

ABSTRACT

Title of dissertation: TENSOR COMPLETION FOR
MULTIDIMENSIONAL INVERSE
PROBLEMS WITH APPLICATIONS
TO MAGNETIC RESONANCE
RELAXOMETRY

Ariel Hafftko, Doctor of Philosophy, 2016

Dissertation directed by: Professor Wojciech Czaja
Department of Mathematics

This thesis deals with tensor completion for the solution of multidimensional inverse problems. We study the problem of reconstructing an approximately low rank tensor from a small number of noisy linear measurements. New recovery guarantees, numerical algorithms, non-uniform sampling strategies, and parameter selection algorithms are developed.

We derive a fixed point continuation algorithm for tensor completion and prove its convergence. A restricted isometry property (RIP) based tensor recovery guarantee is proved. Probabilistic recovery guarantees are obtained for sub-Gaussian measurement operators and for measurements obtained by non-uniform sampling from a Parseval tight frame.

We show how tensor completion can be used to solve multidimensional inverse problems arising in NMR relaxometry. Algorithms are developed for regularization parameter selection, including accelerated k -fold cross-validation and generalized

cross-validation. These methods are validated on experimental and simulated data.

We also derive condition number estimates for nonnegative least squares problems.

Tensor recovery promises to significantly accelerate N -dimensional NMR relaxometry and related experiments, enabling previously impractical experiments.

Our methods could also be applied to other inverse problems arising in machine learning, image processing, signal processing, computer vision, and other fields.

TENSOR COMPLETION FOR MULTIDIMENSIONAL INVERSE
PROBLEMS WITH APPLICATIONS TO MAGNETIC
RESONANCE RELAXOMETRY

by

Ariel Hafftko

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2016

Advisory Committee:
Professor Wojciech Czaja, Chair/Advisor
Dr. Richard G. Spencer
Professor John J. Benedetto
Professor Kasso A. Okoudjou
Professor Amitabh Varshney

© Copyright by
Ariel Hafftko
2016

Acknowledgments

I am incredibly thankful to my advisor, Wojciech Czaja, for his exceptional guidance and unwavering support. Wojciech is a paragon of excellence in academic advising. He gave me the freedom to explore my own mathematical interests, consistently provided deep mathematical insights and guidance, all the while advocating for me. Wojciech never ceased to provide encouragement, particularly during the most difficult times of writing this thesis. Wojciech's guidance helped me become the applied mathematician I am today.

I am deeply grateful to Dr. Richard G. Spencer, who supervised my work on NMR applications at the National Institute on Aging, NIH. I have never met a scientist so knowledgeable and so modest. It is with Dr. Spencer's guidance that I have developed a deeper understanding of how to apply mathematics to study scientific questions. I am very grateful for his continual encouragement and his immense generosity with his time. Dr. Spencer's numerous mathematical and scientific insights were essential to completing this thesis.

I also want to thank Dr. Hasan Celik, a postdoctoral fellow with whom I collaborated on NMR applications at the National Institute on Aging, NIH. Dr. Celik explained a lot of details of NMR physics to me and was always willing to share his extensive knowledge. I am deeply grateful for his numerous insights and feedback. Dr. Celik also acquired several experimental datasets that were used to test the algorithms in this thesis.

I owe thanks to other scientists in Dr. Spencer's group at the National Institute

on Aging, from whom I gained invaluable knowledge: Dr. Kenneth Fishbein, Dr. Mustapha Bouhrara, Dr. Donnie Cameron, Dr. Ariel Zane, Dr. David Reiter, Beth Ashinsky, Dr. Seongjin Choi, and Kyle Sexton.

I am grateful to Professor John J. Benedetto for welcoming me into the Norbert Wiener Center and for teaching several excellent classes on harmonic analysis.

I am thankful to Dr. Alexander Cloninger, with whom I had numerous enlightening conversations.

Special thanks are due to Professor Konstantina Trivisa for welcoming me into the AMSC program, for encouraging my early interest in applied mathematics, and for her overwhelming enthusiasm.

I thank Professor Kasso Okoudjou for providing insightful feedback on my dissertation and I thank Professor Amitabh Varshney for his role in my dissertation defense.

Alverda McCoy, Celeste Regalado, and Linette Berry helped me navigate the administrative aspects of graduate school. Professor Maria K. Cameron mentored and guided me during the early stages of my graduate education. Professor Diane P. O’Leary’s class on scientific computing was exceptional and I thank her for encouraging my interest in computer science.

I am grateful to have had the opportunity to work with Dr. Benjamin Manning, Dr. Daniel Weinberg, Dr. Rongrong Wang, and Dr. Xuemei Chen.

I feel lucky for the friends I made during graduate school, including Jake, Zsolt, Tim, Shoshana, Stefan, Lucia, Viktoria, Mark, Paul, Karamatou, Marie, the Matt’s, Chae, Catie, Franck, Christiana, Scott, Joe, and many others.

I am indebted to the late Professor Paul J. Sally, Jr., and to Professor John Boller, who mentored me during my undergraduate education at the University of Chicago, and whose guidance got me to the University of Maryland. Professor Sally and Professor Boller helped me to see the deep beauty of pure mathematics and inspired me to seek ways to improve the world with mathematics.

I thank my parents for their unending confidence and enthusiasm and for encouraging me since I was a child to pursue my dreams in mathematics. I thank my brother and sister for their encouragement and friendship.

I am most grateful to my beautiful wife Ruirui, the strongest, wisest, and most caring person I know.

Table of Contents

1	Introduction	1
2	Mathematical Preliminaries	6
2.1	Basic Notation	6
2.2	Linear Algebra	7
2.3	Tikhonov Regularization	9
2.4	Tensor Notation	14
2.5	Optimization	18
2.6	Finite Frames	19
2.7	Introduction to Compressed Sensing	20
2.8	Restricted Isometry Property	23
2.9	Recovery Algorithms in Compressed Sensing	25
2.10	Incoherence	27
2.11	Probability Theory	29
2.12	Structured Random Measurement Operators	31
2.13	Matrix Completion	34
2.13.1	Matrix Restricted Isometry Property	38
2.13.2	Singular Value Thresholding (SVT)	39
2.13.3	Fixed Point Continuation	41
2.14	Introduction to Tensor Completion	41
3	Tensor Completion and Recovery Guarantees	46
3.1	Overview	46
3.2	Recovery Guarantees via Restricted Isometry Property (RIP)	48
3.3	Probabilistic Recovery Guarantees	58
3.4	Non-Uniform Sampling from Parseval Tight Frames	63
3.5	Choice of Non-Uniform Sampling Distribution	67
3.6	Tensor Recovery Algorithm	69
3.6.1	Derivation	70
3.6.2	Fixed Point Iteration	76
3.6.3	Convergence	77
3.6.4	Homotopy Path of Regularization Parameters: Version 1	80

3.6.5	Homotopy Path of Regularization Paramaters: Version 2	82
3.7	Choosing the Parameter τ	83
3.8	Choosing the Regularization Parameter μ	84
3.9	Accelerated k -Fold Cross-Validation for Selecting μ	87
4	Nuclear Magnetic Resonance Preliminaries	91
4.1	Relaxometry	91
4.2	1-Dimensional Relaxometry Experiments	92
4.3	Multidimensional Relaxometry and Related Experiments	96
4.4	General Mathematical Setup	98
4.5	Description of the Kernels K_i	99
4.6	Nonnegative Least Squares (NNLS) and Tikhonov Regularization . .	102
4.7	Data Compression	103
5	Compressed Sensing for Nuclear Magnetic Resonance Relaxometry	108
5.1	Overview	108
5.2	Compressed Sensing for 2-Dimensional NMR Relaxometry	109
5.3	Tensor Completion Applied to NMR	111
5.4	A Naive Approach: N -Dimensional NMR Data Recovery by Matrix Completion on 2-Dimensional Slices	112
5.5	Non-Uniform Sampling of N -Dimensional Relaxometry Data	115
5.6	Choosing the Regularization Parameter α in NNLS	118
5.7	3-Dimensional Tensor Recovery on Simulated Data	119
5.7.1	Results at Sampling Ratio 0.0025 and $SNR = 16384$	121
5.7.2	Results at Sampling Ratio 0.01 and $SNR = 16384$	123
5.7.3	Relative Errors vs. Sampling Ratio	126
5.8	Application to $T_1 - D - T_2$ Experimental Data	128
6	Stability Results for Nonnegative Least Squares	133
6.1	Introduction	133
6.2	Stability of Least Squares	134
6.3	Extension to Nonnegative Least Squares	139
	Conclusion	149
	Bibliography	151

Chapter 1: Introduction

This thesis deals with tensor completion and applications to multidimensional inverse problems. Tensor completion, the problem of reconstructing a multidimensional data array from incomplete measurements, is a problem of fundamental importance for high-dimensional data recovery. Applications of tensor completion include machine learning, image processing, signal processing, computer vision, and the efficient solution of multidimensional inverse problems.

Compressed sensing theory, pioneered in 2004-2006 by Candès, Tao, Donoho, and Romberg, enables the recovery of sparse or compressible signals from a small number of incoherent measurements. A fundamental result in compressed sensing is that sparse signals can be approximately recovering by l^1 norm minimization, which is computationally tractable.

Recht, Fazel, and Parrilo [59], Candès [32], and Candès and Plan [31] showed that fundamental results in compressed sensing hold also for matrix recovery. For compressed sensing of matrices, sparsity or approximate sparsity is replaced with a low-rank or approximate low-rank assumption. While the recovery of a compressible signal implicitly involves the identification of a basis (or frame) in which it is sparse and the recovery of its coefficients, the recovery of a compressible matrix requires

learning both its principle singular vectors and its principle singular values.

A number of tensor completion approaches have been studied. J. Liu, Malski, Wonka, and Ye [98] first proposed an algorithm for tensor completion via the minimization of a sum of nuclear norms. The algorithm works by unfolding, or flattening, the tensor into a matrix along each axis. The unfolding approach has been further studied by Tomioka [141, 140, 139], Gandy, Recht and Yamada [64], Krishnamurthy and Singh [89], Yang [148], and Zhang, Yang and Huang [150]. Tensor factorization techniques have also been studied by Y. Liu, Shang, Fan, J. Cheng, and H. Cheng [99, 100, 101]. Other approaches include Bayesian methods by Zhao, Zhang, and Cichocki [151] and Bazerque, Mateos, and Giannakis [8].

Other tensor completion approaches aim to capture more high-dimensional structure than unfolding methods. Rauhut, Schneider, and Stojanac have developed an iterative hard thresholding (IHT) algorithm for tensor completion that uses the full higher-order singular value decomposition (HOSVD) structure [119, 122]. Rauhut and Stojanac have also developed algorithms using θ -norms [121], which are relaxations of the true tensor nuclear norm, not a sum of nuclear norms of unfoldings. Mu, Huang, Wright, and Goldfarb proposed unfolding a high dimensional tensor into a matrix that is as square as possible [111].

Our work was initially motivated by the results of Cloninger and Czaja [42, 41]. Cloninger and Czaja developed an algorithm to solve discrete, separable Fredholm integral equations using matrix completion. They observed that for highly ill-conditioned inverse problems, the solution depends only on the projection of the observed data onto a significantly lower dimensional space. They applied matrix

completion to reconstruct the compressed data. By extending a probabilistic restricted isometry property (RIP) result of Y. Liu [102] to the case of Parseval tight frames, Cloninger and Czaja established a probabilistic recovery guarantee. They showed that the resulting algorithm is highly effective for 2-dimensional nuclear magnetic resonance relaxometry. The algorithm was validated on experimental data in [3].

Here, we extend the results of Cloninger and Czaja to the setting of tensor completion applied to the solution of multidimensional ill-posed inverse problems. We establish a deterministic restricted isometry property (RIP) based recovery guarantee for tensor recovery. We also establish probabilistic recovery guarantees for a more general class of random measurement operators, sub-Gaussian maps.

In nuclear magnetic resonance (NMR) relaxometry and related applications, N -dimensional experiments ($N \geq 2$) promise two primary advantages over 1 or 2-dimensional experiments. First, N -dimensional experiments provide significantly richer information than 1-dimensional experiments [29]. Second, in [38], Celik, Bouhrara, Reiter, Fishbein, and Spencer observed empirically that 2-dimensional relaxometry problems exhibit better stability than similar 1-dimensional problems. By extension, these results suggest that N -dimensional experiments ($N \geq 3$) could provide even better stability than 2-dimensional experiments.

This dissertation is organized as follows. In Chapter 3, we develop tensor recovery guarantees and algorithms for tensor recovery via sum of nuclear norms minimization. In Section 3.2 a restricted isometry property (RIP) recovery guarantee is proved. Section 3.3 develops probabilistic tensor recovery guarantees for a general

class of random measurement operators, sub-Gaussian maps. Section 3.4 obtains recovery guarantees for non-uniform sampling from Parseval tight frames. Section 3.5 shows how the non-uniform sampling distribution can be tailored to minimize the coherence of the random measurement operator, resulting in improved recovery guarantees. Section 3.6 develops a fast fixed point continuation algorithm for tensor completion, improving upon a previous algorithm by Yang, Huang, and Shi [148]. In section 3.7, we develop an algorithm, accelerated k-fold cross-validation, that rapidly estimates the regularization parameter for tensor recovery that minimizes a measure of generalization error.

Chapters 4, 5, and 6, described below, are joint with Dr. Richard G. Spencer and Dr. Hasan Celik.

Chapter 4 introduces nuclear magnetic resonance (NMR) relaxometry and the ill-conditioned inverse problems that arise in NMR.

In Chapter 5, we apply our results to solve ill-conditioned multidimensional inverse problems arising in NMR relaxometry. For 3-dimensional experimental and simulated NMR data, we demonstrate that our algorithm consistently reconstructs distributions of relaxation parameters. The methods described can be applied to T_1 - T_2 , T_2 -store- T_2 , T_1 - D - T_2 , and other experiments. In Section 5.6, we propose a heuristic technique, generalized cross-validation, for the choice of regularization parameter for multidimensional inverse problems.

Chapter 6 develops condition number estimates for Tikhonov regularized non-negative least squares problems. These results can be applied to quantify the stability of multidimensional inverse problems.

This work was supported by the Intramural Research Program, National Institute on Aging, of the National Institutes of Health (NIH). Ariel Hafftko is a participant in the NIH Graduate Partnership Program.

Chapter 2: Mathematical Preliminaries

2.1 Basic Notation

Let \mathbb{R} denote the real numbers. Given sets A and B , let $A \times B = \{(a, b) : a \in A, b \in B\}$ denote the Cartesian product of sets. We denote by $\mathbb{R}^d = \mathbb{R} \times \cdots \times \mathbb{R}$ a d -dimensional Euclidean vector space over \mathbb{R} with the standard basis e_1, \dots, e_d . $\mathbb{R}^{d_1 \times d_2}$ denotes the space of matrices of size $d_1 \times d_2$. In our convention, vectors are denoted by lowercase letters, such as $x, y \in \mathbb{R}^d$, and matrices are denoted by uppercase letters, such as $X, Y \in \mathbb{R}^{d_1 \times d_2}$. The transpose of a matrix X is denoted X^t . For any integer $n \geq 1$, let $[n]$ be the set $\{1, \dots, n\}$. Given two real vector spaces V and W , let $L(V, W)$ denote the space linear maps $T : V \rightarrow W$.

For $N \geq 1$ and $d_1, \dots, d_N \geq 1$, let $\mathbb{R}^{d_1 \times \cdots \times d_N}$ denote the space of N -dimensional tensors, or real N -dimensional arrays, of size $d_1 \times \cdots \times d_N$. Although $\mathbb{R}^{d_1 \times \cdots \times d_N}$ can be viewed as an N -dimensional array, it is a $(d_1 d_2 \cdots d_N)$ -dimensional real vector space. Tensors are denoted by boldface capital letters, such as \mathbf{X}, \mathbf{Y} . Tensors are a natural generalization of vectors and matrices. Individual entries of a vector, matrix, or tensor will be denoted by $x[i_1]$, $X[i_1, i_2]$, or $\mathbf{X}[i_1, \dots, i_N]$.

We define the following l^p norms and inner products for vectors $x \in \mathbb{R}^d$. $\|x\|_p = (\sum_{i=1}^d |x_i|^p)^{1/p}$ for $1 \leq p < \infty$. For $p = \infty$, let $\|x\|_\infty = \max_i |x_i|$. For any set A , let

$|A|$ denote its cardinality. The sparsity of a vector x is defined by $\|x\|_0 = |\{i|x_i \neq 0\}|$ and is equal to the number of nonzero entries in x . Despite the similar notation to the $\|x\|_p$ norms for $p \geq 1$, sparsity $\|x\|_0$ is not a norm. The inner product of two vectors $x, y \in \mathbb{R}^d$ is $\langle x, y \rangle = \sum_{i=1}^d x_i y_i$.

2.2 Linear Algebra

We state several standard facts from linear algebra. Given a matrix $A \in \mathbb{R}^{d_1 \times d_2}$, we define $\ker(A) = \{x|Ax = 0\}$ and $\text{im}(A) = \{Ax|x \in \mathbb{R}^{d_2}\}$. The $n \times n$ identity matrix is denoted by $\text{Id}_{n \times n}$. The rank of a matrix A is denoted $\text{rank}(A)$.

Definition 2.1 (Orthonormal matrix). *A matrix $U \in \mathbb{R}^{d_1 \times d_2}$ is said to have orthonormal columns if $U^t U = \text{Id}_{d_2 \times d_2}$. U is said to have orthonormal rows if $U U^t = \text{Id}_{d_1 \times d_1}$. A square matrix U is orthonormal if it has orthonormal columns and orthonormal rows.*

Definition 2.2 (Singular value decomposition (SVD)). *Let $A \in \mathbb{R}^{d_1 \times d_2}$ with $r = \text{rank}(A) \geq 1$. The SVD of A consists of matrices $U \in \mathbb{R}^{d_1 \times r}$ and $V \in \mathbb{R}^{d_2 \times r}$ with orthonormal columns and a diagonal matrix $S = \text{diag}(\sigma_1, \dots, \sigma_r) \in \mathbb{R}^{r \times r}$, $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r \geq 0$, satisfying*

$$A = U S V^t.$$

Definition 2.3 (Full singular value decomposition (full SVD)). *Let $A \in \mathbb{R}^{d_1 \times d_2}$. The full SVD of A consists of orthonormal matrices $U \in \mathbb{R}^{d_1 \times d_1}$ and $V \in \mathbb{R}^{d_2 \times d_2}$, and a rectangular matrix $S \in \mathbb{R}^{d_1 \times d_2}$ with diagonal entries $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{\min(d_1, d_2)} \geq 0$*

and all off-diagonal entries equal to 0, satisfying

$$A = USV^t.$$

The following is a standard theorem [67].

Theorem 2.4 (Existence and uniqueness of SVD). *Let $A \in \mathbb{R}^{d_1 \times d_2}$ be nonzero. Then, the SVD of A , of the form $A = USV^t$, exists and is unique, up to possibly replacing U and V with $U\Lambda$ and $V\Lambda$, where Λ is a block diagonal matrix with blocks corresponding to the distinct singular values of A , and where each block occurring along the diagonal of Λ is orthonormal.*

In the simple case in which all singular values are distinct, the SVD is unique up to possible sign changes of the columns of U and V . In the more general case, Theorem 2.4 states that if several singular vectors have the same singular value, they can be transformed by an orthonormal transformation.

We define the following inner products and norms for matrices. Let $X, Y \in \mathbb{R}^{d_1 \times d_2}$. We define $\text{vec}(X)$ to be the column vector in $\mathbb{R}^{d_1 d_2}$ obtained by concatenating the columns of X . We define the inner product $\langle X, Y \rangle = \langle \text{vec}(X), \text{vec}(Y) \rangle = \sum_{i_1=1}^{d_1} \sum_{i_2=1}^{d_2} X[i_1, i_2]Y[i_1, i_2]$. The Frobenius norm for matrices is defined to be $\|X\|_F = \sqrt{\langle X, X \rangle}$.

For $p \geq 1$, the Schatten- p norms for matrices are defined as follows. Let $X = U \text{diag}(\sigma_1, \dots, \sigma_r) V^t$ be the SVD of X , and define $\|X\|_{S_p} = (\sum_{i=1}^r \sigma_i^p)^{1/p}$. The nuclear norm for matrices is defined to be $\|X\|_* = \sum_{i=1}^r \sigma_i(X)$. The operator norm for matrices is defined by $\|X\|_2 = \sigma_{\max}(X)$.

We have the following equality of norms. The nuclear norm $\|X\|_*$ is equal to the Schatten-1 norm $\|\mathbf{X}\|_{S_1}$. The operator norm $\|X\|_2$ is equal to the Schatten- ∞ norm $\|X\|_{S_\infty}$. The Frobenius norm $\|X\|_F$ is equal to the Schatten-2 norm $\|X\|_{S_2}$.

Given a nonzero matrix $X \in \mathbb{R}^{d_1 \times d_2}$ of rank $r \geq 1$ with singular value decomposition $U \text{diag}(\sigma_1, \dots, \sigma_r) V^t$, we define $\sigma_{\max}(X) = \sigma_1(X)$ to be the largest singular value of X . We define $\sigma_r(X)$ to be the smallest nonzero singular value of X . We define $\sigma_{\min}(X)$ to be the smallest singular value in the full SVD of X . Hence, if X is not full-rank, $\sigma_{\min}(X) = 0$ while $\sigma_r(X) \neq 0$. We define the condition number of X to be

$$\kappa(X) = \frac{\sigma_1(X)}{\sigma_r(X)},$$

the ratio between the largest and smallest nonzero singular values of X .

Lemma 2.5. *Let X be a matrix with $\text{rank}(X) \leq r$. Then $\|X\|_* \leq \sqrt{r} \|X\|_F$.*

Proof. Let $X = U \text{diag}(\sigma_1, \dots, \sigma_r) V^t$ be the SVD of X . Then by the Cauchy-Schwarz inequality,

$$\|X\|_* = \sum_{k=1}^r \sigma_k \leq \sqrt{r} \sqrt{\sum_{k=1}^r \sigma_k^2} = \sqrt{r} \|X\|_F. \quad (2.1)$$

□

2.3 Tikhonov Regularization

Given a kernel $K \in \mathbb{R}^{m \times n}$ and a data vector $y \in \mathbb{R}^m$, consider the least squares problem

$$\min_{x \in \mathbb{R}^n} \|Kx - y\|_2^2 \quad (2.2)$$

If the kernel K is ill-conditioned, small changes in the data y could result in large changes in the solution x . Tikhonov regularization is a standard method to improve the conditioning of (2.2). Tikhonov regularization solves the following problem instead of (2.2):

$$\min_{x \in \mathbb{R}^n} \|Kx - y\|_2^2 + \alpha^2 \|x\|_2^2 \quad (2.3)$$

The following standard result shows that the Tikhonov regularized least squares problem (2.3) is equivalent to a least squares problem [78, Equation 4.9].

Lemma 2.6. *Let $K \in \mathbb{R}^{m \times n}$ and $y \in \mathbb{R}^m$. Define the augmented kernel*

$$K_{\text{aug}} := \begin{pmatrix} K \\ \alpha \text{Id}_{n \times n} \end{pmatrix} \quad (2.4)$$

and the augmented data

$$y_{\text{aug}} := \begin{pmatrix} y \\ 0_{n \times 1} \end{pmatrix}. \quad (2.5)$$

Then (2.3) is equivalent to:

$$\min_{x \in \mathbb{R}^n} \|K_{\text{aug}}x - y_{\text{aug}}\|_2^2 \quad (2.6)$$

in the sense that $x \in \mathbb{R}^n$ is a solution of (2.3) if and only if x is a solution of (2.6).

The following result characterizes the SVD of K_{aug} .

Lemma 2.7. *Let K be a matrix of size $m \times n$ and of rank $r \geq 1$. Let $K = USV^t$ be any full SVD of K , where U is $m \times m$, V is $n \times n$, U and V have orthonormal columns, and S is an $m \times n$ diagonal matrix with diagonal entries $\sigma_1 \geq \dots \geq$*

$\sigma_{\min(m,n)} \geq 0$. Then an SVD of the augmented matrix K_{aug} for $\alpha > 0$ is given by

$$K_{\text{aug}} = \tilde{U}\tilde{S}\tilde{V}^t,$$

where

$$\tilde{U} = \begin{pmatrix} US \\ \alpha V \end{pmatrix} \left(S^t S + \alpha^2 \text{Id}_{n \times n} \right)^{-1/2}$$

$$\tilde{S} = \left(S^t S + \alpha^2 \text{Id}_{n \times n} \right)^{1/2},$$

and

$$\tilde{V} = V.$$

Proof. We need to show that \tilde{U} and \tilde{V} have orthonormal columns, that \tilde{S} is diagonal with positive, non-increasing diagonal entries, and that $K_{\text{aug}} = \tilde{U}\tilde{S}\tilde{V}^t$. Observe that

$$\begin{pmatrix} S^t U^t & \alpha V^t \end{pmatrix} \begin{pmatrix} US \\ \alpha V \end{pmatrix} = S^t U^t U S + \alpha^2 V^t V = (S^t S + \alpha^2 \text{Id}_{r \times r}),$$

so the columns of the right factor on the left side above are orthogonal with norms $\sqrt{\sigma_i^2 + \alpha^2}$ for $i = 1, \dots, n$, where for notational brevity we let $\sigma_i = 0$ if $i \geq \min(m, n)$. It follows that $\tilde{U}^t \tilde{U} = \text{Id}_{r \times r}$. Since $\tilde{V} = V$, orthonormality of the columns of \tilde{V} follows from the same property for V .

Finally, we have

$$\begin{aligned} \tilde{U}\tilde{S}\tilde{V}^t &= \begin{pmatrix} US \\ \alpha V \end{pmatrix} \left(S^t S + \alpha^2 \text{Id}_{n \times n} \right)^{-1/2} \left(S^t S + \alpha^2 \text{Id}_{n \times n} \right)^{1/2} V^t \\ &= \begin{pmatrix} USV^t \\ \alpha VV^t \end{pmatrix} \end{aligned}$$

$$\begin{aligned}
&= \\
&= \begin{pmatrix} K \\ \alpha \text{Id}_{n \times n} \\ K_{\text{aug}} \end{pmatrix}
\end{aligned}$$

□

The following lemma shows that if a matrix K is replaced by a submatrix consisting of a subset of its columns, its largest singular value cannot increase. Furthermore, if $\ker(K) = \{0\}$, the smallest singular value cannot decrease. This lemma can be used to bound the expressions appearing in the condition number estimates in Chapter 6.

Lemma 2.8. *Let K be a nonzero matrix of size $m \times n$ and of rank r . Let $\Lambda \subset \{1, \dots, n\}$ be nonempty and let K_Λ be the submatrix of K consisting of columns from K with indices in Λ . Let $\sigma_1(K)$ and $\sigma_r(K)$ denote the largest and smallest nonzero singular values of K . Then*

$$\sigma_1(K_\Lambda) \leq \sigma_1(K). \quad (2.7)$$

If, in addition, we have $\ker K = \{0\}$, then

$$\sigma_r(K_\Lambda) \geq \sigma_r(K) \quad (2.8)$$

Proof. Let $\sigma_{\max}(K)$ and $\sigma_{\min}(K)$ denote the largest and smallest singular values of K occurring in the full SVD of K . Then $\sigma_1(K) = \sigma_{\max}(K)$ and $\sigma_{\min}(K) \leq \sigma_r(K)$. In general, $\sigma_r(K)$ is always nonzero, since K is nonzero, but $\sigma_{\min}(K)$ could be zero.

We use the following standard characterization of the maximum and minimum

singular values of the full SVD of a matrix [136, Equation 6.4.7]:

$$\sigma_{\max}(K) = \sup_{\|x\|=1} \|Kx\|_2 \quad (2.9)$$

and

$$\sigma_{\min}(K) = \inf_{\|x\|=1} \|Kx\|_2. \quad (2.10)$$

Hence, by (2.9) we have

$$\sigma_1(K_\Lambda) = \sup_{\|x\|=1} \|K_\Lambda x\|_2 \leq \sup_{\|x\|=1} \|Kx\|_2 = \sigma_1(K).$$

If $\ker(K) = \{0\}$, we have $\sigma_r(K) = \sigma_{\min}(K)$. Since $\ker(K) = \{0\}$, it follows that $\ker(K_\Lambda) = 0$ also. Hence we also have $\sigma_r(K_\Lambda) = \sigma_{\min}(K_\Lambda)$. Hence, by (2.10), we have

$$\begin{aligned} \sigma_r(K_\Lambda) &= \sigma_{\min}(K_\Lambda) \\ &= \inf_{\|x\|=1} \|K_\Lambda x\|_2 \\ &\geq \inf_{\|x\|=1} \|Kx\|_2 \\ &= \sigma_{\min}(K) \\ &= \sigma_r(K). \end{aligned}$$

□

Corollary 2.9. *If K is an $m \times n$ matrix, $\Lambda \subset \{1, \dots, n\}$ is nonempty, and $\ker(K) = \{0\}$, then all the singular values of K_Λ are contained in the interval $[\sigma_{\min}(K), \sigma_{\max}(K)]$ and the condition number of K_Λ satisfies*

$$\kappa(K_\Lambda) \leq \kappa(K). \quad (2.11)$$

2.4 Tensor Notation

We introduce standard tensor notation as in [87, 72]. Tensors are denoted by boldface capital letters, such as \mathbf{X}, \mathbf{Y} . Recall that matrices are denoted by plain capital letters such as X, Y , and vectors are denoted by plain lowercase letters, such as x, y . Individual entries are denoted by $\mathbf{X}[i_1, \dots, i_N]$.

Definition 2.10 (Tensor). *A tensor is any element $\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}$, where $N \geq 1$ and $d_1, \dots, d_N \geq 1$.*

We denote by $\mathbb{R}^{d_1 \times \dots \times d_N}$ the space of tensors of size $d_1 \times \dots \times d_N$. In contrast, $\mathbb{R}^{d_1 d_2 \dots d_N}$ denotes the space of column vectors of length $d_1 d_2 \dots d_N$.

Definition 2.11 (Lexicographical ordering). *The lexicographical ordering on $[d_1] \times \dots \times [d_n]$ is defined by $(i_1, \dots, i_N) < (j_1, \dots, j_N) \iff i_1 < j_1$ or for some $1 \leq k < N, i_1 = j_1, i_2 = j_2, \dots, i_k = j_k$ and $i_{k+1} < j_{k+1}$.*

Recall that for any integer $n \geq 1$, $[n]$ denotes the set $\{1, \dots, n\}$.

Definition 2.12 (Vectorization). *Given a tensor $\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}$, its vectorization,*

$$\text{vec}(\mathbf{X}),$$

is the column-vector of length $d_1 d_2 \dots d_N$ in $\mathbb{R}^{d_1 d_2 \dots d_N}$ obtained by arranging all the entries of \mathbf{X} into a column vector according to the lexicographical ordering on the indices of \mathbf{X} . Here $d_1 \dots d_N$ denotes the product of numbers.

Definition 2.13 (Reshape). *Given vector $x \in \mathbb{R}^{d_1 d_2 \dots d_N}$, $\text{reshape}(d_1, \dots, d_N)$ is the*

result of arranging the entries of x into a tensor $\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}$ according to the lexicographical order on $[d_1] \times \dots \times [d_N]$.

Definition 2.14 (Kronecker product). *The Kronecker product of matrices $X \in \mathbb{R}^{m_1 \times m_2}$ and $Y \in \mathbb{R}^{n_1 \times n_2}$ is the matrix in $X \otimes Y \in \mathbb{R}^{(m_1 n_1) \times (m_2 n_2)}$ defined by*

$$X \otimes Y := \begin{pmatrix} X[1, 1]Y & \cdots & X[1, n_1]Y \\ \vdots & \ddots & \vdots \\ X[m_1, 1]Y & \cdots & X[m_1, n_1]Y \end{pmatrix} \quad (2.12)$$

The following standard result, found in [93] and other textbooks, characterizes the SVD of a Kronecker product.

Theorem 2.15. *Let K_1, \dots, K_N be matrices with SVD's*

$$K_i = U_i S_i V_i^t.$$

The reduced SVD for the Kronecker product, up to a possible permutation of the order of the singular values and singular vectors, is given by

$$K_1 \otimes \dots \otimes K_N = (U_1 \otimes \dots \otimes U_N)(S_1 \otimes \dots \otimes S_N)(V_1 \otimes \dots \otimes V_N)^t.$$

The following definition generalizes the definition of rows and columns to tensors.

Definition 2.16. (*k-columns*) *Given a tensor $X \in \mathbb{R}^{d_1 \times \dots \times d_N}$, its k -columns are the vectors in \mathbb{R}^{d_k} of the form $\mathbf{X}[i_1, \dots, i_{k-1}, :, i_{k+1}, \dots, i_N]$ obtained by fixing all indices except the k -th, with the k -th index ranging from $1, \dots, d_k$.*

For a matrix, 1-columns are columns while 2-columns are rows.

Definition 2.17. (*Unfolding*) Given a tensor $\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}$, its k -th unfolding is the tensor $\mathbf{X}^{(k)} \in \mathbb{R}^{d_k \times (d_1 \dots d_{k-1} d_{k+1} \dots d_N)}$ obtained by arranging all the k -columns of \mathbf{X} into a matrix, according to the lexicographical ordering on $[d_1] \times \dots \times [d_{k-1}] \times [d_{k+1}] \dots \times [d_N]$.

Recall that for an integer n , $[n] = \{1, \dots, n\}$.

Definition 2.18 (*Refolding*). Given a vector $x \in \mathbb{R}^{d_1 \dots d_N}$, define $\text{refold}_{d_1, \dots, d_N}(x)$ to be the tensor in $\mathbb{R}^{d_1 \times \dots \times d_N}$ obtained by arranging the entries of x into a tensor according to the lexicographical ordering on $[d_1] \times \dots \times [d_N]$.

It is straightforward to show that unfolding and refolding are adjoint to each other. This property will be useful later when analyzing first order conditions for tensor optimization.

Given two vectors, $v_1 \in \mathbb{R}^{d_1}$ and $v_2 \in \mathbb{R}^{d_2}$, their *outer product* $v_1 \circ v_2 = v_1 v_2^t \in \mathbb{R}^{d_1 \times d_2}$ is the matrix whose (i_1, i_2) entry is given by $v_1[i_1]v_2[i_2]$. The following definition generalizes the outer product to more than two vectors.

Definition 2.19 (*Outer product*). Let $v_i \in \mathbb{R}^{d_i}$ for $i = 1, \dots, N$. The outer product of v_1, \dots, v_N , denoted by $v_1 \circ \dots \circ v_N \in \mathbb{R}^{d_1 \times \dots \times d_N}$, is the tensor with entries

$$(v_1 \circ \dots \circ v_N)[i_1, \dots, i_N] = v_1[i_1]v_2[i_2] \dots v_N[i_N].$$

The following definition provides a concise way to extract a subset of the entries of a tensor.

Definition 2.20. (*Tensor sampling*) Let $\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}$. For any subset $\Omega \subset [d_1] \times \dots \times [d_N]$ of the indices of \mathbf{X} , let $\mathbf{X}[\Omega]$ denote the vector in $\mathbb{R}^{|\Omega|}$ obtained by arranging

the entries $\{\mathbf{X}[i_1, \dots, i_N] : (i_1, \dots, i_N) \in \Omega\}$ into a column vector, according to the ordering lexicographical ordering inherited by Ω from $[d_1] \times \dots \times [d_N]$.

The following definition generalizes the action of a matrix on a vector to the action of a matrix on a tensor.

Definition 2.21. (*k-mode product*) The *k-mode product* of a tensor $\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}$ with a matrix $U \in \mathbb{R}^{n_k \times d_k}$ is the tensor $\mathbf{X} \otimes_k U \in \mathbb{R}^{d_1 \times \dots \times d_{k-1} \times n_k \times d_{k+1} \times \dots \times d_N}$ obtained by multiplying all the *k*-columns of \mathbf{X} by U .

If X is a matrix, $X \otimes_1 U = UX$ and $X \otimes_2 U = XU'$.

As shown in [92], the *k*-mode product can be described in terms of unfolding and matrix multiplication by

$$(\mathbf{X} \otimes_k U)^{(k)} = U\mathbf{X}^{(k)}. \quad (2.13)$$

Definition 2.22. (*Tucker decomposition*) A *Tucker decomposition* of a tensor $\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}$ consists of a core-tensor $\mathbf{C} \in \mathbb{R}^{r_1 \times \dots \times r_N}$ and collection of matrices $U_i \in \mathbb{R}^{d_i \times r_i}$, $i = 1, \dots, N$, satisfying

$$\mathbf{X} = \mathbf{C} \otimes_1 U_1 \otimes_2 U_2 \cdots \otimes_N U_N. \quad (2.14)$$

\mathbf{C} is called the core-tensor and U_i are called the Tucker factors.

The Tucker decomposition induces structure on vectorizations and unfoldings. The following is a standard result [92].

Lemma 2.23. Suppose $\mathbf{X} = \mathbf{C} \otimes_1 U_1 \otimes_2 U_2 \cdots \otimes_N U_N$ is a Tucker decomposition of \mathbf{X} . Then the following properties hold.

- (Vectorization) We have

$$\text{vec}(\mathbf{X}) = (U_N \otimes \cdots \otimes U_1)\text{vec}(\mathbf{C}). \quad (2.15)$$

- (Unfolding) For all $k = 1, \dots, N$,

$$\mathbf{X}^{(k)} = U_k \text{vec}(\mathbf{C})(U_N \otimes \cdots \otimes U_{k+1} \otimes U_{k-1} \otimes \cdots \otimes U_1). \quad (2.16)$$

Definition 2.24 (Tucker rank). A tensor $\mathbf{X} \in \mathbb{R}^{d_1 \times \cdots \times d_N}$ has Tucker rank (r_1, \dots, r_N) if $\text{rank}(\mathbf{X}^{(i)}) = r_i$ for $i = 1, \dots, N$.

The Tucker rank is equivalent to the size of the core tensor in any higher-order singular value decomposition (HOSVD) of \mathbf{X} [92].

2.5 Optimization

Definition 2.25 (convex function). A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex if for all $x, y \in \mathbb{R}^d$, and for all $\alpha \in [0, 1]$ we have

$$f(\alpha x + (1 - \alpha)y) \leq \alpha f(x) + (1 - \alpha)f(y).$$

Definition 2.26 (Subdifferential of convex function). Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a convex function. Let $x \in \mathbb{R}^d$. The subdifferential of f at x , denoted $\partial f(x)$, is the set

$$\partial f(x) = \{v \in \mathbb{R}^d : f(y) \geq f(x) + \langle y - x, v \rangle \forall y \in \mathbb{R}^d\}.$$

The following result is standard.

Lemma 2.27. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a convex function and let $x \in \mathbb{R}^d$. Then f attains a global minimum value at x if and only if $0 \in \partial f(x)$.

Proof. Suppose that $0 \in \partial f(x)$. Set $v = 0$ in the definition $\partial f(x)$. It follows that $0 \in \partial f(x)$ if and only if for all $y \in \mathbb{R}^d$,

$$f(y) \geq f(x) + \langle y - x, 0 \rangle = f(x).$$

Hence the claim follows. □

2.6 Finite Frames

Definition 2.28 ([17]). *A frame for a d -dimensional real Hilbert space \mathcal{H} is a set of vectors $\{u_1, \dots, u_n\} \subset \mathcal{H}$ such that there exist constants $A, B > 0$ for which*

$$A\|x\|_{\mathcal{H}}^2 \leq \sum_{i=1}^n |\langle u_i, x \rangle|^2 \leq B\|x\|_{\mathcal{H}}^2 \quad \forall x \in \mathcal{H}.$$

A frame is tight if $A = B$ and a frame is a Parseval tight if $A = B = 1$.

Frames have significantly increased the scope and applicability of results in harmonic analysis and compressed sensing [12, 11, 13, 14, 15, 16, 48, 47, 10, 49]. The redundancy and non-uniqueness in frame representations makes it easier to find stable sparse representations [4, 117]. A topic of recent interest has been scalable frames [40, 90].

