
$\oplus$
$\oplus$

# Tomographic Reconstruction <br> using Anatomy <br> and <br> Regularization 

| Title | Tomographic Reconstruction using Anatomy and Regularization |
| :--- | :--- |
| Author | Esben Høgh-Rasmussen |
| Keywords | Tomographic reconstruction |
|  | LaRGE-SCALE COMPUTATIONS |
|  | Krybov, SVD, LANCZOS, Bidiagonalization. |

Supervisors Lars Kai Hansen, IMM
Per Christian Hansen, IMM
Claus Svarer, NRU

| Typesetting | $\mathrm{ETT}_{\mathrm{E}} \mathrm{X}$ and $\mathrm{BiBT}_{\mathrm{E}} \mathrm{X}$ |
| :--- | :--- |
|  | Main text: Times-Roman <br>  <br>  <br> Mathematics: Computer-Modern <br> Graphs |
| Headlines and notes: Helvetica |  |
| MATLAB |  |


$\overline{\overline{\underline{\| M M}} \quad$|  Technical University of Denmark  |
| :--- |
|  Informatics and Mathematical Modelling  |
|  Building 321, DK-2800 Kongens Lyngby, Denmark  |
|  Phone (+45) 45253351, Fax (+45) 45882673  |
|  reception@imm.dtu.dk  |
|  www.imm.dtu.dk  |$}$

Neurobiology Research Unit
Department of Neurology
The Neuroscience Centre, Rigshospitalet Blegdamsvej 9, Copenhagen University Hospital
Building 9201, DK-2100 Copenhagen, Denmark
Phone (+45) 35456711, Fax (+45) 35456713
www.nru.dk

Copyright © 2005, Esben Høgh-Rasmussen.

## Reference IMM-PHD-2005-156

ISSN Unassigned
ISBN Unassigned

## Abstract

## Abstract (English)

This thesis develops a method for reconstructing images in emission tomography. Rooted in user requirements, it is decided to go forward with Tikhonov regularization followed by post-filtering with a Gaussian kernel. This is followed by an analysis of the theoretical loss of performance implied by linear methods.

After this, the focus switches to development of algorithms for large-scale SVDproblems and regularized solutions of linear systems. This culminates in a toolbox for MATLAB ${ }^{\circledR}$, BBTools, which implements these algorithms in a familiar and practical setting.

Finally, we show the advantages of the method by comparing several images reconstructed using both the new and a classical algorithm. We also show that anatomical structures can be incorporated by means of restrictions.

## Resumé (Danish)

## Tomografisk rekonstrution med brug af anatomi og regularisering.

Denne afhandling udvikler en metode til rekonstruktion af billeder inden for emissionstomografi. Med udgangspunkt i brugerkrav besluttes det at fortsætte med Tikhonov regularisering efterfulgt af en filtrering med en Gausskerne. Dette følges op med en analyse af de teoretiske ulemper ved lineære metoder.

Herefter skifter fokus til at udvikle algoritmer, der kan håndtere beregninger af store SVD-problemer og bestemme regulariserede løsninger af lineære systemer. Dette kulminerer i en pakke til MATLAB ${ }^{\circledR}$, BBTools, der implementerer disse algoritmer i et velkendt og anvendeligt miljø.

Endeligt demonstreres fordelene ved den nye metode ved at sammenligne et antal billeder, der er rekonstrueret med både den nye metode og en klassisk metode. Det vises at anatomiske strukturer kan respekteres gennem restriktioner.

## Preface

This thesis was prepared at the Neurobiology Research Unit (NRU), Copenhagen University Hospital, and Informatics and Mathematical Modeling (IMM), Technical University of Denmark, in partial fulfillment of the requirements for acquiring the PhD degree in engineering.

