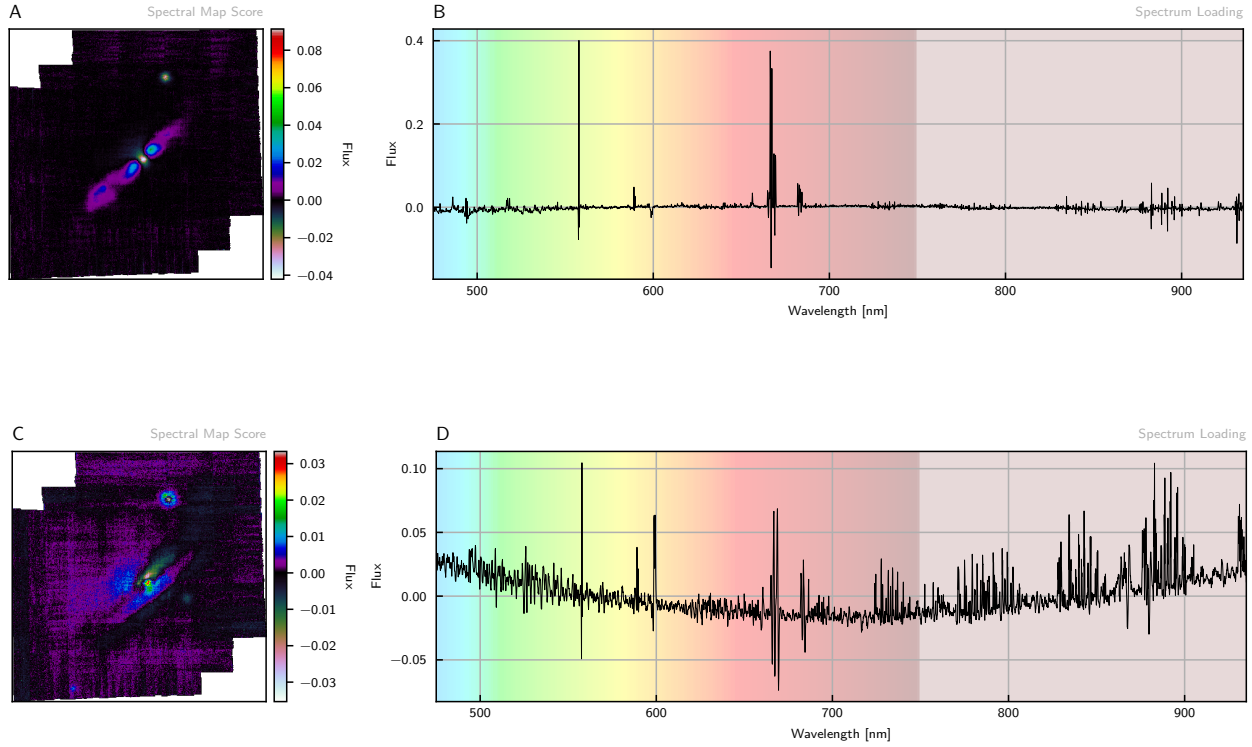


Figure 9.17 | Scores and loadings for the tenth component of PCA (A, B) and fifth component PCP (C, D) for the spiral galaxy AM2210-262. Flux is given in normalized units.

9.5.3 Discussion

In this case study, we reviewed the use of PCP and PCA for exploratory purposes. For the spiral galaxy AM2210-262. We observed that the approach of PCP in the first two components is similar to PCA, but that both methods diverge as singular values become smaller. PCA, specifically using the ℓ_2 -norm, captures energetic line-emissions and corresponding observable galaxy kinematics in further components. This is not the case for PCP, with its ℓ_1 -norm use. The retrieved components for PCP present entangled physical phenomena. For this case study and its data, PCA seems to be more appropriate.

9.6 Case Study D: Continuum and Line Emission Extraction with Improved SNR

In this case study, we research the use of PCP for the extraction of the continuum spectrum of individual galaxies in combination with obtaining line emissions with increased SNR. These line emissions are the result of local regions with (strongly) increased star formation and can be used to derive velocities between receding and approaching parts of the galaxy. From that information along with the galaxy distance, estimated from its redshift, the physical distance between those receding and approaching parts can be calculated and from that the mass of the galaxy. Hence, observing them with higher SNR increases the reliability of such measures. The use of PCP in this case study is similar to the detection of exoplanets for high contrast imaging¹⁴, which focuses on the residual term instead of the low-rank term. Concurrent to the extraction of continuum, we identify line emissions through simple spatial and spectral correlation measures. Nevertheless, dedicated line emission extraction tools such as LSDcat¹³ can be used.

9.6.1 Additional Experimental Design

In this section, we provide additional experimental design settings.

Preprocessing No preprocessing was performed after the data reduction.

Method We use PCP with randomized SVD, with QR iteration and including recycling step. The number of iterations was set to 2 and the number of oversamples to 10. The λ was set to $0.6\lambda_{\text{candes}}$, through a “maximal compression” setting of our developed framework.

Assumptions The main assumption in this study is that line emissions are spatially as well as spectrally sparse. Furthermore, for detection purposes we as-

sume that these line emissions correlated on short wavelength and small spatial distances.

9.6.2 Results

In this section, we present our results. We first show the white light images of A , i.e. the retrieved low-rank matrix, and B , i.e. the retrieved residual term. A will contain our continuum spectra associated to grid-like structures and object, i.e. galaxy, while B will contain sparse instrumental and pipeline artifacts, but also the extracted line emissions. Secondly, we show the working of our simple detector for the identification of line emissions and finally discuss the improvement of SNR with respect to the reduced data.

White Light Images and Mean Spectra The white light images and mean spectra associated to the retrieved matrices A and B from PCP are depicted in **Fig. 9.18**. Some line emissions are annotated for an heuristically matched redshift $z = 0.153$. From the residual term, we observe that the “structural” part of the OH forest residuals are already drastically reduced, making the line emissions and other features more apparent. A second observation is the gradient in the “background” of the white light image of B , probably caused by sparse instrumental artifacts.

Spectral Correlation To detect the line emissions, we applied two separate auto-correlation measures in respectively the spatial and spectral dimension. For the detection of even fainter features a combination of these measures can be used, such as implemented in LSDcat¹³. Our methods provide a spatial and spectral view, as depicted in **Fig. 9.19**. In the spatial view we obtain indications of possible line-emitting regions, while in the spectral view possible indications of wavelengths associated to such emissions can be seen. As a side observation, we see that the background correlates on short spectral scales, but less on long scales. This is one of the properties that ZAP³ exploits when dividing the spectra in different parts for the removal of OH forest residuals.

Line Emissions Having extracted the line emissions and identified their wavelengths, we can visualize them by plotting these particular residual monochromatic layers in B . The result is depicted in **Fig. 9.20**. We observe the $H\alpha$ at around 755 nm, the strongest detectable emission line in most spiral and irregular galaxies. It is therefore often used as a reliable quantitative trace of massive star formation¹⁵. We believe to have observed the following line emissions:

- $H\gamma$ at 500.402 nm
- $H\beta$ at 560.508 nm
- [OIII] at 570.620 nm

- [OIII] at 577.284 nm
- [NII] at 755.200 nm
- H α at 756.691 nm
- S2 at 774.355 nm

SNR Improvement The SNR calculations for the line emissions found by our simple detector are given in **Fig. 9.21**. Comparing their spatial SNR with that of reduced data, shows that our method increases the SNR on median by a factor 1.98. In the worst case, an improvement of 1.01 is observed. The largest improvement is a 4 fold increase in SNR. For these calculations we used spatial correlation and a simple threshold value to define the spaxels associated to line emissions and galaxy continuum. To increase the reliability of our study, we changed the threshold for both values to obtain a total of 400 SNR measurements per feature. The SNR is calculated by

$$\text{SNR} = \frac{\mu_{\text{signal}}}{\sigma_{\text{background}}}, \quad (9.5)$$

where μ represents the mean over all spaxels in the monochromatic layer associated to signal, and σ is the standard deviation associated to all spaxels in the monochromatic layer associated to background.