Definition 2.29 (Analysis operator). *The analysis operator $T : \mathcal{H} \rightarrow \mathbb{R}^n$ for a frame $\{u_i\}_{i=1}^n$ is given by*

$$T(x)[i] = \langle x, u_j \rangle,$$

for $x \in \mathbb{R}^d$ and $i = 1, \dots, n$.

Lemma 2.30. *Let $\{u_i\}_{i=1}^n$ be a Parseval tight frame. Then the analysis operator $T : \mathcal{H} \rightarrow \mathbb{R}^n$ is an isometry.*

Proof. The result follows from the definition of a Parseval tight frame, since

$$\|T(x)\|_2^2 = \sum_{i=1}^n |\langle u_i, x \rangle|^2.$$

□

The following result is well-known.

Lemma 2.31. *Let $\{u_i\}_{i=1}^n$ be a Parseval tight frame for a d -dimensional Hilbert space \mathcal{H} . Then*

$$\sum_{i=1}^n \|u_i\|_{\mathcal{H}}^2 = d.$$

Proof. Without loss of generality, assume that $\mathcal{H} = \mathbb{R}^d$. Let M_T be the matrix for T with respect to the standard basis for \mathbb{R}^d . Then M_T has rows given by u_i^t , $i = 1, \dots, n$. Since M_T is an isometry, it must have exactly d nonzero singular values all equal to 1. Since $\sum_{i=1}^n \|u_i\|_2^2$ is equal to the Frobenius norm of M_T and since the Frobenius norm is equal to the sum of the squares of the singular values of M_T , we obtain

$$\sum_{i=1}^n \|u_i\|_2^2 = \|M_T\|_F^2 = \sum_{i=1}^d \sigma_i(M_T)^2 = \sum_{i=1}^d 1 = d.$$

□

In [137], geometric optimization over finite frames is studied.

2.7 Introduction to Compressed Sensing

Compressed sensing (CS) was pioneered by Candès, Tao, Romberg and Donoho. In a series of papers [34, 36, 35, 55, 33, 33], they proved that under certain assumptions on a signal and on a linear measurement operator, signals can be recovered

from a very small number of measurements. In Fourier analysis, the Shannon sampling theorem sets a fundamental lower limit on the number of measurements needed to reconstruct a band-limited function. Compressed sensing enables the recovery of signals from fewer measurements than would be required for traditional methods. Major additional contributions from DeVore, Baraniuk, Davenport, Rauhut, Wakin, Cohen, and Dahmen quickly followed [5, 43, 52, 44, 116].

We briefly list several pioneering papers before describing central results in more detail. In [34], Candès, Romberg, and Tao proved that it is possible to reconstruct a discrete signal $x \in \mathbb{C}^d$ from random observations of its Fourier coefficients. In [36], Candès and Tao proved recovery results for random Gaussian measurements under the assumption that the observed signals $x \in \mathbb{R}^d$ obey a power-law decay. In [35], Candès and Tao show that sparse signals can be exactly recovered by l_1 minimization.

In [55], Donoho proved recovery results for signals that are sparse with respect to an orthonormal basis or a tight frame. Donoho derived conditions on the sampling operator under which these results hold and showed that good CS measurement operators give what look like almost random linear combinations of the the signal entries.

In [33], Candès and Romberg introduced a fundamental concept, incoherence, and showed that incoherent measurements can be used to recover sparse signals from a small number of measurements. Candès and Romberg observed that incoherence can be used to ensure recovery from sampling operators constructed by randomly choosing a subset of the rows of an orthonormal matrix.

In compressed sensing, the goal is to recover a signal $x_0 \in \mathbb{R}^d$ from linear measurements of the form

$$y = \mathcal{A}(x_0) \tag{2.17}$$

in the noise-free case, or

$$y = \mathcal{A}(x_0) + z \tag{2.18}$$

in the noisy case, where $\mathcal{A} : \mathbb{R}^d \rightarrow \mathbb{R}^m$ is a linear operator, $y \in \mathbb{R}^m$ are the observed measurements, m is the number of measurements observed, and $z \in \mathbb{R}^m$ is a noise vector. In the noise free setting, $z = 0$.

If \mathcal{A} is described by a matrix A , we define

$$\mathcal{A}(x) = Ax,$$

so that each observation y_i is given by $y_i = \langle A[i, :], x \rangle$, where $A[i, :]$ is the i -th row of A .

Compressed sensing is concerned with the case in which the dimension of the signal to be recovered is much higher than the number of observations, i.e., $m \ll d$, in which case the linear system (2.17) is undetermined. Hence, recovery cannot be guaranteed for all signals x . A main theme in CS is that recovery can be guaranteed provided two general assumptions hold: (1) the signal is sparse in some respect, meaning that it contains a small amount of ‘information’, and (2) the operator \mathcal{A} approximately preserves the information in such sparse signals. There are several ways to quantify these conditions and each leads to different recovery theorems.

The following is perhaps the simplest notion of sparsity [51, 6].

Definition 2.32 (*r*-sparsity). A vector $x \in \mathbb{R}^d$ is *r*-sparse if $\|x\|_0 \leq r$, i.e., at most *r* entries $x[i]$ are nonzero.

Definition 2.33 (compressibility). A vector $x \in \mathbb{R}^d$ is said to be compressible if it is approximately equal to an *r*-sparse vector.

While there are many ways to quantify compressibility, one is a power-law decay.

Definition 2.34. (*Power-law decay*) [36] A vector $x \in \mathbb{R}^d$ obeys a power-law decay with constants $C > 0$ and $p > 0$ if $x_{\text{descending}}[i] \leq Ci^{-1/p}$, where $x_{\text{descending}}$ denotes the result of arranging the entries of x in order of descending absolute value.

In [43], Cohen, Dahmen, and DeVore derived necessary and sufficient conditions on the operator \mathcal{A} to approximately recover best *r*-term approximations of a given (possibly non-sparse) vector x in terms of null space properties.

In [113], applications of compressed sensing to image representation and compression are studied. The degree of sparsity, in biological and other datasets, plays an important role in determining the applicability of compressed sensing methods [69].

2.8 Restricted Isometry Property

The following property, introduced by Candès and Tao in [35], quantifies the notion that the operator \mathcal{A} should preserve information in sparse vectors.

Definition 2.35 (Restricted isometry property (RIP)). A linear operator $\mathcal{A} : \mathbb{R}^d \rightarrow$

\mathbb{R}^m satisfies the RIP of order r and constant $\delta_r > 0$ if

$$(1 - \delta_r)\|x\|_2^2 \leq \|\mathcal{A}(x)\|_2^2 \leq (1 + \delta_r)\|x\|_2^2$$

holds for all r -sparse vectors $x \in \mathbb{R}^d$.

If \mathcal{A} satisfies the RIP with $\delta_r < 1$, \mathcal{A} acts as an approximate isometry on all r -sparse vectors, thus preserving ‘most’ of the information in sparse vectors.

An alternate definition of RIP is sometimes used, without the squares.

Definition 2.36 (Restricted isometry property (RIP), square-free version). *A linear operator $\mathcal{A} : \mathbb{R}^d \rightarrow \mathbb{R}^m$ satisfies the square-free RIP of order r and constant $\bar{\delta}_r > 0$ if*

$$(1 - \bar{\delta}_r)\|x\|_2 \leq \|\mathcal{A}(x)\|_2 \leq (1 + \bar{\delta}_r)\|x\|_2$$

holds for all r -sparse vectors $x \in \mathbb{R}^d$.

We will consistently use $\bar{\delta}_r$ to denote the square-free RIP constant and δ_r to denote the standard RIP constant.

Lemma 2.37. *If \mathcal{A} satisfies the square-free RIP with constant $\bar{\delta}$, then \mathcal{A} satisfies the standard RIP with constant $\delta = 2\bar{\delta} + \bar{\delta}^2$. If \mathcal{A} satisfies the standard RIP with constant δ , then \mathcal{A} satisfies the square-free RIP with constant $\bar{\delta}$ for all $\bar{\delta} > 0$ satisfying $\delta \leq 2\bar{\delta} - \bar{\delta}^2$.*

Proof. Assume that \mathcal{A} satisfies the square-free RIP with constant $\bar{\delta}$. We could deduce that \mathcal{A} satisfies the standard RIP with constant δ if $1 - \delta \leq (1 - \bar{\delta})^2$ and $(1 + \bar{\delta})^2 \leq 1 + \delta$. These inequalities are equivalent to $\delta \geq 2\bar{\delta} + \bar{\delta}^2$.

Next assume that \mathcal{A} satisfies the standard RIP with constant δ . We could deduce the square-free RIP with constant $\bar{\delta}$ if $(1 - \bar{\delta})^2 \leq 1 - \delta$ and $(1 + \bar{\delta})^2 \leq 1 + \delta$. These inequalities are equivalent to $\delta \leq 2\bar{\delta} - \bar{\delta}^2$. \square

Alternative versions of the RIP, such as the statistical RIP, have been proposed. In [146], the stability of recovery algorithms under the statistical restricted isometry properties is studied. There are connections between compressed sensing and randomized dimensionality reduction [75].

2.9 Recovery Algorithms in Compressed Sensing

A naive approach to solve (2.17) would be to search for the lowest rank solution, i.e., to solve

$$\min_{x \in \mathbb{R}^d : \mathcal{A}x=y} \|x\|_0, \tag{2.19}$$

in the noise-free case, or

$$\min_{x \in \mathbb{R}^d : \|\mathcal{A}x-y\|_2 \leq \epsilon} \|x\|_0, \tag{2.20}$$

in the noisy case.

Unfortunately, (2.19) and (2.20) are NP-hard [112]. The non-convexity of the rank function creates significant difficulties. The set of vectors of sparsity r in \mathbb{R}^d is the union of $\frac{d!}{r!(d-r)!}$ subspaces, so (2.19) could be solved by searching for solutions to a large number of distinct systems of linear equations, but doing so is impractical for all but very low dimensional problems.

The problems (2.19) and (2.20) can be replaced by their convex relaxations

$$\min_{x \in \mathbb{R}^d : \mathcal{A}x=y} \|x\|_1, \quad (2.21)$$

in the noise-free case, or

$$\min_{x \in \mathbb{R}^d : \|\mathcal{A}x-y\|_2 \leq \epsilon} \|x\|_1, \quad (2.22)$$

in the noisy case. Since the $\|\cdot\|_1$ norm is convex, standard convex optimization and linear programming techniques can be applied.

A fundamental result in CS is that under certain conditions, problems (2.21) and (2.22) exactly or approximately recover the solutions of (2.19) and (2.20). The following theorem was proven by Cohen, Dahmen, and DeVore [43, Theorem 4.3].

Theorem 2.38. *Let $x_0 \in \mathbb{R}^d$, assume that measurements $y = \mathcal{A}(x_0)$ are obtained according to (2.17), and assume that the measurement operator \mathcal{A} satisfies the square-free RIP of order $2r$ with $\delta_{2r} \leq \delta < 1/3$. Then the solution x^* of the minimization problem (2.21) satisfies the error estimate*

$$\|x_0 - x^*\|_1 \leq C \|x_0 - x_{\text{best}}\|_1, \quad (2.23)$$

where x_{best} is the best r -sparse approximation of x_0 in the $\|\cdot\|_1$ norm and $C = \frac{2+2\delta}{1-3\delta}$.

Hence, the RIP of order $2r$ is sufficient to guarantee that $\|\cdot\|_1$ norm minimization produces the best r -sparse approximation, up to a constant factor. If x_0 is r -sparse, the right hand side in (2.23) vanishes, so recovery is exact.

Given a measurement operator $\mathcal{A} : \mathbb{R}^d \rightarrow \mathbb{R}^m$ with associated matrix $A \in \mathbb{R}^{m \times d}$, it is possible to directly check whether the RIP of rank r holds by computing the largest and smallest singular values of every r -column submatrix, i.e., every

submatrix of the form $A[:, \Lambda]$, where $\Lambda \subset \{1, \dots, d\}$ and $|\Lambda| = r$. However, the number of such submatrices is $\frac{d!}{r!(d-r)!}$, so this computation is impractical.

2.10 Incoherence

Since the RIP is difficult to verify deterministically, a large body of research in CS has focused on proving that the RIP holds with high probability for various random measurement operators \mathcal{A} .

Probabilistic recovery guarantees have been proved for random measurement operators \mathcal{A} given by random Fourier submatrices, random Gaussian ensembles (matrices in which each entry is independent, identically distributed (i.i.d.) Gaussian), and Bernoulli ensembles, as described by Candès and Tao in [30]. Rudelson and Vershynin provided recovery proofs for random Fourier and Gaussian measurements [131].

More generally, Baraniuk, Davenport, DeVore, and Wakin proved that the RIP holds with high probability for random Gaussian and Bernoulli measurements [5]. Their arguments elucidate a fundamental connection between the restricted isometry property and covering numbers for finite dimensional unit balls in Euclidean space.

In [110], Mendelson, Pajor, and Tomczak-Jaegermann proved RIP-like results for isotropic sub-Gaussian random measurements. Their argument involved bounding the supremum of a Gaussian process and revealed fundamental connections to the γ_2 functional and the majorizing measures theorem of Guédon, Mendelson, Pajor, and Tomczak-Jaegermann in [70].

Intuitively, recovery of a sparse signal from a small number of measurements is only possible if each measurement provides information about ‘most’ of the entries in the signal. The necessity of such a condition is apparent by considering an operator \mathcal{A} defined by randomly selecting a subset $\Omega \subset \{1, \dots, d\}$ and setting $\mathcal{A}(x) = x[\Omega]$, i.e., only observing the entries $x[i]$ for $i \in \Omega$. If x is r -sparse with $r \ll d$, with very high probability most of the observed entries will be 0, so $\mathcal{A}(x)$ captures almost none of the information in the signal. In contrast, for a random Gaussian ensemble, each measurement is a dense random linear combination of the entries of x , so each measurement can be expected to capture a nontrivial amount of information about the signal.

In [33], Candès and Romberg defined coherence, which quantifies the idea that each measurement provide information about ‘most’ of the signal.

Consider a random measurement operator defined as follows. Let $A \in \mathbb{R}^{d \times d}$ be an orthonormal matrix. Choose a subset $\Omega \subset \{1, \dots, d\}$, and define

$$\mathcal{A}(x) = A[\Omega, :]x, \tag{2.24}$$

where $A[\Omega, :]$ is the submatrix of A consisting of the rows of A indexed by Ω . Hence, if a_1^t, \dots, a_d^t are the rows of A , the measurements take the form of inner products $\langle a_i, x \rangle$ for $i \in \Omega$.

The following definition is from [33].

Definition 2.39 (Mutual coherence).

- *The mutual coherence of an orthonormal matrix $A \in \mathbb{R}^{d \times d}$ with respect to the*

standard basis is

$$\nu(A) = \max_{i,j} |A[i, j]|$$

- If $A, B \in \mathbb{R}^{d \times d}$ are two orthonormal matrices, the mutual coherence of A and B is

$$\nu(A, B) = \max_{i,j} |\langle a_i, b_j \rangle|, \quad (2.25)$$

where a_i and b_j are the rows of A and B , viewed as column-vectors.

Observe that if $B = \text{Id}_{d \times d}$, the second definition reduces to the first. Mutual coherence measures the degree of similarity between the two bases. Since A and B are orthonormal, mutual coherence always satisfies $\nu \geq 1/\sqrt{d}$. Values of ν near $1/\sqrt{d}$ indicate that measurements of the form $\langle a_i, x \rangle$ capture robust information about the representation of x with respect to the basis b_1, \dots, b_d . The discrete Fourier transform matrix $\mathcal{F} \in \mathbb{C}^{d \times d}$ has coherence $1/\sqrt{d}$ while the identity matrix $\text{Id}_{d \times d}$ has coherence 1.

In [117], Rauhut, Schnass, and Vendergheynst extended compressed sensing results to the case in which signals are sparse with respect to redundant dictionaries, not just orthonormal bases.

2.11 Probability Theory

The following are standard definitions from probability theory [88].

Definition 2.40. *Let Ω be a set. A collection \mathcal{F} of subsets of Ω is a σ -algebra if the following holds:*

- $\Omega \in \mathcal{F}$.
- If $F \in \mathcal{F}$, then $\Omega \setminus F \in \mathcal{F}$, where $A \setminus B$ denotes the points in A that are not in B .
- If $\{F_i\}_{i=1}^{\infty} \subset \mathcal{F}$, then $\cup_{i=1}^{\infty} F_i \in \mathcal{F}$, where \cup denotes the set union.

Definition 2.41 (Probability measure space). A probability measure space is a triple (Ω, \mathcal{F}, P) where Ω is a set, \mathcal{F} is a σ -algebra on Ω , and $P : \mathcal{F} \rightarrow [0, 1]$ is a function satisfying:

- $P(\Omega) = 1$.
- If $F_1, F_2, \dots \in \mathcal{F}$ are disjoint, then $P(\cup_{i=1}^{\infty} F_i) = \sum_{i=1}^{\infty} P(F_i)$.

Definition 2.42 (Random variable). Let D be a subset of \mathbb{R}^d , for some $d \geq 1$. A random variable (RV) on a probability space (Ω, \mathcal{F}, P) with values in D is any function $f : \Omega \rightarrow D$ that is measurable with respect to \mathcal{F} . In other words, for all open sets $U \subset D$, the set $f^{-1}(U) \in \mathcal{F}$, where $f^{-1}(U) = \{\omega \in \Omega : f(\omega) \in U\}$.

When there is no ambiguity, we will describe a random variable as taking values in a given set, without explicitly identifying the measure space (Ω, \mathcal{F}, P) .

Definition 2.43 (Expectation). Let (Ω, \mathcal{F}, P) be a probability measure space and let $x : \Omega \rightarrow \mathbb{R}$ be a real-valued random variable. The expectation of x is defined to be

$$\mathbb{E}[x] = \int_{\Omega} xP,$$

where the integral is a Lebesgue integral with respect to the measure P .

The expectation of a random variable is in general not guaranteed to exist.

Definition 2.44 (Rademacher random variable). *A Rademacher random variable is a random variable ϵ with values in $\{-1, 1\}$ that satisfies $P(\epsilon = 1) = 1/2$ and $P(\epsilon = -1) = 1/2$.*

2.12 Structured Random Measurement Operators

In the book [60], Foucart and Rauhut show that the RIP holds with high probability for measurement operators with independent, isotropic, sub-Gaussian rows, a significant generalization of the setting of orthonormal measurements or Gaussian or Bernoulli ensembles. In this setting, each row of the measurement matrix A is an iid drawing from a random vector with values in \mathbb{R}^d ; however, the entries in any given row need not be independent.

The following definitions appear in [57].

Definition 2.45 (sub-Gaussian random variable). *A real-valued random variable x is sub-Gaussian with constant c if for all $p \geq 1$,*

$$(\mathbb{E}|x|^p)^{1/p} \leq c\sqrt{p}.$$

Definition 2.46. *If x is a sub-Gaussian random variable, its sub-Gaussian norm, denoted $\|x\|_{\psi_2}$, is defined to be the smallest constant c satisfying $(\mathbb{E}|x|^p)^{1/p} \leq c\sqrt{p}$ for all $p \geq 1$.*

The following appears in [57, Example 5.8].

Lemma 2.47. *If x is a bounded real-valued random variable satisfying $\mathbb{E}[x] = 0$ and $|x| \leq c$, then x is sub-Gaussian with $\|x\|_{\psi_2} \leq c$.*

The following definition appears in [57].

Definition 2.48 (sub-Gaussian and isotropic random vectors). *Let x be a random vector taking values in \mathbb{R}^d .*

- *x is sub-Gaussian with constant c if for all $z \in \mathbb{R}^d$ with $\|z\|_2 = 1$, the random variable $\langle x, z \rangle$ is sub-Gaussian with $\|\langle x, z \rangle\|_{\psi_2} \leq c$. The sub-Gaussian norm of x is defined to be*

$$\|x\|_{\psi_2} = \sup_{z \in \mathbb{R}^d : \|z\|_2=1} \|\langle x, z \rangle\|_{\psi_2}.$$

- *x is **isotropic** if for all $z \in \mathbb{R}^d$,*

$$\mathbb{E} [|\langle x, z \rangle|^2] = \|z\|_2^2.$$

The following is [60, Theorem 9.6]

Theorem 2.49. *Let A be a random $m \times d$ matrix with independent, isotropic, and sub-Gaussian rows with sub-Gaussian parameter c . There exists a constant $C > 0$ depending only on the sub-Gaussian constant c such that if the number of measurements satisfies*

$$m \geq C\delta^{-2} \left(r \log \left(\frac{ed}{r} \right) + \log \left(\frac{2}{\epsilon} \right) \right)$$

then the measurement operator given by

$$\mathcal{A}(x) = \frac{1}{\sqrt{m}} Ax$$

satisfies the RIP of rank r with constant $\delta_r \leq \delta$ with probability at least $1 - \epsilon$.

The following example is well known; see for example [60].

Example 2.50 (Bounded orthonormal system). *Let U be an $n \times d$ matrix with orthonormal columns. Let A be an $m \times d$ random matrix with iid rows of the form $\sqrt{n}\epsilon_i u_i$, where u_1, \dots, u_m are m rows of U , chosen uniformly at random with replacement and $\epsilon_1, \dots, \epsilon_m$ are iid Rademacher random variables with values in $\{-1, 1\}$. Then A is a random matrix with independent, isotropic, sub-Gaussian rows. Moreover, the sub-Gaussian constant is given by*

$$c = \max_{i=1, \dots, n} \sqrt{n} \|U[i, :]\|_2.$$

Proof. By construction, the columns of A are iid.

Let x be the random vector obtained by drawing a row uniformly at random from the rows U . To prove that $z = \sqrt{n}\epsilon x$ is sub-Gaussian, observe that because of the Rademacher variable ϵ , we have $\mathbb{E}[z] = 0$. Also, we have $\|z\|_2 \leq c$ by the definition of c . Hence by Lemma 2.47, z is sub-Gaussian.

To prove that z is isotropic, observe that for any $w \in \mathbb{R}^d$,

$$\begin{aligned} \mathbb{E}[|\langle z, w \rangle|^2] &= \frac{1}{n} \sum_{i=1}^n |\sqrt{n}\epsilon_i U[i, :]\langle w \rangle|^2 \\ &= \|Uw\|_2^2 \\ &= \|w\|_2^2 \end{aligned} \tag{2.26}$$

where (2.26) follows because U has orthonormal columns.

□

The assumption that each vector is multiplied by a Rademacher random variable does not create any loss of generality. Indeed, given observations without

the Rademacher random variables, one could reduce to the above case by multiplying the rows of A and the observations y by the Rademacher random variables.

Since the constant appearing in Theorem 2.49 depends on the sub-Gaussian constant c , the sub-Gaussian constant affects the number of measurements required to guarantee RIP with high probability. In the above example, the columns of U form a Parseval tight frame for \mathbb{R}^d . The sub-Gaussian constant c can be interpreted as a measure of coherence, as will be discussed in Chapter 3.

2.13 Matrix Completion

Initial results in CS sensing sought to recover discrete signals $x \in \mathbb{R}^d$. Although such an x is a d -dimensional vector, it can also be thought of as 1-dimensional, in the sense that its entries $x[i]$ are indexed by a single integer i . Similarly, while the set of matrices $\mathbb{R}^{d_1 \times d_2}$ is a $(d_1 d_2)$ -dimensional vector space, matrices $X \in \mathbb{R}^{d_1 \times d_2}$ can be thought of as a 2-dimensional, since its entries $X[i_1, i_2]$ are indexed by two integers i_1 and i_2 .

The results of CS can be directly applied to recover sparse matrices $X \in \mathbb{R}^{d_1 \times d_2}$. Indeed, given a linear measurement operator $\mathcal{A} : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^m$, one could reformulate the problem in terms of standard CS recovery by vectorizing the matrices and the corresponding operator. Indeed, after defining $x = \text{vec}(X)$ and $\tilde{\mathcal{A}}(x) = \mathcal{A}(\text{refold}_{(d_1, d_2)}(x))$, one could directly apply CS with the operator $\tilde{\mathcal{A}}$ to recover x . After recovering the vectorization x , one recovers X by $X = \text{refold}_{(d_1, d_2)}(x)$. Unfortunately, this procedure ignores the substantial additional structure in matri-

ces, or 2-dimensional arrays.

Often, matrices have additional structure that can only be described by using the full matrix structure, such as low rank or approximate low rank properties. Such matrices arise in numerous applications; one example is recommender systems. A famous problem in this area is the Netflix problem [9], in which a matrix is used to store user preferences for movies, with rows corresponding to users and columns to movies. Since it can be expected that only a relatively small number of factors determine a user’s preferences for movies, the resulting matrix can be expected to be approximately low-rank.

Consider the general matrix completion problem in which $X_0 \in \mathbb{R}^{d_1 \times d_2}$ is a low-rank matrix to be recovered, $\mathcal{A} : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^m$ is a linear measurement operator, the observed data $y \in \mathbb{R}^m$ is given by

$$y = \mathcal{A}(X_0) + z, \tag{2.27}$$

and $z \in \mathbb{R}^m$ is a noise vector satisfying $\|z\|_2 \leq \epsilon$. In the noise-free case, we assume $z = 0$. The goal is to recover X_0 from the observations y .

In the simplest case, the measurement operator \mathcal{A} is defined by directly observing a subset of the entries of X . In this case, let $\Omega \subset [d_1] \times [d_2]$ be a subset of the indices of X chosen uniformly at random and satisfying $|\Omega| = m$. Define

$$\mathcal{A}_\Omega(X) = X[\Omega],$$

where $X[\Omega]$ denotes the entries $X[i_1, i_2]$ for which $(i_1, i_2) \in \Omega$, arranged into a column-vector in \mathbb{R}^m according to the lexicographical ordering on Ω inherited from $[d_1] \times [d_2]$.

Candès and Recht showed in [32] that (2.27) can be solved, in the noise-free case with direct observations $\mathcal{A} = \mathcal{A}_\Omega$, via a convex relaxation similar to the l_1 relaxation in CS. A naive approach would be to solve

$$\min_{X : X[\Omega]=y} \text{rank}(X) \quad (2.28)$$

Unfortunately, as for l_0 minimization in the vector case, rank-minimization is NP-hard. In the case of observations of matrix entries, Candès and Recht suggested the following convex relaxation:

$$\min_{X : X[\Omega]=y} \|X\|_*, \quad (2.29)$$

where the nuclear norm is defined by $\|X\|_* = \sum_{i=1}^{\text{rank}(X)} \sigma_i(X)$.

As in 1-dimensional CS, not all low-rank matrices are recoverable from random observations. Indeed, consider the low rank matrix $e_1 e_1^t \in \mathbb{R}^{d \times d}$, with all entries 0 except for a 1 in position (1, 1). Unless the number of observations m is on the order of d^2 , with high probability all the observations will miss the nonzero entry, resulting in $y = 0$. In [32], Candès and Recht proved that the solutions of (2.28) and (2.29) are equal with high probability over the choice of Ω , provided that the singular vectors of X_0 satisfy certain incoherence properties.

In [68], Gross significantly generalized previous results in matrix completion, showing that a $d \times d$ matrix with rank r can be recovered from $O(dr\mu_{\text{Gross}} \log^2(d))$ measurements drawn randomly from any basis of $\mathbb{R}^{d \times d}$, where μ_{Gross} is defined below in Definition 2.51.

Gross considered the following random measurement operator. Let A_1, \dots, A_{d^2} be an orthonormal basis for $\mathbb{R}^{d \times d}$ with respect to the inner product $\langle X, Y \rangle =$

$\langle \text{vec}(X), \text{vec}(Y) \rangle$. For a subset $\Omega \subset [d^2]$ satisfying $|\Omega| = m$ and chosen uniformly at random, define the operator $\mathcal{A}_\Omega : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^m$ by letting

$$\mathcal{A}_\Omega(X)[i] = (X, A_{\omega_i}), \quad (2.30)$$

where $\omega_1, \dots, \omega_m$ are the entries of Ω ordered lexicographically.

The following definitions are from [68]. Let X_0 be the matrix to be recovered and assume $X_0 = U_0 S_0 V_0^t$ is the SVD of X_0 . Define

$$T = \{X | (\text{Id}_{d \times d} - U_0 U_0^t) X (\text{Id}_{d \times d} - V_0 V_0^t) = 0\}$$

and define

$$\mathcal{P}_T(X) = U_0 U_0^t X + X V_0 V_0^t - U_0 U_0^t X V_0 V_0^t,$$

so that \mathcal{P}_T defines an orthogonal projection onto T . The space T and the projection \mathcal{P}_T depend on the true matrix X_0 ; however, the dependence is suppressed from the notation.

For $x \in \mathbb{R}$ define $\text{sgn}(x) = x/|x|$ for $x \neq 0$ and $\text{sgn}(0) = 0$. Similarly, if X is a matrix, define $\text{sgn}(X) = U \text{sgn}(S) V^t$, where $X = U S V^t$ is the SVD of X and the sgn function is applied to the singular values of X .

The following definition is from [68].

Definition 2.51 (Gross coherence). *A $d \times d$ matrix X_0 of rank r has Gross-coherence μ_{Gross} with respect to a basis A_1, \dots, A_{d^2} of $\mathbb{R}^{d \times d}$ if either*

$$\max_{i=1, \dots, d^2} d \|A_i\|_2^2 \leq \mu_{\text{Gross}}$$

or both of the following hold:

$$\max_{i=1, \dots, d^2} d \|\mathcal{P}_T(X)\|_2 \leq 2r \mu_{\text{Gross}}$$

and

$$\max_{i=1,\dots,d^2} d^2(X, \text{sgn}(X_0))^2 \leq 2r\mu_{\text{Gross}},$$

where $r = \text{rank}(X_0)$.

The following recovery theorem applies to self-adjoint matrices. However, by a simple argument any matrix completion problem can be reduced to the self-adjoint case [68].

Theorem 2.52. ([68, Theorem 3]) *Let $X_0 \in \mathbb{R}^{d \times d}$ be a self-adjoint matrix of rank r with Gross-coherence μ_{Gross} with respect to a basis A_1, \dots, A_{d^2} of $\mathbb{R}^{d \times d}$. Let \mathcal{A}_Ω be defined as in (2.30) and assume that noise-free observations of the form*

$$y = \mathcal{A}_\Omega(X_0)$$

are obtained. There exists a constant C such that if the number of measurements $m \geq Cdr\mu_{\text{Gross}}(1 + \beta) \log^2(d)$, then the solution X^ of the problem*

$$\min_{X : \mathcal{A}_\Omega(X)=y} \|X\|_{S_1}$$

is unique and equal to X_0 with probability at least $1 - d^{-\beta}$.

This recovery guarantee applies in very general settings, since there is no requirement that the basis A_1, \dots, A_{d^2} be orthonormal.

2.13.1 Matrix Restricted Isometry Property

Definition 2.53 (square-free RIP for matrix operators). *Let $\mathcal{A} : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^m$ be a measurement operator. Let $U \subset \mathbb{R}^{d_1 \times d_2}$ be a fixed set. Then \mathcal{A} satisfies the*

square-free RIP with constant $\bar{\delta}$ over the set U if for all matrices $X \in U$, we have

$$(1 - \bar{\delta}_r)\|X\|_F \leq \|\mathcal{A}(X)\|_2 \leq (1 + \bar{\delta})\|X\|_F.$$

Yi-Kai Liu proved that the square-free RIP holds for random matrix measurements of the same form as (2.30) with respect to an orthonormal basis A_1, \dots, A_{d^2} for $\mathbb{C}^{d \times d}$. The following theorem is from [102, Theorem 2.1].

Theorem 2.54. *Let $0 \leq \delta < 1$ and let A_1, \dots, A_{d^2} be an orthonormal basis of $\mathbb{R}^{d \times d}$ satisfying the first part of the definition of Gross-coherence (2.51), i.e.*

$$\max_{i=1, \dots, d^2} d\|A_i\|_2^2 \leq \mu.$$

Let \mathcal{A} be defined as in (2.30). There exists a constant C such that if the number of measurements satisfies

$$m \geq C\mu r d \log^6(d)/\delta^2$$

then the operator \mathcal{A}_Ω satisfies the RIP with constant δ over the set of matrices $\{X : \|X\|_ \leq \sqrt{r}\|X\|_F\}$ with probability at least $1 - \exp(-C)$.*

Observe that the set $\{X : \|X\|_* \leq \sqrt{r}\|X\|_F\}$ contains all matrices $\mathbf{X} \in \mathbb{R}^{d_1 \times d_2}$ of rank at most r , so Liu's result implies that with high probability the square-free RIP holds over the set of matrices of rank at most r .

2.13.2 Singular Value Thresholding (SVT)

In [26], Cai, Candès, and Shen developed a simple algorithm to solve the problem

$$\min_{X : \mathcal{A}(X)=y} \|X\|_* \tag{2.31}$$

Definition 2.55 (Shrinkage). Let $\tau \geq 0$. For $X \in \mathbb{R}^{d_1 \times d_2}$ with SVD given by $X = USV^t$, define

$$\text{shrink}_\tau(X) = U \max(0, S - \tau)V^t,$$

where the subtraction and maximum operations are applied pointwise to the singular values of X .

For a fixed $\tau > 0$, consider the modified problem

$$\min_{X : \mathcal{A}(X)=y} \tau \|X\|_* + \frac{1}{2} \|X\|_{\mathbb{F}}^2 \quad (2.32)$$

It is shown in [26] that the solution of (2.32) converges to the solution of (2.31) as $\tau \rightarrow \infty$. For a fixed value of $\tau > 0$, the SVT algorithm solves (2.32).

Algorithm 2.56 ([26]). 1: *procedure* SVT($X_0, \mathcal{A}, y, \{\delta_i\}_{i \geq 1}, \tau$)

2: *while not converged do*

3: $X_k = \text{shrink}_\tau(\mathcal{A}^*(z_{k-1}))$

4: $z_k = z_{k-1} + \delta_k(y - \mathcal{A}(X_k))$

In [26], convergence of the above algorithm is proved under assumptions on the step sizes δ_k . The following lemma is important in proving the correctness of singular value thresholding algorithms.

The following is [26, Theorem 2.1].

Lemma 2.57. Let $X_0 \in \mathbb{R}^{d_1 \times d_2}$ and let $\tau \geq 0$. The solution of

$$\min_{X \in \mathbb{R}^{d_1 \times d_2}} \tau \|X\|_* + \frac{1}{2} \|X - X_0\|_{\mathbb{F}}^2$$

is given by $\text{shrink}_\tau(X_0)$.

2.13.3 Fixed Point Continuation

An alternative, but related, algorithm is fixed point continuation, as developed by Ma, Goldfarb, and Chen in [106]. The algorithm works by solving a fixed point equation that holds at the optimal solution of

$$\|X\|_* + \frac{1}{2\mu} \|\mathcal{A}(X) - y\|_2^2.$$

Fix a decreasing sequence of parameters $\mu_{\text{init}} = \mu_1 > \mu_2 > \mu_3 \cdots > \mu_l = \mu_{\text{final}} > 0$ and initialize $X = X_0$.

Algorithm 2.58 ([106]). 1: *procedure* FPC($\mathbf{X}_0, \mathcal{A}, y, \{\mu_i\}_{i=1,\dots,l}$)

2: *while not converged do*

3: *Select* τ .

4: $Y = X - \tau \mathcal{A}^*(\mathcal{A}(X) - y)$

5: $X_k = \text{shrink}_{\tau\mu}(Y)$

2.14 Introduction to Tensor Completion

Consider the following general tensor completion problem. Let $\mathcal{A} : \mathbb{R}^{d_1 \times \cdots \times d_N} \rightarrow \mathbb{R}^m$ be a linear measurement operator and let $\mathbf{X}_0 \in \mathbb{R}^{d_1 \times \cdots \times d_N}$ be an unknown tensor. We aim to recovery \mathbf{X}_0 from observations of the form

$$y = \mathcal{A}(\mathbf{X}_0) + z,$$

where $z \in \mathbb{R}^m$ is a noise vector. In the literature, various operators \mathcal{A} have been considered. One special case is when $\Omega \subset [d_1] \times \cdots \times [d_N]$ is a random subset of the

indices of \mathbf{X} of fixed cardinality m , and the operator \mathcal{A} is given by

$$\mathcal{A}(\mathbf{X}) = \mathbf{X}[\Omega].$$

Recall that $\mathbf{X}[\Omega]$ denotes the vector in \mathbb{R}^m obtained by listing the entries $\mathbf{X}[i_1, \dots, i_N]$ for which $(i_1, \dots, i_N) \in \Omega$ in lexicographical order.

In [98], Ji Liu, Musialski, Wonka, and Ye proposed solving the following convex optimization problem in the case of direct observations $\mathbf{X}[\Omega]$:

$$\min_{\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N} : \mathbf{X}[\Omega] = y} \sum_{k=1}^N \|\mathbf{X}^{(i)}\|_*,$$

where y is the vector of observations of $\mathbf{X}[\Omega]$ and $\mathbf{X}^{(i)}$ is the i -th unfolding of \mathbf{X} . Liu et al. introduced variables Z_i to represent the unfoldings $\mathbf{X}^{(i)}$ to reduce the interdependence between the summands in the objective function. The resulting problem, after replacing the constraints $Z_i = \mathbf{X}^{(i)}$ with penalty terms in the objective function, is:

$$\min_{\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}, Z_1, \dots, Z_N} \left(\sum_{k=1}^N \|Z_i\|_* + \frac{1}{2\mu} \sum_{k=1}^N \|\mathbf{X}^{(i)} - Z_i\|_F^2 \right)$$

Sum of nuclear norms minimization is a standard approach for tensor completion and has been studied by Tomioka et al. [139, 140, 141], Signoretto, Lathauwer, and Suykens [134], Signoretto, Plas, Moor, and Suykens [135], Gandy, Recht, and Yamada [64], Romera-Paredes and Pontil [130], Yang, Huang, and Shi [148], Al-Qizwini and Radha [1], and Cao et al. [37].

One potential drawback of minimizing the sum of nuclear norms of unfoldings is the loss of additional higher-order structure [111]. Mu, Huang, Wright, and Goldfarb proposed an alternative strategy to preserve additional structure in the

case of tensors of dimension $N \geq 4$, unfolding a tensor by grouping its axes into two approximately equally sized sets, and unfolding the tensor into a maximally square-shaped matrix [111].

In the case of a matrices $X \in \mathbb{R}^{d_1 \times d_2}$, the operator norm $\|X\|_2$ is dual to the nuclear norm $\|X\|_*$, in the sense that

$$\|X\|_* = \sup_{Y \in \mathbb{R}^{d_1 \times d_2} : \|Y\|_2 \leq 1} \langle X, Y \rangle.$$

where recall that $\langle X, Y \rangle$ denotes the Euclidean inner product of $\text{vec}(X)$ and $\text{vec}(Y)$.

Based on this duality, a natural generalization of the nuclear norm to tensors $\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}$ can be defined as follows. First, for $\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}$, define the operator norm by

$$\|\mathbf{X}\|_{\text{op}} = \max_{u_i \in \mathbb{R}^{d_i} : \|u_i\|_2 \leq 1} \langle \mathbf{X}, u_1 \circ u_2 \circ \dots \circ u_N \rangle,$$

where \circ denotes outer product. Then, the duality-based tensor nuclear norm can be defined by

$$\|\mathbf{X}\|_{\text{duality},*} = \sup_{\mathbf{Y} \in \mathbb{R}^{d_1 \times \dots \times d_N} : \|\mathbf{Y}\|_{\text{op}} \leq 1} \langle \mathbf{X}, \mathbf{Y} \rangle.$$

In the case of 3-dimensional tensors, Yuan and Zhang [149] proved a recovery guarantee for the recovery of incoherent tensors via the minimization of $\|\mathbf{X}\|_{\text{duality},*}$. Unfortunately, it is computationally intractable to minimize $\|\mathbf{X}\|_{\text{duality},*}$ in practice, since doing so is NP-hard, as shown by Friedland and Lim [63].