The aim of the work is to improve the quality of the images produced in emission tomography, i.e. PET (Positron Emission Tomography) and SPECT (Single Photon Emission Computed Tomography). Improved image quality would have wide-spread diagnostic applications, both in medical research and clinically. For instance, better images would facilitate diagnosis of cancer, Alzheimer's disease and its early stage known as MCI (mild cognitive impairment), epilepsy, and many others.

In medicine, almost all images are produced using one of two algorithms: Filtered Back-Projection (FBP) or Expectation Maximization (EM). They have been used for many years, although relatively primitive.

Many technologically oriented minds have stumbled upon this fact. I am by no means the first to have attempted to make a better solution; indeed it appears to be a frequently presented topic for a PhD work. Many of these projects have successfully shown that they can produce better images than either of the two classical methods.

So why do images continue to be produced using these two algorithms?
One major reason is that medical usage have a number of requirements that are difficult, or impossible, to quantify. These requirements tends to be ignored, because it is discouraging to be unable to measure whether a progress really has been made.

If the success of a method is measured by the number of doctors using a new method after, say, 10 years, then the goals to be pursued changes radically. One example is backwards compatibility, i.e. the ability to imitate the previous methods. So far only the EM algorithm have been able to overcome this obstacle. Unlike previous projects, I have devoted a chapter to describe these requirements.

This project is an attempt to make evolutionary, rather than revolutionary, progress in image formation. There are a number of reasons that made me believe I could succeed when I started this PhD.

- When I started it seemed easy to improve on the current methods. This was only true because I was naive and unaware of prior art.
- The method chosen for the computational work is a text-book solution used in many scientific fields. I therefore assumed that off-the-shelf software was available to tackle the problem. As it turned out, this was not the case


## Computational

 powerEngineering spirit

User knowledge
and, consequently, an excessive fraction of my time as a PhD student had to be devoted to software development.

- The power of computers increases at a breath-taking pace. A problem as important as the one addressed in this thesis, which relies heavily on computations, should continuously try to take advantage of this.
- As an engineer, I have been trained to seek solutions that will work for the intended user. This is, in my opinion, the weakest point in most previous work and the real reason why many previous attempts to improve currently used methods have failed to be adopted clinically.
- I have spent a great deal of time at NRU listening to doctors who use the images. Even though they are top researchers in their fields, hence willing to adopt new methods faster than most medical staff, they have comparable needs.


## Acknowledgments

First, I wish to thank the organizations and private funds who funded this research:

- EU 5th Framework Programme (QLK6-CT-2000-00502, NCI-MCI)

Neuroreceptor Changes in Mild Cognitive Impairment

- EU 6th Framework Programme (LSHB-CT-2005-512146, DiMI)

Diagnostic Molecular Imaging

- H:S - Copenhagen Hospital Corporation
- M.L. Jørgensens og Gunnar Hansens Fond (J. nr. 2956)
- Savværksejer Jeppe Juhl og hustru Ovita Juhls Mindelegat
- Otto Mønsteds Fond (J. nr. 03-71-898)

I would also like to thank Misha Kilmer, whom I visited 6 months at Tuft's University in Boston, USA. She providing a great large-scale test-problem in diffuision optical tomography. Such test-problems are invaluable when developing numerical software. I also thank her for for pointing me at [53] which includes a proof that is fundamental to the computation routines in chapter 7.

Furthermore, I thank my brother, Nis Høgh-Ramussen, for encourage- and discouragement throughout, for helping reduce a 3D-problem to 2 D in chapter 6 , and the maxim "if you can not express what you want, then you will not get it".

Finally, I thank NRU and everyone there for supporting the project, not least financially. In particular, everyone there has been very patient during the long periods where I have been buried in a dark mathematical pit, where only I could (occasionally) glimpse where the project was heading.