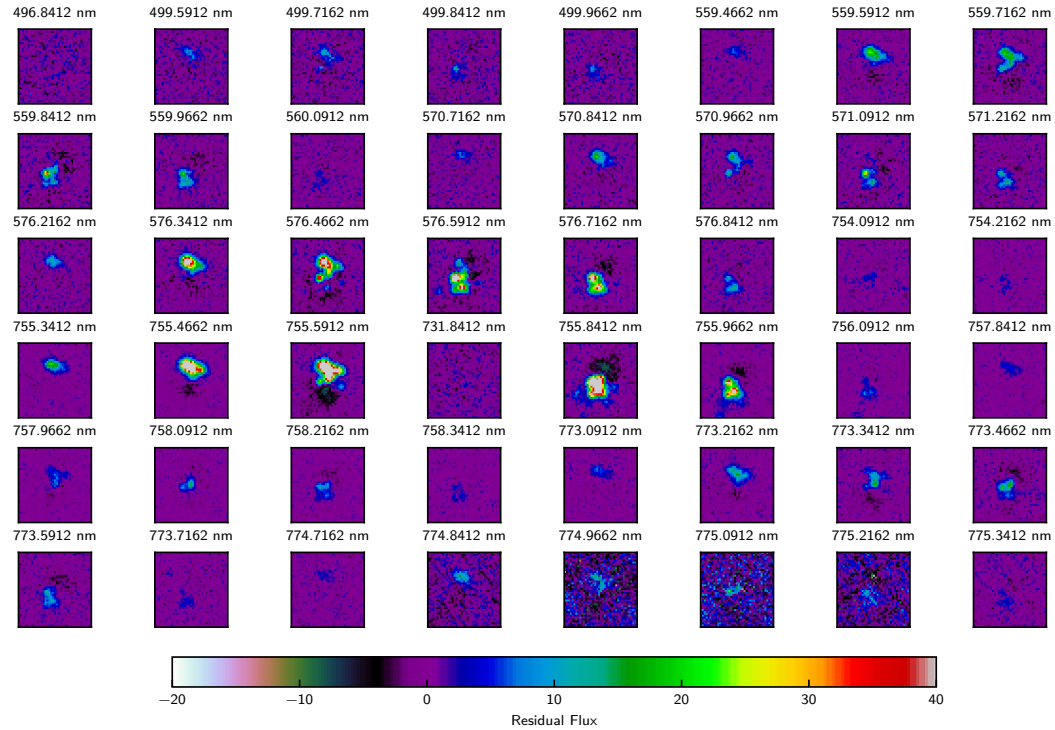


Figure 9.20 | Visualization of line emissions in the residual term B, they consists of area of under and over subtraction with respect to the continuum spectrum of the underlying galaxy. Flux is given in units of 10^{-20} erg/s/cm²/Angstrom.

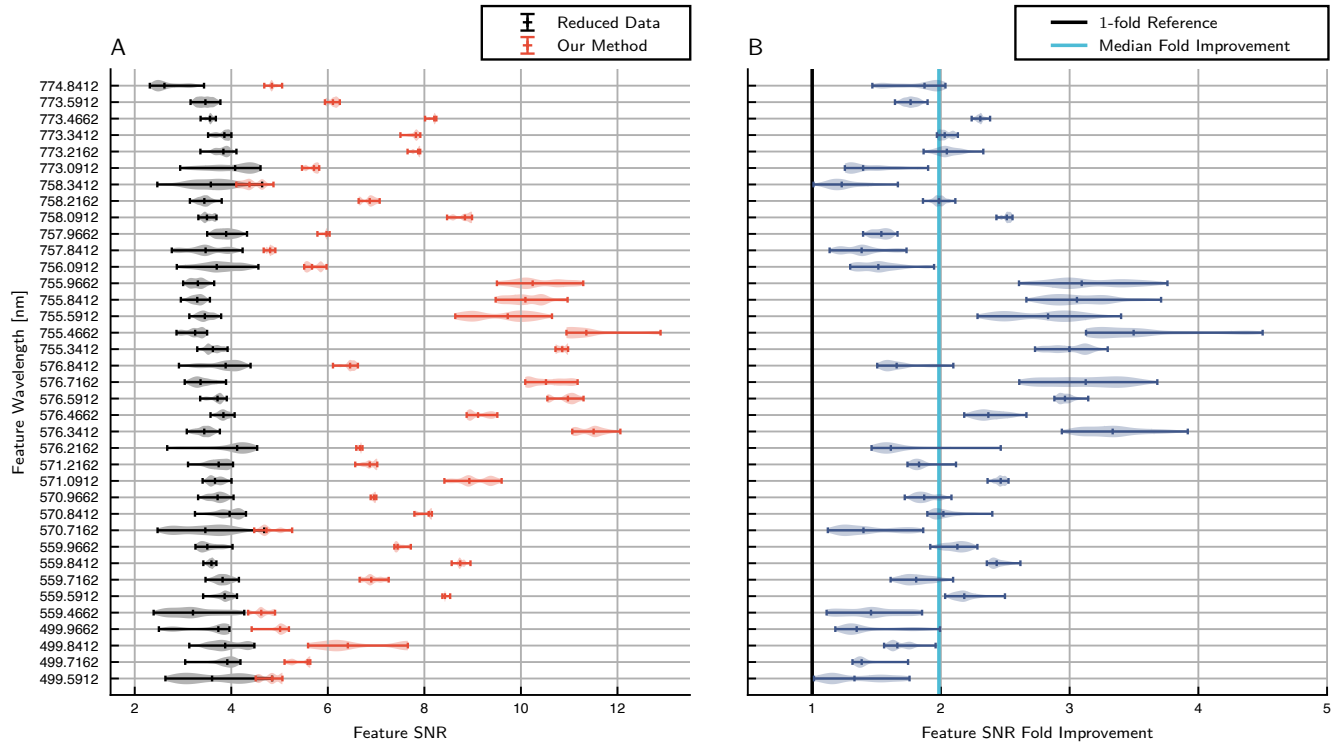


Figure 9.21 | Spatial signal-to-noise ratio measurement of all features as presented in **Fig. 9.20**. In subplot (A) both SNR on the reduced data and the B term (our method) are given. In subplot (B) we present the relative improvement with respect to the reduced data in terms of SNR.

9.6.3 Discussion

In this chapter, we have shown the use of PCP for the extraction of continuum and line emissions, both obtained simultaneously. Secondly, we showed the usefulness of simple correlation measures to spatially as well as spectrally indicate regions of possible line emissions. Finally, we proved that our method provides an improvement in signal-to-noise ratio 2-fold on average (median) and in best case up to 4-fold, with no loss of SNR in worst case. Finally, our method is preferred to running median approaches as no widths or cut-off frequencies need to be predefined, leading to less flux loss in continuum and increasing the SNR of the recovered line emissions.

9.7 Conclusions

In this chapter, we explored the versatility and applicability of our developed strategies on different levels ranging from low-level reduction, to high-level data exploration and visualization, and domain-specific applications.

In the first case study, we presented the use of PCP for co-adding purposes and showed its ability to reduce instrumental noise and sky background residuals. This is valuable with an eye towards the detection and quantization of faint continuum spectra associated to detected line emissions of line-emitting galaxies. It is also valuable for the detection of line-emitting galaxies that appear overlapping with OH forest and possible related residuals after subtraction. Finally, our method shows the ability to achieve lower variability while using less exposures. When using all 54 available exposures a 1.2 – 2.4 fold reduction in variability is observed. In the second case study, we presented the use of PCP for the concurrent spectral and spatial identification of instrumental and pipeline artifacts and argue its advantage with respect to proposed running median solutions to reduce the introduction of new subtraction artifacts. We also provide failed strategies that we pursued to remove those artifacts from objects that were masked out. In the third case study, we presented the use of PCP for exploratory and visualization purposes on the well extended spiral galaxy AM2210-262. We observe that PCP and PCA provide similar results in first few components, but that PCA is more appropriate for this data set to identify galaxy kinematics. In a final case study, we have shown the use of PCP for the extraction of continuum and line emissions, both obtained simultaneously. We proved that our methods provide an improvement in signal-to-noise ratio of 2-fold on average (median) and in the best case up to 4-fold, with no loss of SNR in the worst case.

Hence, we have shown on different levels how the developed methodology can be applied to measurement matrices of the MUSE instrument picturing a spiral galaxy (AM2210-262), the Hubble Ultra Deep-Field (UDF-06), and different exposures of the Hubble Deep Field South (HDFS), for noise reduction purposes.

Bibliography

- [1] Bacon, R. *et al.* The muse second-generation vlt instrument. In *Ground-based and Airborne Instrumentation for Astronomy III*, vol. 7735, 773508 (International Society for Optics and Photonics, 2010).
- [2] Weilbacher, P. M. *et al.* The data processing pipeline for the muse instrument. *arXiv preprint arXiv:2006.08638* (2020).
- [3] Soto, K. T., Lilly, S. J., Bacon, R., Richard, J. & Conseil, S. Zap-enhanced pca sky subtraction for integral field spectroscopy. *Monthly Notices of the Royal Astronomical Society* **458**, 3210–3220 (2016).
- [4] Bro, R. & Smilde, A. K. Centering and scaling in component analysis. *Journal of Chemometrics* **17**, 16–33 (2003).
- [5] Lin, Z., Chen, M. & Ma, Y. The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. *arXiv preprint arXiv:1009.5055* (2010).
- [6] Savitzky, A. & Golay, M. J. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry* **36**, 1627–1639 (1964).
- [7] Virtanen, P. *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17**, 261–272 (2020).
- [8] Van Rossum, G. & Drake, F. L. *Python 3 Reference Manual* (CreateSpace, Scotts Valley, CA, 2009).
- [9] Herenz, E. C. *et al.* The muse-wide survey: A first catalogue of 831 emission line galaxies. *Astronomy & Astrophysics* **606**, A12 (2017).
- [10] Borisova, E. *et al.* Ubiquitous giant $\text{Ly}\alpha$ nebulae around the brightest quasars at $z \sim 3.5$ revealed with muse. *The Astrophysical Journal* **831**, 39 (2016).
- [11] Vidal, R., Ma, Y. & Sastry, S. Generalized principal component analysis (gpca). *IEEE transactions on pattern analysis and machine intelligence* **27**, 1945–1959 (2005).
- [12] Bacon, R. *et al.* The muse 3d view of the hubble deep field south. *Astronomy & Astrophysics* **575**, A75 (2015).
- [13] Herenz, E. C. & Wisotzki, L. Lsdcat: Detection and cataloguing of emission-line sources in integral-field spectroscopy datacubes. *Astronomy & Astrophysics* **602**, A111 (2017).
- [14] Gonzalez, C. G. *et al.* Low-rank plus sparse decomposition for exoplanet detection in direct-imaging adi sequences-the llsg algorithm. *Astronomy & Astrophysics* **589**, A54 (2016).