As a consequence of the intractability of minimizing the norm $\|\mathbf{X}\|_{\text{duality},*}$, a number of researchers have attempted to replace the norm $\|\mathbf{X}\|_{\text{duality},*}$ with convex relaxations that enable tractable minimization. In [120], Rauhut and Stojanac de-

veloped a family of convex relaxations of $\|\mathbf{X}\|_{\text{duality},*}$ in the case of 3-dimensional tensors, called θ -bodies, using results from algebraic geometry.

In [122, 119], Rauhut, Schneider, and Stojanac developed an algorithm called *iterative hard threshold* (IHT). The algorithm is the following:

Algorithm 2.59 ([119]). 1: *procedure* IHT($\mathbf{X}_0, \mathcal{A}, y, \{\mu_i\}_{i \geq 0}, (r_1, \dots, r_N)$)

2: *while not converged do*

3: $Y = X - \mu_j \mathcal{A}^*(\mathcal{A}(X) - y)$

4: $X_k = \mathcal{H}_{(r_1, \dots, r_N)}(Y)$

In the above, the operator $\mathcal{H}_{(r_1, \dots, r_N)} : \mathbb{R}^{d_1 \times \dots \times d_N} \rightarrow \mathbb{R}^{d_1 \times \dots \times d_N}$ acts by computing the HOSVD of \mathbf{Y} , $\mathbf{Y} = \mathbf{C} \otimes_1 U_1 \cdots \otimes_N U_N$ and then replacing \mathbf{C} with $\mathbf{C} \begin{bmatrix} r_1 \\ \vdots \\ r_N \end{bmatrix}$ and replacing each U_i with $U_i[:, 1 : r_i]$ for $i = 1, \dots, N$. Under restricted isometry and other assumptions, Rauhut et al. proved convergence of IHT.

Krishnamurthy and Singh have studied adaptive sampling for low Tucker rank tensor completion [89] in the case in which the Tucker factors have low coherence. Krishnamurthy et al. show that their algorithm recovers a low Tucker rank tensor \mathbf{X} with high probability under incoherence assumptions. Their algorithm, however, differs significantly from the non-uniform sampling we introduce in Chapter 3, since they are not solving a nuclear norm minimization problem and are considering only noise-free direct observations of tensor entries, as in the operator $\mathcal{A}(\mathbf{X}) = \mathbf{X}[\Omega]$ discussed earlier.

Other tensor completion and tensor decomposition approaches include Bayesian

methods, as in the work of Zhao, Zhang, and Cichocki [151], Zhao, Zhou, Zhang, Cichoki, and Amari [152], and Bazerque, Mateos, and Giannakis [8]. These methods impose probabilistic prior assumptions on the Tucker or CP factorization of a tensor to predict its true value.

Extensions of compressed sensing to multiple dimensions have also been developed, as in the work of Duarte and Baraniuk [56], Qun Li, Schonfeld, and Friedland [95, 62], and Caiafa and Cichoki [28]. An overview of multidimensional compressed sensing with applications is found in [27].

Chapter 3: Tensor Completion and Recovery Guarantees

3.1 Overview

Given the success of nuclear norm minimization for low-rank matrix completion, a natural extension to the tensor case is obtained by defining the nuclear norm of a tensor to be the sum of the nuclear norms of its unfoldings. This approach, first suggested in [98] and subsequently studied in [141, 111, 133] has proven highly successful in applications. Since the unfolding operation loses some of the tensor structure, this approach has suboptimal recovery guarantees in the case of Gaussian measurement ensembles, in comparison to using a true tensor nuclear norm [111, 149] incorporating the full structure; unfortunately, it is NP-hard to minimize the true tensor nuclear norm when $N \geq 3$ [79].

Recall that the nuclear norm $\|X\|_*$ of a matrix is defined to be the sum of its singular values. We consider the following standard extension of the nuclear norm to tensors via unfolding:

Definition 3.1 (Nuclear norm of tensor). *Given a tensor $\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}$, its nuclear norm,*

$$\|\mathbf{X}\|_* = \frac{1}{N} \sum_{k=1}^N \|\mathbf{X}^{(k)}\|_*, \quad (3.1)$$

is defined to be the average of the nuclear norms of all its mode- k unfoldings, $k = 1, \dots, N$.

We consider the following general setting for approximately low-rank tensor completion. Let $\mathbf{X}_0 \in \mathbb{R}^{d_1 \times \dots \times d_N}$ be the tensor we wish to recover, where \mathbf{X}_0 need not be low-rank. Assume that $\mathcal{A} : \mathbb{R}^{d_1 \times \dots \times d_N} \rightarrow \mathbb{R}^m$ is a given linear measurement operator. We are given observations of the form

$$y = \mathcal{A}(\mathbf{X}_0) + z, \quad (3.2)$$

where z is a noise vector satisfying

$$\|z\|_2 < \epsilon \quad (3.3)$$

for some $\epsilon > 0$.

We aim to recover \mathbf{X}_0 by solving the constrained convex optimization problem:

$$\min_{\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N} : \|\mathcal{A}(\mathbf{X}) - y\|_2 \leq \epsilon} \|\mathbf{X}\|_*. \quad (3.4)$$

This problem can be reformulated to its unconstrained Lagrangian form:

$$\min_{\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}} \|\mathbf{X}\|_* + \frac{1}{2\mu} \|\mathcal{A}(\mathbf{X}) - y\|_2^2. \quad (3.5)$$

The regularization parameter μ determines how much to weight the nuclear norm. Larger values of μ tend to produce lower-rank solutions, as measured by $\|\mathbf{X}\|_*$, with larger residuals.

3.2 Recovery Guarantees via Restricted Isometry Property (RIP)

In compressed sensing and in low rank matrix completion, an important ingredient of many recovery guarantees is the restricted isometry property, which essentially guarantees that \mathcal{A} acts as an approximate isometry on sparse vectors or low-rank matrices. Several authors have extended the definition of RIP to the tensor case. In [133] and [122], the RIP is defined using Tucker rank, while in [118] a version of RIP is defined for hierarchical tensor decompositions.

We adopt the definition in [133], based on Tucker rank.

Definition 3.2 ([133]). *The measurement operator \mathcal{A} satisfies the square-free restricted isometry property (RIP) of Tucker rank (r_1, \dots, r_N) with constant $\bar{\delta}_{(r_1, \dots, r_N)}$ if for all tensors $\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}$ with $\text{rank}(\mathbf{X}) \leq (r_1, \dots, r_N)$, the following bound holds:*

$$(1 - \bar{\delta}_{(r_1, \dots, r_N)}) \|\mathbf{X}\|_F \leq \|\mathcal{A}(\mathbf{X})\|_2 \leq (1 + \bar{\delta}_{(r_1, \dots, r_N)}) \|\mathbf{X}\|_2$$

Definition 3.3. *\mathcal{A} satisfies the square-free strong RIP of rank (r_1, \dots, r_N) with constant $\bar{\delta}_{(r_1, \dots, r_N)}$ if, for each $k = 1, \dots, N$, \mathcal{A} satisfies the RIP of rank $(d_1, \dots, d_{k-1}, r_k, d_{k+1}, \dots, d_N)$ with constant $\bar{\delta}_{(r_1, \dots, r_N)}$.*

The RIP and strong RIP can also be defined with squared norms; in that case, the RIP constant is denoted $\delta_{(r_1, \dots, r_N)}$.

In [133], an RIP-based recovery guarantee is developed for (3.4) under the assumption of noise-free measurements, i.e. with $z = 0$. However, as far as we are aware, there is currently no known RIP-based measurement guarantee for (3.4)

under inexact or noisy measurements. A guarantee has been shown for 3.5 in [141] under a different assumption, called *restricted strong convexity* (RSC); however, the RSC assumption is data dependent. In other words, whether or not the guarantee is applicable is determined by the true tensor \mathbf{X}_0 , which is unknown. In contrast, our RIP-based guarantee will hold for arbitrary tensors \mathbf{X}_0 .

We now state our recovery guarantee for (3.4), which is an extension of Theorem 4 in [59] from the matrix case to the tensor case.

Theorem 3.4 (Tensor RIP recovery guarantee from noisy measurements). *Assume that the linear operator \mathcal{A} satisfies the square-free strong RIP of rank $(5r_1, \dots, 5r_N)$ with RIP constants $\bar{\delta}_{(5r_1, \dots, 5r_N)} < \bar{\delta}_5$ and $\bar{\delta}_{(3r_1, \dots, 3r_N)} < \bar{\delta}_3$, where $\bar{\delta}_5 \geq 0$, $\bar{\delta}_3 \geq 0$, and $\bar{\delta}_5 + (1 + \bar{\delta}_3)\sqrt{\frac{2}{3}} < 1$. Let $\epsilon > 0$. Then, for all $\mathbf{X}_0 \in \mathbb{R}^{d_1 \times \dots \times d_N}$ and $y \in \mathbb{R}^m$ satisfying $\|\mathcal{A}(\mathbf{X}_0) - y\|_2 \leq \epsilon$, the solution \mathbf{X}^* of (3.4) satisfies the error estimate*

$$\|\mathbf{X}^* - \mathbf{X}_0\|_F \leq C_0 \epsilon + C_1 \left(\frac{1}{\frac{1}{N} \sum_{k=1}^N \sqrt{r_k}} \right) \|\mathbf{X}_0 - \mathbf{X}_{\text{best}}\|_*$$

where \mathbf{X}_{best} is the best approximation to \mathbf{X}_0 in the $\|\cdot\|_*$ norm among all tensors \mathbf{X} satisfying $\text{rank}(\mathbf{X}) \leq (r_1, \dots, r_N)$. The constants are given by

$$C_0 = \frac{2 \left(1 + \sqrt{\frac{2}{3}} \right)}{1 - \bar{\delta}_5 - (1 + \bar{\delta}_3)\sqrt{\frac{2}{3}}}$$

and

$$C_1 = \frac{2}{\sqrt{3}} \left(\frac{2 - \bar{\delta}_5 + \bar{\delta}_3}{1 - \bar{\delta}_5 - (1 + \bar{\delta}_3)\sqrt{\frac{2}{3}}} \right)$$

Remark 3.5. *This condition of Theorem 3.4 can be simplified slightly. Indeed, the inequality $\bar{\delta}_{(d_1, \dots, d_{k-1}, 3r_k, d_{k+1}, \dots, d_N)} \leq \bar{\delta}_{(d_1, \dots, d_{k-1}, 5r_k, d_{k+1}, \dots, d_N)}$ follows from the definition*

of RIP, so it is always possible to choose $\bar{\delta}_3 \leq \bar{\delta}_5$ (by decreasing $\bar{\delta}_3$ if necessary).

Hence, the condition $\bar{\delta}_5 + (1 + \bar{\delta}_3)\sqrt{\frac{2}{3}} < 1$ is implied by

$$\bar{\delta}_5 < \bar{\delta}_{\text{critical}} =: \frac{1 - \sqrt{\frac{2}{3}}}{1 + \sqrt{\frac{2}{3}}} \approx 0.101. \quad (3.6)$$

The proof of Theorem 3.4 uses a similar argument to the one in [59] for the matrix case. We use the construction from [59] separately along each axis of the tensor and then compute a weighted average of the results, with weights given by $\sqrt{r_k}/N$.

Before starting the proof, we briefly outline the strategy used in [59] for matrices. First, the difference between the solution to problem (3.4) and the true matrix, $E = X^* - X_0$, is decomposed into the sum of two matrices, $E = E_0 + E_c$ with certain orthogonality properties, where E_0 is rank at most $2r$. The second matrix, E_c , is split up into a finite sum of rank $3r$ matrices, $E_c = E_1 + E_2 + \dots$, of descending norm. Finally, the norm of $E_0 + E_1$ and the norm of $E_2 + E_3 + \dots$ are each approximated. The RIP properties are required because $E_0 + E_1$ is of rank at most $5r$ while each E_i is of rank at most $3r$. Our proof essentially repeats this construction individually along each axis $k = 1, \dots, N$, and combines the results.

We now state two useful lemmas. The following is [123, Lemma 2.3].

Lemma 3.6. *If X, Y are matrices of size $d_1 \times d_2$ satisfying $XY^t = 0$ and $X^tY = 0$, then*

$$\|X + Y\|_* = \|X\|_* + \|Y\|_*.$$

The following is [123, Lemma 3.4].

Lemma 3.7. *Let X, Y be matrices of size $d_1 \times d_2$. Assume that the SVD decomposition of X is $X = USV^t$. Define*

$$Y_2 = (\text{Id}_{d_1 \times d_1} - UU^t)Y(\text{Id}_{d_2 \times d_2} - VV^t)$$

and

$$Y_1 = Y - Y_2$$

Then the following holds:

1. $Y = Y_1 + Y_2$;
2. $Y_2^t X = 0$ and $Y_2 X^t = 0$;
3. $\text{rank}(Y_1) \leq 2\text{rank}(X)$;
4. $\langle Y_1, Y_2 \rangle = 0$.

A similar construction as in Lemma 3.7 is used [114] for the analysis of matrix completion and in [141] for the analysis of tensor completion.

The following Lemma is from [123, Theorem 3.3].

Lemma 3.8. *Let $k \geq 1$. Let $X \in \mathbb{R}^{d_1 \times d_2}$ be a nonzero matrix with singular value decomposition $X = U \text{diag}(\sigma_1, \dots, \sigma_r) V^t$, where $U \in \mathbb{R}^{d_1 \times r}$ and $V \in \mathbb{R}^{d_2 \times r}$. For $i = 1, \dots, \lceil r/k \rceil$, define $I_i = \{(i-1)k + 1, \dots, \min(ik, r)\}$ and*

$$X_i = U \sigma_{I_i} V^t \tag{3.7}$$

Then $X = X_1 + X_2 + \dots + X_{\lceil r/k \rceil}$, $\text{rank}(X_i) \leq k$ for all i , $\langle X_i, X_j \rangle = \delta_{(i,j)}$ all i, j , and the following estimate holds:

$$\sum_{i \geq 2} \|X_i\|_{\text{F}} \leq \frac{1}{\sqrt{k}} \|X\|_* \tag{3.8}$$

Proof. The first three statements follow from the definition of $X_1, X_2, \dots, X_{\lceil r/k \rceil}$. For the last statement, observe that for all $l \in I_i$ and $m \in I_{i-1}$, $\sigma_l \leq \sigma_m$, since the singular values are in non-increasing order. It follows that for all $l \in I_i$, $\sigma_l \leq \frac{1}{k} \sum_{m \in I_{i-1}} \sigma_m$. Hence

$$\begin{aligned} \sum_{i \geq 2} \|X_i\|_F &= \sum_{i \geq 2} \sqrt{\sum_{l \in I_i} \sigma_l^2} \\ &\leq \sum_{i \geq 1} \sqrt{k \frac{1}{k^2} (\sum_{m \in I_i} \sigma_m)^2} \\ &= \frac{1}{\sqrt{k}} \sum_{i \geq 1} \sum_{m \in I_i} \sigma_m \\ &= \frac{1}{\sqrt{k}} \|X\|_* \end{aligned}$$

□

Proof of theorem 3.4.

(Part 1: Setup.) We following the argument in [59], but extend from the matrix to the tensor case. Let

$$R = \sum_{k=1}^N \sqrt{r_k}.$$

Define

$$\mathbf{E} = \mathbf{X}^* - \mathbf{X}_0$$

and

$$\mathbf{\Delta} = \mathbf{X}_0 - \mathbf{X}_{\text{best}}.$$

\mathbf{E} is the prediction error we aim to bound and $\mathbf{\Delta}$ is the amount by which the tensor \mathbf{X}_0 differs from its best rank (r_1, \dots, r_N) approximation.

For each $k = 1, \dots, N$, consider the unfolding $\mathbf{E}^{(k)}$. By Lemma 3.7 we can

decompose $\mathbf{E}^{(k)}$ into a sum of matrices

$$\mathbf{E}^{(k)} = E_{k,0} + E_{k,c}$$

satisfying

$$\text{rank}(E_{k,0}) \leq 2 \text{rank}(\mathbf{X}_{\text{best}}^{(k)}) \leq 2r_k,$$

$$E_{k,c}(\mathbf{X}_{\text{best}}^{(k)})^t = 0, E_{k,c}^t \mathbf{X}_{\text{best}}^{(k)} = 0, \text{ and } (E_{k,0}, E_{k,c}) = 0.$$

Using the construction in Lemma 3.8, with $3r_k$ instead of k , we decompose each $E_{k,c}$ into a sum of matrices $E_{k,c} = E_{k,0} + E_{k,1} + \dots + E_{k,L}$, where $L = \lceil \frac{\text{rank}(E_{k,c})}{3r_k} \rceil$ and each $E_{k,l}$ is of rank at most $3r_k$. It follows immediately that

$$\mathbf{E}^{(k)} = E_{k,0} + E_{k,1} + E_{k,2} + \dots + E_{k,L},$$

where $\text{rank}(E_{k,0}) \leq 2r_k$ and $\text{rank}(E_{k,l}) \leq 3r_k$ for $l \geq 1$.

Our goal is to bound the prediction error $\|\mathbf{E}\|_{\text{F}}$. Since $\|\mathbf{E}\|_{\text{F}} = \|\mathbf{E}^{(k)}\|_{\text{F}}$ for all k , we have

$$\begin{aligned} \|\mathbf{E}\|_{\text{F}} &= \|\mathbf{E}^{(k)}\|_{\text{F}} \\ &= \|E_{k,0} + E_{k,c}\|_{\text{F}} \\ &= \|E_{k,0} + E_{k,1} + \sum_{l \geq 2} E_{k,l}\|_{\text{F}} \\ &\leq \|E_{k,0} + E_{k,1}\|_{\text{F}} + \sum_{l \geq 2} \|E_{k,l}\|_{\text{F}} \end{aligned} \tag{3.9}$$

Multiplying (3.9) by $\sqrt{r_k}$ and computing the average over $k = 1, \dots, N$, we obtain

$$R\|\mathbf{E}\|_{\text{F}} \leq \frac{1}{N} \sum_{k=1}^N \sqrt{r_k} \|E_{k,0} + E_{k,1}\|_{\text{F}} + \frac{1}{N} \sum_{k=1}^N \sum_{l \geq 2} \sqrt{r_k} \|E_{k,l}\|_{\text{F}}, \tag{3.10}$$

where $R = \sum_{k=1}^N \sqrt{r_k}$.

Hence, in order to bound $\|\mathbf{E}\|_{\text{F}}$, it suffices to bound the two terms on the right hand side of (3.10) separately.

(Part 2: bounding $\frac{1}{N} \sum_{k=1}^N \sum_{l \geq 2} \sqrt{r_k} \|E_{k,l}\|_{\text{F}}$.) By Lemma 3.8,

$$\sum_{l \geq 2} \|E_{k,l}\|_{\text{F}} \leq \frac{1}{\sqrt{3r_k}} \|E_{k,c}\|_*$$

for all $k = 1, \dots, N$.

Multiplying by $\sqrt{r_k}$ and computing the average over $k = 1, \dots, N$, we obtain

$$\frac{1}{N} \sum_{k=1}^N \sum_{l \geq 2} \sqrt{r_k} \|E_{k,l}\|_{\text{F}} \leq \frac{1}{\sqrt{3}} \frac{1}{N} \sum_{k=1}^N \|E_{k,c}\|_* \quad (3.11)$$

We also have

$$\|\mathbf{X}_0\|_* \geq \|\mathbf{X}^*\|_* \quad (3.12)$$

$$= \|\mathbf{X}_0 + \mathbf{E}\|_*$$

$$= \|\mathbf{X}_{\text{best}} + \mathbf{\Delta} + \mathbf{E}\|_*$$

$$= \frac{1}{N} \sum_{k=1}^N \|\mathbf{X}_{\text{best}}^{(k)} + \mathbf{\Delta}^{(k)} + \mathbf{E}^{(k)}\|_*$$

$$= \frac{1}{N} \sum_{k=1}^N \|\mathbf{X}_{\text{best}}^{(k)} + \mathbf{\Delta}^{(k)} + E_{k,c} + E_{k,0}\|_*$$

$$\geq \frac{1}{N} \sum_{k=1}^N \left(\|\mathbf{X}_{\text{best}}^{(k)} + E_{k,c}\|_* - \|\mathbf{\Delta}^{(k)}\|_* - \|E_{k,0}\|_* \right)$$

$$= \frac{1}{N} \sum_{k=1}^N \left(\|\mathbf{X}_{\text{best}}^{(k)}\|_* + \|E_{k,c}\|_* - \|\mathbf{\Delta}^{(k)}\|_* - \|E_{k,0}\|_* \right) \quad (3.13)$$

$$= \frac{1}{N} \sum_{k=1}^N \left(\|\mathbf{X}_0^{(k)} - \mathbf{\Delta}^{(k)}\|_* + \|E_{k,c}\|_* - \|\mathbf{\Delta}^{(k)}\|_* - \|E_{k,0}\|_* \right)$$

$$\geq \frac{1}{N} \sum_{k=1}^N \left(\|\mathbf{X}_0^{(k)}\|_* + \|E_{k,c}\|_* - 2\|\mathbf{\Delta}^{(k)}\|_* - \|E_{k,0}\|_* \right)$$

$$= \|\mathbf{X}_0\|_* - 2\|\mathbf{\Delta}\|_* + \frac{1}{N} \sum_{k=1}^N \|E_{k,c}\|_* - \frac{1}{N} \sum_{k=1}^N \|E_{k,0}\|_* \quad (3.14)$$

Inequality (3.12) follows because \mathbf{X}^* is an optimal solution of problem (3.4).

Equation (3.13) follows from $\|\mathbf{X}_{\text{best}}^{(k)} + \mathbf{E}^{(k)}\|_* = \|\mathbf{X}_{\text{best}}^{(k)}\|_* + \|\mathbf{E}^{(k)}\|_*$, which is a consequence of Lemma 3.6, $E_{k,c} \left(\mathbf{X}_{\text{best}}^{(k)} \right)^t = 0$ and $(E_{k,c})^t \mathbf{X}_{\text{best}} = 0$.

Rearranging (3.14), we obtain

$$\frac{1}{N} \sum_{k=1}^N \|E_{k,c}\|_* \leq 2\|\Delta\|_* + \frac{1}{N} \sum_{k=1}^N \|E_{k,0}\|_* \quad (3.15)$$

Combining inequalities (3.11) and (3.15) we obtain

$$\begin{aligned} \frac{1}{N} \sum_{k=1}^N \sum_{l \geq 2} \sqrt{r_k} \|E_{k,l}\|_{\text{F}} &\leq \frac{1}{\sqrt{3}} \frac{1}{N} \sum_{k=1}^N \|E_{k,c}\|_* \\ &\leq \frac{1}{\sqrt{3}} \left(2\|\Delta\|_* + \frac{1}{N} \sum_{k=1}^N \|E_{k,0}\|_* \right) \\ &\leq \frac{1}{\sqrt{3}} \left(2\|\Delta\|_* + \frac{1}{N} \sum_{k=1}^N \sqrt{2r_k} \|E_{k,0}\|_{\text{F}} \right) \end{aligned} \quad (3.16)$$

$$= \frac{2}{\sqrt{3}} \|\Delta\|_* + \sqrt{\frac{2}{3}} \frac{1}{N} \sum_{k=1}^N \sqrt{r_k} \|E_{k,0}\|_{\text{F}}, \quad (3.17)$$

where (3.16) follows from $\text{rank}(E_{k,0}) \leq 2r_k$ and Lemma 2.5.

(Part 3: bounding $\frac{1}{N} \sum_{k=1}^N \sqrt{r_k} \|E_{k,0} + E_{k,1}\|_{\text{F}}$.)

For brevity, let $\text{refold}(\mathbf{Z})$ denote $\text{refold}_{(d_1, \dots, d_N)}(\mathbf{Z})$ for any \mathbf{Z} . For each k , we have

$$\mathbf{E} = \text{refold}(E_{k,0} + E_{k,1}) + \sum_{l \geq 2} \text{refold}(E_{k,l})$$

The first summand has rank less than or equal to $(d_1, \dots, d_{k-1}5r_k, d_{k+1}, \dots, d_N)$ while all the remaining summands have rank less than or equal to $(d_1, \dots, d_{k-1}3r_k, d_{k+1}, \dots, d_N)$. The square-free strong RIP property of \mathcal{A} implies that the RIP of rank $(d_1, \dots, d_{k-1}5r_k, d_{k+1}, \dots, d_N)$ holds for each $k = 1, \dots, N$. Hence, by these RIP properties,

$$\begin{aligned}
\|E_{k,0} + E_{k,1}\|_{\text{F}} &= \|\text{refold}(E_{k,0} + E_{k,1})\|_{\text{F}} \\
&\leq \frac{1}{1-\bar{\delta}_5} \|\mathcal{A}(\text{refold}(E_{k,0} + E_{k,1}))\|_2 \\
&= \frac{1}{1-\bar{\delta}_5} \|\mathcal{A}(\mathbf{E}) - \sum_{l \geq 2} \mathcal{A}(\text{refold}(E_{k,l}))\|_2 \\
&\leq \frac{1}{1-\bar{\delta}_5} \|\mathcal{A}(\mathbf{E})\|_2 + \frac{1}{1-\bar{\delta}_5} \sum_{l \geq 2} \|\mathcal{A}(\text{refold}(E_{k,l}))\|_2 \\
&\leq \frac{1}{1-\bar{\delta}_5} \|\mathcal{A}(\mathbf{E})\|_2 + \frac{1+\bar{\delta}_3}{1-\bar{\delta}_5} \sum_{l \geq 2} \|E_{k,l}\|_{\text{F}} \tag{3.18}
\end{aligned}$$

Since \mathbf{X}^* and \mathbf{X}_0 are both feasible solutions of problem (3.4),

$$\begin{aligned}
\|\mathcal{A}(\mathbf{E})\|_2 &= \|\mathcal{A}(\mathbf{X}^* - \mathbf{X}_0)\|_2 \\
&= \|\mathcal{A}(\mathbf{X}^*) - y + y - \mathcal{A}(\mathbf{X}_0)\|_2 \\
&\leq \|\mathcal{A}(\mathbf{X}^*) - y\|_2 + \|\mathcal{A}(\mathbf{X}_0) - y\|_2 \\
&\leq 2\epsilon \tag{3.19}
\end{aligned}$$

Multiplying inequality (3.18) by $\sqrt{r_k}$, averaging over $k = 1, \dots, N$, and combining with (3.19), and (3.17), we obtain

$$\frac{1}{N} \sum_{k=1}^N \sqrt{r_k} \|E_{k,0} + E_{k,1}\|_{\text{F}} \leq \frac{1}{1-\bar{\delta}_5} R(2\epsilon) + \left(\frac{1+\bar{\delta}_3}{1-\bar{\delta}_5}\right) \left(\frac{2}{\sqrt{3}} \|\Delta\|_* + \sqrt{\frac{2}{3}} \frac{1}{N} \sum_{k=1}^N \sqrt{r_k} \|E_{k,0}\|_{\text{F}}\right).$$

(Recall that $R = \frac{1}{N} \sum_{k=1}^N \sqrt{r_k}$.) Since $\langle E_{k,0}, E_{k,1} \rangle = 0$, we have

$$\|E_{k,0}\|_{\text{F}} = \sqrt{\|E_{k,0} + E_{k,1}\|_{\text{F}}^2 - \|E_{k,1}\|_{\text{F}}^2} \leq \|E_{k,0} + E_{k,1}\|_{\text{F}}.$$

Hence, by collecting the term $\frac{1}{N} \sum_{k=1}^N \sqrt{r_k} \|E_{k,0} + E_{k,1}\|_{\text{F}}$ onto one side, we obtain

$$\left(1 - \frac{1+\bar{\delta}_3}{1-\bar{\delta}_5} \sqrt{\frac{2}{3}}\right) \frac{1}{N} \sum_{k=1}^N \sqrt{r_k} \|E_{k,0} + E_{k,1}\|_{\text{F}} \leq \frac{2}{1-\bar{\delta}_5} R\epsilon + \frac{1+\bar{\delta}_3}{1-\bar{\delta}_5} \frac{2}{\sqrt{3}} \|\Delta\|_*$$

It follows that

$$\frac{1}{N} \sum_{k=1}^N \sqrt{r_k} \|E_{k,0} + E_{k,1}\|_{\mathbb{F}} \leq \alpha R \epsilon + \beta \|\Delta\|_* \quad (3.20)$$

where

$$\alpha = \frac{2}{1 - \bar{\delta}_5 - (1 + \bar{\delta}_3) \sqrt{\frac{2}{3}}}$$

and

$$\beta = \frac{2(1 + \bar{\delta}_3)}{\sqrt{3} \left(1 - \bar{\delta}_5 - (1 + \bar{\delta}_3) \sqrt{\frac{2}{3}}\right)}$$

(Part 3: Collecting results.) Finally, by (3.10), (3.20), and (3.17) we have

$$\begin{aligned} R \|\mathbf{E}\|_{\mathbb{F}} &\leq \frac{1}{N} \sum_{k=1}^N \sqrt{r_k} \|E_{k,0} + E_{k,1}\|_{\mathbb{F}} + \frac{1}{N} \sum_{l \geq 2} \sqrt{r_k} \|E_{k,l}\|_{\mathbb{F}} \\ &\leq \alpha R \epsilon + \beta \|\Delta\|_* + \frac{2}{\sqrt{3}} \|\Delta\|_* + \sqrt{\frac{2}{3}} \frac{1}{N} \sum_{k=1}^N \sqrt{r_k} \|E_{k,0}\|_{\mathbb{F}} \end{aligned}$$

Again, using $\|E_{k,0}\|_{\mathbb{F}} \leq \|E_{k,0} + E_{k,1}\|_{\mathbb{F}}$ and (3.20) we have

$$\begin{aligned} R \|\mathbf{E}\|_{\mathbb{F}} &\leq \alpha R \epsilon + \beta \|\Delta\|_* + \frac{2}{\sqrt{3}} \|\Delta\|_* + \sqrt{\frac{2}{3}} (\alpha R \epsilon + \beta \|\Delta\|_*) \\ &= \alpha \left(1 + \sqrt{\frac{2}{3}}\right) R \epsilon + \left(\beta \left(1 + \frac{2}{\sqrt{3}}\right) + \sqrt{\frac{2}{3}}\right) \|\Delta\|_* \end{aligned}$$

Dividing by $R = \frac{1}{N} \sum_{k=1}^N \sqrt{r_k}$, we arrive at the claim

$$\|\mathbf{E}\|_{\mathbb{F}} \leq C_0 \epsilon + \frac{C_1 \|\Delta\|_*}{R},$$

where

$$C_0 = \frac{2 \left(1 + \sqrt{\frac{2}{3}}\right)}{1 - \bar{\delta}_5 - (1 + \bar{\delta}_3) \sqrt{\frac{2}{3}}}$$

and

$$C_1 = \frac{2(1 + \bar{\delta}_3) \left(1 + \frac{2}{\sqrt{3}}\right)}{\sqrt{3} \left(1 - \bar{\delta}_5 - (1 + \bar{\delta}_3) \sqrt{\frac{2}{3}}\right)} + \sqrt{\frac{2}{3}}.$$

□

Remark 3.9. *The requirement to split each $E_{k,c}$ into a sum of rank 3 matrices may seem artificial. However, if we instead we split of $E_{k,c}$ into a sum of rank 2 matrices, the expression $\sqrt{3}$ appearing in the proof would change to $\sqrt{2}$. Hence, the condition on the RIP constants would reduce to $\bar{\delta}_5 + 1 + \bar{\delta}_3 < 1$, which is always false, since RIP constants are nonnegative.*

3.3 Probabilistic Recovery Guarantees

We have shown in Theorem 3.4 that recovery of a tensor \mathbf{X} , with error controlled by the difference between \mathbf{X} and its best Tucker rank (r_1, \dots, r_N) approximation, is guaranteed by solving the minimization problem (3.4), provided that the random measurement operator \mathcal{A} satisfies the strong square-free RIP of rank $(5r_1, \dots, 5r_N)$ with $\bar{\delta}_5 + (1 + \bar{\delta}_3)\sqrt{\frac{2}{3}} < 1$. As remarked previously, this condition is implied by $\bar{\delta}_5 < \bar{\delta}_{\text{critical}}$. Hence, in order to prove recovery with high probability, it suffices to prove that the strong square-free RIP of rank $(5r_1, \dots, 5r_N)$ holds with constant $\bar{\delta}_5 < \bar{\delta}_{\text{critical}}$ with high probability.

We consider random measurement operators, as described in the following definition.

Definition 3.10 (Random measurement operator). *Let (Ω, \mathcal{F}, P) be a probability measure space. A random measurement operator \mathcal{A} is a measurable function*

$$\mathcal{A} : \Omega \longrightarrow L(\mathbb{R}^{d_1 \times \dots \times d_N}, \mathbb{R}^m),$$

where $L(\mathbb{R}^{d_1 \times \dots \times d_N}, \mathbb{R}^m)$ is the space of linear maps from $\mathbb{R}^{d_1 \times \dots \times d_N}$ to \mathbb{R}^m .

In [102], Yi-Kai Liu first proved that the RIP holds with high probability, in the case of matrices, for a measurement operator drawn from random entries in an orthonormal basis. In [42], Cloninger and Czaja extended Liu’s result to the case of measurements drawn from a Parseval tight frame. Cloninger and Czaja combined the resulting RIP guarantee with a recovery result of Fazel, Candès, Recht, and Parillo [59] to obtain a general probabilistic recovery guarantee for measurement operators obtained from a Parseval tight frame. Here, we extend the results of Cloninger and Czaja in two respects. First we extend from the matrix to the tensor case. Second, we obtain a recovery result for a more general class of random measurement operators, isotropic, sub-Gaussian measurement maps.

For tensors, Rauhut, Schneider, and Stojanac proved in [122, 119], that Tucker rank (r_1, \dots, r_N) RIP holds for random sub-Gaussian ensembles.

Definition 3.11 (Random sub-Gaussian ensemble). *Let ω be a random tensor with values in $\mathbb{R}^{d_1 \times \dots \times d_N}$ such that the entries $\omega[i_1, \dots, i_N]$ for $(i_1, \dots, i_N) \in [d_1] \times \dots \times [d_N]$ are iid sub-Gaussian with the same sub-Gaussian norm. Then the measurement operator $\mathcal{A} : \mathbb{R}^{d_1 \times \dots \times d_N} \rightarrow \mathbb{R}^m$ defined by*

$$\mathcal{A}(\mathbf{X}) = (\mathbf{X}, \omega_i),$$

where $\omega_1, \dots, \omega_m$ are iid drawings of ω is a random sub-Gaussian ensemble.

A limitation of sub-Gaussian ensembles is that all the entries of the resulting matrix must be i.i.d., which excludes important examples such as when each row is drawn from a Parseval tight frame. We are thus interested in a more general class of random measurement operators. The following definition is from [54].

Definition 3.12 (sub-Gaussian map). *Let (Ω, \mathcal{F}, P) be a probability measure space. Let \mathcal{H} be a d -dimensional real Hilbert space and let $\mathcal{S} = \{x \in \mathcal{H} : \|x\|_{\mathcal{H}} = 1\}$. Let $m \geq 1$. Let $\mathcal{A} : \Omega \rightarrow L(\mathcal{H}, \mathbb{R}^m)$ be a random measurement operator. Then \mathcal{A} is a sub-Gaussian map with parameter ν if the following conditions hold:*

- (Linearity) \mathcal{A} is linear.
- (Independence) For all $x \in \mathcal{S}$, the entries $\mathcal{A}(x)[1], \dots, \mathcal{A}(x)[m]$ are independent.
- (Isotropy) For all $x \in \mathcal{S}$, $\mathbb{E}\|\mathcal{A}(x)\|_2^2 = \|x\|_2^2$.
- (sub-Gaussianity) For all $x, y \in \mathcal{S}$ and for all $i = 1, \dots, m$,

$$\|\mathcal{A}(x - y)[i]\|_{\psi_2} \leq \sqrt{\frac{\nu}{m}} \|x - y\|_{\mathcal{H}}.$$

An important ingredient in the proof of the RIP is a covering number computation. The following result was proved by Rauhut, Schneider, and Stojanac [122, 119].

Theorem 3.13 ([119]). *Let $\mathcal{B}_{(r_1, \dots, r_N)} = \{\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N} : \|\mathbf{X}\|_F = 1, \text{rank}(\mathbf{X}) \leq (r_1, \dots, r_N)\}$. Then for any $0 < \lambda < 1$, the following condition holds:*

$$N(\mathcal{B}_{(r_1, \dots, r_N)}, d_{\text{Euclidean}}, \lambda) < \left(\frac{3(N+1)}{\lambda} \right)^{r_1 r_2 \dots r_N + \sum_{k=1}^N d_k r_k},$$

where $N(\mathcal{B}_{(r_1, \dots, r_N)}, d_{\text{Euclidean}}, \lambda)$ is the covering number of $\mathcal{B}_{(r_1, \dots, r_N)}$ with balls of radius λ with respect to the Euclidean metric $d_{\text{Euclidean}}(\mathbf{X}) = \sqrt{(\mathbf{X}, \mathbf{X})}$.

Using the above covering number result, Rauhut, Schneider, and Stojanac proved that the RIP holds with high probability for sub-Gaussian ensembles.

In [54, Example 5.8], Dirksen proved the following theorem, which establishes that the RIP holds with high probability for a sub-Gaussian maps. Dirksen's proof uses the covering number result Theorem 3.13. In the following theorem, Dirksen uses the standard RIP (not the square-free RIP).

Theorem 3.14 ([54]). *Consider a sub-Gaussian measurement operator $\mathcal{A} : \Omega \rightarrow L(\mathbb{R}^{d_1 \times \dots \times d_N}, \mathbb{R}^m)$ with parameter μ . There exists an absolute constant $C > 0$, such that for any $0 < \delta < 1$ and $0 < p < 1$, we have $P(\delta_{(r_1, \dots, r_N)} < \delta) > 1 - p$, provided that the number of measurements m satisfies*

$$m \geq C \frac{\mu^2}{\delta^2} \max \left\{ \log\left(\frac{1}{p}\right), \left(r_1 r_2 \cdots r_N + \sum_{k=1}^N d_k r_k \right) \log(N) \right\}.$$

Dirksen's proof of the above result uses generic chaining, a theory developed by Talagrand [138] to bound the suprema of stochastic processes.

We now combine Theorem 3.14 with our previous recovery guarantee, Theorem 3.4, to obtain a general probabilistic recovery guarantee for sub-Gaussian maps.

Theorem 3.15. *There exists a numerical constant $C > 0$ such that the following holds. Define*

$$\delta_{\text{critical}} = 2\bar{\delta}_{\text{critical}} - \bar{\delta}_{\text{critical}}^2 \approx 0.1918, \quad (3.21)$$

where $\bar{\delta}_{\text{critical}}$ is defined by (3.6). Assume that $\delta > 0$ satisfies $\delta < \delta_{\text{critical}}$. If $\mathcal{A} : \Omega \rightarrow L(\mathbb{R}^{d_1 \times \dots \times d_N}, \mathbb{R}^m)$ is a sub-Gaussian map with parameter ν and if the number of measurements satisfies

$$m \geq C \frac{\nu^2}{\delta^2} \max \left\{ \log\left(\frac{N}{p}\right), \max_{k=1, \dots, N} \left\{ r_k \prod_{j \neq k} d_j + r_k d_k + \sum_{j \neq k} d_j^2 \right\} \log(N) \right\}, \quad (3.22)$$

then with probability greater than $1 - p$ over the choice of \mathcal{A} , the following holds:

For all $\epsilon > 0$ and for all $\mathbf{X}_0 \in \mathbb{R}^{d_1 \times \dots \times d_N}$ and $y \in \mathbb{R}^m$ satisfying $\|\mathcal{A}(\mathbf{X}_0) - y\|_2 \leq \epsilon$,

the solution \mathbf{X}^* of the minimization problem

$$\min_{\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N} : \|\mathcal{A}(\mathbf{X}) - y\|_2 \leq \epsilon} \|\mathbf{X}\|_* \quad (3.23)$$

satisfies

$$\|\mathbf{X}^* - \mathbf{X}_0\|_F \leq C_0 \epsilon + C_1 \left(\frac{1}{\frac{1}{N} \sum_{k=1}^N \sqrt{r_k}} \right) \|\mathbf{X}^* - \mathbf{X}_{\text{best}}\|_*, \quad (3.24)$$

where \mathbf{X}_{best} is the best rank (r_1, \dots, r_N) approximation of \mathbf{X}_0 in the $\|\cdot\|_*$ norm.