Copenhagen, September, 2005

## Contents

Abstract ..... iii
Abstract (English) ..... iii
Resumé (Danish) ..... iii
Preface ..... iv
Acknowledgments ..... V
Part I: Propaedeutic
1 Introduction ..... 3
1.1 Tomographic reconstruction ..... 3
1.2 Thesis outline ..... 3
1.3 On the writing style ..... 4
1.4 Coordinates and directions ..... 4
1.5 Linear algebra ..... 5
1.6 Software and programs ..... 6
1.7 Literate programming ..... 6
1.8 Numerical notation ..... 7
2 Scanner instrumentation ..... 9
2.1 Scanners and their abbreviations ..... 9
2.2 Resolution ..... 10
2.3 The CT scanner ..... 12
2.4 Attenuation ..... 12
2.5 The SPECT scanner ..... 14
2.6 The PET scanner ..... 16
2.7 Intrinsic resolution in PET ..... 18
2.8 2D and 3D acquisitions ..... 19
2.9 References ..... 19
3 Scanner models ..... 21
3.1 Phantoms ..... 21
3.2 The Radon transform ..... 22
3.3 The discrete Radon transform ..... 23
3.4 Basic SPECT model ..... 24
3.5 Basic PET model ..... 27
3.6 Attenuation model ..... 29
3.7 General linear models ..... 29
3.8 Noise model ..... 30
3.9 Software ..... 32
3.10 References ..... 33
4 Classical reconstruction ..... 35
4.1 Methods and abbreviations ..... 35
4.2 Back-projection ..... 35
4.3 Examples of back-projection ..... 37
4.4 Filtered back-projection ..... 38
4.5 Examples of FBP ..... 41
4.6 EM reconstruction ..... 42
4.7 Examples of EM ..... 45
4.8 Other algorithms ..... 46
4.9 References ..... 47
Part II: Improved Reconstruction
5 Reconstruction goals and strategies ..... 51
5.1 User requirements ..... 51
5.2 Technical requirements ..... 53
5.3 Statistical reconstruction ..... 55
5.4 Linear reconstruction ..... 57
5.5 Regularized reconstruction ..... 58
5.6 Incorporating anatomy ..... 61
5.7 References ..... 62
6 Weighted least-squares ..... 65
6.1 Chapter overview ..... 65
6.2 Likelihood functions ..... 66
6.3 Weighted least-squares ..... 67
6.4 Model-weighted least-squares ..... 67
6.5 General reconstruction ..... 68
6.6 Hausdorff distance ..... 69
6.7 Log-likelihood range ..... 71
6.8 Analysis of weighted least-squares ..... 72
6.9 Analysis of model-weighted least-squares ..... 73
6.10 Summary of the analysis ..... 74

## Part III: Numerical Algorithms

7 Iterative framework ..... 79
7.1 Partial SVD ..... 79
7.2 Diagonalizing the NPSVD ..... 82
7.3 Updating the NPSVD ..... 82
7.4 Updates with reorthogonalization ..... 85
7.5 Krylov algorithms ..... 86
7.6 Reorthogonalization ..... 89
7.7 Restarting Lanczos ..... 90
7.8 Restarting the NPSVD ..... 92
7.9 References ..... 93
8 Computing the SVD ..... 95
8.1 On the meaning of "large-scale" ..... 95
8.2 The memory hierarchy ..... 96
8.3 Left residual factorization ..... 97
8.4 Accuracy of matrix-vector product ..... 98
8.5 Basic iteration ..... 99
8.6 Accuracy and aliasing ..... 100
8.7 Singular value decomposition ..... 102
8.8 Explicit deflation ..... 102
8.9 Implicit deflation ..... 104
8.10 Sliding window ..... 104
8.11 References ..... 105
9 Solving linear systems ..... 107
9.1 Prelude to iterative solvers ..... 107
9.2 Bidiagonalizing solver ..... 108
9.3 Unregularized solver ..... 109
9.4 Regularized solver ..... 110
9.5 Hybrid solver ..... 111
9.6 References ..... 112
10 BBTools ..... 113
10.1 Introduction to BBTools ..... 113
10.2 Central features ..... 114
10.3 Working with black-box operators ..... 117
10.4 Compatibility ..... 120
10.5 References ..... 123