- [15] Kennicutt Jr, R. C. The integrated spectra of nearby galaxies-general properties and emission-line spectra. *The Astrophysical Journal* **388**, 310–327 (1992).

Chapter 10

Conclusions

In this chapter, we review the conclusions of each research chapter and formulate the answers on the research questions and subquestions. First, we conclude the chapters related to the developed methodology, i.e. chapters 5, 6 and 7. Second, we provide the conclusions for the case study chapters, i.e. chapters 8 and 9, and include a revision of the conclusions of our literature review. Thereafter, we provide our specific contributions and recommendations for future work. Finally, we summarize our experiences on the collaborative work presented throughout this thesis as track record for future collaborations.

10.1 Methodology

In this section, we present our conclusions related to chapters 5, 6, and 7. First, we formulate a summarizing answer to research question A, where-after the subquestions related to research question A are reviewed.

A. Can matrix decomposition methods based on the ℓ_1 -robustness principles efficiently reduce massively multivariate and high-dimensional spectral imaging measurement matrices to a lower-dimensional representation while incurring minimal information loss?

We have demonstrated that such methodology is capable of running efficiently in time complexity $\mathcal{O}(mnr)$, both utilizing CPU and GPU. Secondly, we have shown that such methodology is capable of scaling to massively multivariate and high-dimensional matrices, practically up to the range of >1 terabyte memory requiring matrices, via an out-of-memory implementation. Thirdly, we have shown that such methodology is capable of handling noise cases often associated to spectral imaging. And finally, we have demonstrated that it is capable,

via a novel application-driven framework, to obtain a lower-dimensional representation while incurring minimal information loss. Hence, we can affirmatively answer to research question A.

Chapter 5: A Comparative Study under Inexact Recovery Conditions

A.1. How do these methods perform with respect to classical dimensionality reduction methods under inexact recovery conditions: (a) quantitatively, i.e. in terms of efficiency, compression ratio and information loss; (b) qualitatively, i.e. what break down phenomena appear in the atoms, dictionaries, and matrix terms?

In this chapter we reviewed a selection of robust methods under inexact recovery conditions, i.e. with respect to different noise types, and compared them to our yardstick, PCA. We compared PCP and SPCP to PCA on a total of 20 noise cases: 5 Gaussian, 5 Gaussian with outliers, 5 speckle, and 5 speckle with outliers.

The recovered rank suggests that both PCP and SPCP might benefit from an improved λ setting, as the rank is consistently overestimated in most non-outlier noise scenarios. Secondly, the low-rank relative offset suggests that SPCP could further benefit from an adjusted δ setting. The latter is manifesting itself through a flattening trend in performance towards smaller intensity noise cases. On the other hand, the incoherence results suggest that the singular vectors of PCP and SPCP are consistently incoherent. The latter probably positively affects the ability to compress measurement matrices efficiently. The same incoherence results revealed a consistent outlying value for PCP in an intermediate intensity case of the Gaussian noise with outliers. Whether this is due to a slight overestimation in rank or noise related phenomenon, is not clear. Nevertheless, in contrast to PCA, the spread on the incoherence values is small for both PCP and SPCP. Finally, the running time shows that SPCP and PCP are several factors slower than PCA. This can be partially alleviated by setting less restrictive stopping criteria, by using a warm start, or using a randomized SVD.

The qualitative performance analysis visually confirms the flattening trend of SPCP for the cases of low intensity Gaussian noise with outliers. It also suggests that PCP and SPCP improve the retrieval of low SNR patterns, but do fail under high intensity noise. Finally, the absolute cosine similarity suggests that the left singular vectors retrieved by PCP and SPCP are highly affected by high intensity speckle noise, while lowering this intensity increases their ability to align their left singular vectors with the ground truth.

What does it enable? Firstly, it enables a deeper understanding of the working principles of these types of methods in a practical setting. Secondly, it shows

that the Hadamard matrix provides a manner to disjointly study the effects of the incoherence principles and noise. Finally, it enables further research on real spectral measurement matrices, as it proves the viability for noise types often related to imaging applications.

Chapter 6: An Application-Driven Framework for Parameter Optimization

A.2. Can a framework be built to efficiently tune the regularization parameters of the above-mentioned methods, but customized for a specific application?

We reviewed the use of several ground truth and heuristic information loss measures and suggested other measures that can be used concomitantly. In second instance, we proposed a framework for parameter optimization along with a grid search as solution to our framework's non-linear optimization problem. The associated computational cost is estimated to be $\mathcal{O}(qmn^2/qmn)$ flops, where q is the number of samples. We have shown its use through different case studies and analyzed both the parameter setting for those and their performance with respect to five ground truth information loss measures. The latter was performed on a total of 20 noise cases: 5 Gaussian, 5 Gaussian with outliers, 5 speckle, and 5 speckle with outliers. From this analysis we conclude that the measures are not only application, but also noise type dependent. Furthermore, we have shown that certain measures can be sensitive to the noise. This has to be taken into account when using the framework. We conclude with the fact that this framework is a first step towards a truly data-driven (and thus automated) parameter setting.

What does it enable? Our framework enables the setting of hyperparameters by domain experts, e.g. analytical chemists or astronomers, using domain-specific knowledge in an application-optimized fashion. Hence, it enables achieving improved results for domain experts by providing them a framework to translate their preferred measures and weights into an optimization.

Chapter 7: Acceleration and Scaling of PCP

A.3. Can a memory and compute resources restricted version of such approaches be developed, and how does it perform with respect to the original batch version that assumes the entire data set can be held in memory?

We have reviewed the theoretical computational cost for IALM and provided different improvements with respect to the standard implementation. Additionally, we have addressed these improvements, time- as well as memory-related. In a second phase, we have practically implemented these solutions and compared them over different settings and criteria.

For acceleration purposes, we provided results for different SVD methods in both 32 and 64 bit precision. We shortly reviewed a randomized approach and analyzed the different iterators, along with proposing a simplified recycling step. The recycling step is heuristically found to be contributing to the iterative stability while leaving the time cost unchanged and remaining simple to implement. To conclude the acceleration section, we advise a 32 bit precision randomized implementation with recycling step and appropriate iterator, depending on the singular value spectrum. For smaller matrices a truncated SVD in the form of ARPACK or PROPACK is also possible, but one has to closely monitor the convergence process for those methods. For larger matrices, GPU-based approaches based on randomized SVD methods show to be promising.

For scaling purposes, we have assessed different truncated and randomized methods and provided a comparison with respect to the standard in-memory GESVD. We conclude that scaling is possible with both ARPACK and randomized methods, but that the latter are preferred as they easily allow to trade off speed and accuracy. It was found that the ρ and μ parameter are of importance to safeguard convergence for those cases. For out-of-memory implementations, enabling scaling to the terabyte range, it was shown that randomized methods are a viable solution. Our analysis indicates that such implementations have reasonable overhead when used in-memory with respect to the in-memory LAPACK GESVD and thus can be considered reasonably optimally implemented.

What does it enable? This chapter contributed to the practical use of ℓ_1 -robustness methodology in real-life applications by proposing PCP implementations enabling analysis on measurement matrices in reasonable time, achieving equal or even better low-rank relative offsets. This ultimately leads to overall reduced time and computational costs for any project using PCP in its pipeline. Furthermore, we have shown that out-of-memory implementations are viable and perform equally well with respect to low-rank relative offset and running time in comparison to in-memory versions, but are able to perform in compute resource restricted environments and scale to terabyte memory requiring matrices.