The constants C_0 and C_1 are the same as in Theorem 3.4.

Proof. For each $k = 1, \dots, N$, we replace (r_1, \dots, r_N) with $(d_1, \dots, d_{k-1}, r_k, d_{k+1}, \dots, d_N)$

and p with $\frac{p}{N}$ in the statement of Theorem 3.14. The assumption on m implies that

for all $k = 1, \dots, N$

$$m \geq C \frac{\nu^2}{\delta^2} \max \left\{ \log\left(\frac{N}{p}\right), \left(r_k \prod_{j \neq k} d_j + \sum_{j \neq k} d_j^2 + d_k r_k \right) \log(N) \right\}.$$

Hence, by Theorem 3.14 the random measurement operator \mathcal{A} satisfies the RIP of rank $(d_1, \dots, d_{k-1}, r_k, d_{k+1}, \dots, d_N)$ with constant δ with probability greater than

$1 - p/N$ provided that the number of measurements m satisfies

$$m \geq C \frac{\mu^2}{\delta^2} \max \left\{ \log\left(\frac{N}{p}\right), \left(r_1 r_2 \dots r_N + \sum_{k=1}^N d_k r_k \right) \log(N) \right\}.$$

By the union bound, the RIP of rank $(d_1, \dots, d_{k-1}, r_k, d_{k+1}, \dots, d_N)$ holds for all

$k = 1, \dots, N$ with probability greater than $1 - p$. Hence, with probability greater

than $1 - p$, the strong RIP of rank $(5r_1, \dots, 5r_N)$ holds with constant δ . By Lemma

2.37, the square-free strong RIP holds with constant $\bar{\delta}$ satisfying $\delta = 2\bar{\delta} - \bar{\delta}^2$. Since

$$2\bar{\delta} - \bar{\delta}^2 = \delta < \delta_{\text{critical}} = 2\bar{\delta}_{\text{critical}} - \bar{\delta}_{\text{critical}}^2$$

and since $x \mapsto 2x - x^2$ is strictly increasing on $[0, 1]$, it follows that $\bar{\delta} < \bar{\delta}_{\text{critical}}$. Hence, by Theorem 3.4, the estimate (3.24) holds with probability at least $1 - p$. \square

3.4 Non-Uniform Sampling from Parseval Tight Frames

Here we show that sub-Gaussian measurement maps arise from non-uniform sampling from Parseval tight frames, after appropriate rescaling of the frame elements. These results establish a theoretical basis for tensor completion via non-uniform sampling from Parseval tight frames.

The example of sub-Gaussian maps obtained from uniform sampling from frame elements is given by Eldar in [57]. We consider a more general setting in which the entries of a frame $\{u_i\}$ are sampled with a non-uniform probability distribution p_i .

Definition 3.16 (Non-uniform sampling from a Parseval tight frame). *Let \mathcal{H} be a d -dimensional Hilbert space. Let $\{u_j\}_{j=1}^n$ be a Parseval tight frame for \mathcal{H} . Let $\{p_j\}_{j=1}^n$ satisfy $\sum_{j=1}^n p_j = 1$ and $p_j > 0$ for $j = 1, \dots, n$. Let ω be a random vector such that $P(\omega = \frac{u_j}{\sqrt{p_j}}) = p_j$, where $\sum_{j=1}^n p_j = 1$ and $p_j > 0$ for $j = 1, \dots, n$. Let $\omega_1, \dots, \omega_m$ be i.i.d. realizations of the random variable ω and let $\epsilon_1, \dots, \epsilon_m$ be i.i.d. Rademacher random variables with values in $\{-1, 1\}$, independent of the ω_i random variables. The random measurement operator \mathcal{A} with values in $L(\mathcal{H}, \mathbb{R}^m)$ generated by $\{u_j\}_{j=1}^n$ and $\{p_j\}_{j=1}^n$ is defined for all $x \in \mathcal{H}$ by*

$$\mathcal{A}(x)[i] = \langle x, \frac{1}{\sqrt{m}} \epsilon_i \omega_i \rangle. \quad (3.25)$$

The following result establishes that \mathcal{A} , as defined above, is a sub-Gaussian

map. The result also provides a theoretical basis for adjusting sampling strategies to reduce coherence.

Theorem 3.17 (sub-Gaussian measurements via non-uniform sampling from a tight frame). *Let \mathcal{H} be a d -dimensional real Hilbert space. Let $\{u_j\}_{j=1}^n$ be a Parseval tight frame for \mathcal{H} and let $\{p_j\}_{j=1}^n$ satisfy $\sum_{j=1}^n p_j = 1$ and $p_j > 0$ for $j = 1, \dots, n$. Then the operator \mathcal{A} defined by (3.25) is a sub-Gaussian map with parameter*

$$\nu = \max_{i=1, \dots, n} \frac{\|u_i\|_{\mathcal{H}}^2}{p_i}.$$

Proof. Observe that \mathcal{A} is clearly linear and for all $x \in \mathcal{S}$, $\mathcal{A}(x)[1], \dots, \mathcal{A}(x)[m]$ are independent by construction. (Recall that $\mathcal{S} = \{x \in \mathcal{H} : \|x\|_{\mathcal{H}} = 1\}$.)

To show isotropy, let $x \in \mathcal{S}$ and observe that

$$\begin{aligned} \mathbb{E}\|\mathcal{A}(x)\|_2^2 &= \sum_{i=1}^m \left| \langle x, \frac{\epsilon_i \omega_i}{\sqrt{m}} \rangle \right|^2 \\ &= \sum_{i=1}^m \sum_{j=1}^n p_j \left| \frac{\langle x, u_j \rangle}{\sqrt{p_j m}} \right|^2 \\ &= \frac{1}{m} \sum_{i=1}^m \sum_{j=1}^n |\langle x, u_j \rangle|^2 \\ &= \frac{1}{m} \sum_{i=1}^m \|x\|_{\mathcal{H}}^2 \\ &= \|x\|_{\mathcal{H}}^2, \end{aligned}$$

where the second to last equality follows from the fact that $\{u_i\}_{i=1}^n$ is a Parseval tight frame.

Next, we show the sub-Gaussian property. Let $x, y \in \mathcal{S}$. Observe that

$$\mathcal{A}(x - y)[i] = \langle x - y, \frac{\epsilon_i \omega_i}{\sqrt{m}} \rangle.$$

Because of the Rademacher random variable ϵ_i , we have $\mathbb{E}[\mathcal{A}(x - y)[i]] = 0$. Also,

we have

$$\|\mathcal{A}(x - y)[i]\|_2 \leq \max_{j=1, \dots, n} \frac{\|x - y\|_{\mathcal{H}} \|u_j\|_{\mathcal{H}}}{\sqrt{p_j m}}.$$

Since $\mathcal{A}(x - y)[i]$ is bounded and has mean 0, by Lemma 2.47 it is sub-Gaussian, with sub-Gaussian norm satisfying

$$\|\mathcal{A}(x - y)[i]\|_{\psi_2} \leq \max_{j=1, \dots, n} \frac{\|x - y\|_{\mathcal{H}} \|u_j\|_{\mathcal{H}}}{\sqrt{p_j m}}.$$

Let $\nu = \max_{j=1}^n \frac{\|u_j\|_{\mathcal{H}}^2}{p_j}$. Then the above equation becomes:

$$\|\mathcal{A}(x - y)[i]\|_{\psi_2} \leq \frac{\sqrt{\nu}}{\sqrt{m}} \|x - y\|_{\mathcal{H}},$$

which is the condition required in Definition 3.12. Hence \mathcal{A} is a sub-Gaussian random operator with parameter ν . \square

Remark 3.18. *If in the previous result we sample the entries of $\{u_i\}_{i=1}^n$ uniformly, then $p_i = \frac{1}{n}$. Hence, the formula for ν reduces to*

$$\nu = \max_{i=1}^n n \|u_i\|_{\mathcal{H}}^2.$$

This formula for ν agrees up to a constant with the definition of incoherence used by Cloninger and Czaja in [42]. Hence, the parameter ν of a sub-Gaussian map is a natural generalization of the incoherence of a Parseval tight frame. These observations motivate the following definition.

Definition 3.19 (Incoherence of sub-Gaussian map). *The incoherence ν of a sub-Gaussian map \mathcal{A} with values in $L(\mathcal{H}, \mathbb{R}^m)$ is defined to be the smallest parameter ν satisfying the last condition of Definition 3.12.*

We now combine Theorem 3.17 on non-uniform sampling with the recovery guarantee Theorem 3.15 to obtain a general tensor recovery result for non-uniform sampling from Parseval tight frames.

Theorem 3.20 (Tensor recovery guarantee for non-uniform sampling from a Parseval tight frame). *There exists $C > 0$ such that the following holds. Let $\{u_j\}_{j=1}^n$ be a Parseval tight frame for $\mathbb{R}^{d_1 \times \dots \times d_N}$ and let $\{p_j\}_{j=1}^n$ satisfy $\sum_{j=1}^n p_j = 1$ and $p_j > 0$ for $j = 1, \dots, n$. Let \mathcal{A} be the random measurement operator with values in $L(\mathbb{R}^{d_1 \times \dots \times d_N}, \mathbb{R}^m)$ be defined by (3.25). Let $p \in (0, 1)$ and let $\delta > 0$ satisfy $\delta < \delta_{\text{critical}}$, where δ_{critical} is defined by (3.21). If the number of measurements m satisfies*

$$m \geq C \frac{\nu^2}{\delta^2} \max \left\{ \log\left(\frac{N}{p}\right), \max_{k=1, \dots, N} \left\{ r_k \prod_{j \neq k} d_j + r_k d_k + \sum_{j \neq k} d_j^2 \right\} \log(N) \right\},$$

then with probability at least $1 - p$ over \mathcal{A} , for all $\epsilon > 0$, $\mathbf{X}_0 \in \mathbb{R}^{d_1 \times \dots \times d_N}$ and $y \in \mathbb{R}^m$ satisfying $\|\mathcal{A}(\mathbf{X}_0) - y\|_2 \leq \epsilon$, the solution \mathbf{X}^* of the minimization problem

$$\min_{\mathbf{X} : \|\mathcal{A}(\mathbf{X}) - y\|_2 \leq \epsilon} \|\mathbf{X}\|_* \tag{3.26}$$

satisfies the error estimate

$$\|\mathbf{X}^* - \mathbf{X}_0\|_{\text{F}} \leq C_0 \epsilon + C_1 \left(\frac{1}{\frac{1}{N} \sum_{k=1}^N \sqrt{r_k}} \right) \|\mathbf{X}^* - \mathbf{X}_{\text{best}}\|_*,$$

where \mathbf{X}_{best} is the best rank (r_1, \dots, r_N) approximation of \mathbf{X}_0 in the $\|\cdot\|_*$ norm.

The constants C_0 and C_1 are the same as in Theorem 3.4.

Proof. By Theorem 3.17, the measurement operator \mathcal{A} is a sub-Gaussian map with parameter $\nu = \max_{i=1, \dots, n} \frac{\|u_i\|_{\mathcal{H}}^2}{p_i}$. Hence, by Theorem 3.15, we obtain the required probabilistic recovery guarantee. □

3.5 Choice of Non-Uniform Sampling Distribution

We describe three possible choices for the sampling probabilities. One choice is *uniform sampling*, in which $p_j = \frac{1}{n}$ for $j = 1, \dots, n$. Two other possible choices are *minimally coherent non-uniform sampling* and *nearly minimally coherent non-uniform sampling*, which are described below.

In minimally coherent non-uniform sampling, the sampling probabilities are chosen to minimize the coherence ν .

Theorem 3.21 (Minimally coherent non-uniform sampling). *Assume that $\{u_j\}_{j=1}^n$ is a Parseval tight frame for a d -dimensional real Hilbert space \mathcal{H} and assume that $u_j \neq 0$ for all $j = 1, \dots, n$. Consider the non-uniform sampling operator (3.25). The choice of probabilities $\{p_j\}_{j=1}^n$ that minimizes the coherence*

$$\nu = \max_{j=1, \dots, n} \frac{\|u_j\|_{\mathcal{H}}^2}{p_j}$$

is given by

$$p_j = \frac{\|u_j\|_{\mathcal{H}}^2}{d}.$$

Furthermore, the optimal coherence is given by $\nu = d$.

Proof. We aim to solve the following problem for (p_1, \dots, p_n) :

$$\min_{p_1, \dots, p_n : \sum_{i=1}^n p_i = 1, p_i \geq 0} \left(\max_{j=1, \dots, n} \frac{\|u_j\|_{\mathcal{H}}^2}{p_j} \right)$$

The objective function converges to $+\infty$ as any p_j converges to 0, since each frame element u_j is assumed to be nonzero. Hence we can consider only (p_1, \dots, p_n) outside

of a neighborhood \mathcal{B} of 0. Without loss of generality, we may assume that \mathcal{B} is an open ball of radius r in the l^1 norm. On the region $[0, 1] \times \cdots \times [0, 1] - \mathcal{B}$, the objective function is continuous and defined on a compact domain, so an optimal solution exists. Let (p_1, \dots, p_n) denote an optimal solution and let ν denote the optimal objective value.

We claim that $\nu = \frac{\|u_j\|_{\mathcal{H}}^2}{p_j}$ for all $j = 1, \dots, n$. Indeed, suppose that equality does not hold for some k . Then, by replacing p_k with $p_k - \epsilon$ and replacing p_j with $p_j + \frac{\epsilon}{n-1}$ for all $j \neq k$, the objective value would decrease. Furthermore, since the l^1 norm of (p_1, \dots, p_n) remains unchanged, the new point cannot enter \mathcal{B} . Hence (p_1, \dots, p_n) would remain feasible but would have a smaller objective value, a contradiction.

Since $\nu = \frac{\|u_j\|_{\mathcal{H}}^2}{p_j}$ for all $j = 1, \dots, n$, it follows that $p_j = C\|u_j\|_{\mathcal{H}}^2$ for some C .

The constraint that $\sum_{j=1}^n p_j = 1$ implies that $C = \frac{1}{\sum_{i=1}^n \|u_i\|_{\mathcal{H}}^2}$.

Finally, since $\{u_j\}_{j=1}^n$ is a Parseval tight frame, by Lemma 2.31 we have $\sum_{i=1}^n \|u_i\|_{\mathcal{H}}^2 = d$. Hence $C = \frac{1}{d}$ and we conclude that each $p_j = \frac{\|u_j\|_{\mathcal{H}}^2}{d}$.

Finally, to compute the optimal objective value, we have for any $j = 1, \dots, n$,

$$\nu = \frac{\|u_j\|_{\mathcal{H}}^2}{p_j} = \frac{\|u_j\|_{\mathcal{H}}^2}{\left(\frac{\|u_j\|_{\mathcal{H}}^2}{d}\right)} = d.$$

□

The above theorem shows that the coherence can be minimized by sampling Parseval tight frame entries with probabilities proportional to $\|u_j\|_{\mathcal{H}}^2$. The resulting sampling strategy minimizes the ν^2 factor occurring in the measurement bound given by Theorem 3.20.

One potential disadvantage of minimally coherent sampling is that if the ratio between the largest $\|u_j\|_{\mathcal{H}}^2$ and the smallest $\|u_j\|_{\mathcal{H}}^2$ is very large, some probabilities p_j could be very small, resulting in very large scaling factors $\frac{1}{\sqrt{p_j}}$. One way to avoid this problem is to add a small quantity to all the sampling probabilities, to prevent any p_j from being too small.

Definition 3.22 (Nearly minimally coherent non-uniform sampling). *Assume that $\{u_j\}_{j=1}^n$ is a Parseval tight frame for a d -dimensional Hilbert space \mathcal{H} and assume that $u_j \neq 0$ for all $j = 1, \dots, n$. Consider the non-uniform sampling operator (3.25). Let $p_0 > 0$. Nearly minimally coherent non-uniform sampling is defined by setting*

$$p_j = \frac{\frac{\|u_j\|_{\mathcal{H}}^2}{d} + \frac{p_0}{n}}{1 + p_0}.$$

As was observed by Cloninger and Czaja in [42], in relaxometry applications it is often the case that only a very small fraction of the Parseval frame elements have large norm. In such examples, nearly minimally coherent non-uniform sampling can significantly reduce the coherence ν while avoiding excessively sparse sampling of the low-norm frame elements.

3.6 Tensor Recovery Algorithm

We now describe a fixed point continuation technique for solving the regularized tensor completion problem

$$\min_{\mathbf{X}} \frac{1}{N} \sum_{k=1}^N \|\mathbf{X}^{(k)}\|_* + \frac{1}{2\mu_1} \|\mathcal{A}(\mathbf{X}) - y\|_2^2 \quad (3.27)$$

In [148], Yang, Huang, and Shi first developed a fixed point continuation

algorithm for (3.27). Their algorithm uses operator splitting, a technique in which a quantity is added and subtracted to the first order equation, resulting in a new, equivalent, optimization problem. We develop a similar algorithm, using similar techniques; however, we generalize it to apply to settings with arbitrary noise levels in the data.

3.6.1 Derivation

The derivation of our algorithm is essentially the same as in [148]. The primary difference is that in [148], it is assumed that a single regularization constant μ converges to 0. In [148], it is proved that as μ tends to 0, the resulting solution, under mild assumptions, converges to a solution of

$$\min_{\mathbf{X} : \mathcal{A}(\mathbf{X})=y} \|\mathbf{X}\|_*.$$

Hence, as $\mu \rightarrow 0$, the algorithm in [148] converges to a solution satisfying the constraint $\mathcal{A}(\mathbf{X}) = y$ exactly. We show that our algorithm converges, under mild assumptions, to the solution of (3.27). This will enable the tuning of the parameter μ_1 . The choice of μ_1 will be discussed later.

We start with a derivation of the algorithm, which follows [148] closely. Assume that \mathbf{X}_0 is a solution of (3.27). Since the objective function in (3.27) is convex, by Lemma 2.27, \mathbf{X}_0 is a solution if and only if $\mathbf{0}$ is in the subdifferential of the objective function at \mathbf{X}_0 , where $\mathbf{0}$ denotes the zero tensor in $\mathbb{R}^{d_1 \times \dots \times d_N}$. Hence,

$$\mathbf{0} \in \frac{1}{N} \sum_{k=1}^N \left(\partial \|\mathbf{X}_0^{(k)}\|_* + \frac{1}{\mu_1} \mathcal{A}^* (\mathcal{A}(\mathbf{X}_0) - y) \right) \quad (3.28)$$

Recall that ∂ denotes the subdifferential, as defined in Definition 2.26. For brevity, define $g(\mathbf{X}) = \mathcal{A}^*(\mathcal{A}(\mathbf{X}_0) - y)$ for all \mathbf{X} . Let $\tau > 0$ and let $\mathbf{Y}_0 = \mathbf{X}_0 - \tau g(\mathbf{X}_0)$. Multiplying (3.29) by $\tau\mu_1$ and then adding and subtracting \mathbf{X}_0 , we obtain

$$\mathbf{0} \in \frac{1}{N} \sum_{k=1}^N \left(\tau\mu_1 \partial \|\mathbf{X}_0^{(k)}\|_* + \mathbf{X}_0 - (\mathbf{X}_0 - \tau g(\mathbf{X}_0)) \right) \quad (3.29)$$

which is equivalent to

$$\mathbf{0} \in \frac{1}{N} \sum_{k=1}^N \left(\tau\mu_1 \partial \|\mathbf{X}_0^{(k)}\|_* + \mathbf{X}_0 - \mathbf{Y}_0 \right)$$

Hence, by Lemma 2.27, \mathbf{X}_0 is an optimal solution of

$$\min_{\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}} \frac{1}{N} \sum_{k=1}^N \left(\tau \|\mathbf{X}^{(k)}\|_* + \frac{1}{2\mu_1} \|\mathbf{X}^{(k)} - \mathbf{Y}_0^{(k)}\|_{\text{F}}^2 \right) \quad (3.30)$$

if and only if X_0 is an optimal solution of (3.27).

To reduce the interdependency between the two terms in the objective function, we introduce matrices Z_i to approximate the unfoldings $\mathbf{X}^{(i)}$, and solve the problem

$$\min_{\mathbf{x}, Z_1, \dots, Z_N} \frac{1}{N} \sum_{k=1}^N \left(\tau \|Z_i\|_* + \frac{1}{2\mu_1} \|Z_i - \mathbf{Y}_0^{(k)}\|_{\text{F}}^2 \right)$$

under the constraint $Z_i = \mathbf{X}^{(i)}$ for $i = 1, \dots, N$. This problem can be relaxed to

$$\min_{\mathbf{x}, Z_1, \dots, Z_N} \frac{1}{N} \sum_{k=1}^N \left(\tau \|Z_i\|_* + \frac{1}{2\mu_1} \|Z_i - \mathbf{Y}_0^{(k)}\|_{\text{F}}^2 + \frac{1}{2\mu_2} \|Z_i - \mathbf{X}^{(i)}\|_{\text{F}}^2 \right) \quad (3.31)$$

In [148], Yang et al. proceed under the assumption $\mu_1 = \mu_2$. In the case in which $\mu_1 \rightarrow 0$ and $\mu_2 \rightarrow 0$, the resulting solution converges to a tensor \mathbf{X} satisfying the constraint $\mathcal{A}(\mathbf{X}) = y$ exactly, under mild assumptions, so no generality is lost by the simplification $\mu_1 = \mu_2$. However, in applications in which the data is highly noisy, as can arise in magnetic resonance relaxometry applications,

the regularization parameter μ_1 must be chosen to appropriately balance bias and variance in the solution. However, it is still appropriate to let $\mu_2 \rightarrow 0$, to enforce approximate equality $X^{(i)} \approx Z_i$. Hence, we proceed with a similar derivation as in [148], but without the simplifying assumption that μ_1 and μ_2 are equal.

We will analyze problem (3.31) using block-coordinate descent (BCD), in which we fix all the variables $\mathbf{X}, Z_1, \dots, Z_N$ except for one, and then optimize over the only non-fixed variable. The following useful result, also cited in [148], provides a condition under which BCD converges.

We first state several definitions from [143].

Definition 3.23. *A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is quasiconvex if for all $x, y \in \mathbb{R}^n$ and $\lambda \in [0, 1]$, we have $f(x + \lambda y) \leq \max\{f(x), f(x + y)\}$.*

Definition 3.24. *A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is hemivariate if it is not constant on any line segment.*

Let $X_0 = (x_{1,0}, \dots, x_{N,0}) \in \mathbb{R}^{d_1} \times \dots \times \mathbb{R}^{d_N}$ be an arbitrary initial starting point. The block coordinate descent (BCD) algorithm minimizes the function $f(x_1, \dots, x_N)$ by, at each stage, fixing all but one variables and minimizing with respect to the non-fixed variable. The cyclic BCD proceeds by minimizing with respect to x_1 , then x_2 , then x_3, \dots , then x_N , and then repeats cyclically starting again at x_1 .

Definition 3.25. *A coordinate wise minimum of a function $f : \mathbb{R}^{d_1} \times \dots \times \mathbb{R}^{d_N} \rightarrow \mathbb{R}$ is a point $(x_1, \dots, x_N) \in \mathbb{R}^{d_1} \times \dots \times \mathbb{R}^{d_N}$ such that for all k , we have $f\left((x_1, \dots, x_N) + (0, \dots, z_k, \dots, 0)\right) \geq f(x_1, \dots, x_N)$ for all $z_k \in \mathbb{R}^{d_k}$.*

Definition 3.26. A point x is a stationary point of f if for all z , we have

$$\liminf_{\lambda \rightarrow 0^+} \frac{f(x + \lambda z) - f(x)}{\lambda} \geq 0.$$

The following is [143, Lemma 3.1].

Lemma 3.27. Let $f : \mathbb{R}^{d_1} \times \cdots \times \mathbb{R}^{d_N} \rightarrow \mathbb{R}$ be a function of the following form.

Suppose that there are functions $f_0 : \mathbb{R}^{d_1} \times \cdots \times \mathbb{R}^{d_N} \rightarrow \mathbb{R}$ and functions $f_i : \mathbb{R}^{d_i} \rightarrow$

\mathbb{R} such that for all $x = (x_1, \dots, x_N) \in \mathbb{R}^{d_1} \times \cdots \times \mathbb{R}^{d_N}$, we have

$$f(x_1, \dots, x_N) = f_0(x_1, \dots, x_N) + \sum_{i=1}^N f_i(x_i)$$

and assume that f_0 is differentiable. If (x_1, \dots, x_N) is a coordinate-wise minimum of f , then it follows that f is a stationary point of f .

The following is [143, Proposition 5.1].

Theorem 3.28. Let f be as in Lemma 3.27. Assume that the functions f_0 and f_1, \dots, f_N satisfy the following properties:

- f_0, f_1, \dots, f_N are continuous.
- Let $k \in \{1, \dots, N\}$ and assume that each x_i is fixed for $i \neq k$. Then the function $x_k \mapsto f(x_0, \dots, x_N)$ is quasiconvex and hemivariate.

Starting at any initial point $X_0 \in \mathbb{R}^{d_1} \times \cdots \times \mathbb{R}^{d_N}$, let X_1, X_2, \dots be the sequence of iterates generated by the cyclic BCD algorithm. Then, if the iterates X_i are defined, either $\lim_{k \rightarrow \infty} f(X_k) = -\infty$ or every limit point of the sequence $\{X_k\}$ is a coordinate-wise minimum of f .

The following lemma, which gives a formula for the solution of (3.31), with Y_0 arbitrary, generalizes [148, Theorem 4.2].

Lemma 3.29. *Let $\tau, \mu_1, \mu_2 > 0$ and let $\mathbf{Y}_0 \in \mathbb{R}^{d_1 \times \dots \times d_N}$ be any tensor. Then $(X_{\mu_2}, Z_{1, \mu_2}, \dots, Z_{N, \mu_2})$ is a minimizer of (3.31) if and only if the following two equations hold:*

$$Z_{i, \mu_2} = \text{shrink}_{\left(\frac{\tau \mu_1 \mu_2}{\mu_1 + \mu_2}\right)} \left(\frac{\mu_2 \mathbf{Y}_0^{(i)} + \mu_1 \mathbf{X}_{\mu_2}^{(i)}}{\mu_1 + \mu_2} \right), \text{ for } i = 1, \dots, N, \quad (3.32)$$

and

$$\mathbf{X}_{\mu_2} = \frac{1}{N} \sum_{i=1}^N \text{refold}(Z_{i, \mu_2}). \quad (3.33)$$

The shrink and refold operators were defined in Definition 2.55 and Definition 2.18.

Proof. The solution $(X_{\mu_2}, Z_{1, \mu_2}, \dots, Z_{N, \mu_2})$ depends on τ, μ_1 , and μ_2 ; however, the dependence on τ and μ_1 is suppressed from the notation. This notational choice will be convenient later, when we fix μ_1 and τ and let $\mu_2 \rightarrow 0$.

We aim to show that equations (3.32) and (3.33) are the update steps for cyclic BCD.

First, we derive the update step for Z_{i, μ_2} assuming $\mathbf{X} = \mathbf{X}_{\mu_2}$ and $Z_j = Z_{j, \mu_2}$ are fixed for all $j \neq i$. Hence Z_{i, μ_2} satisfies the first order condition

$$0 \in \tau \partial \|Z_{i, \mu_2}\|_* + \frac{1}{\mu_1} \text{refold}_{(d_1, \dots, d_N)}(Z_{i, \mu_2} - \mathbf{Y}_0^{(i)}) + \frac{1}{\mu_2} \text{refold}_{(d_1, \dots, d_N)}(Z_i - \mathbf{X}_{\mu_2}^{(i)}).$$

Collecting the Z_{i, μ_2} terms and dividing by $\mu_1 + \mu_2$, we have

$$0 \in \left(\frac{\tau \mu_1 \mu_2}{\mu_1 + \mu_2} \right) \partial \|Z_{i, \mu_2}\|_* + \text{refold}_{(d_1, \dots, d_N)}(Z_{i, \mu_2}) - \text{refold}_{(d_1, \dots, d_N)} \left(\frac{\mu_2 \mathbf{Y}_0^{(i)} + \mu_1 \mathbf{X}_{\mu_2}^{(i)}}{\mu_1 + \mu_2} \right)$$

Hence Z_{i,μ_2} is an optimal solution of

$$\min_{Z_i} \left(\frac{\tau\mu_1\mu_2}{\mu_1 + \mu_2} \right) \|Z_i\|_* + \frac{1}{2} \left\| Z_i - \left(\frac{\mu_2 \mathbf{Y}_0^{(i)} + \mu_1 \mathbf{X}_{\mu_2}^{(i)}}{\mu_1 + \mu_2} \right) \right\|_{\mathbb{F}}^2,$$

of which the optimal solution is given by (3.32) by Lemma 2.57. Hence (3.32) holds.

Now, we derive the BCD update formula for \mathbf{X}_{μ_2} . If \mathbf{X}_{μ_2} is optimal for (3.31) when holding all the $Z_i = Z_{i,\mu_2}$ fixed for $i = 1, \dots, N$, then \mathbf{X}_{μ_2} satisfies the first order condition

$$\mathbf{0} = \frac{1}{N} \sum_{i=1}^N \frac{1}{\mu_2} \text{refold}(Z_{i,\mu_2} - \mathbf{X}_{\mu_2}^{(i)})$$

Since $\text{refold}(\mathbf{X}_{\mu_2}^{(i)}) = \mathbf{X}_{\mu_2}$ for all $i = 1, \dots, N$, we obtain (3.33).

Now, we aim to apply Theorem 3.28 to the objective function of (3.31). Let

$$f_i(Z_i) = \frac{1}{N} \left(\tau \|Z_i\|_* + \frac{1}{2\mu_1} \|Z_i - \mathbf{Y}_0^{(k)}\|_{\mathbb{F}}^2 \right)$$

and let

$$f_0(\mathbf{X}, Z_1, \dots, Z_N) = \frac{1}{N} \sum_{k=1}^N \left(\frac{1}{2\mu_2} \|Z_i - \mathbf{X}^{(i)}\|_{\mathbb{F}}^2 \right),$$

so that (3.31) can be rewritten as

$$\min_{\mathbf{X}, Z_1, \dots, Z_N} f_0(\mathbf{X}, Z_1, \dots, Z_N) + \sum_{k=1}^N f_i(Z_i).$$

Now, observe that f_0, \dots, f_N are continuous. Observe that if all but one of the variables $\mathbf{X}, Z_1, \dots, Z_N$ are fixed, the resulting function $g_0 : \mathbf{X} \mapsto f(\mathbf{X}, Z_1, \dots, Z_N)$ or $g_i : Z_i \mapsto f(\mathbf{X}, Z_1, \dots, Z_N)$ is quasiconvex and hemivariate. Since the objective function is nonnegative, it follows by Theorem 3.28 that any limit point of the cyclic BCD iterates is a coordinate-wise minimum of (3.31).

Now, suppose that $(X_{\mu_2}, Z_{1,\mu_2}, \dots, Z_{N,\mu_2})$ satisfies (3.32) and (3.33). Then $(X_{\mu_2}, Z_{1,\mu_2}, \dots, Z_{N,\mu_2})$ is a fixed point of the cyclic BCD iteration, so it must be a

coordinate-wise minimum. Next, by Lemma 3.27, it follows that $(X_{\mu_2}, Z_{1,\mu_2}, \dots, Z_{N,\mu_2})$ is a local minimizer of (3.31).

To prove the opposite implication, suppose that $(X_{\mu_2}, Z_{1,\mu_2}, \dots, Z_{N,\mu_2})$ is a local minimum of (3.31). It follows trivially that $(X_{\mu_2}, Z_{1,\mu_2}, \dots, Z_{N,\mu_2})$ is also a coordinate-wise minimum, so the first order conditions (3.32) and (3.33), as derived above, hold. \square

3.6.2 Fixed Point Iteration

We now state a fixed point algorithm for tensor completion, which is motivated by the derivation in the previous section. This algorithm essentially equivalent to the one in [148].

Algorithm 3.30.

```

1: procedure FIXED POINT ITERATION FOR TC( $\mathbf{X}_0, \mathcal{A}, y, \mu, \tau$ )
2:   while not converged do
3:      $Y_k = \mathbf{X}_k - \tau g(\mathbf{X}_k)$ 
4:     for  $i = 1, \dots, N$  do
5:        $Z_i = \text{shrink}_{\left(\frac{\tau\mu_1\mu_2}{\mu_1+\mu_2}\right)} \left( \frac{\mu_2 \mathbf{Y}_k^{(i)} + \mu_1 \mathbf{X}_k^{(i)}}{\mu_1 + \mu_2} \right)$ 
6:      $\mathbf{X}_{k+1} = \frac{1}{N} \sum_{i=1}^N \text{refold}(Z_i)$ 

```

3.6.3 Convergence

In [148], Yang et al. proved that \mathbf{X}_{μ_2} converges to the minimizer of the constrained problem

$$\min_{\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N} : \mathcal{A}(\mathbf{X})=y} \frac{1}{N} \|\mathbf{X}^{(i)}\|_*$$

as $\mu_1 = \mu_2 \rightarrow 0$. Here, we conduct a similar analysis, but instead we show that \mathbf{X}_{μ_2} converges to the minimizer \mathbf{X}_0 of (3.27) as $\mu_2 \rightarrow 0$ and μ_1 remains fixed. This result is important in cases where the data is noisy, since μ_1 controls the tradeoff between the norm $\|\mathbf{X}\|_*$ and the residual $\|\mathcal{A}(\mathbf{X}) - y\|_2$, as is a standard result of regularization theory [78].

We start by stating a simplified version of a useful lemma from [7, Theorem 9.2.2], which is also cited in [148].

Lemma 3.31 (Convergence of penalty method for constrained optimization). *Let $U \subset \mathbb{R}^d$ and let $f : U \rightarrow \mathbb{R}$ be continuous. Let $g_1, \dots, g_n : U \rightarrow \mathbb{R}$ be continuous and consider the constrained minimization problem*

$$\min_{x \in U : g_i(x)=0, i=1, \dots, n} f(x) \tag{3.34}$$

Let $\phi_i : U \rightarrow \mathbb{R}$ be continuous functions satisfying $\phi_i(0) = 0$ and $\phi_i(x) > 0$ if $x \neq 0$.

For all $\mu > 0$, consider the unconstrained minimization problem

$$\min_{x \in U} f(x) + \mu \sum_{i=1}^N \phi_i(g_i(x)) \tag{3.35}$$

Assume that there exists a compact set $U_0 \subset U$ such for all $\mu > 0$, every solution of x_μ of (3.35) satisfies $x_\mu \in U_0$. Let $\mu_1, \mu_2, \mu_3, \dots$ be any sequence satisfying

$\mu_i > 0$ and $\lim_{i \rightarrow \infty} \mu_i = \infty$, with corresponding solutions x_{μ_i} of (3.35). Then any accumulation point \bar{x} of x_{μ_i} is an optimal solution of (3.34), $\lim_{i \rightarrow \infty} f(x_{\mu_i}) = f(\bar{x})$, and

$$\lim_{i \rightarrow \infty} \mu \sum_{i=1}^n \phi_i(f_i(x_{\mu_i})) = 0.$$

In [148], it is shown that the map described by one iteration of the above algorithm is a contraction, and hence the algorithm converges under mild assumptions, provided that τ is sufficiently small. We briefly summarize these results.

The following is [148, Lemma 5.3].

Lemma 3.32. *Let $\mathcal{A}(\mathbf{X}) = A(\text{vec}(\mathbf{X}))$, where A is the matrix representation of \mathcal{A} . The map $T : \mathbf{X}_k \mapsto \mathbf{X}_{k+1}$ defined by the iteration in Algorithm 3.30 is a contraction, provided that $0 < \tau < \frac{2}{\|A'A\|_2}$. Furthermore, we have $\|\mathbf{X} - \mathbf{X}'\|_F = \|T(\mathbf{X}) - T(\mathbf{X}')\|_F$ if and only if $\mathbf{X} - \mathbf{X}' = T(\mathbf{X}) - T(\mathbf{X}')$.*

It is straightforward to verify that the proof of [148, Lemma 5.3] holds in the case $\mu_1 \neq \mu_2$. In the following theorem we generalize [148, Theorem 5.1] to the case in which $\mu_1 \neq \mu_2$.

Theorem 3.33. *Let $\mathcal{A}(\mathbf{X}) = A(\text{vec}(\mathbf{X}))$, where A is the matrix representation of \mathcal{A} , and assume $0 < \tau < \frac{2}{\|A'A\|_{s_\infty}}$. Assume that the map T defined by one iteration of Algorithm 3.30 has at least one fixed point \mathbf{X}^* satisfying $\mathbf{X}^* = T(\mathbf{X}^*)$. Then, for any starting point \mathbf{X}_0 , the sequence $\mathbf{X}_0, \mathbf{X}_1, \dots$ obtained from Algorithm 3.30 has at least one accumulation point, \mathbf{X}_{lim} . Furthermore, any such accumulation point is a solution of problem (3.31) with $\mathbf{Y}_0 = \mathbf{X}_{\text{lim}} - \tau \mathcal{A}(\mathcal{A}(\mathbf{X}_{\text{lim}}) - y)$.*

Proof. We follow the argument in [148]. Let \mathbf{X}^* be any fixed point of T . It follows that for all k , $\|\mathbf{X}_{k+1} - \mathbf{X}^*\| = \|T(\mathbf{X}_k) - T(\mathbf{X}^*)\|_{\mathbb{F}} \leq \|\mathbf{X}^k - \mathbf{X}^*\|_{\mathbb{F}}$, so the sequence $\|\mathbf{X}_{k+1} - \mathbf{X}^*\|$ is non-increasing. It follows that $\{\mathbf{X}_k\}$ is bounded, so it has a convergent subsequence. Hence, we may assume $\{\mathbf{X}_k\}$ converges to \mathbf{X}_{lim} .

Let $L = \lim_{k \rightarrow \infty} \|\mathbf{X}_{k+1} - \mathbf{X}^*\|$. It follows that $L = \|T(\mathbf{X}_{\text{lim}}) - T(\mathbf{X}^*)\|_{\mathbb{F}} = \|\mathbf{X}_{\text{lim}} - \mathbf{X}^*\|_{\mathbb{F}}$, so by Lemma 3.32 we have $T(\mathbf{X}_{\text{lim}}) - T(\mathbf{X}^*) = \mathbf{X}_{\text{lim}} - \mathbf{X}^*$. Since $T(\mathbf{X}^*) = \mathbf{X}^*$, it follows that $T(\mathbf{X}_{\text{lim}}) = \mathbf{X}_{\text{lim}}$.

Finally, we conclude that \mathbf{X}_{lim} is a solution of (3.31) with $\mathbf{Y}_0 = \mathbf{X}_{\text{lim}} - \tau \mathcal{A}^*(\mathcal{A}(\mathbf{X}_{\text{lim}}) - y)$ by Lemma 3.29. \square

Finally, we prove that in the limit as $\mu_2 \rightarrow 0$, we can obtain from Algorithm 3.30 a solution of (3.27). This result differs from [148, Theorem 5.3] in that we let $\mu_2 \rightarrow 0$ while fixing μ_1 , whereas they let $\mu_1 = \mu_2 \rightarrow 0$.