## Part IV: Results

11 Results ..... 127
11.1 Introduction ..... 127
11.2 Low-count reconstruction ..... 127
11.3 Medium-count reconstruction ..... 132
11.4 Using anatomy in reconstruction ..... 134
Conclusion ..... 136
Part V: Appendix
A Addendum to WLS ..... 141
A1 Image-space bound ..... 141
A2 Noise-perturbation of WLS-surface ..... 142
B Base iteration ..... 143
B1 Stopping criteria ..... 143
C Matrix analysis ..... 147
C1 Perturbation theory ..... 147
C2 Unitarily invariant norms ..... 147
C3 Norm inequalities ..... 149
C4 Matrix approximations ..... 149
C5 Near-orthogonal matrices ..... 150
C6 Triplet orthogonality ..... 151
C7 Accuracy of singular values ..... 153
C8 Triplet accuracy and aliasing ..... 155
C9 One-sided reorthogonalization ..... 155
D BBTools reference ..... 157
Bibliography ..... 160
Index ..... 166
Nomenclature ..... 171
Papers ${ }^{1}$ ..... Attached

A cache-friendly kin to bidiagonalization with orthogonalization and restart (submitted)

[^0]

## Part I

## Propaedeutic

$\oplus$
$\oplus$

## Introduction

This chapter outlines the contents of the thesis, and establishes basic concepts and typographic conventions.

### 1.1 Tomographic reconstruction

In Greek $\pi \rho o$ (pro) means "before" and $\pi \alpha \iota \delta \epsilon v \in \epsilon \nu$ (paideuein) means "to teach", so propaedeutic means a preparatory study. By using this word I hope to maximize the chance of getting the readers attention.

The material covered in this chapter is assumed knowledge throughout the thesis. As a motivation for reading the chapter, it explains how to identify chapters and sections you may skip without missing too much of the bigger picture.

The word $\tau o ́ \mu o \sigma$ (tomos) is Greek and means a slice or a cut and $\gamma \rho \alpha \phi \hat{o}$ (graphô) means a drawing. Thus tomography is the science of drawing slices. By making many cuts through an object, a 3D description of the target can be obtained.

Reconstruction
Of course the cuts needs to be retrieved without actually doing cutting, when we are dealing with medical imaging. Instead we make indirect measurements of some function of interest and tries to "reconstruct" the original function. What is recorded and how it needs to be reconstructed depends on the physical process involved.

### 1.2 Thesis outline

One fact can not be hidden and we may as well get it settled right away. This project was made by an engineer and the text will unavoidably have a certain perspective. However, the intended readership includes doctors and others without extensive formal mathematical training.

The problem is clear: parts of the thesis make heavy use of math to accomplish an objective which is (hopefully) interesting to users who will be unable to read it. I have done my best to encapsulate the mathematics in separate chapters, and make non-technical chapters explaining the rationale independent of these.

A disadvantage of this approach is that it sometimes breaks a natural flow, and occasionally requires forward references. I decided that overall this was preferable to the alternatives.

The thesis is divided into the following parts.

## Propaedeutic

Part I The remaining chapters in this part describes the scanner physics (chapter 2), the data they produce (chapter 3), and the classical reconstruction algorithms (chapter 4).

## Improved reconstruction

Chapter 11 starts by showing that the new algorithm can reproduce the results from the Filtered Back-Projection algorithm. We then show that it may also be used to create images with higher resolution at the same noise suppression, and that any reconstruction in between these two extremes can be achieved in realtime by the operator.

Furthermore, we show that anatomical restrictions can be incorporated directly into the reconstruction. This can be accomplished in a way where the operator is fully in control of the process.

We follow up with a conclusion in chapter 11.4.

### 1.3 On the writing style

As recommended by John Kirkman [41], most of the thesis is written in an active first-person style, to the extent allowed by my ability on English. Most of the thesis in in plural, and singular is reserved mainly for personal observations, opinions, and decisions about the text itself.