HannibalSpecter: Custom Python Library

We implemented, following PEP-8 standards, an installable object-oriented based Python library `HannibalSpecter`, providing 8 different matrix decomposition algorithms optimized for CPU and GPU-based use, built up of 31 files comprising 3619 lines of code and 4564 lines of comments. Furthermore, it en-

compasses different packages that provide functionalities to unify the handling of IMS and MUSE data, providing corresponding templates and preprocessing methods to offer application-specific operations, provide quality measures, generate synthetic measurement matrices, including shape-based and randomized generation methods, a package containing optimization stopping criteria that can interchangeably be used by the algorithms, and a package for plotting figures as provided in this thesis. Finally, we have developed a package enclosing 14 distinct SVD methods. Unit testing has been put in place, however, it still needs to be extended to all modules and functionalities. Currently, the library is available upon request.

What does it enable? The `HannibalSpecter` custom Python library enables further structured research towards robust matrix decomposition methods through its object-oriented based implementation in Python. Secondly, these efforts are also made to promote transparency for this research, enabling easy comparison and analysis. Finally, the package ensures continuity for future BSc, MSc, and PhD students following up on this work, but also researchers interested in applying the same methodology in their own application domains. See Appendix B for a full layout of the package.

10.2 Application Domains

In this section, we present our conclusions related to chapters 8 and 9. First, we formulate a summarizing answer to research question B, where-after the sub-questions related to research question B are reviewed, along with the conclusions of our literature review.

B. How can the application-driven methodology we have developed, be used to address application domain-specific challenges?

For imaging mass spectrometry, we show that our application-driven methodology can be used for the efficient setting of application-driven parameters for MALDI-TOF measurement matrices, improving on the predefined λ_{candes} parameter, in compression as well as in retrieved quality of the corresponding spatial components. Secondly, we show that our methodology can be used for improving over the compression abilities of PCA by 1.73-fold, while obtaining a similar spatial reconstruction score. The latter coincides with the optimal λ -setting defined by our framework. Thirdly, we have shown that our implementations scale linearly with measurement matrix size as theoretically predicted by $\mathcal{O}(mnr)$. Finally, we have demonstrated the processing of a MALDI timsTOF measurement matrix requiring > 1 TB of processing memory, along with its denoising capabilities.

For integral field spectroscopy, we show that our application-driven methodology can be used for co-adding purposes, reducing instrumental noise and sky background residuals, and further decreasing the variability. Secondly, we show that our methodology can be used for the concurrent spectral and spatial identification of instrumental and pipeline artifacts and subtracting of those for improving detectability of faint line-emitting galaxies and other faint object. Finally, we have shown that our developed methodology can be used to increase the spatial signal-to-noise ratio and thus detectability of line emissions.

Literature Review

In the literature review, we introduced the basics of matrix decompositions from the perspective of atoms. We explained the concept of ℓ_1 -robustness with respect to sparse outliers and introduced the family of extended linear mixture models together with its validity for imaging applications. Furthermore, 18 different classical and robust matrix decomposition methods were analyzed in the framework of atoms, along with the presentation of some of their applications in an imaging context, in particular in the fields of imaging mass spectrometry and astronomical imaging.

The novelty of the literature review lies in the inclusion of the framework of atoms into the broader context of matrix decompositions for analysis and comparison of methods. A second accomplishment was the discovery of promising classical methods that have an application legacy in only one of the considered imaging modalities, but not in both. This knowledge can be transferred to establish some mutually beneficial cross-talk between these distinct fields of imaging. Thirdly, the review of ℓ_1 -robust matrix decomposition methodology has indicated a substantial gap between method literature itself and application-driven literature from both fields. This knowledge gap can and should be closed, and we hope that this thesis helps contribute to that interdisciplinary effort. Hence, we concluded this review by focusing on two possible research directions.

Case Studies I: Imaging Mass Spectrometry

B.1. In molecular imaging, and specifically imaging mass spectrometry high-dimensionality and *massive data set footprints* are the most pressing bottleneck. How can we apply the developed methodology to measurement matrices of MALDI TOF and MALDI timsTOF nature, respectively of mouse kidney and mouse pup origin, specifically for *dimensionality reduction and compression purposes*?

In this chapter, we presented four different case studies related to the use of our developed methodology in a molecular imaging context, for use on MALDI

(tims)TOF-type IMS measurement matrices and with a focus on dimensionality reduction and compression.

In a first case study, we reviewed the use of our developed framework for the application-driven hyperparameter setting of PCP for MALDI-TOF IMS data. We introduced three applications and demonstrated some of their quantitative as well as qualitative behavior. We demonstrated its improvement over the common, predefined, and wide-spread λ_{candes} parameter, in compression as well as in retrieved quality of the corresponding spatial components in U . Secondly, we have shown that the three defined application cases result in distinct solutions with specific favorable properties. In a second case study, we have compared PCP to PCA on its ability to compress MALDI-TOF IMS data. We have shown that a 1.73-fold improvement in compression rate can be obtained with PCP in comparison to PCA, while obtaining a similar spatial reconstruction score. This results in a compression ratio of 19.03 for PCP instead of 10.98 for PCA. Concurrently, this observation coincides with our framework's optimal λ for a maximal compression application. In the third case study, we reviewed the in-memory scalability of our methodology in combination with our framework's hyperparameter optimization for the trade-off case. We have shown that our implementations scale linearly with measurement matrix size as theoretically predicted by $\mathcal{O}(mnr)$. We have shown that a fixed λ -setting gives similar results when the measurement matrix is scaled. And, we have shown for our largest in-memory measurement matrix the ability to compress its corresponding 30 GB data set size into a version requiring less than 30 MB on disk, while retaining spatial and spectral reconstruction scores. Finally, in the fourth case study, we presented the proof-of-concept results of a MALDI timsTOF measurement matrix normally requiring > 1 TB of memory for processing via PCP, which as far as we know is a first. We gradually scaled up the measurement matrix to a size of $164,808 \times 568,142$ (otherwise requiring around 1.38 TB of memory). Furthermore, we have shown its denoising capabilities. Finally, we have discussed the current bottleneck for further scaling of the methodology related to disk read/write time.

Hence, we have shown that for molecular imaging, and specifically imaging mass spectrometry with its high-dimensionality and massive data sets, our developed methodology can be applied for dimensionality reduction and compression purposes, while also providing measures to detect any leakage of signals in the process.

What does it enable? For imaging mass spectrometry, these developments will be essential in making full ion mobility (IM) IMS, a powerful generalization of MALDI-TOF IMS that collects two spectral dimensions per pixel (or 4-mode data tensors), and that can yield petabyte measurements using standard techniques, a practical reality.

Case Studies II: Integral-Field Spectroscopy

B.2. In astronomical imaging, and specifically integral field spectroscopy *low signal-to-noise ratios and instrumental artifacts* are the most pressing bottleneck. How can we apply the developed methodology to measurement matrices of the MUSE instrument picturing a spiral galaxy (AM2210-262), the Hubble Ultra Deep-Field (UDF-06), and different exposures of the Hubble Deep Field South (HDFS), *for noise reduction purposes?*

In this chapter, we explored the versatility and applicability of our developed strategies on different levels ranging from low-level reduction, to high-level data exploration and visualization, and domain-specific applications.

In the first case study, we presented the use of PCP for co-adding purposes and showed its ability to reduce instrumental noise and sky background residuals. This is valuable with an eye towards the detection and quantization of faint continuum spectra associated to detected line emissions of line-emitting galaxies. It is also valuable for the detection of line-emitting galaxies that appear overlapping with OH forest and possible related residuals after subtraction. Finally, our method shows the ability to achieve lower variability while using less exposures. When using all 54 available exposures a 1.2 – 2.4 fold reduction in variability is observed. In the second case study, we presented the use of PCP for the concurrent spectral and spatial identification of instrumental and pipeline artifacts and argue its advantage with respect to proposed running median solutions to reduce the introduction of new subtraction artifacts. We also provide failed strategies that we pursued to remove those artifacts from objects that were masked out. In the third case study, we presented the use of PCP for exploratory and visualization purposes on the well extended spiral galaxy AM2210-262. We observe that PCP and PCA provide similar results in first few components, but that PCA is more appropriate for this data set to identify galaxy kinematics. In a final case study, we have shown the use of PCP for the extraction of continuum and line emissions, both obtained simultaneously. We proved that our methods provide an improvement in signal-to-noise ratio of 2-fold on average (median) and in the best case up to 4-fold, with no loss of SNR in the worst case.

Hence, we have shown on different levels how the developed methodology can be applied to measurement matrices of the MUSE instrument picturing a spiral galaxy (AM2210-262), the Hubble Ultra Deep-Field (UDF-06), and different exposures of the Hubble Deep Field South (HDFS), for noise reduction purposes.

What does it enable? For integral field spectroscopy, these developments will be essential in further enabling the investigation of faint and complex astrophysi-

cal signatures, one of the scientific goals of the MUSE instrument¹. Secondly, our methodology will enable achieving higher SNR values with less exposures, subsequently increasing observation time for other studies. Finally, the presented achievable compression ratios are promising for future reduction pipelines and the archiving of scientific data from novel instruments, e.g. the Extremely Large Telescope (ELT)².