Theorem 3.34. *Let $\tau < \frac{2}{\|A^t A\|_2}$ and let $\mu_1 > 0$ be fixed. Assume that for all $\mu_2 > 0$, the map T has at least one fixed point. Let $\mu_{2,i} > 0$ be a sequence of values of μ_2 satisfying $\lim_{i \rightarrow \infty} \mu_{2,i} = 0$. For any set of starting values $\mathbf{X}_{i,0}$, $i = 1, 2, \dots$, let $\mathbf{X}_{i,\text{lim}}$ be an accumulation point of the iterates generated by Algorithm 3.30, starting at $\mathbf{X}_{i,0}$ and with $\mu_2 = \mu_{2,i}$, which exists by Theorem 3.33. Then, any accumulation point of $\{\mathbf{X}_{i,\text{lim}}\}$ is a solution of (3.27).*

Proof. Let \mathbf{X}_{lim} be an accumulation point of $\{\mathbf{X}_{i,\text{lim}}\}$. For all $i = 1, 2, \dots$, by Theorem 3.33 $\mathbf{X}_{i,\text{lim}}$ is a solution of (3.31) with \mathbf{Y}_0 replaced by $\mathbf{Y}_{i,0} = \mathbf{X}_{i,\text{lim}} - \tau g(\mathbf{X}_{i,\text{lim}})$. For convenience, we restate (3.31):

$$\min_{\mathbf{x}, Z_1, \dots, Z_N} \frac{1}{N} \sum_{k=1}^N \left(\tau \|Z_i\|_* + \frac{1}{2\mu_1} \|Z_i - \mathbf{Y}_{i,0}^{(k)}\|_{\mathbb{F}}^2 + \frac{1}{2\mu_{2,i}} \|Z_i - \mathbf{X}^{(i)}\|_{\mathbb{F}}^2 \right)$$

We view the first the first two terms

$$f(\mathbf{X}, Z_1, \dots, Z_N) = \frac{1}{N} \sum_{k=1}^N \left(\tau \|Z_i\|_* + \frac{1}{2\mu_1} \|Z_i - \mathbf{Y}_{i,0}^{(k)}\|_F^2 \right)$$

as an objective function and the last term

$$g(\mathbf{X}, Z_1, \dots, Z_N) = \frac{1}{N} \frac{1}{2\mu_{2,i}} \sum_{k=1}^N \|Z_i - \mathbf{X}^{(i)}\|_F^2$$

as a penalty function for the constraints $X^{(i)} = Z_i$. Since the objective and penalty functions satisfy the hypotheses of Lemma 3.31, it follows that any accumulation point of $\{X_{i,\text{lim}}\}$ is a solution of

$$\min_{\mathbf{x}, Z_1, \dots, Z_N} \frac{1}{N} \sum_{k=1}^N \left(\tau \|Z_i\|_* + \frac{1}{2\mu_1} \|Z_i - \mathbf{Y}_{i,0}^{(k)}\|_F^2 \right)$$

satisfying the constraints $\mathbf{X}^{(i)} = Z_i$ for $i = 1, \dots, N$. In the argument preceding (3.30), it was shown that (3.27) and (3.30) are equivalent. Since the above problem is the same as (3.30), we conclude that it is equivalent to (3.27). Hence, any accumulation point of $\{X_{i,\text{lim}}\}$ is a solution of (3.27), as was claimed. \square

3.6.4 Homotopy Path of Regularization Paramaters: Version 1

Our goal is to solve (3.27). By Theorem 3.34, the solution of (3.27) can be approximated arbitrarily well by the fixed-point iteration in Algorithm 3.30 for a decreasing sequence of values of μ_2 , while holding μ_1 fixed. However, in practice the algorithm often converges very slowly if initialized with very small values of μ_1 and μ_2 . This leads us to the following homotopy path in the parameters (μ_1, μ_2) .

Let μ_{init} be an initial starting value for both μ_1 and μ_2 . Let $0 < \eta < 1$ be fixed. Let $\mu_{1,\text{final}} > 0$ and $\mu_{2,\text{final}} > 0$ be fixed, where $\mu_{2,\text{final}} < \mu_{1,\text{final}}$.

Algorithm 3.35.

```

1: procedure FPC FOR TC( $\mathbf{X}_0, \mathcal{A}, y, \mu_{\text{init}}, \mu_{1,\text{final}}, \mu_{2,\text{final}}, \tau, \eta$ )
2:   Let  $k = 0$ ,  $\mu_1 \leftarrow \mu_{\text{init}}$ ,  $\mu_2 \leftarrow \mu_{\text{init}}$ , and  $\tau \leftarrow \tau_{\text{init}}$ .
3:   while  $\mu_2 > \mu_{2,\text{final}}$  do
4:      $\mu_1 \leftarrow \max(\mu_1 \eta, \mu_{1,\text{final}})$  ▷ Decrease  $\mu_1$  if above  $\mu_{1,\text{final}}$ 
5:      $\mu_2 \leftarrow \mu_2 \eta$  ▷ Decrease  $\mu_2$ 
6:     while not converged do
7:        $\mathbf{Y}_k = \mathbf{X}_k - \tau \mathcal{A}^* (\mathcal{A}(\mathbf{X}_k) - y)$ 
8:        $\mathbf{X}_{k+1} = \frac{1}{N} \sum_{i=1}^N \text{refold} \left( \text{shrink}_{\left(\frac{\tau \mu_1 \mu_2}{\mu_1 + \mu_2}\right)} \left( \frac{\mu_2 \mathbf{Y}_k^{(i)} + \mu_1 \mathbf{X}_k^{(i)}}{\mu_1 + \mu_2} \right) \right)$ 

```

In Algorithm 3.35, we start with $\mu_1 = \mu_2 = \mu_{\text{init}}$. At each iteration, both μ_1 and μ_2 are decreased by the factor η until $\mu_1 = \mu_2 < \mu_{1,\text{final}}$. Once μ_1 reaches its final value $\mu_{1,\text{final}}$, we continue decreasing μ_2 until μ_2 reaches its final value $\mu_{2,\text{final}}$. Under the assumptions of Theorem 3.34, Algorithm 3.35 is guaranteed to converge to a solution of problem (3.27), provided that $\mu_{2,\text{final}} = 0$.

While $\mu_{2,\text{final}}$ can be chosen to be very small, $\mu_{1,\text{final}}$ often needs to be significantly larger if the initial data y is noisy. Larger values of $\mu_{1,\text{final}}$ result in more noise filtering. A small value of $\mu_{2,\text{final}}$ guarantees that the matrices Z_i described previously agree well with the unfoldings $\mathbf{X}^{(i)}$ of the solution \mathbf{X} .

3.6.5 Homotopy Path of Regularization Paramaters: Version 2

We now aim to simplify the expressions in Algorithm 3.35. Observe that if we define

$$\tilde{\tau} = \frac{\tau\mu_2}{\mu_1 + \mu_2},$$

then

$$\frac{\mu_2 \mathbf{Y}_k^{(i)} + \mu_1 \mathbf{X}_k^{(i)}}{\mu_1 + \mu_2} = \frac{\mu_2(\mathbf{X}_k^{(i)} - \tau \mathcal{A}^*(\mathcal{A}(\mathbf{X}_k) - y)) + \mu_1 \mathbf{X}_k^{(i)}}{\mu_1 + \mu_2} \quad (3.36)$$

$$= \mathbf{X}_k^{(i)} - \tilde{\tau} \mathcal{A}^*(\mathcal{A}(\mathbf{X}_k) - y) \quad (3.37)$$

Also, the shrink amount in Algorithm 3.35 simplifies to:

$$\frac{\tau\mu_1\mu_2}{\mu_1 + \mu_2} = \tilde{\tau}\mu_1.$$

If τ is fixed and we decrease $\mu_1 = \mu_2$, observe that $\tilde{\tau}$ remains fixed. If μ_1 and τ is fixed, observe that $\tilde{\tau}$ decreases as μ_2 decreases. Hence the previous approach, which involves fixing τ first and decreasing both $\mu_1 = \mu_2$ and then decreasing μ_2 only, is equivalent to the following alternative approach. First fix $\tilde{\tau}$ and let μ_1 decrease to $\mu_{1,\text{final}}$. Then fix μ_1 and let $\tilde{\tau}$ decrease to $\tilde{\tau}_{\text{final}}$. This leads to the following algorithm. For brevity of notation, we have replaced $\tilde{\tau}$ with τ and μ_1 with μ .

Let μ_{init} be an initial starting value for μ . Let $0 < \eta < 1$ be fixed. Let $\mu_{\text{final}} > 0$ and $\tau_{\text{final}} > 0$ be fixed.

Algorithm 3.36.

```
1: procedure FPC FOR TC( $\mathbf{X}_0, \mathcal{A}, y, \mu_{\text{init}}, \tau_{\text{init}}, \tau_{\text{final}}, \eta$ )
2:   Let  $k = 0$ ,  $\mu \leftarrow \mu_{\text{init}}$ , and  $\tau \leftarrow \tau_{\text{init}}$ .
3:   while  $\tau > \tau_{\text{final}}$  do
4:      $\mu = \max(\mu\eta, \mu_{\text{final}})$ . ▷ Decrease  $\mu$  until it reaches  $\mu_{\text{final}}$ .
5:     if  $\mu = \mu_{\text{final}}$  then
6:        $\tau = \max(\tau\eta, \tau_{\text{final}})$  ▷ Then decrease  $\tau$  until it reaches  $\tau_{\text{final}}$ .
7:       while not converged do ▷ Run fixed point iteration.
8:          $\mathbf{Y}_k = \mathbf{X}_k - \tau \mathcal{A}^* (\mathcal{A}(\mathbf{X}_k) - y)$ 
9:          $\mathbf{X}_{k+1} = \frac{1}{N} \sum_{i=1}^N \text{refold}(\text{shrink}_{\tau\mu}(\mathbf{Y}_k))$ 
```

Algorithm 3.36 first fixes τ and lets μ decrease. Once μ has reached μ_{final} , μ is fixed and the algorithm decreases τ until it reaches τ_{final} . The above computations show that, up to possible differences in the values of the parameters, both Algorithm 3.35 and Algorithm 3.36 solve the same problem.

Under the assumptions of Theorem 3.34, Algorithm 3.36 is guaranteed to converge to a solution of problem (3.27), provided that $\tau_{\text{final}} = 0$. In our implementation, we use Algorithm 3.36.

3.7 Choosing the Parameter τ

By Theorem 3.33, we must have $\tau_{\text{init}} \leq \frac{2}{\|\mathcal{A}^*\mathcal{A}\|_2}$. If \mathcal{A} is a sub-Gaussian map, then for any \mathbf{X} , we have

$$\mathbb{E}\|\mathcal{A}(\mathbf{X})\|_{\text{F}}^2 = \|\mathbf{X}\|_{\text{F}}^2.$$

Hence, on average, \mathcal{A} behaves like an isometry. Theoretical results on isotropic sub-Gaussian random matrices, such as [145, Theorem 5.39] suggest that the singular values of \mathcal{A} remain near 1 with high probability.

In our simulations, $\tau_{\text{init}} = 0.01$ works well. Larger values of τ_{init} resulted in faster convergence. However, too large a value of τ_{init} causes divergence. We chose $\tau_{\text{final}} = 0.001$. Decreasing τ_{final} significantly does not appear to cause significant improvement in solution accuracy.

3.8 Choosing the Regularization Parameter μ

Recall that Algorithm 3.36 approximately solves (3.27), rewritten here with μ instead of μ_1 :

$$\min_{\mathbf{X}} \frac{1}{N} \sum_{k=1}^N \|\mathbf{X}^{(i)}\|_* + \frac{1}{2\mu} \|\mathcal{A}(\mathbf{X}) - y\|_2^2 \quad (3.38)$$

The parameter μ controls the weight of the regularizer $\|\mathbf{X}\|_*$. The value of μ controls the trade-off between the norm of the residual and the Tucker rank, as measured by $\|\mathbf{X}\|_*$. There are a number of standard techniques for choosing a regularization parameter, such as L -curve and the discrepancy principle [77, 78, 2]. However, such methods need not yield the optimal parameter μ and are inherently qualitative. The discrepancy principle requires a choice of the maximal ratio by which the residual can exceed its minimal value. The L -curve is based upon the intuitive assumption that the optimal parameter lies near the corner of the curve parametrizing the trade-off between $\log(\|\mathbf{X}\|_*)$ and $\log(\|\mathcal{A}(\mathbf{X}) - y\|_2)$. (The L -curve

is usually presented with $\|\mathbf{X}\|_2$ instead of $\|\mathbf{X}\|_*$.)

We propose to use k -fold cross-validation to choose μ . k -fold cross validation is a standard statistical technique for choosing an optimal predictive model for a set of inputs and outputs.

Assume that \mathcal{A} is defined via non-uniform sampling from a Parseval tight frame $\{u_i\}_{i=1}^n$ for a d -dimensional Hilbert space \mathcal{H} with probabilities $\{p_j\}_{j=1}^n$ by (3.25), rewritten here:

$$\mathcal{A}(x)[i] = \langle x, \frac{1}{\sqrt{m}} \epsilon_i \boldsymbol{\omega}_i \rangle,$$

where $P(\boldsymbol{\omega}_i = \frac{u_j}{p_j}) = p_j$ for all $j = 1, \dots, n$ and $\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_m$ are independent.

Consider the modified version of \mathcal{A} , denoted $\mathcal{A}_{\text{all observations}} : \mathbb{R}^{d_1 \times \dots \times d_N} \rightarrow \mathbb{R}^n$, that gives all the possible observations (excluding the Rademacher random variables):

$$\mathcal{A}_{\text{all observations}}(x)[j] = \langle x, \frac{1}{\sqrt{n}} \frac{u_i}{\sqrt{p_i}} \rangle,$$

for $\mathbf{X} \in \mathbb{R}^{d_1 \times \dots \times d_N}$ and $j = 1, \dots, n$. While \mathcal{A} only gives m out of n possible observations, $\mathcal{A}_{\text{all observations}}$ gives all possible observations.

Now consider the following learning problem. Given only one realization of the random operator \mathcal{A} , denoted \mathcal{A}_0 , and observations of the form $y_0 = \mathcal{A}(\mathbf{X}_0) + z$, where z is a noise vector satisfying $\|z\|_2 \leq \epsilon$, we aim to predict all possible observations, i.e., we aim to predict $\mathcal{A}_{\text{all observations}}(\mathbf{X})$.

Let \mathbf{X}_μ denote the solution of

$$\min_{\mathbf{X}} \frac{1}{N} \sum_{k=1}^N \|\mathbf{X}^{(k)}\|_* + \frac{1}{2\mu} \|\mathcal{A}_0(\mathbf{X}) - y_0\|_2^2 \quad (3.39)$$

We can now view $\mathcal{A}_{\text{all observations}}(\mathbf{X}_\mu)$ as a predictor of $\mathcal{A}_{\text{all observations}}(\mathbf{X}_0)$. Hence, we

aim to minimize the following expected squared generalization error:

$$\min_{\mu > 0} \mathbb{E}_{\mathcal{A}} [\|\mathcal{A}_{\text{all observations}}(\mathbf{X}_\mu) - \mathcal{A}_{\text{all observations}}(\mathbf{X}_0)\|_2^2] \quad (3.40)$$

The expectation is taken over all realizations of the random operator \mathcal{A} . In the above, \mathbf{X}_μ is the solution (3.39) and \mathbf{X}_0 is the true tensor. In contrast to the residual $\|y - \mathcal{A}(\mathbf{X}_\mu)\|_2$, which is monotonically decreasing as μ decreases, the above expectation estimates the true generalization error of all observations(\mathbf{X}_μ) as a predictor of $\mathcal{A}_{\text{all observations}}(\mathbf{X}_0)$.

It is not possible to directly solve (3.40), since \mathbf{X}_0 is unknown. However, we can estimate it via k -fold cross-validation as follows. Let $T = \{1, \dots, m\}$ and let $\omega_{i,0}$ and $\epsilon_{i,0}$ be the realizations of the random variables ω_i and ϵ_i occurring in the definition of \mathcal{A}_0 . For a fixed integer $k \geq 1$ (we use $k = 10$), partition T uniformly at random into k sets of almost-equal size, $T = T_1 \sqcup S_2 \sqcup \dots \sqcup T_k$, where \sqcup denotes the disjoint set union. For each $i = 1, \dots, k$, let $S_i = T \setminus T_i$, the set of indices in $\{1, \dots, m\}$ that are not in T_i . Suppose $|T_i| = t_i$ and $|S_i| = s_i$. Now, for each $i = 1, \dots, k$, let $S_i = \{l_1, \dots, l_{s_i}\}$, define $\mathcal{A}_i : \mathbb{R}^{d_1 \times \dots \times d_N} \rightarrow \mathbb{R}^{s_i}$ by

$$\mathcal{A}_i(\mathbf{X})[j] = \langle \mathbf{X}, \frac{1}{\sqrt{s_i}} \epsilon_{l_j,0} \omega_{l_j,0} \rangle \quad (3.41)$$

and define

$$y_i[j] = y_0[l_j] \quad (3.42)$$

for all $i = 1, \dots, k$ and for all $j = 1, \dots, s_i$. The result is to partition the observations occurring in \mathcal{A}_0 into k almost-equally sized sets T_i randomly. Each operator \mathcal{A}_i , $i = 1, \dots, k$, includes the observations corresponding to $k - 1$ of the sets T_i .

Now, for each $i = 1, \dots, k$, solve the problem

$$\mathbf{X}_{\mu,i} = \operatorname{argmin}_{\mathbf{X}} \frac{1}{N} \sum_{k=1}^N \|\mathbf{X}^{(i)}\|_* + \frac{1}{2\mu} \|\mathcal{A}_i(\mathbf{X}) - y_i\|_2^2 \quad (3.43)$$

Now we can view $\mathcal{A}_0(\mathbf{X}_{\mu,i})[T_i]$ as a predictor of $y[T_i]$, where recall that $x[T]$ is the vector with entries $x[i]$ for $i \in T$. For any μ , the estimated squared generalization error appearing inside the expectation in (3.40) can be approximated by

$$\text{GE}_{\text{est}}^2 = \frac{1}{k} \sum_{k=1}^n \|\mathcal{A}_0(\mathbf{X}_{\mu,i})[T_i] - y_0[T_i]\|_2^2. \quad (3.44)$$

Hence, μ can be chosen to minimize GE_{est}^2 .

We are thus led to the following k -fold cross-validation (CV) method for selecting μ . Let $\mu_1 > \mu_2 > \dots > \mu_l > 0$ be decreasing sequence of μ values.

Algorithm 3.37.

- 1: **procedure** SLOW k -FOLD CV FOR $\mu(\mathcal{A}_0, y_0, \{\mu_j\}_{j=1}^l)$
- 2: **for** $\mu = \mu_1, \dots, \mu_l$ **do**
- 3: **for** $i = 1, \dots, k$ **do**
- 4: Solve 3.43 for $\mathbf{X}_{\mu,i}$ using Algorithm 3.36.
- 5: Compute $\text{GE}_{\text{est}}^2(\mu)$ via (3.44).
- 6: $\mu \leftarrow \operatorname{argmin}_{\mu} \text{GE}_{\text{est}}^2(\mu)$

3.9 Accelerated k -Fold Cross-Validation for Selecting μ

Unfortunately, the above algorithm is very slow because it requires solving the tensor completion problem once for each μ and for each fold $i = 1, \dots, k$. Here, we show that it is possible to significantly accelerate k -fold cross-validation.

Recall that algorithm Algorithm 3.36 works by initializing $\mu = \mu_{\text{init}}$ and decreasing μ by a factor η repeatedly. By recording the residuals at each intermediate value of μ , it is possible to compute $\text{GE}_{\text{est}}^2(\mu)$ for $\mu = \mu_{\text{init}}\eta^j$ for $j = 0, \dots, l - 1$, where l is the number of μ values. However, to accelerate the algorithm, when doing k -fold cross-validation, we entirely skip the second stage during which τ is decreased. Based on simulations, decreasing τ does not appear to significantly reduce the generalization error. Hence, we propose the following accelerated k -fold cross validation algorithm.

The following algorithm estimates the best μ among $\{\mu_{\text{init}}\eta^j\}_{j=0}^{\text{num}_\mu-1}$.

Algorithm 3.38.

- 1: **procedure** $\mu = \text{FAST K-FOLD CV}(\mathbf{X}_0, \mathcal{A}_0, y_0, \mu_{\text{init}}, \text{num}_\mu, \eta, \tau)$
- 2: Randomly partition observation indices $\{1, \dots, m\}$ into k sets T_1, \dots, T_k and define \mathcal{A}_i and y_i as in (3.41) and (3.42).
- 3: **for** $i = 1, \dots, k$ **do**
- 4: \triangleright Solve TC problem with observations indexed by $\{1, \dots, m\} - T_i$, as follows:
- 5: Let $k = 0$, $\mu \leftarrow \mu_{\text{init}}$, and $\tau \leftarrow \tau_{\text{init}}$.
- 6: **for** $j = 0, \dots, \text{num}_\mu - 1$ **do**
- 7: $\mu \leftarrow \mu_{\text{init}} \eta^j$
- 8: **while** not converged **do** \triangleright Run fixed point iteration.
- 9: $\mathbf{Y}_k = \mathbf{X}_k - \tau \mathcal{A}_i^*(\mathcal{A}_i(\mathbf{X}_k) - y)$
- 10: $\mathbf{X}_{k+1} = \frac{1}{N} \sum_{i=1}^N \text{refold}(\text{shrink}_{\tau\mu}(\mathbf{Y}_k))$
- 11: \triangleright Estimate squared generalization error on T_i , as follows:
- 12: $\text{GE}_{\text{est}}^2[\mu_j, i] = \|\mathcal{A}_0(\mathbf{X}_{\mu, i})[T_i] - y[T_i]\|_2^2$.
- 13: \triangleright Aggregate the estimated squared generalization errors:
- 14: **for** $j = 0, \dots, \text{num}_\mu - 1$ **do**
- 15: $\text{GE}_{\text{est}}^2[\mu_j] = \frac{1}{k} \sum_{i=1}^k \text{GE}_{\text{est}}^2[\mu_j, i]$.
- 16: \triangleright Choose μ to minimize the estimated generalization error:
- 17: $\mu \leftarrow \text{argmin}_{j=0}^{\text{num}_\mu - 1} \{\text{GE}_{\text{est}}^2[\mu_j]\}$.

Algorithm 3.38 requires less than k times the computation of one tensor completion problem, as solved by Algorithm 3.36. The reason the computation time

is less than k multiples, and not exactly k , is that the second stage, in which τ is decreased, is omitted. Because each iteration of the first outermost loop in Algorithm 3.38 is independent of other iterations, the algorithm is highly parallelizable. When parallelized, the total computation time is less than that of solving one tensor completion problem via Algorithm 3.36, ignoring parallelization overhead. Even taking into account parallelization overhead, the run time of Algorithm 3.38 when parallelized on k cores is usually less than 1 run of Algorithm 3.36, due to the time savings from omitting the τ reduction stage.

Chapter 4: Nuclear Magnetic Resonance Preliminaries

4.1 Relaxometry

Nuclear magnetic resonance (NMR) relaxometry and related experiments can provide useful information about the chemical and physical properties of materials. Relaxometry experiments aim to study the properties of a given sample. Examples of samples of interest in medicine include cartilage [22, 125, 83, 96, 124, 128, 127, 126, 108, 129, 97, 91, 115], muscle [132, 21], and brain tissue [24, 23, 107]. Relaxometry has also been applied in other areas, such as in food science [80] for quality control and in oil logging (the detection of underground oil reserves) [61].

Quantum mechanical spin is associated with many atomic nuclei; these are the ones that we can study with NMR. The most widely-studied is the nucleus of the hydrogen atom, that is, the proton [58, 85], which is found in great abundance in water and hence in water-containing materials. When placed in a magnetic field, a particularly large magnetic moment is induced in water-containing materials [25], which is what we study here.

4.2 1-Dimensional Relaxometry Experiments

One dimensional NMR relaxometry experiments aim to quantify the distribution of a single parameter in a sample, such as the decay constant of longitudinal magnetization (T_1), the decay constant of transverse magnetization (T_2), the decay constant of longitudinal magnetization in a rotating frame ($T_{1,\rho}$), or apparent diffusion coefficient (ADC).

In experiments that measure the longitudinal relaxation time, T_1 , the evolution of longitudinal magnetization is described by the differential equation

$$\frac{d\left(m_{\text{longitudinal}}(\tau)\right)}{d\tau} = \frac{M_0 - m_{\text{longitudinal}}(\tau)}{T_1},$$

where $m_{\text{longitudinal}}(\tau)$ is the longitudinal magnetization, τ is time, and M_0 is the equilibrium value of induced magnetization given by Curie's law. For the case of an inversion recovery experiment, $m_{\text{longitudinal}}(\tau = 0) = -2M_0$, so after the inversion pulse,

$$m_{\text{longitudinal}}(\tau) = M_0 \left(1 - 2 \exp(-\tau/T_1)\right). \quad (4.1)$$

In general, the observed signal consists of the superposition of T_1 components, each described by (4.1); hence, it can be modeled as an integral

$$m_{\text{longitudinal}}(\tau) = \int_{T_1=0}^{\infty} f(T_1) \left(1 - 2 \exp(-\tau/T_1)\right) dT_1, \quad (4.2)$$

where $m_{\text{longitudinal}}(\tau)$ is the magnetization at time τ and $f : (0, \infty) \rightarrow [0, \infty)$ describes the distribution of T_1 constants in the sample. While f need not be normalized, it can be viewed as a distribution describing the relative sizes of the T_1 components in the sample.

In practice, the longitudinal magnetization $m_{\text{longitudinal}}(\tau)$ is sampled at a discrete set of times $\tau_1 < \tau_2, \dots < \tau_m$. The integral in equation (4.2) can be approximated by a discrete sum. Let $0 < T_{1,1} < T_{1,2} < \dots < T_{1,n}$ be a discretization of the T_1 values. Then (4.2) can be approximated as

$$m(\tau_i) = \sum_{j=1}^m f(T_{1,j}) (1 - \exp(-\tau/T_{1,j})) w[j], \quad (4.3)$$

where $w[j]$ are appropriately chosen quadrature weights.

Another 1-dimensional relaxometry experiment aims to measure the distribution of T_2 , the decay constant of transverse magnetization. The evolution of transverse magnetization is described by the differential equation

$$\frac{d(m_{\text{transverse}}(\tau))}{d\tau} = -\frac{m_{\text{transverse}}(\tau)}{T_2},$$

where $m_{\text{transverse}}(\tau)$ is the transverse magnetization and τ is time. In the case of Carr-Purcell-Meiboom-Gill (CPMG) experiments in which spin-echoes [74] are formed through the application of radiofrequency pulses at time interval $TE > 0$, the magnetization measured at the echo times iTE , for $i = 1, \dots, m$, is given by [25]

$$m_{\text{transverse}}(iTE) = M_0 \exp(-iTE/T_2), \quad (4.4)$$

where M_0 is defined in the same way as above. T_2 experiments have been applied in the analysis and quantification of cartilage [129].

As in the case of T_1 , a sample in general has a continuous distribution of T_2 values, so (4.4) can be replaced by

$$m_{\text{transverse}}(iTE) = \int_{T_2=0}^{\infty} f(T_2) \exp(-iTE/T_2) dT_2, \quad (4.5)$$

where $f(T_2)$ describes the distribution of T_2 values in the sample and the acquisition times $TE, 2TE, 3TE, \dots, mTE$ are discrete multiples of TE . As for T_1 , the integral (4.5) can be discretized. Let $0 < T_{2,1} < \dots < T_{2,n}$ be a discrete set of T_2 relaxation times. Then (4.5) can be approximated for $i = 1, \dots, m$ by

$$m_{\text{transverse}}(iTE) = \sum_{j=1}^n f(T_{2,j}) \exp(-iTE/T_{2,j}) w[j] \quad (4.6)$$

where $w[j]$ are quadrature weights.

Another parameter of interest is apparent diffusion coefficient (ADC), which quantifies the diffusion of moving particles, such as water molecules, in a sample [45]. In a diffusion experiment, the external diffusion sensitizing gradient strength g is varied and the resulting signal is observed. For a sample with only one apparent diffusion component equal to D , the observed signal can be modeled by the Stejskal-Tanner formula [81]:

$$m_{\text{diffusion}}(b) = \exp(-bD) \quad (4.7)$$

where

$$b(g) = \gamma^2 g^2 \delta^2 \left(\Delta - \frac{1}{3} \delta \right). \quad (4.8)$$

Here γ is the gyromagnetic ratio, δ is the duration of each gradient pulse, and Δ is the delay between pulses. The exact dependence of b on g depends, in general, on the specific experimental setup. However, (4.7) provides a simple equation relating the experimentally controlled acquisition variable, b , and the unknown ADC value D .

As for the T_1 and T_2 parameters, a given sample may contain a distribution

of ADC values. Hence we obtain the signal equation

$$m_{\text{diffusion}}(b) = \int_0^{\infty} f(D) \exp(-bD) dD \quad (4.9)$$

where $f : (0, \infty) \rightarrow [0, \infty)$ is the unknown distribution of ADC values in the sample, and D are the ADC values. As in the previous cases, this equation is discretized as:

$$m_{\text{diffusion}}(b_i) = \sum_{j=1}^m f(D_j) \exp(-b_i D_j) w[j] \quad (4.10)$$

where $w[j]$ are quadrature weights.

We have described three parameters of interest, T_1 , T_2 , and D . After appropriate substitutions, the distribution of each of these quantities in a sample can in general be obtained by solving the 1-dimensional Laplace transform-type equation

$$m(\tau) = \int_0^{\infty} f(t) \exp(-\tau/t) dt \quad (4.11)$$

or, in practice, its discretization

$$m(\tau_i) = \int_0^{\infty} f(t_j) \exp(-\tau_i/t_j) dt$$

Definition 4.1 (Laplace Transform). *Given an integrable function $f : [0, \infty) \rightarrow \mathbb{R}$, its Laplace transform $\mathcal{L}f$ is defined to be the function*

$$(\mathcal{L}f)(\tau) = \int_0^{\infty} f(\lambda) \exp(-\tau\lambda) d\lambda$$

Let $\lambda = \frac{1}{t}$ and $\tilde{f}(t) = f(\frac{1}{t})\frac{1}{t^2}$. By changing variables

$$m(\tau) = \int_0^{\infty} \tilde{f}(\frac{1}{t})\frac{1}{t^2} \exp(-\tau/t) dt \quad (4.12)$$

$$= \int_0^{\infty} \tilde{f}(\lambda) \exp(-\tau\lambda) d\lambda \quad (4.13)$$

$$= \mathcal{L}(\tilde{f})(\tau) \quad (4.14)$$

Hence (4.11) becomes

$$m = \mathcal{L}\tilde{f}.$$

It follows that f can be obtained indirectly by an inverse Laplace Transform (ILT). Indeed, \tilde{f} can be obtained as an ILT of m , and then f can be obtained via $f(t) = \tilde{f}(\frac{1}{t})\frac{1}{t^2}$.

In the context of NMR, we will primarily be concerned with the equation in the form (4.11), not with the standard definition of the Laplace transform, as in Definition 4.1.

The Laplace transform is an infinitely ill-conditioned operator, meaning that arbitrarily small changes in m can result in arbitrarily large changes in the solution f . As a result, solving for f is non-trivial and in general requires regularization techniques [78]. These problems carry over to the discretized case, where the discrete Laplace transform has rapidly decaying singular values, resulting in very large condition numbers [20].

4.3 Multidimensional Relaxometry and Related Experiments

Multidimensional NMR experiments aim to compute the joint density function f of one or more parameters. For example, in a T_1 - T_2 experiment, $f(t_1, t_2)$ denotes the joint distribution of T_1 and T_2 . In certain such experiments, the data $m(\tau_1, \tau_2)$ satisfies a multidimensional separable Laplace transform-type equation:

$$m(\tau_1, \tau_2) = \int_{T_1=0}^{\infty} \int_{T_2=0}^{\infty} f(T_1, T_2) (1 - 2 \exp(-\tau_1/T_2)) \exp(-\tau_2/T_2) dT_1 dT_2$$

Other 2-dimensional experiments quantify the joint distribution of parameters

such as $D - T_2$ or $T_1 - T_{1,\rho}$. Also, two dimensional T_2 -store- T_2 experiments quantify the joint distribution of T_2 with itself after a delay, which can be used to quantify the exchange between components in a sample. 2-dimensional experiments have seen growing applications in the chemical and biological sciences and permit a more complete description of materials [29] [80]. Applications of T_2 -store- T_2 include the the quantification of pore sizes in cement [109].

Celik, Bouhrara, Reiter, Fishbein, and Spencer have shown empirically that 2-dimensional relaxometry problems exhibit better stability than 1-dimensional problems [38], which provides an additional motivation to pursue higher dimensional experiments. These results suggest that N -dimensional experiments could provide even greater stability than 1 or 2-dimensional experiments. Because of the greater descriptive power and the potentially improved stability, it is of great value to have available higher dimensional NMR experiments for materials and tissue characterization.

A 3-dimensional experiment acquires the joint distribution of 3 parameters, such as T_1 , D , and T_2 . For one such an experiment, the observed data $m(\tau_1, b, \tau_2)$ is related to the distribution $f(T_1, D, T_2)$ of parameters by the equation

$$m(\tau_1, b, \tau_2) = \int_{T_1=0}^{\infty} \int_{D=0}^{\infty} \int_{T_2=0}^{\infty} f(T_1, D, T_2) (1 - 2 \exp(-\tau_1/T_1)) \exp(-bD) \exp(-\tau_2/T_2) dT_1 dD dT_2$$

In such an NMR experiment, the data $m(\tau_1, b, \tau_2)$ is acquired on a grid of values of τ_1 , b , and τ_2 . In many experiments, the acquisition time is proportional to the total number of such points. Hence, high dimensional experiments, in which the dimension N is 3 or greater, can take hours, days, or even weeks to acquire. This

long acquisition time has substantially limited applications of higher dimensional NMR relaxometry. By accelerating these experiments, compressive sensing offers to make possible many biomedical applications that were previously impractical.

4.4 General Mathematical Setup

Let $\mathbf{M} \in \mathbb{R}^{m_1 \times \dots \times m_N}$ be the observed data tensor and let $\mathbf{F} \in \mathbb{R}^{n_1 \times \dots \times n_N}$ be the discretized distribution of parameters to be solved for. We consider the following N -dimensional separable linear equation

$$\mathbf{M} = \mathbf{F} \otimes_1 K_1 \otimes_2 K_2 \cdots \otimes_N K_N + \mathbf{Z}, \quad (4.15)$$

where each kernel $K_i \in \mathbb{R}^{m_i \times n_i}$. The tensor $\mathbf{Z} \in \mathbb{R}^{m_1 \times \dots \times m_N}$ consists of noise. We assume that \mathbf{Z} satisfies $\|\mathbf{Z}\|_F \leq \epsilon$, for some $\epsilon > 0$. Equation (4.15) is equivalent to

$$\text{vec}(\mathbf{M}) = (K_N \otimes K_{N-1} \cdots \otimes K_1) \text{vec}(\mathbf{F}) + \text{vec}(\mathbf{Z}),$$

where $K_N \otimes K_{N-1} \cdots \otimes K_1$ denotes the Kronecker product of matrices.

Equation (4.15) is an example of an N -dimensional *discrete separable Fredholm integral equation of the first kind*.

For notational convenience, we define the operator $\mathcal{K} : \mathbb{R}^{n_1 \times \dots \times n_N} \longrightarrow \mathbb{R}^{m_1 \times \dots \times m_N}$ by

$$\mathcal{K}(\mathbf{X}) = \mathbf{X} \otimes_1 K_1 \otimes_2 K_2 \cdots \otimes_N K_N, \quad (4.16)$$

so that (4.15) can be rewritten as

$$\mathbf{M} = \mathcal{K}(\mathbf{F}) + \mathbf{Z}. \quad (4.17)$$

4.5 Description of the Kernels K_i

Each kernel K_i can be viewed as a discretization of a continuous kernel $\kappa_i : [0, \infty) \times (0, \infty) \rightarrow \mathbb{R}$. We parametrize the domain of κ_i by (τ_i, t_i) , where τ_i is the acquisition time corresponding to the i -th axis of the tensor M and t_i is the value of the parameter corresponding to the i -th axis of the tensor \mathbf{F} .

Let $0 < \tau_i[1] < \dots < \tau_i[m_i]$ be a fixed discretization of the acquisition time τ_i and let $0 < t_i[1] < \dots < t_i[n_i]$ be a fixed discretization of the parameter t_i . Then the kernel $K_i \in \mathbb{R}^{m_i \times n_i}$ is defined by

$$K_i[l, k] = \kappa(\tau_i[l], t_i[k])w_i[k], \quad (4.18)$$

where $1 \leq l \leq m_i$, $1 \leq k \leq n_i$, and $w_i[1], \dots, w_i[n_i]$ are quadrature weights. The quadrature weights can in general be defined by any quadrature rule, such as the trapezoidal rule, Simpson's rule, or Gaussian quadrature.

Standard practice in NMR is to discretize the points τ_i and t_i by either linear or logarithmic spacing.

Definition 4.2 (Linear spacing). *Given a parameter x , its m -point linearly-spaced discretization on the interval $[x_{\min}, x_{\max}]$ consists of the points*

$$x_j = x_{\min} + (x_{\max} - x_{\min}) \frac{j-1}{m-1},$$

for $j = 1, \dots, m$.

Definition 4.3 (Logarithmic spacing). *Let $0 < x_{\min} < x_{\max}$. Given a parameter x , its m -point logarithmically-spaced discretization on the interval $[x_{\min}, x_{\max}]$ consists*

of the points

$$x_j = x_{\min} \exp\left(\log\left(\frac{x_{\max}}{x_{\min}}\right) \frac{j-1}{m-1}\right),$$

for $j = 1, \dots, m$,

Although more sophisticated quadrature rules are available, for simplicity we use the following two rules. In the case of linear spacing, we let

$$w[j] = x_{j+1} - x_j = \frac{x_{\max} - x_{\min}}{m-1},$$

for $j = 1, \dots, m$. In the case of logarithmic spacing, we let

$$w[j] = x_{j+1} - x_j = x_j \left(\exp\left(\log\left(\frac{x_{\max}}{x_{\min}}\right) \frac{1}{m-1}\right) - 1 \right),$$

for $j = 1, \dots, m$. Although the point x_{m+1} is not defined, the right hand side of the above two equations can still be used to define w_m . The resulting quadrature rules approximate the integral on the slightly expanded interval $[x_{\min}, x_{\max} + \frac{x_{\max} - x_{\min}}{m-1}]$ in the linear spacing case and on a similarly expanded interval in the case of logarithmic spacing, although no generality is lost as an approximation on the interval $[x_{\min}, x_{\max}]$ would be obtained by letting j range from $1, \dots, m-1$, instead of $1, \dots, m$.

Hence, the kernels K_i are defined explicitly by (4.18), where the points $0 < \tau_i[1] < \dots < \tau_i[m_i]$ are linearly or logarithmically spaced on the interval $[\tau_{i,\min}, \tau_{i,\max}]$ and the points $0 < t_i[1] < \dots < t_i[m_i]$ are linearly or logarithmically spaced on the interval $[t_{i,\min}, t_{i,\max}]$. To be clear, each τ_i corresponds to the times at which experimental measurements are obtained while each t_i corresponds to a parametrization of one of the desired parameters (such as T_1 , T_2 , or D).

We will let $\mathcal{D}_i = [\tau_{i,\min}, \tau_{i,\max}] \times [t_{i,\min}, t_{i,\max}]$. It follows that the operator \mathcal{K} defined in (4.16) is a discretization of the continuous function $\kappa_1 \times \cdots \times \kappa_N$ restricted to the subdomain $\mathcal{D}_1 \times \cdots \times \mathcal{D}_N$. (In fact, it is an approximation on a slightly expanded subdomain, due to the choice of the quadrature weights.)

Each kernel κ_i , and hence its discretization K_i , depends on the specific experimental parameter t_i being measured at the observation times τ_i . In the case of a T_1 inversion recovery experiment, we have

$$\kappa_{T_1}(\tau, t) = 1 - 2 \exp(-\tau/t).$$

For a T_2 CPMG experiment, we have

$$\kappa_{T_2}(\tau, t) = \exp(-\tau/t).$$

For an ADC experiment, we have

$$\kappa_D(\tau, t) = \exp(-\tau t).$$

By appropriate transformations, each of these kernels κ_i can be reformulated as the kernel of a Laplace transform. Hence, after an appropriate transformation, the operator \mathcal{K} is a discretization of a multidimensional Laplace transform.

In cases in which the distribution \mathbf{F} exhibits details at very different scales of (t_1, \dots, t_N) , logarithmic spacing is useful to provide resolution at different time scales. Logarithmic spacing has also been observed to provide better conditioning, corresponding to a slower rate of decay of singular values in some problems [20]. For these reasons, the use of logarithmic spacing is standard in NMR relaxometry.