In order to maintain a high level of readability, I have tried to maintain a consistent notation throughout the report, as defined in the nomenclature on p. 171. To further enhance the reader's focus on the contents rather than the text, I hope that the typographical quality provided by the typesetting package IATEX will keep such annoyances at a minimum.

### 1.4 Coordinates and directions

The coordinates of interest in this thesis are defined with respect to the scanner. Usually a patient rests on the back and is looking upwards, and the spine is thought of as the "human axis" and is defined as the z -axis, with the positive directions going from the feet to the head. The $y$-axis goes upwards and $x$-axis is positive in the left direction as seen from the patient.

Unfortunately, the coordinate system resulting from this definition is left-handed, going against deeply grown habits in the engineering/mathematical community. However, using a right-handed system would be just as confusing for the medical society, and thus the definitions above are adopted.

The terms left and right are always relative to the person in the scanner.

## Planes

There are several ways to define planes anatomically, and these may in turn be used to establish a coordinate system [79]. Some are defined with respect to cranial features and others are based on anatomical landmarks. For this thesis we do not need the precise definitions, and the meaning of the following terms can be given roughly:

- Sagittal: yz-plane. These are planes parallel to the median plane, which seperates the two hemispheres of the brain. This coresponds to seing the person from the side.
- Transverse/transaxial: xy-plane. The transverse planes are horizontal for a person standing up. A transaxial plane is like looking at a person from above.
- Coronal: xz-plane. Orthogonal to the two other planes and corresponds to looking at a person from behind.

In radiological practice the brain is often tilted in order to record the whole brain with a fewer number of slices. We will not distinguish between the differences in this thesis.

### 1.5 Linear algebra

$\iiint \quad$ Sections, marked like this one, contains mathematical derivations. It is used in general chapters as a signal to people who prefer to skip the mathematical parts. Chapters full of mathematics can be identified from the introduction.

The thesis assumes a basic knowledge of linear algebra and numerical algorithms. Nonetheless, different sources use slightly different notation and definitions. Although most notation is introduced the first time it is used, the nomenclature given in the back, starting on page 171 , should settle such ambiguities.

In any case, the most important rules are given here:

- Matrices are written in uppercase roman bold, e.g. A.
- All vectors are columns. They are printed in lowercase roman bold, e.g. x.
- We occasionally use the term " $\mathbf{A}$ is orthogonal" as a short-hand for " $\mathbf{A}$ has orthonormal columns".
- The Hermitian of an operator, $\mathbf{A}$, is the Hermitian adjoint operator of $\mathbf{A}$. If $\mathbf{A}$ is a matrix, the Hermitian is the transpose-conjugate of $\mathbf{A}$, i.e. $\mathbf{A}^{\mathrm{H}}$. Thus, the Hermitian of an operator is not to the same as saying that $\mathbf{A}$ is Hermitian or self-adjoint (which, for matrices, say that $\mathbf{A}=\mathbf{A}^{\mathrm{H}}$ ).

If a matrix is "small" it it with respect to the dimensions. Other interpretations will always be specified such as " $\mathbf{A}$ is small in norm" or $\|\mathbf{A}\|$ is small.

When a key-equation is derived or defined, it will either be encapsulated in a box or be given a name (or both). I hope the burden of reading is lightened when an equation is referred to by name rather than merely by a number such as eq. (7.12).

### 1.6 Software and programs

001010
100001 111100

The following quote is attributed to folklore in [14]:

The difference between theory and practice is smaller in theory than it is in practice.

Almost all the theory in the thesis has been fleshed out as working computer programs written in MATLAB ${ }^{\circledR}$, which is a convenient language for numerical algorithms. This is particularly true in linear algebra which forms the back-bone of this thesis. They have been collected in a toolbox, BBTools [33], which is described in chapter 10, and is available under the GNU General Public License (GPL) as published by the Free Software Foundation [24].