10.3 Contributions

Chapter 4:

- Inclusion of the framework of atoms into the broader context of matrix decompositions for analysis and comparison of methods

Chapter 5:

- Object-oriented Python implementations of several methods
- Shown the use of robust methods under inexact recovery conditions
- Suggested the use of the Hadamard matrices for fixing incoherence
- Proposed different quantitative and qualitative measures
- Recognized shortcomings and identified breakdown phenomena

Chapter 6:

- Showcased the use of a framework to define hyperparameter optimality
- Suggested the use of different ground truth and heuristic information loss measures
- Implemented several of those measures and defined different case studies
- Shown some of the vulnerabilities of this approach

Chapter 7:

- Theoretically and practically analyzed IALM and its SVD methods
- Suggested a novel, simplified type of recycling
- Provided an overview of available SVD methods in Python
- Implemented different SVD methods in `HannibalSpecter`
- Fixed and improved the official out-of-memory SVD implementation of the Python toolbox `dask`, and committed this fix back to the open source project^a

^a<https://github.com/dask/dask/issues/6757>

- Implemented different out-of-memory implementations able to scale to the terabyte range

Chapter 8:

- The efficient setting of an application-driven λ -parameter, outperforming the common, predefined, and wide-spread λ_{candes} value
- Demonstrated the ability to improve over PCA's compression ratio
- Demonstrated the ability to process increasingly large in-memory MALDI measurement matrices
- The robust matrix decomposition of a MALDI timsTOF measurement matrix requiring > 1 TB memory (a first proof-of-concept)

Chapter 9:

- Proposal of PCP for co-adding purposes with improved SNR and reduced variability
- Artifact identification and partial removal via PCP and several attempts to directly reduce these artifacts in object spectra
- Proposal of PCP for continuum and line emission extraction
- Identification and annotation of several line emissions of spiral galaxy in UDF-06

10.4 Recommendations

Future research could include:

Chapter 4: A literature study on the use of (robust) matrix decomposition methods in astronomy-based applications.

Chapter 5: The expansion towards fixed noise, variable incoherence case studies and their influence on the optimal λ -setting.

Chapter 6: A review of convex measures to reduce complexity and improve convergence through convex optimization techniques, e.g. as a non-negative weighted sums of convex measures.

Chapter 7: Compare the efficiency of dask with respect to ScaLAPACK and online methods.

Chapter 8: Explore the possibilities of scaling dask on GPU-based systems.

Chapter 9: Research the artifact reduction in masked objects via a custom method optimization.

10.5 Collaborative Perspective

We started with a collaboration with 6 instances. The basis collaboration was between Dr. Ing. Van de Plas of the Delft Center for Systems and Control (DCSC) and Prof. Dr. Brandl of the Faculty of Aerospace Engineering through my double degree. Secondly, we collaborated through Dr. Ing Van de Plas with the Mass Spectrometry Research Center, Vanderbilt University (TN, US) and through a research proposal with the Center for Nanophase Materials Sciences (CNMS) at the U.S. Department of Energy's Oak Ridge National Lab (TN, US). The proposal and its review are given in Appendix C. These institutions provided respectively the imaging mass spectrometry data and technical/chemical input (during the whole thesis), and the computational power and computational input required for the initial phase of this research (my internship). Via Prof. Dr. Brandl, we collaborated with the Leibniz Institute (AIP), Potsdam (GER). They provided the MUSE data, support on the data reduction pipeline, and astronomical input. In first instance, we also collaborated with Rijksuniversiteit Groningen for Omega-CAM data. However, the latter was unsuccessful, but open to be revisited. We believe that the in-person visits to Vanderbilt, Oak Ridge National Laboratory and AIP made a lot of difference in terms of getting up to speed on the different aspects of the various disciplines involved.

10.5.1 Timeline

This timeline briefly summarizes the progress of the thesis: from first contact with supervisors to the final defense.

2017 August: first brainstorm meeting with supervisors Prof. Dr. Brandl and Dr. Ing. Van de Plas

2017 September: start of individual double degree

2019 June: first contact with Leibniz Institute for Astrophysics Potsdam (MUSE) and Rijksuniversiteit Groningen (OmegaCAM)

2019 August 11 – 13: in-person visit to Leibniz Institute for Astrophysics Potsdam (AIP)

2019 August 15: first Skype meeting with Rijksuniversiteit Groningen

2019 September - November: internship at Mass Spectrometry Research Center, Vanderbilt University (TN, US) and five week research visit to Oak Ridge National Lab (TN, US)

2020 February 1: start of graduation thesis

2020 June 5: progress meeting - Oak Ridge National Lab

2020 **August 24**: progress meeting - AIP & Groningen University

2020 **November 25**: DCSC midterm

2021 **April 8**: green light meeting

2021 **April 30**: final defense

Bibliography

- [1] Bacon, R. *et al.* The muse second-generation vlt instrument. In *Ground-based and Airborne Instrumentation for Astronomy III*, vol. 7735, 773508 (International Society for Optics and Photonics, 2010).
- [2] Gilmozzi, R. & Spyromilio, J. The european extremely large telescope (e-elt). *The Messenger* **127**, 3 (2007).

Part III

Appendices

Appendix A

Additional Materials

A.1 A Comparative Study under Inexact Recovery Conditions

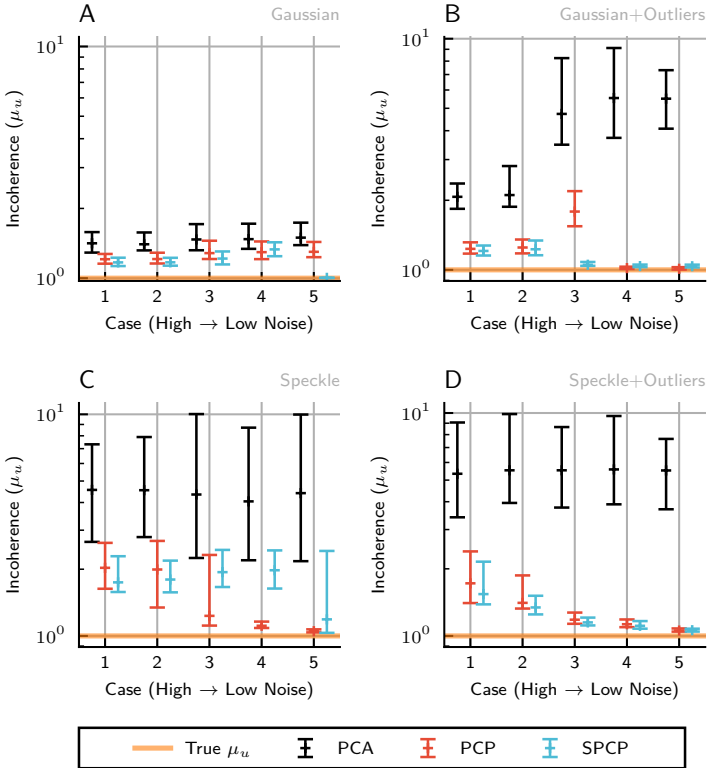


Figure A.1 | The figure depicts the μ_u incoherence for the cases: (A) Gaussian noise, (B) Gaussian noise with additional outliers, (C) Speckle noise, (D) Speckle noise with additional outliers. The + sign represents the mean of 100 simulations, and is given along with maximum and minimum values. The “True μ_u ” reflects the calculated incoherence as given in Eq. 5.7

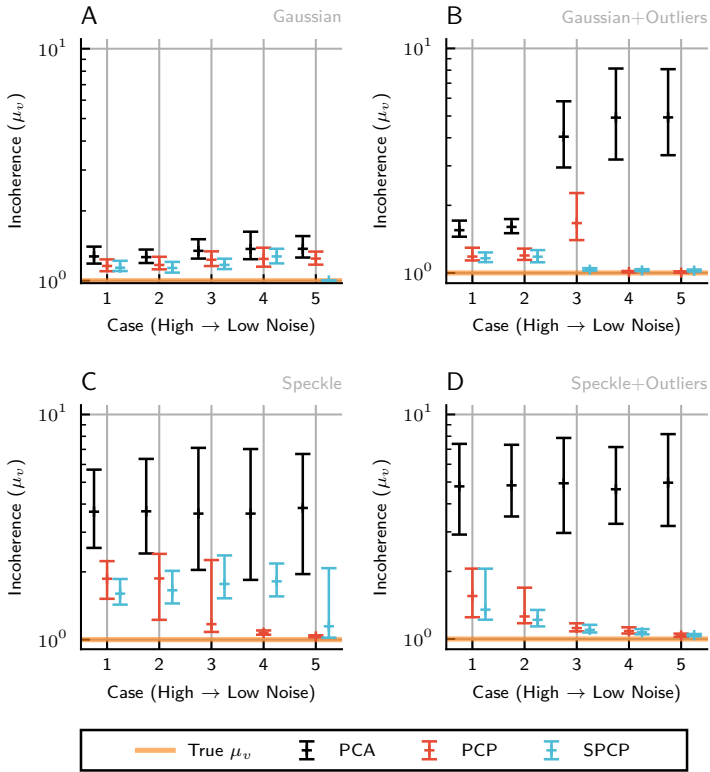


Figure A.2 | The figure depicts the μ_v incoherence for the cases: (A) Gaussian noise, (B) Gaussian noise with additional outliers, (C) Speckle noise, (D) Speckle noise with additional outliers. The + sign represents the mean of 100 simulations, and is given along with maximum and minimum values. The “True μ_v ” reflects the calculated incoherence as given in Eq. 5.7

A.2 An Application-Driven Framework for Parameter Optimization

Table A.1 | Explanation of measure values. This table includes the measure name, its type (GT: ground truth; HEU: heuristic), its natural range of possible values and an interpretation for different of those values, i.e. low, high and optimal.