4.6 Nonnegative Least Squares (NNLS) and Tikhonov Regularization

A solution of the equation (4.17) can be sought by solving the following nonnegative least squares (NNLS) problem:

$$\min_{\mathbf{F} \in \mathbb{R}^{n_1 \times \dots \times n_N} : \mathbf{F} \geq 0} \|\mathbf{M} - \mathcal{K}(\mathbf{F})\|_{\mathbf{F}}^2, \quad (4.19)$$

where the operator \mathcal{K} is defined by (4.16).

The kernels K_i are highly ill-conditioned in relaxometry applications. Since the matrix for \mathcal{K} , when viewed as an operator on vectors $\text{vec}(\mathbf{X})$, takes the form of a Kronecker product $K_N \otimes \dots \otimes K_1$, by Lemma 2.15 the singular values of \mathcal{K} take the form of products $\prod_{i=1}^N \sigma_i[l_i]$, for each $(l_1, \dots, l_N) \in [\text{rank}(K_1)] \times \dots \times [\text{rank}(K_N)]$, where $\sigma_i[1], \dots, \sigma_i[\text{rank}(K_i)]$ are the singular values of K_i . Hence, if the singular values decay rapidly for each kernel K_i , the singular values will also decay rapidly for \mathcal{K} . Moreover, the condition number $\kappa(\mathcal{K}) = \prod_{i=1}^N \kappa(K_i)$. Hence, if each $\kappa(K_i)$ is bounded below by $O(\kappa_0)$, $\kappa(\mathcal{K})$ will be bounded below by $O(\kappa_0^N)$. Thus $\kappa(\mathcal{K})$ is extremely large and (4.19) is highly ill-conditioned. As a result, even small amounts of noise in the observed data \mathbf{M} can result in very large changes in the solution \mathbf{F} , resulting in a physically meaningless solution.

A standard technique to improve the conditioning of ill-posed inverse problems is regularization [78]. While there are a large number of different regularization techniques, a common idea in regularization is to introduce a term that penalizes excessively large solutions, with the goal of reducing the sensitivity of the solution to

small changes in the observed data. We consider the follow Tikhonov regularized nonnegative least squares (Tikhonov-NNLS) problem:

$$\min_{\mathbf{F} \in \mathbb{R}^{n_1 \times \dots \times n_N} : \mathbf{F} \geq 0} \|\mathbf{M} - \mathcal{K}(\mathbf{F})\|_{\mathbb{F}}^2 + \alpha^2 \|\mathbf{F}\|_{\mathbb{F}}^2, \quad (4.20)$$

The parameter $\alpha > 0$ describes the relative weight of the regularizer $\|\mathbf{F}\|_{\mathbb{F}}^2$, with larger values of α corresponding to larger amounts of regularization. Under appropriate statistical assumptions, increasing α results in a solution with smaller variance and larger bias. Hence the choice of α represents a trade-off between bias and variance.

4.7 Data Compression

In [144], Venkataramanan, Song, and Hürlimann introduced an accelerated numerical algorithm, hereafter called the VSH algorithm, to solve (4.20). The VSH algorithm exploits the rapid decay of the singular values of the kernel \mathcal{K} to project the observed data onto a lower-dimensional space. The projected data is then used to solve a smaller NNLS problem, greatly reducing computational time. While the algorithm was initially presented for 2-dimensional kernels of the form $K_2 \otimes K_1$, its derivation remains essentially unchanged for N -dimensional kernels of the form $K_N \otimes \dots \otimes K_1$.

Let $K_i = U_i S_i V_i^t$ be the singular value decomposition of each kernel K_i , where $U_i \in \mathbb{R}^{m_i \times r_i}$ has orthonormal columns, $S_i \in \mathbb{R}^{r_i \times r_i}$ is diagonal with diagonal entries $\sigma_i[1] \geq \sigma_i[2] \geq \dots \sigma_i[r_i] > 0$, $V_i \in \mathbb{R}^{n_i \times r_i}$ has orthonormal columns, and $r_i = \text{rank}(K_i) > 0$.

Let $\mathcal{U} : \mathbb{R}^{r_1 \times \dots \times r_N} \longrightarrow \mathbb{R}^{m_1 \times \dots \times m_N}$ denote the operator

$$\mathcal{U}(\mathbf{X}) = \mathbf{X} \otimes_1 U_1 \otimes_2 U_2 \cdots \otimes_N U_N.$$

and let $\mathcal{U}^t : \mathbb{R}^{m_1 \times \dots \times m_N} \longrightarrow \mathbb{R}^{r_1 \times \dots \times r_N}$ denote the operator

$$\mathcal{U}^t(\mathbf{X}) = \mathbf{X} \otimes_1 U_1^t \otimes_2 U_2^t \cdots \otimes_N U_N^t.$$

Then $\mathcal{U}^t \mathcal{U} = \text{Id}_{\mathbb{R}^{r_1 \times \dots \times r_N}}$ and $\mathcal{U} \mathcal{U}^t$ is a projection onto the range of \mathcal{K} .

Definition 4.4 (Compressed kernel and compressed data). *Define the compressed*

kernel $\tilde{\mathcal{K}} : \mathbb{R}^{r_1 \times \dots \times r_N} \longrightarrow \mathbb{R}^{n_1 \times \dots \times n_N}$ to be the operator defined by

$$\tilde{\mathcal{K}}(\mathbf{X}) = \mathbf{X} \otimes_1 (S_1 V_1^t) \otimes_2 (S_2 V_2^t) \cdots \otimes_N (S_N V_N^t).$$

Define the compressed data $\tilde{\mathbf{M}} \in \mathbb{R}^{r_1 \times \dots \times r_N}$ to be

$$\tilde{\mathbf{M}} = \mathcal{U}^t(\mathbf{M}).$$

The entry $\tilde{\mathbf{M}}[i_1, \dots, i_N] = M \otimes_1 U[:, i_1]^t \cdots \otimes_N U[:, i_N]^t$, and hence are the inner products of \mathbf{M} with the singular vectors of \mathcal{K} .

Lemma 4.5. *The problems (4.20) and*

$$\min_{\mathbf{F} \in \mathbb{R}^{n_1 \times \dots \times n_N} : \mathbf{F} \geq 0} \|\tilde{\mathbf{M}} - \tilde{\mathcal{K}}(\mathbf{F})\|_{\mathbb{F}}^2 + \alpha^2 \|\mathbf{F}\|_{\mathbb{F}}^2, \quad (4.21)$$

are equivalent, in the sense that both problems have the same set of solutions.

Proof. It suffices to verify that (4.20) and (4.21) have the same first order conditions.

First, observe that the first order conditions for (4.20) are

$$(\mathbf{F} \otimes (K_1^t K_1) \cdots \otimes (K_N^t K_N)) [i_1, \dots, i_N] + \alpha^2 \mathbf{F}[i_1, \dots, i_N] \geq 0, \quad (4.22)$$

for all $(i_1, \dots, i_N) \in [n_1] \times \dots \times [n_N]$, with equality whenever $\mathbf{F}[i_1, \dots, i_N] > 0$.

Similarly, the first order conditions for (4.21) are

$$(\mathbf{F} \otimes ((S_1 V_1^t)^t S_1 V_1^t) \cdots \otimes ((S_1 V_1^t)^t S_1 V_1^t)) [i_1, \dots, i_N] + \alpha^2 \mathbf{F}[i_1, \dots, i_N] \geq 0, \quad (4.23)$$

for all $(i_1, \dots, i_N) \in [n_1] \times \dots \times [n_N]$, with equality whenever $\mathbf{F}[i_1, \dots, i_N] > 0$.

Since each U_i has orthonormal columns, $K_i^t K_i = (U_i S_i V_i^t)^t U_i S_i V_i^t = V_i^t S_i U_i^t U_i S_i V_i^t = (S_i V_i^t)^t S_i V_i^t$. Hence the first order conditions (4.22) and (4.23) are equivalent.

□

The VSH algorithm solves the compressed problem (4.21) instead of the full problem (4.20). The primary advantage of solving the compressed problem instead of the full problem is a significant reduction in computational time, if $r_i = \text{rank}(K_i) \ll \min(m_i, n_i)$, where recall that $K_i \in \mathbb{R}^{m_i \times n_i}$.

In relaxometry applications, the kernels K_i are in fact full-rank; however, the singular values $\sigma_i[1], \dots, \sigma_i[r_i]$ decay rapidly. One way to greatly reduce computational time is to fix a threshold $0 < \rho < 1$, and replace each K_i with its low rank approximation $K_{i,\rho}$ obtained by retaining only the singular values $\sigma_i[j]$ satisfying $\sigma_i[j] \geq \rho \sigma_i[1]$. All of the analysis in this chapter holds in the case in which K_i is replaced by such an approximation; however, the rank of $K_{i,\rho}$ is greatly reduced, resulting in a significantly smaller compressed data tensor $\tilde{\mathbf{M}}$ and compressed kernel $\tilde{\mathcal{K}}$, and consequently faster computations. Replacing each K_i with $K_{i,\rho}$ and r_i with the rank of $K_{i,\rho}$, the analysis in this chapter is not changed. Another effect of thresholding, as we will see in Chapter 5, is that the tensor completion algorithm is significantly less computationally expensive, since the tensor being completed will

be the compressed tensor $\tilde{\mathbf{M}}$.

We now describe the VSH algorithm from [144]. We aim to solve the compressed problem (4.21). Let $f = \text{vec}(\mathbf{F})$ and let $m = \text{vec}(\mathbf{M})$. Let $\tilde{K}_i = S_i V_i^t$ and let $\tilde{K} = \tilde{K}_N \otimes \cdots \otimes \tilde{K}_1$. Let $r = r_1 r_2 \cdots r_N$ and let $n = n_1 \cdots n_N$. In [144], it is shown using the first order conditions of (4.21) that if we define

$$c = \frac{\tilde{K}f - m}{-\alpha^2},$$

then the optimal solution satisfies $f = \max(0, \tilde{K}^t c)$. Using this, it is derived in [144] that c solves:

$$\min_{c \in \mathbb{R}^r} \frac{1}{2} c^t (G(c) + \alpha^2 \text{Id}_{n \times n}) c - c^t m, \quad (4.24)$$

where

$$G = \tilde{K} \text{diag}(H(\tilde{K}[:, 1]^t c), \dots, H(\tilde{K}[:, n]^t c)) \tilde{K}^t$$

and $H(x) = 1$ if $x \geq 0$ and $H(x) = 0$ if $x < 0$.

The derivatives of the objective function in (4.24) are easily computed, so (4.24) can be solved rapidly using the inverse Newton method, as suggested in [144] and used by Cloninger and Czaja in [42]. We use a line-search method in our implementation.

Remark 4.6 (N-dimensional singular value thresholding). *When constructing the kernel \tilde{K} , it is not necessary to compute the full Kronecker product $\tilde{K}_i = S_i V_i^t$. Indeed, suppose $\rho > 0$ is a threshold (e.g., $\rho = 10^{-8}$). Since the SVD of \tilde{K} is of the form $(S_N \otimes \cdots \otimes S_1)(V_N \otimes \cdots \otimes V_1)^t$, the columns of \tilde{K} are of the form*

$$\left(\sigma_N [i_N] \cdots \sigma_1 [i_1] \right) v_N[:, i_N] \otimes \cdots \otimes v_N[:, i_1],$$

for $(i_N, \dots, i_1) \in [r_N] \times \dots \times [r_1]$. Let

$$\mathcal{I} = \{(i_N, \dots, i_1) \in [r_N] \times \dots \times [r_1] : \sigma_N[i_N] \cdots \sigma_1[i_1] \geq \rho \sigma_N[1] \cdots \sigma_1[1]\}.$$

Let $\mathcal{I} = \{\iota_1, \dots, \iota_r\}$ be an enumeration of the elements of \mathcal{I} . Now, let \tilde{K}_ρ be the matrix whose j -th column is

$$\left(\prod_{k=1}^N \sigma_k[\iota_j(k)] \right) v_N[:, \iota_j(N)] \otimes \dots \otimes v_1[:, \iota_j(1)].$$

Let $m_\rho = \tilde{M}[\mathcal{I}]$, where $\tilde{M}[\mathcal{I}]$ is the result of listing the entries of \tilde{M} indexed by \mathcal{I} into a column-vector lexicographically. Then, we obtain the truncated SVD problem:

$$\min_{f \geq 0} \|K_\rho f - m_\rho\|_2^2 + \alpha^2 \|f\|_2^2.$$

In practice, the region of $[r_1] \times \dots \times [r_N]$ corresponding to \mathcal{I} corresponds to an approximate triangle in 2 dimensions (or simplex) in higher dimensions. Thresholding the singular values in this manner further reduces computational time, even beyond thresholding the individual kernels K_i . Indeed, it is possible that two singular values $\sigma_1[i]$ of K_1 and $\sigma_2[j]$ of K_2 are above the threshold, but their product is below the threshold. Hence, applying N -dimensional thresholding further reduces the size of the kernel and accelerates the VSH algorithm. An additional advantage is noise-thresholding, although that is less significant since the Tikhonov regularizer also reduces the effect of noise.

Chapter 5: Compressed Sensing for Nuclear Magnetic Resonance Relaxometry

5.1 Overview

While there have been extensive applications of CS to magnetic resonance imaging (MRI) [104, 103, 105, 82] using various types of sparsity, we are not aware of any previous applications of CS to NMR relaxometry or related experiments other than the algorithm developed by Cloninger and Czaja in [42, 41], which was further validated in [3]. Unlike MRI, which requires Fourier methods, relaxometry problems require the solution of discrete Laplace transform-type equations. The work [19] applies sparsity-inducing regularizers to NMR relaxometry, but does not exploit multidimensional tensor structure or address the problem of data reconstruction from incomplete measurements.

In this chapter, we show how our tensor recovery results described in Chapter 3 can be applied to greatly accelerate N -dimensional NMR relaxometry data acquisition. While Cloninger and Czaja’s results apply only when $N = 2$, our results apply for all dimensions $N \geq 2$. Furthermore, our results support non-uniform sampling, while previous results required uniform sampling.

In N -dimensional relaxometry experiments, acquisition time, which can be hours, days, or weeks, has significantly limited previous biomedical applications. By accelerating NMR data acquisition, our new results promise to enable many new biomedical applications.

5.2 Compressed Sensing for 2-Dimensional NMR Relaxometry

In the 2-dimensional case, Cloninger and Czaja first developed an algorithm for solving equations of the form

$$M = F \otimes_1 K_1 \otimes_2 K_2 + Z$$

from incomplete noisy measurements of M using compressed sensing [42, 41]. Cloninger and Czaja observed that the solution of

$$\min_{F \geq 0} \|M - F \otimes_1 K_1 \otimes_2 K_2\|_F^2 + \frac{\alpha^2}{2} \|F\|_2^2$$

depends only on the compressed data $\tilde{M} = M \otimes_1 U_1^t \otimes_2 U_2^t$. Hence, the goal of a compressed sensing recovery algorithm should be to recover \tilde{M} . Since \tilde{M} is much smaller than M , it is reasonable to expect to recover it from a relatively small number of measurements. Moreover, in the noise-free setting, the compressed data \tilde{M} has rapidly decaying singular values. As a result, Cloninger and Czaja suggested the application of matrix completion to recover \tilde{M} from observations of the entries $M[i_1, i_2]$ for $(i_1, i_2) \in \Omega$, where $\Omega \subset [m_1] \times [m_2]$ is a randomly chosen set of entries of fixed cardinality $|\Omega|$.

For any measurement set $\Omega \subset [m_1] \times [m_2]$, define $\mathcal{A}_\Omega : \mathbb{R}^{r_1 \times r_2} \rightarrow \mathbb{R}^{|\Omega|}$ by

$$\mathcal{A}_\Omega(X) = (X \otimes_1 U_1 \otimes_2 U_2) [\Omega].$$

Recall that for a matrix M , $M[\Omega]$ denotes the column-vector obtained by arranging the entries $M[i_1, i_2]$ for which $(i_1, i_2) \in \Omega$ into a vector, according to the lexicographically ordering. Hence, the entries of $\mathcal{A}_\Omega(X)$ can be written as

$$X \otimes_1 U_1[i_1, :] \otimes_2 U_2[i_2, :]$$

for $(i_1, i_2) \in \Omega$.

Observe that in the noise free case, $\mathcal{A}_\Omega[\tilde{M}] = M[\Omega]$, so the operator \mathcal{A} applied to \tilde{M} is equivalent to the observation of M on a random subset Ω of its entries.

Cloninger and Czaja proposed solving the following problem to recover \tilde{M} .

$$\min_{\tilde{M} \in \mathbb{R}^{r_1 \times r_2}} \|\tilde{M}\|_* + \frac{1}{2\mu} \|\mathcal{A}_\Omega(\tilde{M}) - y\|_{\mathbb{F}}^2, \quad (5.1)$$

where $y = M[\Omega]$ and μ is a regularization constant. Cloninger and Czaja showed that this problem can be rapidly solved by a fixed point continuation iterative algorithm developed by Ma, Goldfarb, and Chen [106].

By extending a result by Yi-Kai Liu [102], which established that the RIP holds with high probability for random observations from an orthonormal measurement set, to the case of a Parseval tight Frame, Cloninger and Czaja proved that the solution of (5.1) approximately recovers \tilde{M} with high probability, provided that the number of measurements $|\Omega|$ is sufficiently large [42, 41].

5.3 Tensor Completion Applied to NMR

We now extend Cloninger and Czaja’s algorithm for the solution of discrete 2-dimensional Fredholm integral equations to the tensor case. We will present two algorithms: *reconstruction by slices* and *reconstruction by N -dimensional tensor completion*. The slice algorithm, which only applies when the data is randomly sampled in 2 out of N dimensions and fully sampled in the remaining $N - 2$ dimensions, uses standard matrix completion techniques to recover 2-dimensional slices. The N -dimensional tensor completion algorithm applies for any non-uniform sampling strategy and uses the tensor recovery methods discussed in Chapter 3

In N -dimensional NMR relaxometry and related experiments, measurements of the form

$$\mathbf{M}[i_1, \dots, i_N] = \mathcal{K}(\mathbf{F}) + \mathbf{Z}[i_1, \dots, i_n]$$

are acquired, where \mathbf{Z} is a noise tensor. In N -dimensional experiments, acquisition time, which can be days or weeks, presents a substantial challenge. We aim to accelerate data acquisition by first acquiring only a subset of the entries of \mathbf{M} and second recovering $\tilde{\mathbf{M}}$ using tensor completion. Since problems (4.20) and (4.21) are equivalent, $\tilde{\mathbf{M}}$ contains all the information needed to recover \mathbf{F} . Furthermore, in NMR applications, $\tilde{\mathbf{M}}$ is usually approximately equal to a tensor with low Tucker rank.

In the noise free case, the entry

$$\mathbf{M}[i_1, \dots, i_N] = \tilde{\mathbf{M}} \otimes_1 U[i_1, :] \cdots \otimes_N U_N[i_n, :]$$

$$= (\mathbf{M}, U[i_1, :]^t \circ \cdots \circ U[i_n, :]^t) \quad (5.2)$$

where $(\mathbf{X}, \mathbf{Y}) = (\text{vec}(X), \text{vec}(Y))$ denotes the Euclidean inner product of tensors and $U[i_1, :]^t \circ \cdots \circ U[i_n, :]^t \in \mathbb{R}^{r_1 \times \cdots \times r_N}$ is the outer product of the vectors $U[i_1, :]^t, \dots, U[i_N, :]^t$.

Since the columns of each U_1, \dots, U_N are orthonormal, it follows that the columns of $U_N \otimes \cdots \otimes U_1$ are also orthonormal. Hence, the rows of $U_N \otimes \cdots \otimes U_1$ form a Parseval tight frame of $\mathbb{R}^{r_1 \cdots r_N}$. Translating this result to tensors, it follows that the set of tensors $\{U[i_1, :]^t \circ \cdots \circ U[i_n, :]^t\}_{(i_1, \dots, i_N) \in [m_1] \times \cdots \times [m_N]}$ forms a Parseval tight frame for $\mathbb{R}^{r_1 \times \cdots \times r_N}$.

Hence, observations of individual entries of the raw relaxometry data, of the form $\mathbf{M}[i_1, \dots, i_N]$, are equivalent (in the noise-free case) to inner products of the compressed tensor $\tilde{\mathbf{M}}$ with the entries of a Parseval tight frame, as given by (5.2). This was a key observation of Cloninger and Czaja in [42] in the 2-dimensional case.

5.4 A Naive Approach: N -Dimensional NMR Data Recovery by Matrix Completion on 2-Dimensional Slices

We now present a first algorithm for reconstruction of N -dimensional NMR data via matrix completion on 2-dimensional slices. The algorithm only applies for *sliced* sampling strategies, i.e., random sampling in two axes and full sampling in the remaining axes. While sliced sampling is in general suboptimal, some experiments necessitate such sampling because subsampling in one of the axes does not save time. For example, in a $T_1 - D - T_2$ CPMG experiment, for each TI and b

point acquired, all the points in the TE axis are acquired. Hence, if we randomly subsample the $TI - D - TE$ grid, with high probability we will choose almost all of the TI and b values, leading to very little if any acceleration. For these experiments, we must randomly subsample only the TI and b directions in order to substantially reduce acquisition time.

While the slice algorithm uses less N -dimensional structure than our main algorithm presented in the next section, we present it as an example of how 2-dimensional methods, such as those introduced previously by Cloninger and Czaja, can be applied to higher dimensional problems.

The 2-dimensional slice sampling operator, which subsamples in axes 1 and 2 and fully samples in axes $3, \dots, N$, is defined as follows. Let $p \in \mathbb{R}^{m_1 \times m_2}$ be a probability distribution on $[m_1] \times [m_2]$ satisfying $p[i_1, i_2] > 0$ for all $(i_1, i_2) \in [m_1] \times [m_2]$. Let ι be a random variable with values in $[m_1] \times [m_2]$ satisfying $P(\iota = (i_1, i_2)) = p[i_1, i_2]$. For ι_1, \dots, ι_m i.i.d. drawings of ι and for $\epsilon_1, \dots, \epsilon_m$ i.i.d. Rademacher random variables, define $\mathcal{A}_{\text{slice}} : \mathbb{R}^{r_1 \times r_2} \rightarrow \mathbb{R}^m$ by

$$\mathcal{A}_{\text{slice}}(\mathbf{X})[i] = \langle \mathbf{X}, \frac{\epsilon_i}{\sqrt{p[\iota_i(1), \iota_i(2)]m}} U[\iota_i(1), :]^t \circ U[\iota_i(2), :]^t \rangle.$$

Assume that one realization of $\mathcal{A}_{\text{slice}}$ is fixed, once and for all. Assume that for each $\iota_1, \dots, \iota_m \in [m_1] \times [m_2]$ occurring in the the realization of $\mathcal{A}_{\text{slice}}$, the observations $\mathbf{M}[\iota_i(1), \iota_i(2), :, \dots, :]$ are acquired, for $i = 1, \dots, m$. Since \mathbf{M} is randomly sampled in axes 1 and 2 and fully sampled in axes $3, \dots, N$, the observations can be compressed along axes $3, \dots, N$. Define $\tilde{M}^{3, \dots, N} \in \mathbb{R}^{m_1 \times m_2 \times r_3 \times \dots \times r_N}$ by

$$\tilde{M}^{3, \dots, N} = \mathbf{M} \otimes_3 U_3^t \cdots \otimes_N U_N^t.$$

For each fixed $(j_3, \dots, j_N) \in [r_3] \times \dots \times [r_N]$, the given observations can be used to construct observations of $\tilde{M}^{3, \dots, N}$. The following algorithm works by apply matrix completion to each slice $\tilde{M}^{3, \dots, N}[:, :, j_3, \dots, j_N]$ for $(j_3, \dots, j_N) \in [r_3] \times \dots \times [r_N]$.

Algorithm 5.1.

- 1: **procedure** $\tilde{M} = \text{Tensor Completion for ND NMR by 2D Slices}(p, m, \epsilon)$
- 2: Choose $\epsilon_i \in \{0, 1\}$ and $\iota_i \in [m_1] \times [m_2]$, where $P(\iota_i = (i_1, i_2) = p[i_1, i_2])$.
- 3: Collect experimental observations $\mathbf{M}[\iota_i(1), \iota_i(2), :, \dots, :]$, for $i = 1, \dots, m$.
- 4: **for** $(j_3, \dots, j_N) \in [r_3] \times \dots \times [r_N]$ **do**
- 5: Define

$$\begin{aligned} y[i, j_3, \dots, j_N] &= \mathcal{A}_{\text{slice}}(\tilde{\mathbf{M}}[:, :, j_3, \dots, j_N])[i] \\ &= \frac{\epsilon_i}{\sqrt{p[\iota_i(1), \iota_i(2)]m}} \tilde{\mathbf{M}}^{3, \dots, N}[\iota_i(1), \iota_i(2), j_3, \dots, j_N] \end{aligned}$$

- 6: Reconstruct the slice $\tilde{\mathbf{M}}[:, :, j_3, \dots, j_N]$ by solving

$$\min_{X \in \mathbb{R}^{r_1 \times r_2} : \|\mathcal{A}_{\text{slice}}(X) - y[:, j_3, \dots, j_N]\|_2 \leq \epsilon} (\|X\|_*)$$

We first presented a slightly modified version of the slice algorithm in [73]. Since the observations are initially compressed along axes $3, \dots, N$, the slice algorithm only requires the solution of a relatively small number of matrix completion subproblems. A drawback of the slice algorithm is that some multidimensional structure is lost, since the problem is split into independent 2-dimensional problems. Furthermore, the sliced sampling strategy often exhibits suboptimal recovery compared to full N -dimensional random sampling. In the following section, we de-

scribe a full N -dimensional tensor completion algorithm for NMR, which uses the results introduced in Chapter 3.

5.5 Non-Uniform Sampling of N -Dimensional Relaxometry Data

With the same notation as in the previous section, assume that $p \in \mathbb{R}^{m_1 \times \dots \times m_N}$ defines a probability distribution on $[m_1] \times \dots \times [m_N]$, i.e., on the indices of the possible observations of the relaxometry data \mathbf{M} . Assume that $p[i_1, \dots, i_N] > 0$ for all $(i_1, \dots, i_N) \in [m_1] \times \dots \times [m_N]$. We consider the following non-uniform sampling strategy. Let $m \geq 1$ be a fixed number of measurements. Let ι be a random variable with values in $[m_1] \times \dots \times [m_N]$ satisfying

$$P(\iota = (i_1, \dots, i_N) = p[i_1, \dots, i_N]).$$

Let $\epsilon_1, \dots, \epsilon_m$ be i.i.d. Rademacher random variables and let ι_1, \dots, ι_m be m i.i.d. drawings of ι . that are independent of the ϵ_i 's. Now define the random sampling operator $\mathcal{A} : \mathbb{R}^{m_1 \times \dots \times m_N} \rightarrow \mathbb{R}^m$ by

$$\mathcal{A}(\mathbf{X})[i] = \left\langle \mathbf{X}, \frac{\epsilon_i U[\iota_i(1), :]^t \circ \dots \circ U[\iota_i(N), :]^t}{\sqrt{p[(\iota_i(1), \dots, \iota_i(N))m]}} \right\rangle. \quad (5.3)$$

We have thus arrived at the following algorithm for non-uniform sampling for N -dimensional relaxometry problems.

Algorithm 5.2.

- 1: **procedure** $\tilde{M} = \text{TENSOR COMPLETION FOR ND NMR}(p, m)$
- 2: Let $\epsilon_1, \dots, \epsilon_m$ be i.i.d. Rademacher random variables in $\{-1, 1\}$ and ι_1, \dots, ι_m be i.i.d. realizations of ι , as described above.
- 3: **for** $i = 1, \dots, m$ **do**
- 4: Acquire experimental observation $\mathbf{M}[\iota_i(1), \dots, \iota_i(N)]$.
- 5: $y[i] \leftarrow \mathcal{A}(\tilde{\mathbf{M}})[i] = \frac{\epsilon_i \mathbf{M}[\iota_i(1), \dots, \iota_i(N)]}{\sqrt{p[\iota_i(1), \dots, \iota_i(N)]m}}$.
- 6: Solve

$$\tilde{\mathbf{M}} = \operatorname{argmin}_{\mathbf{X} \in \mathbb{R}^{r_1 \times \dots \times r_N}} \left(\|\mathbf{X}\|_* + \frac{1}{2\mu} \|\mathcal{A}(\mathbf{X}) - y\|_2^2 \right). \quad (5.4)$$

In the last step of the above algorithm, the parameter μ can be chosen using accelerated k -fold cross-validation (Algorithm 3.38). The problem can then be solved using fixed point continuation (Algorithm 3.36).

Theorem 5.3 (Non-uniform sampling tensor recovery guarantee for N -dimensional NMR relaxometry). *There exists a constant C such that the following holds. Let*

$$\nu = \max_{(i_1, \dots, i_N) \in [m_1] \times \dots \times [m_N]} \frac{\prod_{k=1}^N \|U[i_k, :]\|_2^2}{p_i}.$$

Let $1 \leq s_j \leq r_j$, for $j = 1, \dots, N$. Define \mathcal{A} as above and let $\epsilon > 0$ and $p \in (0, 1)$.

Let $\delta < \delta_{\text{critical}}$, where δ_{critical} is the constant appearing in Theorem 3.20. If the number of measurements satisfies

$$m \geq C \frac{\mu^2}{\delta^2} \max \left\{ \log\left(\frac{N}{p}\right), \max_{k=1, \dots, N} \left\{ s_k \prod_{j \neq k} r_j + s_k r_k + \sum_{j \neq k} r_j^2 \right\} \log(N) \right\},$$

then with probability at least $1 - p$ over \mathcal{A} , for every $\tilde{\mathbf{M}}_0 \in \mathbb{R}^{r_1 \times \dots \times r_N}$ and $y \in \mathbb{R}^m$ satisfying $\|\mathcal{A}(\tilde{\mathbf{M}}_0) - y\|_2 \leq \epsilon$, the solution $\tilde{\mathbf{M}}^*$ of (5.4) satisfies

$$\|\tilde{\mathbf{M}}^* - \tilde{\mathbf{M}}_0\|_{\text{F}} \leq C_0 \epsilon + C_1 \left(\frac{1}{\frac{1}{N} \sum_{k=1}^N \sqrt{s_k}} \right) \|\tilde{\mathbf{M}}_0 - \tilde{\mathbf{M}}_{\text{best}}\|_*,$$

where $\tilde{\mathbf{M}}_{\text{best}}$ is the best rank (s_1, \dots, s_N) approximation of $\tilde{\mathbf{M}}_0$ in the $\|\cdot\|_*$ norm.

Proof. Since $U[i_1, :]^t \circ \dots \circ U[i_n, :]^t$ forms a Parseval tight frame, the result follows from Theorem 3.20. The formula for ν follows because for a rank one tensor $v_1 \circ \dots \circ v_N$, we have $\|v_1 \circ \dots \circ v_N\|_{\text{F}} = \prod_{k=1}^N \|v_k\|_{\text{F}}$. Applying this result, we have $\|U[i_1, :]^t \circ \dots \circ U[i_n, :]^t\|_{\text{F}}^2 = \prod_{k=1}^N \|U[i_k, :]\|_2^2$. The formula above for ν then follows from the formula given in Theorem 3.17. \square

Theorem 5.4 (Minimally coherent non-uniform sampling for NMR relaxometry).

The minimally coherent choice of sampling probabilities for relaxometry are given by

$$p[i_1, \dots, i_N] = \frac{\prod_{k=1}^N \|U[i_k, :]\|_2^2}{r_1 r_2 \dots r_N}.$$

Proof. The result follows directly from Theorem 3.21. Observe that d in that theorem is $r_1 r_2 \dots r_N$ here. \square

For $p_0 > 0$, an alternative sampling distribution is given by nearly minimally coherent random sampling, as in Definition 3.22:

$$p[i_1, \dots, i_N] = \frac{\frac{\prod_{k=1}^N \|U[i_k, :]\|_2^2}{r_1 r_2 \dots r_N} + \frac{p_0}{m_1 m_2 \dots m_N}}{1 + p_0}.$$

5.6 Choosing the Regularization Parameter α in NNLS

Recall that for relaxometry and related applications, we solve the Tikhonov regularized nonnegative least squares (NNLS) problem (4.20), rewritten here as

$$\min_{\mathbf{F} \in \mathbb{R}^{n_1 \times \dots \times n_N} : \mathbf{F} \geq 0} \|\mathbf{M} - \mathcal{K}(\mathbf{F})\|_{\mathbb{F}}^2 + \alpha^2 \|\mathbf{F}\|_{\mathbb{F}}^2,$$

using the VSH algorithm, as described in Chapter 4.

Of central importance in obtaining meaningful solutions is the appropriate selection of the regularization parameter α . Intuitively, the parameter α determines the degree to which noise is filtered out of the solution. Too small a value of α results in a solution with extremely high variance while too small a value of α results in extremely high bias, under appropriate statistical assumptions. Standard methods for the selection of α include the L -curve and the discrepancy principle [78].

Here, we use a different method, called generalized cross-validation. Generalized cross-validation (GCV) was first introduced by Craven and Wahba [46] and further developed by Golub, Heath, and Wahba [66]. The method aims to estimate the leave-one-out cross-validation error in the residual. We will apply GCV not to the full problem, but rather to the compressed problem

$$\min_f \|m - \tilde{K}f\|_2^2 + \alpha \|f\|_2^2, \quad (5.5)$$

where $f = \text{vec}(\mathbf{F})$, $m = \text{vec}(\mathbf{M})$, $\tilde{K}_i = S_i V_i^t$, and $\tilde{K} = \tilde{K}_N \otimes \dots \otimes \tilde{K}_1$. Let $r = r_1 r_2 \dots r_N$ and $n = n_1 \dots n_N$.

GCV aims to estimate the leave-one-out CV error, which is obtained by solving the problem with one $\tilde{m}[i]$ omitted and then measuring the squared error on that

point. By repeating r times, where $\tilde{m} \in \mathbb{R}^r$, and then averaging the r errors, we obtain an estimate of the leave-one-out error.

It might seem that estimating the GCV error would be computationally cumbersome, but for linear problems it can in fact be done by only solving one inversion problem for each α (rather than r problems for each α). It is shown in [2] that the GCV error for the compressed problem is approximated by

$$\text{GCV}_{\text{error}} \approx \frac{r \|Kf_\alpha - m\|_2^2}{\text{trace}(\text{Id}_{r \times r} - \tilde{K}\tilde{K}^\#)},$$

where

$$\tilde{K}^\# = (\tilde{K}^t \tilde{K} + \alpha^2 \text{Id}_{n \times n})^{-1} \tilde{K}^t.$$


We implement VSH inversion for a fixed set of values of α , and choose the value that minimizes the above approximation of $\text{GCV}_{\text{error}}$. By exploiting the simple SVD structure \tilde{K} , the above approximation of $\text{GCV}_{\text{error}}$ can be computed rapidly.

5.7 3-Dimensional Tensor Recovery on Simulated Data

We consider simulated 3-dimensional data. The kernels are of T_2 -type, i.e., $\kappa_i(\tau, t) = \exp(-\tau, t)$. In each dimension, we set the times τ to be 64 points logarithmically spaced on $[0.1, 1]$ and we set t to be 64 logarithmically spaced points on $[0.1, 1]$. Hence, \mathbf{M} and \mathbf{F} are tensors of size $64 \times 64 \times 64$. We threshold the kernels at $\rho = 1^{-8}$ according to Remark 4.6. times the largest singular value. The resulting compressed tensor is of size $8 \times 8 \times 8$. The true distribution \mathbf{F} consists of 3 Gaussian peaks with standard deviation 0.01. The peak positions are $(0.2, 0.3, 0.2)$, $(0.7, 0.2, 0.4)$, and $(0.5, 0.5, 0.6)$. The sampling ratios used are:

0.00015625, 0.0003125, 0.000625, 0.00125, 0.0025, 0.005, 0.01, 0.025, 0.05, 0.1, 0.25, and 0.5. For k -fold CV we set $k = 10$.

Figure 5.1: True simulated distribution F with peaks at $(0.2, 0.3, 0.2)$, $(0.7, 0.2, 0.4)$, and $(0.5, 0.5, 0.6)$



SNR values used are 256, 4096, and 16384. Data of given SNR is constructed as follows. Given the true uncompressed data \mathbf{M}_0 , we define $\mathbf{M} = \mathbf{M}_0 + \mathbf{E}$, where \mathbf{E} contains i.i.d. mean 0 Gaussian random variables with standard deviation $\sigma = \|\mathbf{M}\|_\infty / \text{SNR}$.

The true value of F is plotted in Figure 5.1.

5.7.1 Results at Sampling Ratio 0.0025 and $SNR = 16384$

The k -fold cross-validation estimated generalization error, at sampling ratio 0.0025 and $SNR = 16384$, is plotted in Figure 5.2. The GCV curve for α selection is shown in Figure 5.3. The inversion results for various values of α are plotted in Figure 5.4.

Figure 5.2: k -fold cross-validation curve at sampling ratio 0.0025 and $SNR = 16384$. Results from 5 out of the 10 folds are displayed to illustrate the variation in the solution. Each row shows inversion results for a different fold. Optimal $\mu = 3.10 \times 10^{-9}$. Optimal error = 9.7×10^{-4} .




Figure 5.3: Generalized CV curve for selection of α at sampling ratio 0.0025 and $SNR = 16384$. Optimal $\alpha = 4.64 \times 10^{-6}$.


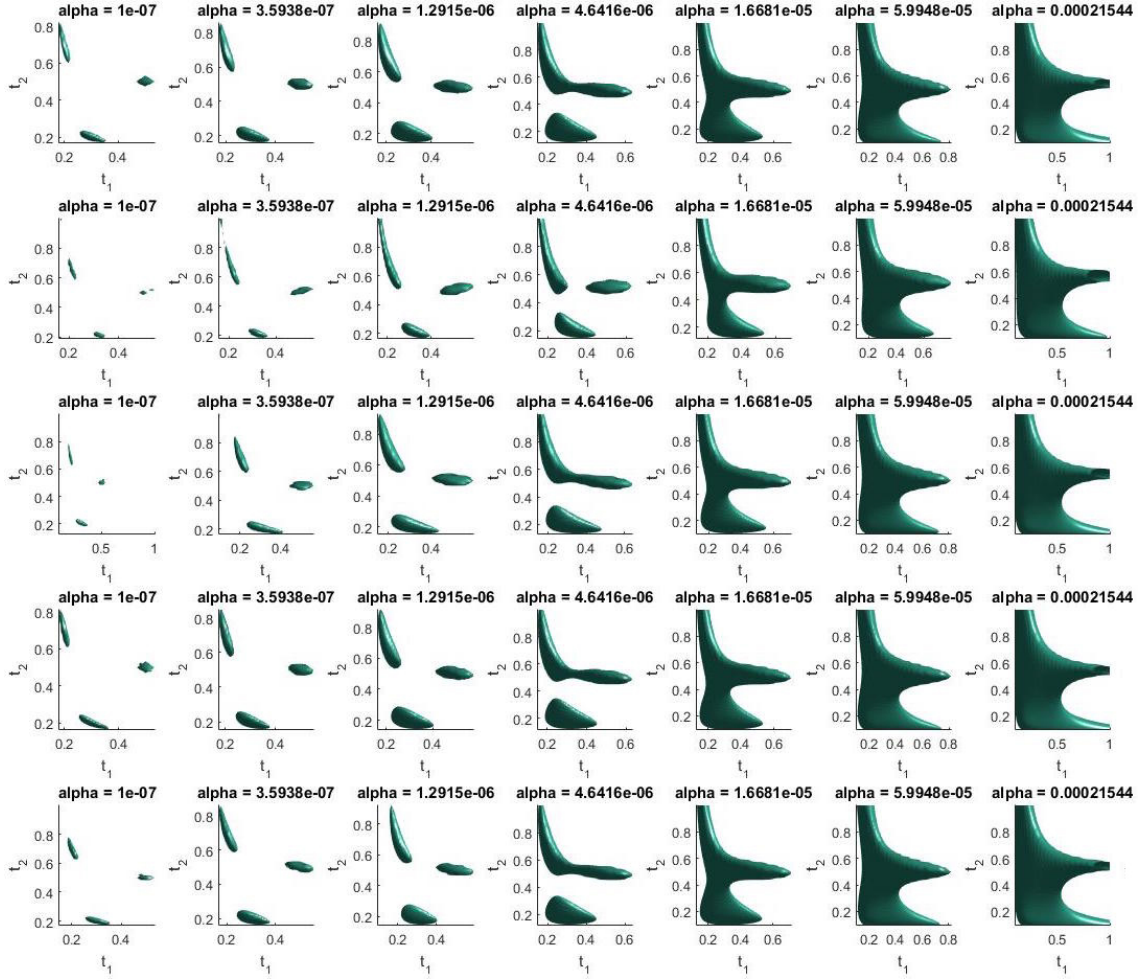


Figure 5.4: Inversion results for various α at ratio 0.0025 and $SNR = 16384$. Each row corresponds to a different fold in k -fold CV. The optimal α selected by GCV corresponds to the 4th column.