Several examples in the text assumes that the BBTools is installed. This allows much of the code to be expressed in a form where the algorithms are undistorted by the algorithmic details of the underlying linear function.

Sections mainly concerned with computer implementations are marked as this one. I personally prefer texts that put mathematical derivations into usage at the earliest possible opportunity, but I realize that this view is not universal.

Programming sections may include code or details that may appear inappropriate for those who are more mathematically oriented. Persons who feel that way should easily be able to skip these sections, possibly leaving them for a second reading.

### 1.7 Literate programming

001010
100001 111100

When it comes to algorithms this thesis differ from most texts. To put it in a nutshell: the algorithms in this thesis works. If the reader can not get an algorithm to work, it must be because the author made a bug. This is in contrast to the use of pseudo-code where it could be a misunderstanding from the reader. It should also be a comfort to know that the author actually ran each algorithm.

Although this approach makes the algorithms more verbose than pseudo-code, it does have the advantage of telling the truth, the whole truth and nothing but the truth. Most people who have tried to implement a program from pseudo-code should appreciate this quality.

Nevertheless, most source-code is sadly quite undocumented. Most examples in this thesis takes its onset on literate programming, as described by its inventor Donald Ervin Knuth [42]:

Let us change our traditional attitude to the construction of programs: Instead of imagining that our main task is to instruct a computer what to do, let us concentrate rather on explaining to human beings what we want a computer to do.

Unfortunately, the existing tools for literate programming are little but advanced pretty-printers. In particular, line-comments are not supported well. The value of good line-comments are evident from a quick glance at respected codes such as [58].

As a result, the code presented in the thesis was produced by a custom set of tools. The most important feature is that scopes has been replaced with connector-lines. The idea is that everything inside a scope shares a common line on the left, except for out-dented expressions. Apart from that, the most important feature is that any identifier can be replaced with an arbitrary command. Thus it is possible to match the notation in the program with the remaining parts of the thesis.

In MATLAB-scripts (see alg. 1.1) every closing scope can be replaced by an end except for functions where nothing should be inserted. The branching construct shows the scoping rules. Since the whole if-line is an out-dented expression, it should be considered to be in the scope of the whole function. The same argument goes for the else-statement but not the assignment on the same line.

## Efficiency

Examples included in the thesis are meant to be illustrative, not efficient. Operations may be repeated, zero by construction, or never used. In particular, language-dependent optimizations are intentionally avoided, although they often have tremendous impact in speed. Examples include preallocating matrices and reformulations such as $\mathbf{A}^{\mathrm{H}} \mathbf{x} \rightarrow\left(\mathbf{x}^{\mathrm{H}} \mathbf{A}\right)^{\mathrm{H}}$ and $\mathbf{A}^{-1} \mathbf{x} \rightarrow \mathbf{A} \backslash \mathbf{x}$.

For efficient code, please consult the code in BBTools [33].

### 1.8 Numerical notation

In order to analyze numerical algorithms, we need some notation. As most texts dealing with rounding errors in floating point arithmetics, we assume there is a constant, $\epsilon_{M}$, so that $a \odot b=(a \circ b)(1+\eta)$ where $|\eta| \leq \epsilon_{M}$. Here $\circ$ is one of the basic arithmetic functions addition, subtraction, multiplication, division, or square root.