Measure	Type	Range	Value	Interpretation
Rank (rank)	GT	$[0, \min(m, n)]$	Low	High compression, low variability
			High	Low compression, high variability
			Opt.	10, perfect match with the true low-rank matrix rank
Low-Rank Relative Offset (lr)	GT	$[0, \infty[$	Low	Good fit in ℓ_2 -sense of the low-rank matrix
			High	A value of 100% represents a recovered low-rank matrix with all zero entries
			Opt.	Perfect retrieval when 0%
Maximum Absolute Offset (mao)	GT	$[0, \infty[$	Low	Values close to 0 represent good entry-wise fits
			High	Entry-wise fit deteriorates for large values
			Opt.	0, in that case every entry is perfectly matched
Median ℓ_1 Absolute Offset (ml1)	GT	$[0, \infty[$	Low	Values close to 0 represent good column-wise fit in ℓ_1 -sense
			High	Column-wise fit deteriorates for large values
			Opt.	0, in that case every column is fitted perfectly. Note that this collides with an lr being equal to 0%

Table A.1 (bis) | Explanation of measure values. This table includes the measure name, its type (GT: ground truth; HEU: heuristic), its natural range of possible values and an interpretation for different of those values, i.e. low, high and optimal.

Measure	Type	Range	Value	Interpretation
Mean Absolute Cosine Similarity (macs)	GT	[0, 1]	Low	Values close to 0 represent perfectly matching direction-wise
			High	Values close to 1 represent orthogonal directions, hence no information overlap
			Opt.	0, represent a perfect match in direction
Structure in B	HEU	[-1, 1]	Low	Values close to -1 correspond to perfect dispersed patterns
			High	Values close to 1 correspond to perfect clustered patterns
			Opt.	0, such values correspond to randomness in patterns
Trace Residual	HEU] -∞, ∞[Low	Negative values reflect the addition of energy to the signal term
			High	reflect large amounts of overlap between signal and noise
			Opt.	0 reflect orthogonal spaces of signal and noise
Cardinality of B	HEU	[0, m.n]	Low	Only a few entries are considered noisy
			High	Almost all entries are considered as noisy
			Opt.	N.A.

A.3 Acceleration and Scaling of PCP

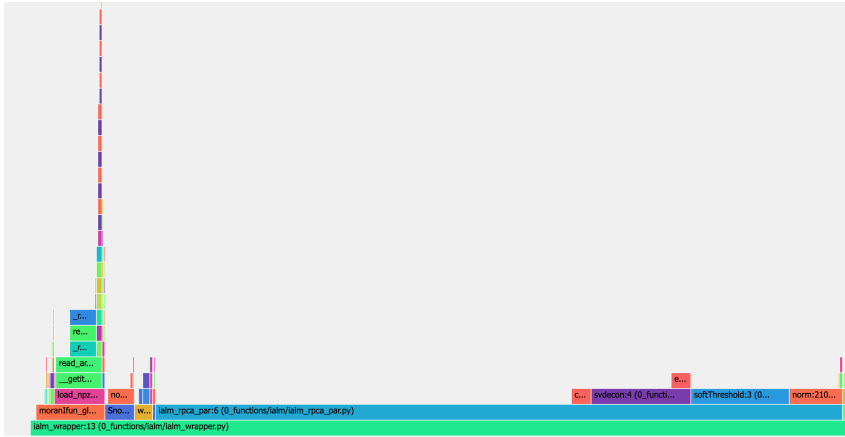


Figure A.3 | CPU profiling of principal component pursuit. Horizontally the relative computational time per function is given, while vertically the embedded functions are given. Hence, one obtains a breakdown on what function the CPU spends most of its time.

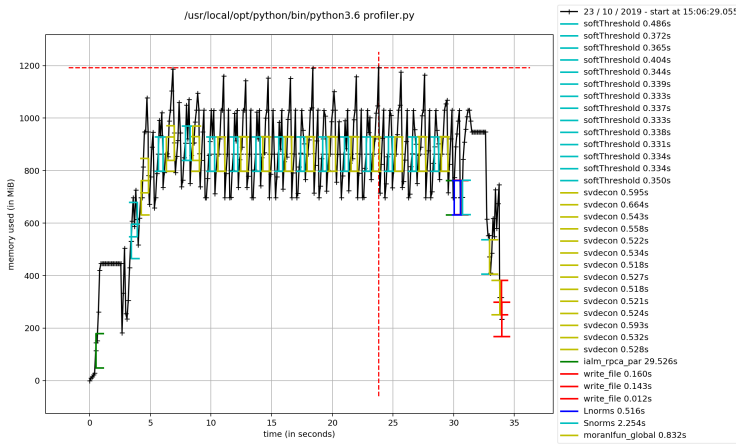


Figure A.4 | Memory profiling of principal component pursuit. The memory usage over time is depicted. The step increase in memory at early times corresponds to the initialization and is due to the creation of a Lagrangian multipliers matrix and residual noise matrices. Thereafter, one sees the different iterations and finally obtaining the converged result, freeing memory from superfluous helper matrices.

```

(base) [rogermoens@or-slurm-login01 script_output]$ cat run-job.error
Traceback (most recent call last):
  File "/home/rogermoens/RF_V0/run_306B.py", line 77, in <module>
    U, V, S, Lrank, Scard, k, crit1, crit2, Sigma = pcp32l(D, lambda_search, tol, maxiter, svd_opt)
  File "/home/rogermoens/RF_V0/functions/ialm/pcp32l.py", line 76, in pcp32l
    U, diagA, Vt = scipy_svd(D-E_hat+(1/Mu)*Y, lapack_driver='gesdd', full_matrices=False)
  File "/home/rogermoens/.local/lib/python3.7/site-packages/scipy/linalg/decomp_svd.py", line 125, in svd
    compute_uv=compute_uv, full_matrices=full_matrices)
  File "/home/rogermoens/.local/lib/python3.7/site-packages/scipy/linalg/lapack.py", line 771, in _compute_lwork
    raise ValueError("Too large work array required -- computation cannot ")
ValueError: Too large work array required -- computation cannot be performed with standard 32-bit LAPACK.
(base) [rogermoens@or-slurm-login01 script_output]$

```

Figure A.5 | Error for the Python GESDD wrapper when too large matrices are provided. This behavior is due to the creation of large additional matrices that are used in its process to calculate the singular values decomposition.

```

Loaded
2 1 0.7738487497859775 0.1575661164134874
2 1 0.47878998038270676 0.283122609035282
46 2 0.40271974873613925 0.15977263253432988
6 5 0.33437302792629503 0.12524670316307615
6 5 0.3161862776032277 0.11721985796591292
50 6 0.2952121034640337 0.08754839375060627
15 14 0.268819471237407 0.095565445508135
59 15 0.23647469250014538 0.10349880206761362
21 20 0.19320720343782902 0.11125442406003333
65 21 0.1437519010814624 0.10827341916316656
28 27 0.0957213579118069 0.09203506616135515
72 28 0.05375035929142282 0.07041179795904004
38 37 0.025200402619438404 0.04447792375583724
82 38 0.011494143483496451 0.022006784725898836
54 53 0.00570636134564752 0.009608580479795733
98 54 0.0032064677560552798 0.004608715769201562
78 77 0.001648132110096903 0.0025002775600763785
122 78 0.0006320856248029003 0.001412188977926934
104 103 0.000495563497563977 0.0006569943470428197
148 104 0.0003032412018030974 0.00039766885947752604
138 137 0.00019128722280105988 0.00023252054846479358
182 138 0.0001200726307231705 0.0001540104527830715
178 177 7.69304425424355e-05 9.243364610953033e-05
222 178 4.8484230321660976e-05 6.19331810348593e-05
184 183 3.154780403526505e-05 3.791401987249416e-05
170 169 2.6325477598727256e-05 2.4699439761958157e-05
158 157 2.7240840218819484e-05 2.6324806861580147e-05
151 150 2.58542056846178e-05 3.459746074638219e-05
147 146 2.2032498749561967e-05 3.429428415011722e-05
145 144 1.7161912830439692e-05 2.9676121860860197e-05
145 144 8.304522315167827e-06 2.313125733681421e-05
143 142 1.4107939861958135e-05 8.237539462594770e-06
143 142 8.44227343019417e-06 2.2227306493123615e-05
143 142 7.378755230048668e-07 8.44388930623667e-06

```

Figure A.6 | Break down phenomena observed for ARPACK svds. From top to bottom the iterations of IALM are given with ARPACK's svd. The first column represents the estimated rank, the second column is the actual calculated rank and the third and fourth column are the convergence criteria. One sees that at the 25th row from the top the rank starts decreasing.