5.7.2 Results at Sampling Ratio 0.01 and $SNR = 16384$

We again plot the k -fold CV error, GCV curve, and inversion results for the simulated data, this time at sampling ratio 0.01. A comparison of Figures 5.4 and 5.7 shows that for fixed α , inversion results are more stable at higher sampling ratios.

Figure 5.5: k -fold cross-validation curve at sampling ratio 0.01 and $SNR = 16384$. Results from 5 out of the 10 folds are displayed to illustrate the variation in the solution. Each row shows inversion results for a different fold. Optimal $\mu = 5.06 \times 10^{-9}$. Optimal error = 2.53×10^{-4} .




Figure 5.6: Generalized CV curve for selection of α at sampling ratio 0.01 and $SNR = 16384$. Optimal $\alpha = 3.59 \times 10^{-7}$.


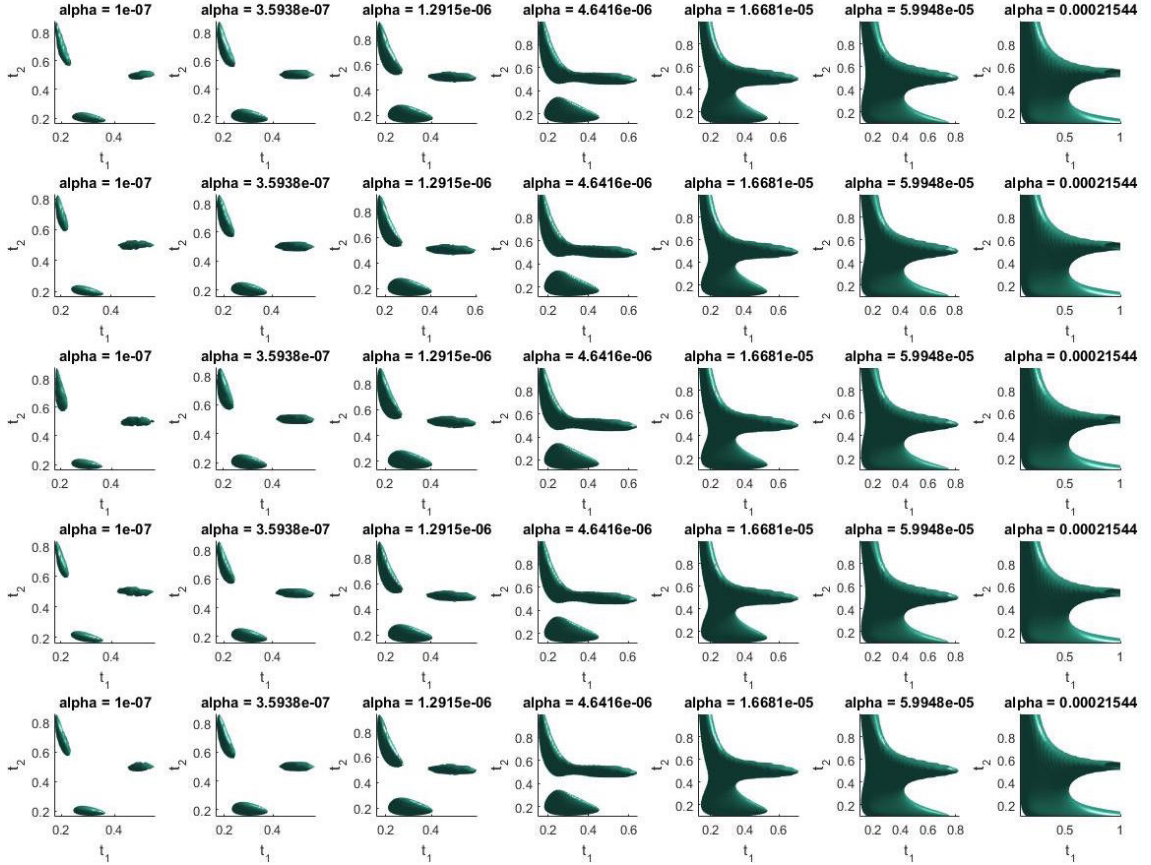


Figure 5.7: Inversion results for various α at ratio 0.01 and $SNR = 16384$. Each column corresponds to a different choice of α . Each row corresponds to a different fold in k -fold CV. The optimal α selected by GCV corresponds to the 2nd column.



5.7.3 Relative Errors vs. Sampling Ratio

We now plot the relative errors versus sampling ratio for tensor recovery at $SNR = 256, 4096, \text{ and } 16384$. The relative errors are calculated as $\|\tilde{\mathbf{M}} - \tilde{\mathbf{M}}_0\|_2 / \|\tilde{\mathbf{M}}_0\|_2$, where $\tilde{\mathbf{M}}_0$ is the true compressed tensor and $\tilde{\mathbf{M}}$ is the recovered compressed tensor.

A comparison of figures 5.8, 5.9, and 5.10 shows that relative error tends to decrease as SNR increases.

Figure 5.8: Relative tensor completion error versus sampling ratio at $SNR = 16384$.




Figure 5.9: Relative tensor completion error versus sampling ratio at $SNR = 4096$.





Figure 5.10: Relative tensor completion error versus sampling ratio at $SNR = 256$.



5.8 Application to $T_1 - D - T_2$ Experimental Data

We apply compressed sensing at sampling ratio 0.025 to $T_1 - D - T_2$ data acquired on an olive oil sample. The experiment was performed by Hasan Celik.

The data \mathbf{M} is $64 \times 64 \times 64$, the solution \mathbf{F} is $32 \times 32 \times 32$, and the compressed data \tilde{M} is $6 \times 4 \times 8$.

Figure 5.11: $T_1 - D - T_2$ experimental inversion result without CS and with CS at sampling ratio 0.025


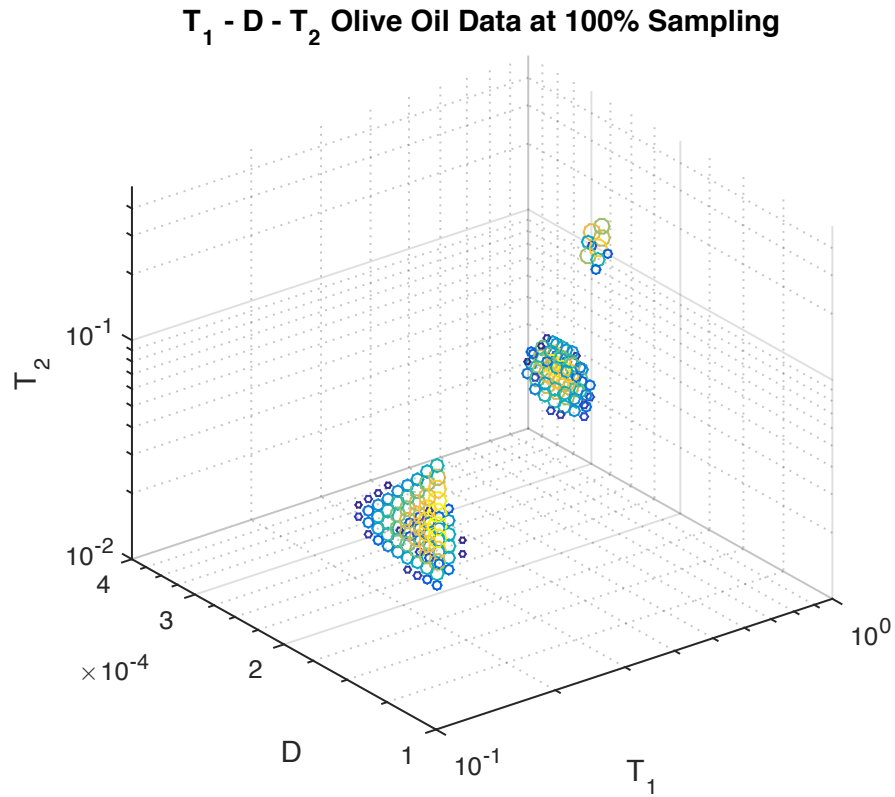


Figure 5.12: $T_1 - D - T_2$ experimental inversion result at sampling ratio 0.025.

Results from 5 out of the 10 folds are displayed to illustrate the variation in the solution. Each row shows inversion results for a different fold.

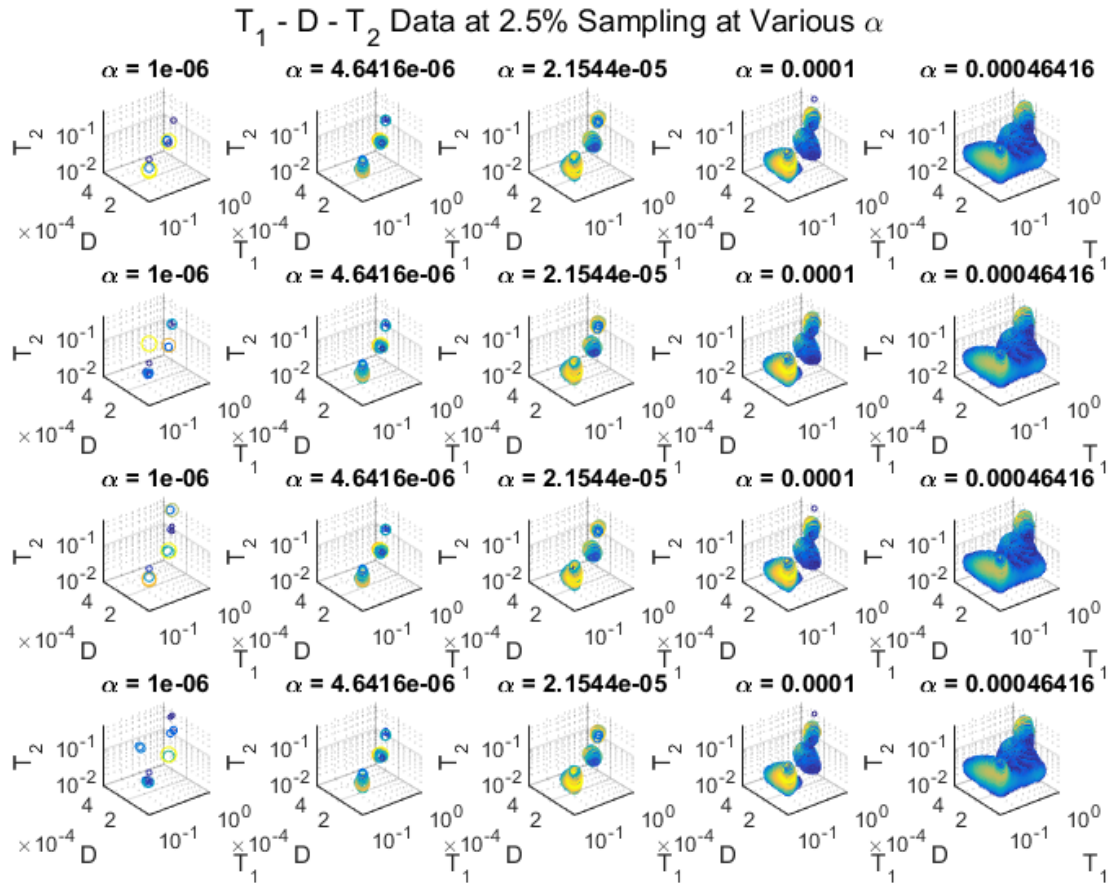


Figure 5.13: k -fold cross-validation curve at sampling ratio 0.025 and $SNR = 16384$ for $T_1 - D - T_2$ experimental data. Results from 5 out of the 10 folds are displayed to illustrate the variation in the solution. Each row shows inversion results for a different fold. The plateau on the left side suggests that the recovery accuracy is limited by noise in the data. This suggests that the number of samples acquired could likely be reduced without reducing recovery accuracy.

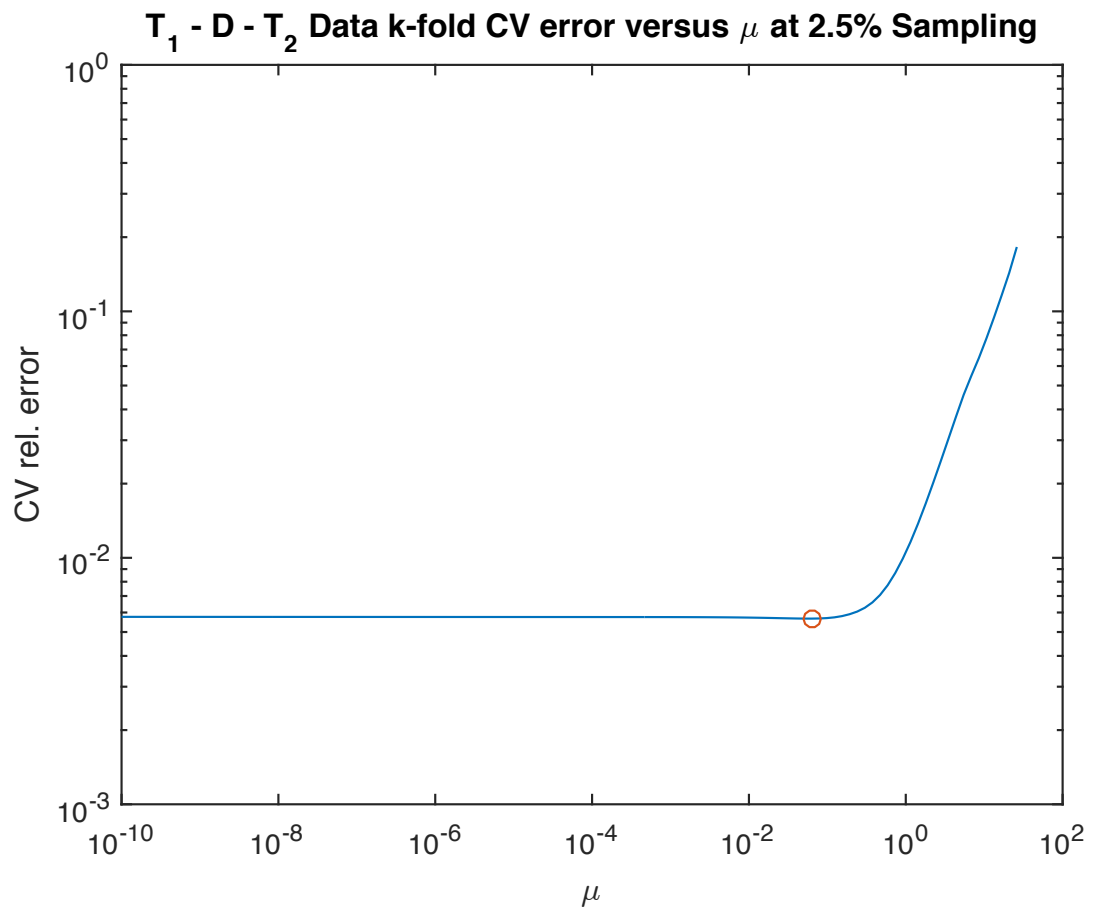
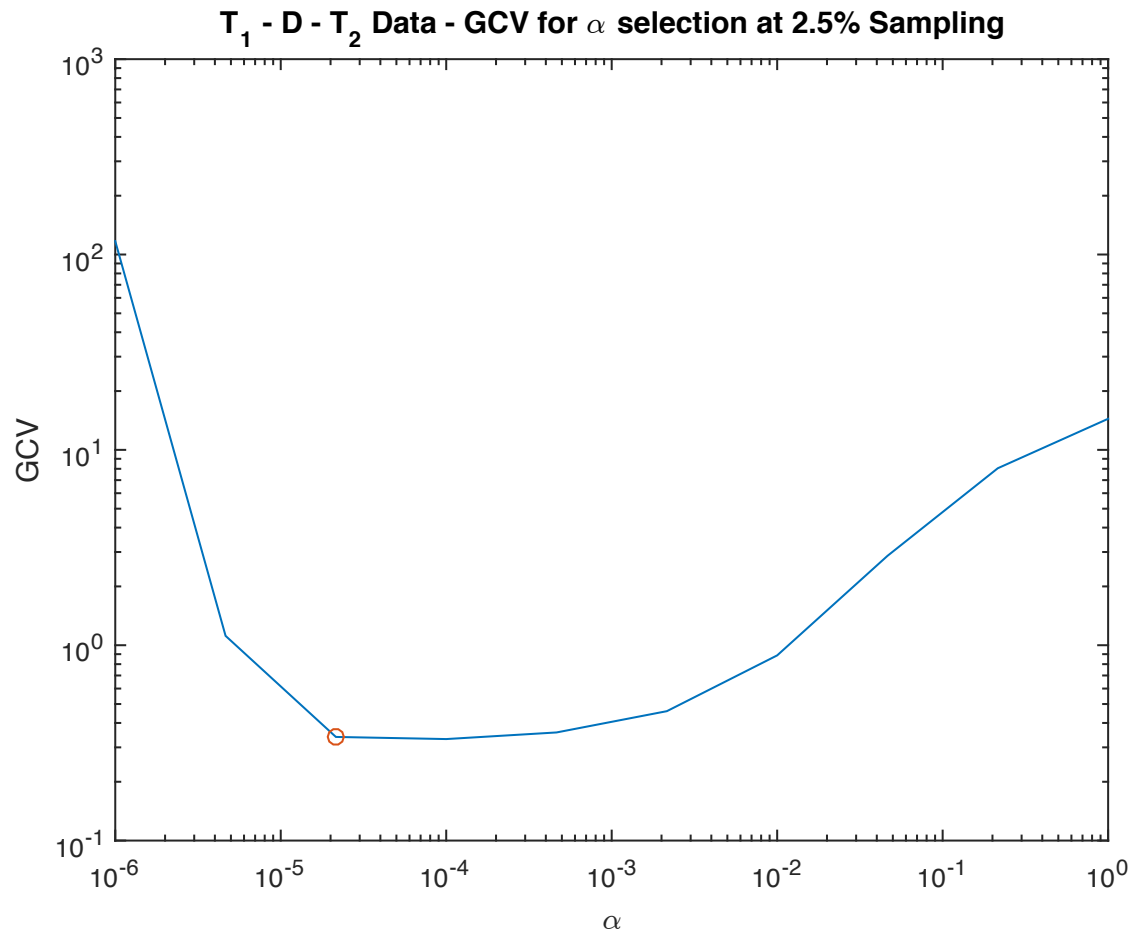


Figure 5.14: Generalized CV curve for $T_1 - D - T_2$ experimental data for selection of α at sampling ratio 0.025 and $SNR = 16384$



Chapter 6: Stability Results for Nonnegative Least Squares

6.1 Introduction

Recall that NMR relaxometry and related experiments require the solution of the Tikhonov regularized nonnegative least squares problem (4.20), rewritten here:

$$\min_{\mathbf{F} \in \mathbb{R}^{n_1 \times \cdots \times n_N} : \mathbf{F} \geq 0} \|\mathbf{M} - \mathcal{K}(\mathbf{F})\|_{\mathbb{F}}^2 + \frac{\alpha^2}{2} \|\mathbf{F}\|_{\mathbb{F}}^2,$$

In NMR relaxometry and related experiments, the kernels K_i are, after a suitable transformation, of Laplace transform type. In [38], Celik, Bouhrara, Reiter, Fishbein, and Spencer observed empirically that the solution of (4.20), in cases in which the solution has two distinct peaks, exhibits improved stability and resolution properties in 2 dimensions than in 1 dimension.

Let $y = \text{vec}(\mathbf{M}) \in \mathbb{R}^m$, $x = \text{vec}(\mathbf{F}) \in \mathbb{R}^n$, and $K = K_N \otimes \cdots \otimes K_N \in \mathbb{R}^{m \times n}$.

Then (4.20) becomes

$$\min_{x \in \mathbb{R}^n : x \geq 0} \|y - Kx\|_{\mathbb{F}}^2 + \frac{\alpha^2}{2} \|x\|_{\mathbb{F}}^2 \quad (6.1)$$

The form of problem (6.1) is independent of the dimension N of the problem (4.20). Hence, we will analyze the stability of (6.1). The resulting analysis will then apply to (4.20) for any value of $N \geq 1$.

We will analyze the stability of (6.1) with respect to perturbations in the data

y and in the regularization α . First, we show how standard estimates can be applied in the absence of the non-negativity constraint $x \geq 0$. Second, we show how those results can be extended to problem (6.1) with non-negativity constraints.

6.2 Stability of Least Squares

We quote a theorem, [142][Theorem 18.1], that describes the stability of the solution to least squares problems with respect to changes in the solution and the kernel. Recall that for a non-zero matrix K , its condition number is defined to be $\kappa(K) = \sigma_1(K)/\sigma_r(K)$, where $r = \text{rank}(K)$.

We now state a more general definition of condition number.

Definition 6.1 (Relative condition number). *Let X and Y be real finite-dimensional Hilbert spaces with norms $\|\cdot\|_X$ and $\|\cdot\|_Y$. The condition number of a function $f : X \rightarrow Y$ at $x \in X$ is given by*

$$\kappa_{f(x):x} = \lim_{\delta \rightarrow 0^+} \sup_{\tilde{x} : \|\tilde{x} - x\|_X \leq \delta} \left(\frac{\|f(\tilde{x}) - f(x)\|_Y}{\|f(x)\|_Y} \bigg/ \frac{\|\tilde{x} - x\|_X}{\|x\|_X} \right)$$

The relative condition number provides an upper bound on the amount by which small relative changes in x are scaled by the transformation f , as measured by the $\|\cdot\|_X$ and $\|\cdot\|_Y$ norms.

Theorem 6.2. *Let K be an $m \times n$ matrix of full rank and let $y \in \mathbb{R}^m$. Consider the least squares problem*

$$\min_{x \in \mathbb{R}^n} \|Kx - y\|_2^2 \tag{6.2}$$

with solution $x^* \in \mathbb{R}^n$. We define the following useful quantities:

$$\eta_K = \frac{\sigma_{\max}(K) \|x^*\|_2}{\|Kx^*\|_2} \quad (6.3)$$

and

$$\cos \theta_K = \frac{\|Kx^*\|_2}{\|y\|_2}. \quad (6.4)$$

Assume that the space of solutions $x \in \mathbb{R}^n$ is equipped with the $\|\cdot\|_2$ norm, the space containing the initial data $y \in \mathbb{R}^m$ is equipped with the $\|\cdot\|_2$ norm, and the space of kernels $K \in \mathbb{R}^{m \times n}$ is equipped with the matrix norm $\|\cdot\|_2$. Then the following hold:

1. The relative condition number of the solution $x^* \in \mathbb{R}^n$ as a function the initial data $y \in \mathbb{R}^m$ satisfies

$$\kappa_{x^*:y} \leq \frac{\kappa(K)}{\eta_K \cos \theta_K}. \quad (6.5)$$

2. The relative condition number of the solution $x^* \in \mathbb{R}^n$ as a function of the kernel $K \in \mathbb{R}^{m \times n}$ satisfies

$$\kappa_{x^*:K} \leq \kappa(K) + \frac{\kappa(K)^2 \tan \theta_K}{\eta_K}. \quad (6.6)$$

Recall that for a non-zero matrix $K \in \mathbb{R}^{m \times n}$, $\sigma_{\min}(K)$ denotes the smallest nonzero singular value of K and $\sigma_{\max}(K)$ denotes the largest singular value of K . The matrix norm $\|K\|_2$ is equivalent to the largest singular value of K , i.e., $\|K\|_2 = \sigma_{\max}(K)$.

Theorem 6.2 can be applied to quantify the stability of unconstrained least squares problems with Tikhonov regularization, as shown in the following corollary. Recall that the augmented kernel K_{aug} and the augmented data y_{aug} , as defined in

Lemma 2.6, are given by

$$K_{\text{aug}} := \begin{pmatrix} K \\ \alpha \text{Id}_{n \times n} \end{pmatrix} \quad (6.7)$$

and

$$y_{\text{aug}} := \begin{pmatrix} y \\ 0_{n \times 1} \end{pmatrix}. \quad (6.8)$$

Corollary 6.3. *Let K be an $m \times n$ matrix of full rank. Let $y \in \mathbb{R}^m$ and $\alpha > 0$.*

Consider the Tikhonov-regularized least squares problem

$$\min_{x \in \mathbb{R}^n} \|Kx - y\|_2^2 + \alpha^2 \|x\|_2^2. \quad (6.9)$$

Assume that the space of solutions $x \in \mathbb{R}^n$ is equipped with the $\|\cdot\|_2$ norm, the space containing the initial data $y \in \mathbb{R}^m$ is equipped with the $\|\cdot\|_2$ norm, and the space of regularization constants $\alpha \in \mathbb{R}$ is equipped with the absolute value $|\cdot|$ norm. Then the following hold:


1. *The relative condition number of the solution $x^* \in \mathbb{R}^n$ of (6.9) as a function of the initial data $y \in \mathbb{R}^m$ satisfies*

$$\kappa_{x^*:y} \leq \frac{\|y\|_2}{\sigma_{\min}(K_{\text{aug}})\|x^*\|_2}. \quad (6.10)$$

2. *The relative condition number of the solution $x^* \in \mathbb{R}^n$ of (6.9) as a function of the regularization constant $\alpha > 0$ satisfies*

$$\kappa_{x^*:\alpha} \leq \frac{\alpha}{\sigma_{\min}(K_{\text{aug}})} + \frac{\alpha\|y_{\text{aug}} - K_{\text{aug}}x^*\|_2}{\sigma_{\min}^2(K_{\text{aug}})\|x^*\|_2} \quad (6.11)$$

Figure 6.1: θ_K as defined in Theorem 6.2 is the angle between Kx^* and y .



See [142, Figure 18.1].

Proof. By Lemma 2.6, (6.9) is equivalent to:

$$\min_{x \in \mathbb{R}^n} \|K_{\text{aug}}x - y_{\text{aug}}\|_2^2. \quad (6.12)$$

For both parts of the proof, we apply Theorem 6.2 to the problem (6.12).

We first prove (6.10). The transformation $y \mapsto y_{\text{aug}}$ is an isometry. Hence, if y is perturbed to \tilde{y} , we have $\frac{\|\tilde{y}-y\|_2}{\|y\|_2} = \frac{\|\tilde{y}_{\text{aug}}-y_{\text{aug}}\|_2}{\|y_{\text{aug}}\|_2}$. It thus follows from Definition 6.1, of relative condition number, that $\kappa_{x^*:y} \leq \kappa_{x^*:y_{\text{aug}}}$, where $\kappa_{x^*:y}$ is the relative condition number of x^* with respect to changes in y and $\kappa_{x^*:y_{\text{aug}}}$ is the relative condition number of x^* with respect to changes in y_{aug} . By Theorem 6.2,

$$\begin{aligned} \kappa_{x^*:y} &\leq \kappa_{x^*:y_{\text{aug}}} \\ &\leq \frac{\kappa(K_{\text{aug}})}{\eta(K_{\text{aug}}) \cos \theta(K_{\text{aug}})} \\ &= \frac{\sigma_{\max}(K_{\text{aug}})}{\sigma_{\min}(K_{\text{aug}})} \frac{\|K_{\text{aug}}x^*\|}{\|K_{\text{aug}}\| \|x^*\|} \frac{\|y\|}{\|K_{\text{aug}}x^*\|} \\ &= \frac{\|y\|_2}{\sigma_{\min}(K_{\text{aug}}) \|x^*\|_2}, \end{aligned}$$

where $\kappa(K_{\text{aug}})$, $\eta_{K_{\text{aug}}}$, and $\theta_{K_{\text{aug}}}$ are defined in the same way as was defined for K

in Theorem 6.2. This proves (6.10).

Next, we prove (6.11). If α is perturbed to $\tilde{\alpha}$, the augmented kernel

$$K_{\text{aug}} = \begin{pmatrix} K \\ \alpha \text{ Id}_{n \times n} \end{pmatrix}$$

is perturbed to

$$\tilde{K}_{\text{aug}} := \begin{pmatrix} K \\ \tilde{\alpha} \text{ Id}_{n \times n} \end{pmatrix}.$$

Hence, we can view a perturbation of α as a perturbation of the kernel K_{aug} . Since

$$\tilde{K}_{\text{aug}} - K_{\text{aug}} = \begin{pmatrix} \mathbb{0}_{m \times n} \\ (\tilde{\alpha} - \alpha) \text{ Id}_{n \times n} \end{pmatrix}, \quad (6.13)$$

we have $\frac{\|\tilde{K}_{\text{aug}} - K_{\text{aug}}\|_2}{\|K_{\text{aug}}\|_2} = \frac{|\tilde{\alpha} - \alpha|}{\|K_{\text{aug}}\|_2} = \frac{|\tilde{\alpha} - \alpha|}{\alpha} \frac{\alpha}{\|K_{\text{aug}}\|_2}$. It thus follows from Definition 6.1, of

relative condition number, that $\kappa_{x^*:\alpha} \leq (\kappa_{x^*:K_{\text{aug}}}) \frac{\alpha}{\|K_{\text{aug}}\|_2}$, where $\kappa_{x^*:\alpha}$ is the relative

condition number of x^* with respect to changes in α and $\kappa_{x^*:K_{\text{aug}}}$ is the relative

condition number of x^* with respect to changes in K_{aug} . By Theorem 6.2,

$$\begin{aligned} \kappa_{x^*:\alpha} &\leq \kappa_{x^*:K_{\text{aug}}} \left(\frac{\alpha}{\|K_{\text{aug}}\|_2} \right) \\ &\leq \left(\kappa(K_{\text{aug}}) + \frac{\kappa(K_{\text{aug}})^2 \tan(\theta_{K_{\text{aug}}})}{\eta K_{\text{aug}}} \right) \left(\frac{\alpha}{\|K_{\text{aug}}\|_2} \right) \\ &= \left(\frac{\sigma_{\max}(K_{\text{aug}})}{\sigma_{\min}(K_{\text{aug}})} + \frac{\sigma_{\max}(K_{\text{aug}})^2 \|y_{\text{aug}} - K_{\text{aug}}x^*\|_2}{\sigma_{\min}(K_{\text{aug}})^2 \|K_{\text{aug}}x^*\|_2} \frac{\|K_{\text{aug}}x^*\|_2}{\sigma_{\max}(K_{\text{aug}})\|x^*\|_2} \right) \left(\frac{\alpha}{\|K_{\text{aug}}\|_2} \right) \\ &= \frac{\alpha}{\sigma_{\min}(K_{\text{aug}})} + \frac{\alpha \|y_{\text{aug}} - K_{\text{aug}}x^*\|_2}{\sigma_{\min}^2(K_{\text{aug}})\|x^*\|_2} \end{aligned}$$

To obtain the second to last line, we used that $\tan(\theta_{K_{\text{aug}}}) = \frac{\|y_{\text{aug}} - K_{\text{aug}}x^*\|}{\|K_{\text{aug}}x^*\|}$, since

$\cos(\theta_{K_{\text{aug}}}) = \frac{\|K_{\text{aug}}x^*\|}{\|y_{\text{aug}}\|}$ and the residual $y_{\text{aug}} - K_{\text{aug}}x^*$ is orthogonal to $K_{\text{aug}}x^*$ by the

first order conditions on x^* , as illustrated in Figure 6.1 (with K_{aug} instead of K and y_{aug} instead of y). We also used $\eta_{K_{\text{aug}}} = \frac{\|K_{\text{aug}}\| \|x^*\|}{\|K_{\text{aug}} x^*\|}$. \square

Corollary 6.3 shows the solution x^* of the Tikhonov regularized least squares problem (6.9) is stable with respect to small perturbations in the data y or the regularization constant α . Furthermore, the theorem provides explicit upper bounds on the relative condition numbers of x^* as a function of y or α . The corollary shows that Tikhonov regularization can provide a physically meaningful solution that is stable under small perturbations of y or α and provides a starting point for a quantitative stability analysis of NMR relaxometry problems.

6.3 Extension to Nonnegative Least Squares

In NMR relaxometry and related applications, the solution $x = \text{vec}(\mathbf{F})$ represents a non-negative distribution. Hence, in order to prevent the solution of (4.20) from having negative entries, the non-negativity constraint $x \geq 0$ must be imposed. Hence, for a fixed kernel $K \in \mathbb{R}^{m \times n}$ and data $y \in \mathbb{R}^m$, we solve the following Tikhonov regularized non-negative least squares problem:

$$\min_{x \geq 0} \|Kx - y\|_2^2 + \alpha^2 \|x\|_2^2. \quad (6.14)$$

In this section, we will extend the condition number results of Corollary 6.3, which applied to the unconstrained Tikhonov regularized least squares problem (6.9), to the Tikhonov regularized non-negative least squares problem (6.14). Our main lemma, Lemma 6.6, shows that the problem with non-negativity constraints

can be reduced, within a neighborhood of the initial data y and regularization parameter α , to an equivalent unconstrained least squares problem. Consequently, similar bounds to those previously derived will hold for the problem with non-negativity constraints.

We will use the idea of active-sets, which is used in the Lawson-Hanson algorithm for nonnegative least squares [94]. Active sets describe the set of variables x_i that must be actively constrained to 0 and would otherwise result in a smaller objective function value if allowed to be negative. We will prove that for $(\tilde{y}, \tilde{\alpha})$ in a neighborhood of (y, α) , the active set remains unchanged. Hence, locally the non-negative least squares problem is equivalent to a least squares problem.

There is substantial existing theory on the sensitivity analysis of constrained quadratic programs [65, 18, 86, 50, 84, 53, 71]. The stability of Tikhonov regularized least squares problems has been considered in [39, 76, 147].

We start by stating a result by Daniel [50, Lemma 2.1] on the stability of quadratic semidefinite minimization problems with constraints.

Lemma 6.4. *Let A be a symmetric positive definite (SPD) matrix of size $n \times n$ with smallest eigenvalue λ . Let $a \in \mathbb{R}^n$. Let $B \in \mathbb{R}^{b \times n}$ and $C \in \mathbb{R}^{c \times n}$ be matrices. Consider the problem*

$$\min_{z \in \mathbb{R}^n : Bz \geq 0, Cz = 0} \left(\frac{1}{2} z^t A z + z^t a \right) \quad (6.15)$$

with solution $z^ \in \mathbb{R}^n$. If the matrix A is perturbed to $\tilde{A} \in \mathbb{R}^{n \times n}$ and the vector a is perturbed to $\tilde{a} \in \mathbb{R}^n$, and if*

$$\epsilon := \max\{\|\tilde{A} - A\|_2, \|\tilde{a} - a\|_2\} < \lambda, \quad (6.16)$$

then the solution \tilde{z}^* to the perturbed problem

$$\min_{z \in \mathbb{R}^n : Bz \geq 0, Cz=0} \left(\frac{1}{2} z^t \tilde{A} z + z^t \tilde{a} \right) \quad (6.17)$$

satisfies

$$\|\tilde{z}^* - z^*\|_2 \leq \frac{\epsilon}{\lambda - \epsilon} (1 + \|z\|_2). \quad (6.18)$$

Corollary 6.5. *The solution of (6.14) depends jointly continuously y and α , for $\alpha > 0$.*

Proof. The objective function in (6.14) can be rewritten as

$$\begin{aligned} \|Kx - y\|_2^2 + \alpha^2 \|x\|_2^2 &= (Kx - y)^t (Kx - y) + \alpha^2 x^t x \\ &= x^t K^t K x - x^t K^t y - y^t K x + y^t y + \alpha^2 x^t x \\ &= \frac{1}{2} x^t A x + x^t a + \text{constant} \end{aligned}$$

where $A = 2(K^t K + \alpha^2 \text{Id}_{n \times n})$ and $a = -2K^t y$.

By Lemma 2.7, A is SPD with smallest eigenvalue $\lambda := 2(\sigma_{\min}(K)^2 + \alpha^2)$.

Observe that if α is perturbed to $\tilde{\alpha}$, then A is perturbed to $\tilde{A} = 2(K^t K + \tilde{\alpha}^2 \text{Id}_{n \times n})$,

so

$$\|\tilde{A} - A\| = \|2(\tilde{\alpha} - \alpha) \text{Id}_{n \times n}\| = 2|\tilde{\alpha} - \alpha| \quad (6.19)$$

If y is perturbed to \tilde{y} , a is perturbed to $\tilde{a} = -2K^t \tilde{y}$, so

$$\|\tilde{a} - a\| = \|-2K^t(\tilde{y} - y)\| \quad (6.20)$$

$$\leq 2\sigma_{\max}(K) \|\tilde{y} - y\| \quad (6.21)$$

Assume that $|\tilde{\alpha} - \alpha| < \frac{\epsilon}{2}$ and $\|\tilde{y} - y\| < \frac{\epsilon}{2\sigma_{\max}(K)}$. It follows that $\|\tilde{A} - A\| \leq \epsilon$

and $|\tilde{a} - a| < \epsilon$. Since $\epsilon < \lambda$, by Lemma 6.4 we have

$$\|\tilde{x}^* - x^*\| \leq \frac{\epsilon}{\lambda - \epsilon}(1 + \|x^*\|) = \frac{\epsilon}{2(\sigma_{\min}(K)^2 + \alpha^2) - \epsilon}(1 + \|x^*\|).$$

Hence, the solution x^* is a continuous function of (y, α) . \square

Our next lemma shows that for $\tilde{\alpha}$ and \tilde{y} in an open neighborhood of (α, y) , problem 6.14 is equivalent to a problem without non-negativity constraints.

Lemma 6.6. *Consider the problem 6.14 with $\alpha > 0$. There exists a nonempty set $\Lambda \subset \{1, \dots, n\}$ and a neighborhood $\mathcal{N} \subset \mathbb{R}^m \times \mathbb{R}$, containing (y, α) , such that for all $(\tilde{y}, \tilde{\alpha}) \in \mathcal{N}$, the two minimizations problems*

$$\min_{x \geq 0} \|Kx - \tilde{y}\|_2^2 + \tilde{\alpha}^2 \|x\|_2^2 \tag{6.22}$$

and

$$\min_{x : x_{\Lambda^c} = 0} \|Kx - \tilde{y}\|_2^2 + \tilde{\alpha}^2 \|x\|_2^2 \tag{6.23}$$

have the same (unique) solution.

Recall that if $\Lambda \subset \{1, \dots, n\}$ and $x \in \mathbb{R}^n$, $x_\Lambda \in \mathbb{R}^{|\Lambda|}$ is the vector with entries x_i for $i \in \Lambda$. Similarly, $x_{\Lambda^c} \in \mathbb{R}^{n-|\Lambda|}$ is the vector with entries x_i for $i \notin \Lambda$.

Proof. The proof involves an analysis of the first order conditions for (6.22) and (6.23).

Let x^* denote the solution to (6.14), let \tilde{x}^* denote the solution to (6.22), and let \tilde{x}^{**} denote the solution to (6.23).