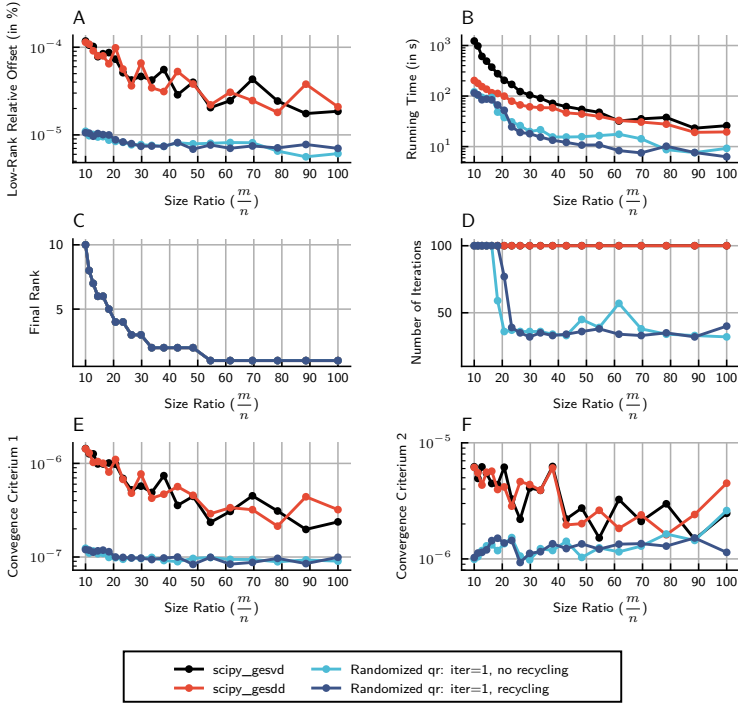


Figure A.7 | Comparison between standard GESVD, GESDD and randomized methods, with and without recycling for different size ratios. A fixed number of iterations 100 was set along with a matrix row size m of 10000 and rank ratio $1/100$, for $\lambda = 1\lambda_{\text{candes}}$. The subplots depict: (A) the low-rank relative offset; (B) the running time; (C) the final rank at convergence; (D) The number of iterations required for convergence; (E) the first convergence criterion, and (F) the second convergence criterion.




Figure A.8 | Comparison between standard GESVD, GESDD and randomized methods, with and without recycling for different size ratios. A fixed number of iterations 50 was set along with a matrix row size m of 10000 and rank ratio $1/100$, for $\lambda = 1\lambda_{\text{candes}}$. The subplots depict: (A) the low-rank relative offset, and (B) the running time.

A.4 Case Studies II: Integral-Field Spectroscopy

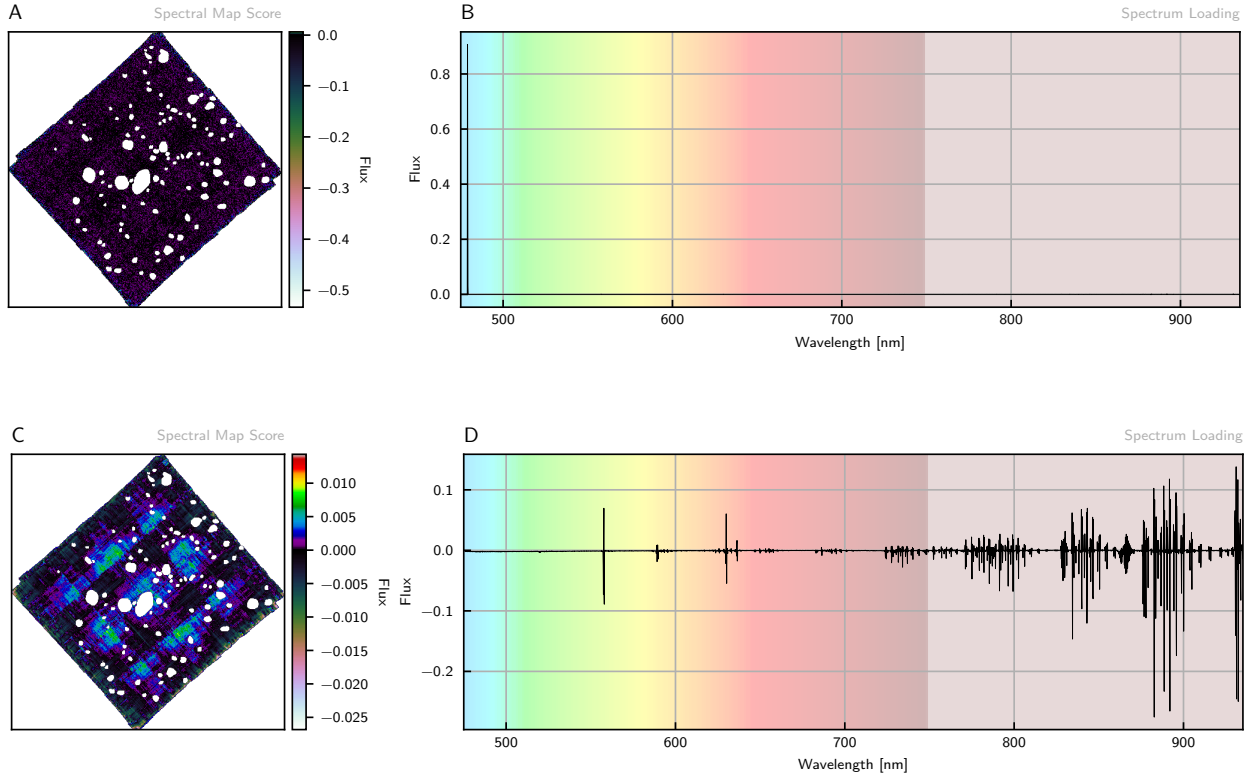


Figure A.9 | Scores and loadings for the first component of PCA (A, B) and PCP (C, D) for UDF-06. Flux is given in normalized units.

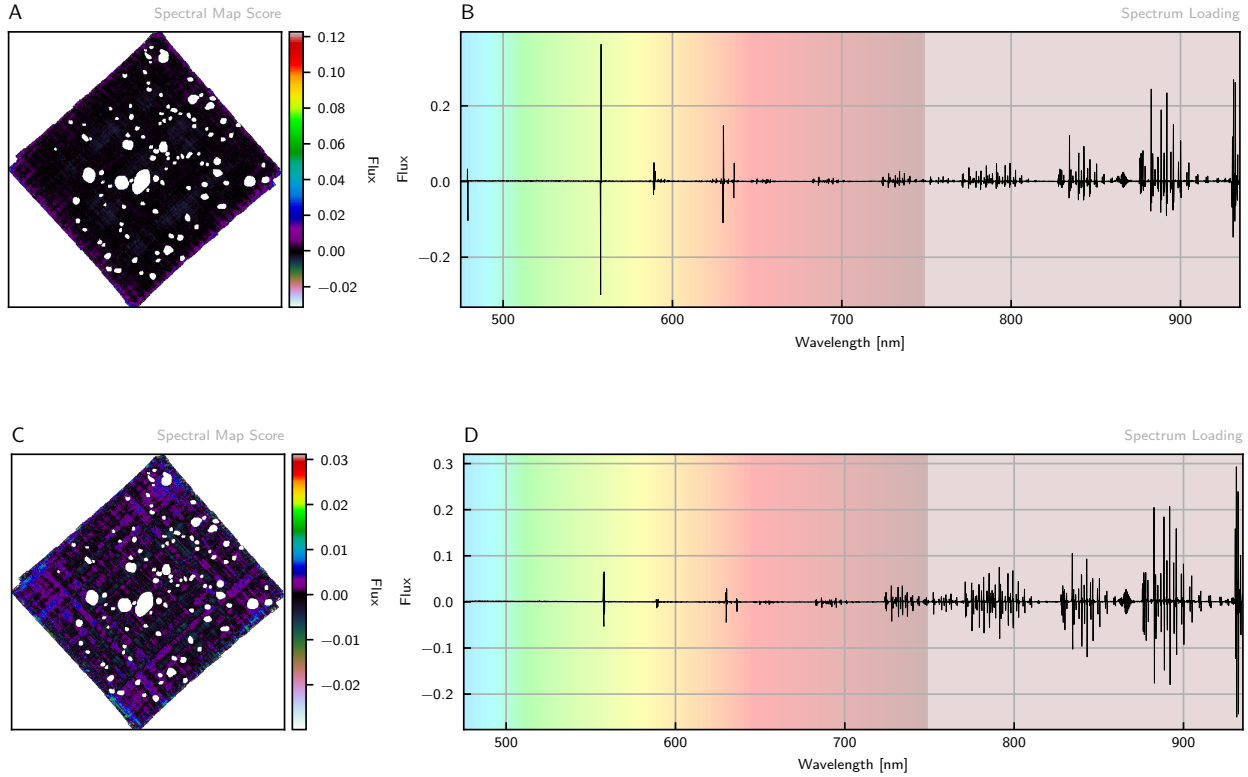


Figure A.10 | Scores and loadings for the second component of PCA (A, B) and PCP (C, D) for UDF-06. Flux is given in normalized units.




Figure A.11 | Score and loading for the third component of PCP (A, B) for UDF-06. Flux is given in normalized units.

Appendix B

HannibalSpecter

HannibalSpecter is fully object-oriented implemented following SOLID principles and the PEP8 standard. An overview of its object oriented structure is given in **Fig. B.1**.




Figure B.1 | Object-Oriented structure of HannibalSpecter.

Appendix C

Oak Ridge National Lab: Research Proposal

DESCRIPTION OF PROPOSED RESEARCH

The description must be limited to a maximum of 2 pages, including text and figures. PIs are encouraged to consult the Review Criteria for CNMS Research Proposals at <https://www.ornl.gov/facility/cnms/subpage/user-policies>.

Note: If you plan to use figures, it is best to copy/paste pre-formatted figures with text into this form.

ADDRESS EACH OF THE FOLLOWING QUESTIONS IN A SEPARATE SECTION.

A maximum of 2 pages can be used to respond to the 6 numbered questions below; Proposers may determine the amount of space used for each question.

1) **What is/are the main scientific or technological question(s) that you plan to address?**

Spectral imaging modalities, such as imaging mass spectrometry, can deliver massively multivariate and extremely high-dimensional measurement sets. The high-dimensional nature of such data makes direct human interpretation of spectral imaging data often impractical, and it raises challenges (e.g. the “curse of dimensionality”) for the efficient and robust application of statistical analyses and machine learning approaches to such datasets.

The questions at the center of this project are:

(a) Can a structured and regularized matrix factorization method (e.g. building on the methods behind RPCA) be developed to efficiently reduce massively multivariate and high-dimensional spectral imaging datasets to a lower-dimensional representation, with minimal information loss and yielding an increased signal-to-noise ratio?

(b) Can an online (i.e. serial) version of this approach be developed for a memory and compute resource restricted environment, that performs equally well as the original batch (i.e. parallel) version that assumes the entire dataset can be held in memory?

2) **Outline the overall technical approach that you plan to use to address the above questions. This section should provide the context for research tasks described below in sections (3), home institution activities, and (4), CNMS research.**

The analysis and modeling of imaging systems and their data has traditionally focused on image types that measure a single feature per pixel (e.g. gray-level image) or a small number of scalar features per pixel (e.g. color image). Recent developments have led to vastly improved spatial, but also spectral, resolution per pixel. This has paved the way for spectral imaging, a growing class of novel imaging modalities that instead of a few features per pixel, record long vectors of hundreds to thousands of variables per pixel, often acquiring a full electromagnetic or molecular mass spectrum per picture element. These improvements in spectral and spatial resolution are concurrently driving an exponential increase in data output, making direct human interpretation of spectral imaging data often impractical. Those novel modalities, hence, have to go hand in hand with advanced mathematical analysis and data mining methods to make the data fully interpretable for the end-user.

In this project, we are working specifically on the development of novel matrix factorization approaches for spectral imaging data, such as imaging mass spectrometry measurements. The goal is to obtain low-dimensional representations of these measurements that are more manageable, require less memory and compute resources, exhibit minimal information loss compared to the original high-dimensional dataset, and that remove sparse noise perturbations from the measurements in the process. The approaches will build on the insights obtained from structured factorization approaches such as Robust Principal Component Analysis [1], in which a low-rank matrix of underlying signals can be retrieved from highly corrupted measurements using a regularization approach that separates out any sparse perturbations.

The low-dimensional representation of an imaging mass spectrometry measurement set will immediately deliver to the analytical chemist a set of trends underlying the data. Each trend consists of a spectral signature showing which ions are correlated or anti-correlated in this dataset, and a spatial signature showing where in the tissue these ions tend to correlate.

Besides the benefit for human interpretation of spectral imaging data, the dimensionality reduction for large IMS datasets will also be beneficial for reducing the “curse of dimensionality” for subsequent computational and statistical analyses, and will yield more robust results in unsupervised and supervised machine learning approaches applied to IMS data. One example in which the dimensionality reduction can mean the difference between feasible and infeasible is the area of data-driven multimodal image fusion (Van de Plas et al., *Nature Methods*, 2015 [2]).

3) **What research tasks will be carried out at the users’ home institution or elsewhere outside of the CNMS? Include any preliminary syntheses, measurements, or tests that have been/will be performed in preparation for the proposed research at the CNMS.**

Imaging mass spectrometry measurements will be acquired at the Mass Spectrometry Research Center (MSRC) in the Vanderbilt University in Nashville, TN, under guidance of Prof. Dr. Richard Caprioli and Prof. Dr. Jeff Spraggins. Initial development of the algorithms will take place at the Delft Center for Systems and Control in the Delft University of

Technology in Delft, the Netherlands, under guidance of Prof. Dr. Raf Van de Plas, and will continue at the MSRC while the student is in the United States.

- 4) **Describe very clearly and specifically the research tasks to be carried out at the CNMS and the expected outcomes from the CNMS tasks. Include any technical milestones that must be met and the need for specialized capabilities and/or expertise at the CNMS for the research to be successful. (This should be the longest and most detailed section in the proposal)**

The algebraic matrix factorization of datasets such as imaging mass spectrometry data is effectively a large mathematical optimization problem with several regularization terms and constraints. For very large datasets (e.g. >1TB), the resources required to calculate this optimization problem lie beyond the capabilities of regular compute infrastructure. The initial goal is therefore to make use of the supercomputing infrastructure at ORNL to further develop and customize the algorithm for very large IMS datasets, and to establish a baseline for performance in a non-memory and compute resource restricted environment such as a supercomputer. Subsequently, the goal will be to develop a serial version of the factorization algorithm that can calculate the solution to the optimization problem by serially reading parts of the data into memory, and to compare its performance against the performance of the unconstrained version developed earlier. After proper characterization and benchmarking, the serial version of the algorithm will then be available for reducing very large IMS datasets even in more restricted compute environment such as regular computer hardware.

For the research to be successful, the primary needs are (i) access to the expertise on molecular imaging-oriented spectral data available in the Ovchinnikova lab, and (ii) access to the supercomputing infrastructure at ORNL.

- 5) **Provide an overall timeline for the CNMS tasks and describe how each facility/instrument that is checked on p. 2 will be used, including estimates of the number/quantities of samples, instrument time, CPU time, etc.**

First two weeks of visit to ORNL: focus on developing unconstrained version of factorization algorithm for very large IMS datasets. Second set of two weeks of visit: focus on developing serial version of factorization algorithm, and benchmarking its performance against unconstrained version.

- 6) **What is your team's specific experience and expertise relevant to this research project?**

The Van de Plas lab focuses on the interface between (i) mathematical engineering and machine learning; (ii) analytical chemistry and instrumentation; and (iii) life sciences and medicine.

Raf Van de Plas is assistant professor at the Delft University of Technology (TU Delft) in the Netherlands. At the Delft Center for Systems and Control, he works on computational analysis of molecular imaging modalities such as imaging mass spectrometry and microscopy. He holds an adjunct assistant professor position in biochemistry at Vanderbilt University School of Medicine in Nashville, TN. He holds a PhD in Engineering (KU Leuven, 2010), a degree in Industrial Engineering (Group T, 2002), and an MSc in Artificial Intelligence (KU Leuven, 2003). Before his appointment at TU Delft, he held a research faculty position at the Mass Spectrometry Research Center at Vanderbilt University with Dr. Richard Caprioli.

PUBLICATION RECORD: Have you had any previous CNMS project(s), including SHaRE? YES ___NO X (response required) If yes, list publications resulting from your past project(s) – maximum of 10: (not included in the 2 page limit)

LITERATURE CITED ABOVE - if any (not included in the 2 page limit)

[1] Candès, E. J., et al. "Robust principal component analysis?." Journal of the ACM (JACM) 58.3 (2011): 11.

[2] Van de Plas, R., Yang, J., Spraggins, J., and Caprioli, R. M., "Image fusion of mass spectrometry and microscopy: a multimodality paradigm for molecular tissue mapping," Nature Methods, vol. 12, no. 4, pp. 366–372, Apr. 2015.

The section below is for PARTNER USER proposals only (half page or less - not included in the 2 page limit)

PARTNER USER proposals only: What unique, new capabilities will be developed at the CNMS as a result of this approach? How will these contribute to future research by other CNMS users?