Step 1: Construction of \mathcal{N}

Define

$$\Gamma := \{i | K_i^t \left(\frac{Kx^* - y}{-\alpha^2} \right) \neq 0\} \quad (6.24)$$

and

$$\tilde{\Gamma} := \{i | K_i^t \left(\frac{K\tilde{x}^* - \tilde{y}}{-\alpha^2} \right) \neq 0\} \quad (6.25)$$

By Corollary 6.5, the solution \tilde{x}^* depends continuously on \tilde{y} and on $\tilde{\alpha}$, so $K_i^t \left(\frac{Kx^* - y}{-\alpha^2} \right)$ also depends continuously on (y, α) . Hence, there exists neighborhood $\mathcal{N} \subset \mathbb{R}^m \times \mathbb{R}$ containing (y, α) such that for all $(\tilde{y}, \tilde{\alpha}) \in \mathcal{N}$, we have $\Gamma \subset \tilde{\Gamma}$.

Step 2: Construction of index set Λ

Define

$$\Lambda := \{i | K_i^t \left(\frac{Kx^* - y}{-\alpha^2} \right) \geq 0\} \quad (6.26)$$

and

$$\tilde{\Lambda} := \{i | K_i^t \left(\frac{K\tilde{x}^* - \tilde{y}}{-\alpha^2} \right) \geq 0\}. \quad (6.27)$$

Step 3: Proof that $\tilde{\Lambda} \subset \Lambda$.

Suppose for contradiction that $i \in \tilde{\Lambda} \setminus \Lambda$. Then, by the definitions of Λ and $\tilde{\Lambda}$, we have

$$K_i^t \left(\frac{Kx^* - y}{-\alpha^2} \right) < 0 \quad (6.28)$$

and

$$K_i^t \left(\frac{K\tilde{x}^* - \tilde{y}}{-\alpha^2} \right) \geq 0 \quad (6.29)$$

By (6.28), we have $i \in \Gamma$. By the construction of \mathcal{N} , we have $\Gamma \subset \tilde{\Gamma}$, so $i \in \tilde{\Gamma}$.

By the definition of $\tilde{\Gamma}$, we then have

$$K_i^t \left(\frac{K\tilde{x}^* - \tilde{y}}{-\alpha^2} \right) \neq 0. \quad (6.30)$$

Combining (6.29) and (6.30) we have

$$K_i^t \left(\frac{K\tilde{x}^* - \tilde{y}}{-\alpha^2} \right) > 0. \quad (6.31)$$

As argued above, $K_i^t \left(\frac{K\tilde{x}^* - \tilde{y}}{-\alpha^2} \right)$ is a jointly continuous function of $(\tilde{y}, \tilde{\alpha})$. Hence, by (6.28) and (6.31) and the intermediate value theorem, there exists a point $\tilde{\tilde{y}}$ on the line from y to \tilde{y} such that, if $\tilde{\tilde{x}}^*$ is the solution of the intermediate problem

$$\min_{x \geq 0} \|Kx - \tilde{\tilde{y}}\|_2^2 + \tilde{\alpha}^2 \|x\|_2^2,$$

we have

$$K_i^t \left(\frac{K\tilde{\tilde{x}}^* - \tilde{\tilde{y}}}{-\alpha^2} \right) = 0. \quad (6.32)$$

Hence $i \notin \tilde{\tilde{\Gamma}}$, where $\tilde{\tilde{\Gamma}} = \{i | K_i^t \left(\frac{K\tilde{\tilde{x}}^* - \tilde{\tilde{y}}}{-\alpha^2} \right) \neq 0\}$ is defined in the same way as Γ was defined in (6.25), but with $\tilde{\tilde{y}}$ instead of \tilde{y} .

By refining the neighborhood \mathcal{N} if needed, may assume \mathcal{N} is convex. Hence $(\tilde{\tilde{y}}, \alpha) \in \mathcal{N}$. By the construction of \mathcal{N} , we have $\Gamma \subset \tilde{\tilde{\Gamma}}$, so $i \notin \Gamma$. This contradicts the initial assumption that $i \in \Gamma$. Hence, we conclude that the assumption was false and it must be true that $\tilde{\Lambda} \subset \Lambda$.

The first order conditions for (6.22) are

$$K_i^t(Kx - \tilde{y}) + \alpha^2 x_i \geq 0 \quad (6.33)$$

with equality whenever $x_i \neq 0$ and $x_i \geq 0$ for all i . The first order conditions for (6.23) are

$$K_i^t(Kx - \tilde{y}) + \alpha^2 x_i = 0 \quad (6.34)$$

for all $i \in \Lambda$ and $x_i = 0$ for all $i \in \Lambda^c$.

We aim to show that the two systems of inequalities (6.33) and (6.34) are equivalent.

Step 4: Proof that (6.33) implies (6.34)

Suppose that \tilde{x}^* satisfies equations (6.33), with \tilde{x}^* instead of x . Since $\tilde{\Lambda} \subset \Lambda$, we consider separately three cases: $i \notin \Lambda$, $i \in \tilde{\Lambda}$, and $i \in \Lambda \setminus \tilde{\Lambda}$.

Case 1: Suppose $i \notin \Lambda$. By Step 3 above, $\tilde{\Lambda} \subset \Lambda$, hence $i \notin \tilde{\Lambda}$. Hence by the definition of $\tilde{\Lambda}$, we have

$$K_i^t \left(\frac{K\tilde{x}^* - \tilde{y}}{-\alpha^2} \right) < 0. \quad (6.35)$$

If $\tilde{x}_i^* \neq 0$, then by (6.33) we have $\tilde{x}_i^* > 0$ and

$$K_i^t \left(\frac{K\tilde{x}^* - \tilde{y}}{-\alpha^2} \right) = \tilde{x}_i^* > 0, \quad (6.36)$$

which contradicts (6.35). Hence $\tilde{x}_i^* = 0$. This proves (6.34) in the case $i \notin \Lambda$.

Case 2: Suppose that $i \in \tilde{\Lambda}$. Then $K_i^t \left(\frac{K\tilde{x}^* - \tilde{y}}{-\alpha^2} \right) \geq 0$. Thus, $K_i^t(K\tilde{x}^* - \tilde{y}) \leq 0$. If $\tilde{x}_i = 0$, we conclude from (6.33) that $K_i^t(K\tilde{x}^* - \tilde{y}) \geq 0$, so $K_i^t(K\tilde{x}^* - \tilde{y}) = 0$. Hence 6.34 holds. If $\tilde{x}_i \neq 0$, then (6.33) holds with equality, so we again conclude that (6.34) holds.

Case 3: Now suppose that $i \in \Lambda \setminus \tilde{\Lambda}$. Then by the definitions of Λ and $\tilde{\Lambda}$ we have

$$K_i^t \left(\frac{Kx^* - y}{-\alpha^2} \right) \geq 0 \quad (6.37)$$

and

$$K_i^t \left(\frac{K\tilde{x}^* - \tilde{y}}{-\alpha^2} \right) < 0. \quad (6.38)$$

Suppose for contradiction that $K_i^t \left(\frac{Kx^* - y}{-\alpha^2} \right) \neq 0$. By (6.37), we have $K_i^t \left(\frac{Kx^* - y}{-\alpha^2} \right) > 0$. By continuity and the intermediate value theorem, as in a previous argument,

there exists an intermediate point $(\tilde{y}, \tilde{\alpha})$ between (y, α) and $(\tilde{y}, \tilde{\alpha})$ such that $K_i^t \left(\frac{K\tilde{x}^* - \tilde{y}}{-\tilde{\alpha}^2} \right) = 0$. Hence $i \in \Gamma$ and $i \notin \tilde{\Gamma}$, which contradicts that $\Gamma \subset \tilde{\Gamma}$ by the construction of \mathcal{N} .

Hence our assumption was false, so

$$K_i^t \left(\frac{K\tilde{x}^* - \tilde{y}}{-\alpha^2} \right) = 0. \quad (6.39)$$

Using (6.39), equation (6.33) simplifies to $\alpha^2 x_i \geq 0$ with equality whenever $x_i > 0$. If $x_i > 0$, this would give a contradiction, hence $x_i = 0$. Hence (6.33) holds with equality, so (6.34) holds.

This completes the proof that (6.33) implies (6.34).

Step 5: Proving the equivalence of (6.22) and (6.23)

Recall that (6.33) is the first order condition for problem (6.22) and (6.34) is the first order condition for problem (6.23). We already showed that (6.33) implies (6.34), hence an optimal solution of (6.22) is an optimal solution of (6.23). To prove the reverse implication, it suffices to show that problem (6.23) has a unique solution.

Let \tilde{x}^{**} be a solution of (6.23). Then \tilde{x}^{**} satisfies the first order conditions (6.34), with \tilde{x}^{**} instead of x . Hence, $\tilde{x}_{\Lambda^c}^{**} = 0$ and \tilde{x}_{Λ}^{**} satisfies

$$(K_{\Lambda}^t K_{\Lambda} + \alpha^2 \text{Id}_{|\Lambda| \times |\Lambda|}) \tilde{x}_{\Lambda}^{**} = K_{\Lambda}^t \tilde{y} \quad (6.40)$$

Since the matrix $K_{\Lambda}^t K_{\Lambda} + \alpha^2 \text{Id}_{|\Lambda| \times |\Lambda|}$ is symmetric positive definite, it is invertible, and hence \tilde{x}_{Λ}^{**} is uniquely determined. We conclude that (6.23) has a unique solution, and hence (6.22) and (6.23) are equivalent. \square

Lemma 6.6 shows that the Tikhonov regularized non-negative least squares problem is equivalent to a Tikhonov regularized unconstrained least squares problem. We can thus extend the previous stability results.

Theorem 6.7. Consider the non-negative least squares problem with Tikhonov regularization (6.14), with $\alpha > 0$. Assume that the space of feasible solutions $\{x \in \mathbb{R}^n : x \geq 0\}$ is equipped with the $\|\cdot\|_2$ norm, the space containing the initial data $y \in \mathbb{R}^m$ is equipped with the $\|\cdot\|_2$ norm, and the space of regularization constants $\alpha \in (0, \infty)$ is equipped with the absolute value $|\cdot|$ norm. Let x^* be the optimal solution of (6.14) and let $\Lambda := \{i | K_i^t (\frac{Kx^* - y}{-\alpha^2}) \geq 0\}$. Then the following holds:

1. The relative condition number of the solution $\tilde{x}^* \in \mathbb{R}^n$ of (6.14) as a function of the initial data $y \in \mathbb{R}^m$ satisfies

$$\kappa_{x^*:y} \leq \frac{\|y\|_2}{\sigma_{\min}((K_\Lambda)_{\text{aug}})\|x^*\|_2}. \quad (6.41)$$

2. The relative condition number of the solution $x^* \in \mathbb{R}^n$ of (6.14) as a function of the regularization constant $\alpha > 0$ satisfies

$$\kappa_{x^*:\alpha} \leq \frac{\alpha}{\sigma_{\min}((K_\Lambda)_{\text{aug}})} + \frac{\alpha\|y_{\text{aug}} - (K_\Lambda)_{\text{aug}}x^*\|_2}{\sigma_{\min}^2((K_\Lambda)_{\text{aug}})\|x^*\|_2} \quad (6.42)$$

Recall that K_Λ is the submatrix of K obtained by extracting the columns $K[:, i]$ for which $i \in \Lambda$. The matrix $(K_\Lambda)_{\text{aug}}$ is defined by

$$(K_\Lambda)_{\text{aug}} := \begin{pmatrix} K_\Lambda \\ \alpha \text{Id}_{n \times n} \end{pmatrix}$$

$$y_{\text{aug}} := \begin{pmatrix} y \\ 0_{n \times 1} \end{pmatrix}.$$

Proof. By Lemma 6.6, there exists a neighborhood \mathcal{N} containing (y, α) such that for all $(\tilde{y}, \tilde{\alpha}) \in \mathcal{N}$, the unconstrained problem (6.22) and the constrained problem (6.23)

are equivalent. The conclusion follows by applying Lemma 6.6 to the unconstrained problem (6.23). □

Theorem 6.7 shows that standard stability results for least squares problems can be extended to nonnegative least squares. The theorem provides a starting point for the quantification of stability of N -dimensional NMR relaxometry problems.

Conclusion

We have developed tensor recovery algorithms and recovery guarantees. These methods enable the solution of N -dimensional inverse problems from a small number of non-uniformly sampled noisy measurements. While the previous methods of Cloninger and Czaja [41, 42] applied only to 2-dimensional inverse problems, our methods apply to N -dimensional problems for all $N \geq 2$.

In NMR relaxometry applications, N -dimensional experiments ($N \geq 3$) provide significantly richer information than 1 or 2-dimensional experiments [29]. However, previous biomedical applications of these high-dimensional experiments have been limited by prohibitively long acquisition times. Our new methods enable significant acceleration of these sophisticated experiments and may thus lead the way to many new biomedical applications.

Our new contributions include theoretical recovery guarantees for sub-Gaussian maps and non-uniform random sampling, heuristic techniques for regularization parameter selection (including accelerated k -fold cross-validation and generalized cross-validation), and the derivation of minimally coherent non-uniform random sampling. These methods promise to significantly accelerate high-dimensional NMR relaxometry data acquisition and improve the NMR data processing pipeline. These

methods could also be applied to other data recovery and inverse problems.

Bibliography

- [1] M. Al-Qizwini and H. Radha. Fast smooth rank approximation for tensor completion. In *Information Sciences and Systems (CISS), 2014 48th Annual Conference on*, pages 1–5, March 2014.
- [2] R.C. Aster, B. Borchers, and C.H. Thurber. *Parameter estimation and inverse problems*, volume 90. Academic Press, 2011.
- [3] R. Bai, A. Cloninger, W. Czaja, and P.J. Basser. Efficient 2d mri relaxometry using compressed sensing. *Journal of Magnetic Resonance*, 255:88–99, 6 2015.
- [4] R. Balan, P. G. Casazza, C. Heil, and Z. Landau. Density, overcompleteness, and localization of frames. i. theory. *Journal of Fourier Analysis and Applications*, 12(2):105–143, 2006.
- [5] R. Baraniuk, M. Davenport, R. A. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- [6] R. Baraniuk, M.A. Davenport, M.F. Duarte, and C. Hegde. An introduction to compressive sensing. *OpenStax CNX*, 8 2014.
- [7] M.S. Bazaraa, H.D. Sherali, and C. M. Shetty. *Nonlinear Programming Theory and Applications*. Wiley, 3 edition, 2006.
- [8] J.A. Bazerque, G. Mateos, and G.B. Giannakis. Rank regularization and bayesian inference for tensor completion and extrapolation. *Signal Processing, IEEE Transactions on*, 61(22):5689–5703, Nov 2013.
- [9] R. M. Bell and Y. Koren. Lessons from the netflix prize challenge. *ACM SIGKDD Explorations Newsletter*, 9(2):75–79, 2007.
- [10] J. J. Benedetto, W Czaja, P. Gadziński, and A. M. Powell. The balian-low theorem and regularity of gabor systems. *The Journal of Geometric Analysis*, 13(2):239–254, 2003.

- [11] J. J. Benedetto and S. Li. Multiresolution analysis frames with applications. In *Acoustics, Speech, and Signal Processing, 1993. ICASSP-93., 1993 IEEE International Conference on*, volume 3, pages 304–307 vol.3, April 1993.
- [12] J. J. Benedetto and H. C. Wu. A multidimensional irregular sampling algorithm and applications. In *Acoustics, Speech, and Signal Processing, 1999. Proceedings., 1999 IEEE International Conference on*, volume 4, pages 2039–2042 vol.4, Mar 1999.
- [13] J.J. Benedetto. Frame decompositions, sampling, and uncertainty principle inequalities. *Wavelets: mathematics and applications*, pages 247–304, 1994.
- [14] J.J. Benedetto and M. Fickus. Finite normalized tight frames. *Advances in Computational Mathematics*, 18(2):357–385.
- [15] J.J. Benedetto and S. Li. The theory of multiresolution analysis frames and applications to filter banks. *Applied and Computational Harmonic Analysis*, 5(4):389 – 427, 1998.
- [16] J.J. Benedetto and D.F. Walnut. Gabor frames for l_2 and related spaces. *Wavelets: mathematics and applications*, pages 97–162, 1993.
- [17] John J Benedetto and Wojciech Czaja. *Integration and modern analysis*. Springer Science & Business Media, 2010.
- [18] A.B. Berkelaar, K. Roos, and T. Terlaky. *Advances in Sensitivity Analysis and Parametric Programming*. Springer US, Boston, MA, 1997.
- [19] Paula Berman, Ofer Levi, Yisrael Parmet, Michael Saunders, and Zeev Wiesman. Laplace inversion of low-resolution nmr relaxometry data using sparse representation methods. *Concepts in Magnetic Resonance Part A*, 42(3):72–88, 2013.
- [20] M. Bertero, P. Boccacci, and E.R. Pike. On the recovery and resolution of exponential relaxation rates from experimental data: a singular-value analysis of the laplace transform inversion in the presence of noise. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 383, pages 15–29. The Royal Society, 1982.
- [21] M. Bouhrara, S. Clerjon, J.-L. Damez, C. Chevarin, S. Portanguen, A. Kondjoyan, and J.-M. Bonny. Dynamic mri and thermal simulation to interpret deformation and water transfer in meat during heating. *Journal of Agricultural and Food Chemistry*, 59(4):1229–1235, 2011.
- [22] M. Bouhrara, D. A. Reiter, H. Celik, J.-M. Bonny, V. Lukas, K. W. Fishbein, and R. G. Spencer. Incorporation of rician noise in the analysis of biexponential transverse relaxation in cartilage using a multiple gradient echo sequence at 3 and 7 tesla. *Magnetic Resonance in Medicine*, 73(1):352–366, 2015.

- [23] M. Bouhrara, D. A. Reiter, and R. G. Spencer. Bayesian analysis of transverse signal decay with application to human brain. *Magnetic Resonance in Medicine*, 74(3):785–802, 9 2015.
- [24] M. Bouhrara and R. G. Spencer. Improved determination of the myelin water fraction in human brain using magnetic resonance imaging through bayesian analysis of mcdespot. *NeuroImage*, In Press, 10 2015.
- [25] R. W. Brown, Y.-C. N. Cheng, E. M. Haacke, M. R. Thompson, and R. Venkatesan. *Magnetic resonance imaging: physical principles and sequence design*. John Wiley & Sons, 2014.
- [26] J.F. Cai, E.J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *Siam J. Optimization*, 20(4):1956–1982, 3 2010.
- [27] C.F. Caiafa and A. Cichocki. Multidimensional compressed sensing and their applications. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 3(6):355–380, 2013.
- [28] C.F. Caiafa and A. Cichocki. Stable, robust, and super fast reconstruction of tensors using multi-way projections. *Signal Processing, IEEE Transactions on*, 63(3):780–793, Feb 2015.
- [29] P.T. Callaghan, C.H. Arns, P. Galvosas, M.W. Hunter, Y. Qiao, and K.E. Washburn. Recent fourier and laplace perspectives for multidimensional nmr in porous media. *Magnetic Resonance Imaging*, 25:441–444, 2007.
- [30] E. J. Candès. Compressive sampling. In *Proceedings of the international congress of mathematicians*, volume 3, pages 1433–1452. Madrid, Spain, 2006.
- [31] E. J. Candès and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 6 2010.
- [32] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 12 2009.
- [33] E. J. Candès and J. K. Romberg. Sparsity and incoherence in compressive sampling. *Inverse Problems*, 23(3):969, 2007.
- [34] E. J. Candès, J. K. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, 6 2004.
- [35] E. J. Candès and T. Tao. Decoding by linear programming. *Information Theory, IEEE Transactions on*, 51(12):4203–4215, 2005.
- [36] E. J. Candès and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *Information Theory, IEEE Transactions on*, 52(12):5406–5425, 2006.

- [37] W. Cao, Y. Wang, C. Yang, X. Chang, Z. Han, and Z. Xu. Folded-concave penalization approaches to tensor completion. *Neurocomputing*, 152:261 – 273, 2015.
- [38] H. Celik, M. Bouhrara, D. Reiter, K.W. Fishbein, and R.G. Spencer. Stabilization of the inverse laplace transform of multiexponential decay through introduction of a second dimension. *Journal of Magnetic Resonance*, 236(2):134–139, 11 2013.
- [39] D. Chu, L. Lin, R. C. E. Tan, and Y. Wei. Condition numbers and perturbation analysis for the tikhonov regularization of discrete ill-posed problems. *Numerical Linear Algebra with Applications*, 18(1):87–103, 2011.
- [40] C.A. Clark. *Spectral Frame Analysis and Learning through Graph Structure*. PhD thesis, University of Maryland, College Park, 2016.
- [41] A. Cloninger. *Exploiting Data-Dependent Structure for Improving Sensor Acquisition and Integration*. PhD thesis, University of Maryland, College Park, 2014.
- [42] A. Cloninger, W. Czaja, R. Bai, and P.J. Basser. Solving 2d fredholm integral from incomplete measurements using compressive sensing. *Siam J. Imaging Sciences*, 7(3):1775–1798, 2014.
- [43] A. Cohen, W. Dahmen, and R. A. DeVore. Compressed sensing and best k-term approximation. *Journal of the American mathematical society*, 22(1):211–231, 2009.
- [44] A. Cohen, W. Dahmen, and R. A. DeVore. Orthogonal matching pursuit under the restricted isometry property. *arXiv:1506.04779*, 6 2015.
- [45] J. Cohen-Adad. and C. Wheeler-Kingshott. Quantitative mri of the spinal cord. In Julien Cohen-AdadClaudia A.M. Wheeler-Kingshott, editor, *Quantitative {MRI} of the Spinal Cord*. Academic Press, San Diego, 2014.
- [46] P. Craven and G. Wahba. Smoothing noisy data with spline functions. *Numerische Mathematik*, 31(4):377–403, 1978.
- [47] W. Czaja, G. Kutyniok, and D. Speegle. The geometry of sets of parameters of wave packet frames. *Applied and Computational Harmonic Analysis*, 20(1):108–125, 2006.
- [48] W. Czaja, G. Kutyniok, and D. Speegle. Beurling dimension of gabor pseudoframes for affine subspaces. *Journal of Fourier Analysis and Applications*, 14(4):514–537, 2008.
- [49] W. Czaja and Z. Rzeszutnik. Pseudodifferential operators and gabor frames: spectral asymptotics. *Mathematische Nachrichten*, 233:77–88, 2002.

- [50] J. W. Daniel. Stability of the solution of definite quadratic programs. *Mathematical Programming*, 5(1):41–53, 2 1973.
- [51] M. A. Davenport, M. F. Duarte, Y. C. Eldar, and G. Kutyniok. Introduction to compressed sensing. *Preprint*, 93(1):2, 2011.
- [52] R. A. DeVore. Deterministic constructions of compressed sensing matrices. *Journal of Complexity*, 23(4–6):918 – 925, 2007.
- [53] J. Ding and W. Hang. New perturbation results for equality-constrained least squares problems. *Linear Algebra and its Applications*, 272(1–3):181 – 192, 1998.
- [54] S. Dirksen. Dimensionality reduction with subgaussian matrices: a unified theory. arXiv:1402.3973, 2014.
- [55] D.L. Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, 4 2006.
- [56] M.F. Duarte and R.G. Baraniuk. Kronecker compressive sensing. *Image Processing, IEEE Transactions on*, 21(2):494–504, Feb 2012.
- [57] Y. Eldar and G. Kutyniok. *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2012.
- [58] R. R. Ernst, G. Bodenhausen, and A. Wokaun. *Principles of Nuclear Magnetic Resonance in One and Two Dimensions*. Oxford University Press, 1988.
- [59] M. Fazel, E.J. Candès, B. Recht, and P. Parrilo. Compressed sensing and robust recovery of low rank matrices. In *Signals, Systems and Computers, 2008 42nd Asilomar Conference on*, pages 1043–1047, 10 2008.
- [60] S. Foucart and H. Rauhut. *A mathematical introduction to compressive sensing*, volume 1. Springer, 2013.
- [61] R. Freedman and N. Heaton. Fluid characterization using nuclear magnetic resonance logging. *Petrophysics*, 45(03), 2004.
- [62] S. Friedland, Q. Li, and D. Schonfeld. Compressive sensing of sparse tensors. *Image Processing, IEEE Transactions on*, 23(10):4438–4447, 10 2014.
- [63] S. Friedland and L.-H. Lim. Computational complexity of tensor nuclear norm. *arXiv preprint arXiv:1410.6072*, 2014.
- [64] S. Gandy, B. Recht, and I. Yamada. Tensor completion and low-n-rank tensor recovery via convex optimization. *Inverse Problems*, 27(2), 1 2011.
- [65] A. J. Geurts. A contribution to the theory of condition. *Numerische Mathematik*, 39(1):85–96.

- [66] G. H. Golub, M. Heath, and G. Wahba. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21(2):215–223, 1979.
- [67] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, 4 edition, 12 2012.
- [68] D. Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, 3 2011.
- [69] M. Guay. *Sparse signal representation in digital and biological systems*. PhD thesis, University of Maryland, College Park, 2016.
- [70] O. Guédon, S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann. Majorizing measures and proportional subsets of bounded orthonormal systems. *Revista matemática iberoamericana*, 24(3):1075–1095, 2008.
- [71] M. Gulliksson, X.-Q. Jin, and Y.-M. Wei. Perturbation bounds for constrained and weighted least squares problems. *Linear Algebra and its Applications*, 349(1–3):221 – 232, 2002.
- [72] W. Hackbusch. *Tensor Spaces and Numerical Tensor Calculus*. Springer, 2012.
- [73] A. Hafftkka, H. Celik, A. Cloninger, W. Czaja, and R. G. Spencer. 2d sparse sampling algorithm for nd fredholm equations with applications to nmr relaxometry. In *Sampling Theory and Applications (SampTA), International Conference on*, pages 367–371, 5 2015.
- [74] E. L. Hahn. Spin echoes. *Physical Review*, 80:580–594, Nov 1950.
- [75] A. Halevy. *Extensions of laplacian eigenmaps for manifold learning*. PhD thesis, University of Maryland, College Park, 2011.
- [76] P. C. Hansen. Perturbation bounds for discrete tikhonov regularisation. *Inverse Problems*, 5(4):L41, 1989.
- [77] P. C. Hansen. Analysis of discrete ill-posed problems by means of the l-curve. *SIAM Review*, 34(4):561–580, 1992.
- [78] P. C. Hansen. *Discrete Inverse Problems*. Society for Industrial and Applied Mathematics, 2010.
- [79] C. J. Hillar and L.-H. Lim. Most tensor problems are np-hard. *Journal of the ACM (JACM)*, 60(6):45, 2013.
- [80] B.P. Hills. Relaxometry: Two-dimensional methods. *Encyclopedia of Magnetic Resonance, Online*, 2009.

- [81] J. Hrabe, G. Kaur, and D. N. Guilfoyle. Principles and limitations of nmr diffusion measurements. *Journal of medical physics/Association of Medical Physicists of India*, 32(1):34, 2007.
- [82] S. Hu, M. Lustig, A. P. Chen, J. Crane, A. Kerr, D. A. C. Kelley, R. Hurd, J. Kurhanewicz, S. J. Nelson, J. M. Pauly, and D. B. Vigneron. Compressed sensing for resolution enhancement of hyperpolarized ^{13}C flyback 3d-mrsi. *J Magn Reson*, 192(2):258–64, Jun 2008.
- [83] O. N. Irrechukwu, S. V. Thaer, E. H. Frank, P.-C. Lin, D. A. Reiter, A. J. Grodzinsky, and R. G. Spencer. Prediction of cartilage compressive modulus using multiexponential analysis of t_2 relaxation data and support vector regression. *NMR in Biomedicine*, 27(4):468–477, 2014.
- [84] W. Kang and H. Xiang. Condition numbers with their condition numbers for the weighted moore-penrose inverse and the weighted least squares solution. *Journal of Applied Mathematics and Computing*, 22(1):95–112.
- [85] J. Keeler. *Understand NMR Spectroscopy*. Wiley, 2010.
- [86] D.S. Kim, N.N. Tam, and N.D. Yen. Solution existence and stability of quadratically constrained convex quadratic programs. *Springer Optimization Letters*, 6(2):363–373, 2 2012.
- [87] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.
- [88] L. Koralov and Y. G. Sinai. *Theory of probability and random processes*. Springer Science & Business Media, 2007.
- [89] A. Krishnamurthy and A. Singh. Low-rank matrix and tensor completion via adaptive sampling. *Advances in Neural Information Processing Systems*, 26:836–844, 2013.
- [90] Gitta Kutyniok, Kasso A. Okoudjou, Friedrich Philipp, and Elizabeth K. Tuley. Scalable frames. *Linear Algebra and its Applications*, 438(5):2225 – 2238, 2013.
- [91] L. Laouar, K. Fishbein, L. E. McGann, W. E. Horton, R. G. Spencer, and N. M. Jomha. Cryopreservation of porcine articular cartilage: mri and biochemical results after different freezing protocols. *Cryobiology*, 54(1):36 – 43, 2007.
- [92] L. D. Lathauwer, B. D. Moor, and J. Vandewalle. A multilinear singular value decomposition. *SIAM Journal on Matrix Analysis and Applications*, 21(4):1253–1278, 2000.
- [93] A. J. Laub. *Matrix Analysis for Scientists and Engineers*. Society for Industrial and Applied Mathematics, 2005.

- [94] C. L. Lawson and R. J. Hanson. *Solving Least Squares Problems*. Society for Industrial and Applied Mathematics, 1995.
- [95] Qun Li, D. Schonfeld, and S. Friedland. Generalized tensor compressive sensing. In *Multimedia and Expo (ICME), 2013 IEEE International Conference on*, pages 1–6, July 2013.
- [96] P.-C. Lin, O. Irrechukwu, R. Roque, B. Hancock, K. W. Fishbein, and R. G. Spencer. Multivariate analysis of cartilage degradation using the support vector machine algorithm. *Magnetic Resonance in Medicine*, 67(6):1815–1826, 2012.
- [97] P.-C. Lin, D. A. Reiter, and R. G. Spencer. Classification of degraded cartilage through multiparametric mri analysis. *Journal of Magnetic Resonance*, 201(1):61 – 71, 2009.
- [98] J. Liu, P. Musialski, P. Wonka, and J. Ye. Tensor completion for estimating missing values in visual data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1):208–220, 1 2013.
- [99] Y. Liu and F. Shang. An efficient matrix factorization method for tensor completion. *Signal Processing Letters, IEEE*, 20(4):307–310, 4 2013.
- [100] Y. Liu, F. Shang, W. Fan, J. Cheng, and H. Cheng. Generalized higher order orthogonal iteration for tensor learning and decomposition. *Neural Networks and Learning Systems, IEEE Transactions on*, PP(99):1–13, 2015.
- [101] Y. Liu, F. Shang, L. Jiao, J. Cheng, and H. Cheng. Trace norm regularized candecomp/parafac decomposition with missing data. *Cybernetics, IEEE Transactions on*, 45(11):2437–2448, Nov 2015.
- [102] Y.-K. Liu. Universal low-rank matrix recovery from pauli measurements. In *Advances in Neural Information Processing Systems*, volume 24, pages 1638–1646, 2011.
- [103] M. Lustig. *SPARSE MRI*. PhD thesis, Stanford University, 2008.
- [104] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly. Compressed sensing mri. *Signal Processing Magazine, IEEE*, 25(2):72–82, 2008.
- [105] M. Lustig, D.L. Donoho, and J.M. Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 12 2007.
- [106] S. Ma, D. Goldfarb, and L. Chen. Fixed point and bregman iterative methods for matrix rank minimization. *Mathematical Programming*, 128(1):321–353, 9 2009.

- [107] A. Mackay, K. Whittall, J. Adler, D. Li, D. Paty, and D. Graeb. In vivo visualization of myelin water in brain by magnetic resonance. *Issue Magnetic Resonance in Medicine Magnetic Resonance in Medicine*, 31(6):673–677, 11 2005.
- [108] R. L. Magin, W. Li, M. P. Velasco, J. Trujillo, D. A. Reiter, A. Morgenstern, and R. G. Spencer. Anomalous nmr relaxation in cartilage matrix components and native cartilage: Fractional-order models. *Journal of Magnetic Resonance*, 210(2):184 – 191, 2011.
- [109] P. J. McDonald, J. Mitchell, M. Mulheron, P. S. Aptaker, J.-P. Korb, and L. Monteilhet. Two-dimensional correlation relaxometry studies of cement pastes performed using a new one-sided nmr magnet. *Cement and Concrete Research*, 37(3):303 – 309, 2007.
- [110] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann. Reconstruction and subgaussian operators in asymptotic geometric analysis. *Geometric and Functional Analysis*, 17(4):1248–1282, 2007.
- [111] C. Mu, B. Huang, J. Wright, and D. Goldfarb. Square deal: Lower bounds and improved relaxations for tensor recovery. *arXiv:1307.5870*, 8 2013.
- [112] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24(2):227–234, 1995.
- [113] A. Nava-Tudela. *Image representation and compression via sparse solutions of systems of linear equations*. PhD thesis, University of Maryland, College Park, 2012.
- [114] S. Negahban and M. J. Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *The Annals of Statistics*, pages 1069–1097, 2011.
- [115] S. Ramaswamy, D.-A. Wang, K. W. Fishbein, J. H. Elisseeff, and R. G. Spencer. An analysis of the integration between articular cartilage and non-degradable hydrogel using magnetic resonance imaging. *Journal of Biomedical Materials Research Part B: Applied Biomaterials*, 77B(1):144–148, 2006.
- [116] H. Rauhut. Compressive sensing and structured random matrices. *Radon Series on Computational and Applied Mathematics*, 9:1–92, 2010.
- [117] H. Rauhut, K. Schnass, and P. Vandergheynst. Compressed sensing and redundant dictionaries. *Information Theory, IEEE Transactions on*, 54(5):2210–2219, 2008.
- [118] H. Rauhut, R. Schneider, and Z. Stojanac. Tensor completion in hierarchical tensor representations. *arXiv:1404.3905*, 4 2014.

- [119] H. Rauhut, R. Schneider, and Z. Stojanac. Low rank tensor recovery via iterative hard thresholding. *arXiv:1602.05217*, 2 2016.
- [120] H. Rauhut and Z. Stojanac. Recovery of third order tensors via convex optimization. In *Sampling Theory and Applications (SampTA), 2015 International Conference on*, pages 397–401, 5 2015.
- [121] H. Rauhut and Z. Stojanac. Tensor theta norms and low rank recovery. *arXiv:1505.05175*, 2015.
- [122] Holger Rauhut, Reinhold Schneider, and Z Stojanac. Low rank tensor recovery via iterative hard thresholding. In *Proc. 10th International Conference on Sampling Theory and Applications*, 2013.
- [123] B. Recht, M. Fazel, and P.A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.
- [124] D. A. Reiter, O. Irrechukwu, P.-C. Lin, S. Moghadam, S. V. Thaer, N. Pleshko, and R. G. Spencer. Improved mr-based characterization of engineered cartilage using multiexponential t2 relaxation and multivariate analysis. *NMR in Biomedicine*, 25(3):476–488, 2012.
- [125] D. A. Reiter, R. L. Magin, W. Li, J. J. Trujillo, M. P. Velasco, and R. G. Spencer. Anomalous t2 relaxation in normal and degraded cartilage. *Magnetic Resonance in Medicine*, pages n/a–n/a, 2015.
- [126] D. A. Reiter, A. Peacock, and R. G. Spencer. Effects of frozen storage and sample temperature on water compartmentation and multiexponential transverse relaxation in cartilage. *Magnetic resonance imaging*, 29(4):561–567, 2011.
- [127] D. A. Reiter, R. A. Roque, P.-C. Lin, S. B. Doty, N. Pleshko, and R. G. Spencer. Improved specificity of cartilage matrix evaluation using multiexponential transverse relaxation analysis applied to pathomimetically degraded cartilage. *NMR in Biomedicine*, 24(10):1286–1294, 2011.
- [128] D. A. Reiter, R. A. Roque, P.-C. Lin, O. Irrechukwu, S. Doty, D. Longo, N. Pleshko, and R. G. Spencer. Mapping proteoglycan-bound water in cartilage: improved specificity of matrix assessment using multiexponential transverse relaxation analysis. *Magnetic Resonance in Medicine*, 65(2):377–384, 2 2011.
- [129] D.A. Reiter, P.C. Lin, K.W. Fishbein, and R.G. Spencer. Multicomponent t2 relaxation analysis in cartilage. *Magnetic Resonance in Medicine*, 61:803–809, 2 2009.
- [130] B. Romera-Paredes and M. Pontil. A new convex relaxation for tensor completion. In *Advances in Neural Information Processing Systems 26*, pages 2967–2975. 2013.

- [131] M. Rudelson and R. Vershynin. Sparse reconstruction by convex relaxation: Fourier and gaussian measurements. In *Information Sciences and Systems, 2006 40th Annual Conference on*, pages 207–212. IEEE, 2006.
- [132] George Saab, R. Terry Thompson, and Greg D. Marsh. Multicomponent t2 relaxation of in vivo skeletal muscle. *Magnetic Resonance in Medicine*, 42(1):150–157, 7 1999.
- [133] Z. Shi, J. Han, T. Zheng, S. Deng, and J. Li. Guarantees of augmented trace norm models in tensor recovery. *arXiv:1207.5326*, 2012.
- [134] M. Signoretto, D.L. Lathauwer, and J. A. K. Suykens. Nuclear norms for tensors and their use for convex multilinear estimation. *Submitted to Linear Algebra and Its Applications*, 43, 2010.
- [135] M. Signoretto, R. Van de Plas, B. De Moor, and J.A.K. Suykens. Tensor versus matrix completion: A comparison with application to spectral data. *Signal Processing Letters, IEEE*, 18(7):403–406, July 2011.
- [136] J. Stoer and R. Bulirsch. *Introduction to numerical analysis*, volume 12 of *Texts in Applied Mathematics*. Springer, 2 edition, 1992.
- [137] N.K. Strawn. *Geometric structures and optimization on spaces of finite frames*. PhD thesis, University of Maryland, College Park, 2011.
- [138] M. Talagrand. *The generic chaining: upper and lower bounds of stochastic processes*. Springer Science & Business Media, 2006.
- [139] R. Tomioka. Convex tensor decomposition via structured schatten norm regularization. In *Advances in neural information processing systems*, pages 1331–1339, 2013.
- [140] R. Tomioka, K. Hayashi, and H. Kashima. Estimation of low-rank tensors via convex optimization. *arXiv:1010.0789*, 2010.
- [141] R. Tomioka, T. Suzuki, K. Hayashi, and H. Kashima. Statistical performance of convex tensor decomposition. In J. Shawe-Taylor, R.S. Zemel, P.L. Bartlett, F. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 972–980. Curran Associates, Inc., 2011.
- [142] L. N. Trefethen and D. Bau, III. *Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, 1997.
- [143] P. Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications*, 109(3):475–494.
- [144] L. Venkataramanan, Y.Q. Song., and M.D. Hürlimann. Solving fredholm integrals of the first kind with tensor product structure in 2 and 2.5 dimensions. *IEEE Transactions on Signal Processing*, 50(5), 5 2002.

- [145] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- [146] R. Wang. *Global geometric conditions on sensing matrices for the success of $L1$ minimization algorithm*. PhD thesis, University of Maryland, College Park, 2013.
- [147] M. Wei. Perturbation theory for the rank-deficient equality constrained least squares problem. *SIAM Journal on Numerical Analysis*, 29(5):1462–1481, 1992.
- [148] L. Yang, Z.-H. Huang, and X. Shi. A fixed point iterative method for low n-rank tensor pursuit. *Signal Processing, IEEE Transactions on*, 61(11):2952–2962, June 2013.
- [149] M. Yuan and C.-H. Zhang. On tensor completion via nuclear norm minimization. *arXiv:1405.1773*, 5 2014.
- [150] M. Zhang, L. Yang, and Z.-H. Huang. Minimum n-rank approximation via iterative hard thresholding. *Applied Mathematics and Computation*, 256:860–875, 2015.
- [151] Q. Zhao, L. Zhang, and A. Cichocki. Bayesian cp factorization of incomplete tensors with automatic rank determination. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 37(9):1751–1763, Sept 2015.
- [152] Q. Zhao, G. Zhou, L. Zhang, A. Cichocki, and S.-I. Amari. Bayesian robust tensor factorization for incomplete multiway data. *Neural Networks and Learning Systems, IEEE Transactions on*, PP(99):1–1, 2016.