Algorithm 1.1: Example of MATLAB-script
A nonsense demonstration showing how a MATLAB-script is typeset in the thesis. The source is not shown in actual programs. The end that closes the if-structure is not shown since it collapses into an end-scope symbol.

| function $a=$ format_demo( $\mathbf{A}$ ) | Source | Demonstrates |
| :---: | :---: | :---: |
| $\mathbf{x}_{1}=\mathbf{A}(:$, end $) ;$ | $\mathrm{x} 1=\mathrm{A}(\mathrm{l}$, end) ; | Vector and matrix |
| $\beta=\\|\mathbf{A}\\|_{2} ;$ | beta=norm(A) ; | Scalar, matrix and standard function |
| $\boldsymbol{Y}$ if $\beta>$ length $\left(\mathbf{x}_{1}\right)$ | if beta>length(x1) | Start of scope |
| $a=$ length $\left(\mathbf{x}_{1}\right) ;$ | $a=1$ ength (x1) ; | Normal identifier |
| $\square_{4}$ else,$\quad a=\beta$; | else, $\mathrm{a}=$ beta; | Outdenting and multiple commands |

The constant $\epsilon_{M}$ is known as machine precision or machine epsilon. The relation is only valid as long as the result stays within the range of the arithmetic. While we may occasionally consider underflow (e.g. results flushed to zero or quantized to lower accuracy), it will generally be assumed that overflow does not occur.

For completeness it is usually mentioned that a few old computers (like the Cray C90 [15]) does not satisfy the relation for addition and subtraction. Such machines will be ignored, and we note that common software such as MATLAB ${ }^{\circledR}$ have abandoned such machines.

For linear algebra it is practical to have some bounds on a higher level. Specifically, we will consider the precision of an inner product of two vectors of length $l$. It can be proved (barring overflow and underflow) that accumulating the sum is stable, and computes the sum of the elements perturbed by a relative error smaller than $\nu_{l}=l \epsilon_{M} /\left(1-l \epsilon_{M}\right) \approx l \epsilon_{M}$, where $l$ is the number of entries added ${ }^{1}$ [69].

[^1]
## Chapter

## Scanner instrumentation

Here the basic principles behind the scanners are introduced. We show example images from each scanner type, and explain the physics behind the CT, PET, and SPECT scanners. The MR scanner is only mentioned briefly, as the principles differ significantly from the others.

This chapter serves as a primer for those who wants to understand the principles behind the scanners. It may also be beneficial for people who just want to know a bit about medical scanners in general.

### 2.1 Scanners and their abbreviations

There are 4 kinds of tomographic scanners used in medical applications that are capable of imaging any part of the body. They are all identified by abbreviations as given in table 2.1. Some synonyms found in the literature are given in table 2.2, but they will not be used further in this thesis. For completeness it should be mentioned that other methods exist, with ultra-sound probably being the most well-known.

Table 2.1:
Acronyms for scanners and modalities used throughout the thesis.

| Acronym | Meaning | Type |
| :--- | :--- | :--- |
| CT | Computer Tomography | Transmission |
| PET | Positron Emission Tomography | Emission |
| SPECT | Single Photon Emission Computerized | Emission |
|  | Tomography |  |
| MR | Magnetic Resonance | Magnetic |

The first scanner to emerge was the CT-scanner invented independently by Allan MacLeod Cormack and Sir Godfrey Newbold Hounsfield who shared the 1979 "Nobel Prize in Physiology or Medicine". It is based on x-rays but instead of taking a single 2D photo, a large number of images are taken from different angles. These are then fed into a computer (hence the name) and the final images generated as explained in sec. 2.3. Since the x-rays are sent through the object, a CT-scan is known as a transmission scan.

This is in contrast to PET and SPECT, where a small dose of a radioactive substance is administered to the subject. Common examples are injections of water with a radioactive isotope of oxygen atom or inhalation of radioactive xenon.

Table 2.2:
Other acronyms found in the literature. These will not be used in this thesis.

| Acronym | Meaning | Explanation |
| :--- | :--- | :--- |
| CAT | Computer Aided Tomography | Same as CT |
|  | or | Computerized Axial Tomography |

When the radioactive atom decays a photon (in SPECT) or a positron (in PET) is expelled. In both cases the scanner tracks the direction of the decay, ultimately leading to a map of the concentration of the administered drug. Since both methods rely on emission from the substance itself, the methods are known as emission tomography.

Finally MR uses a completely different principle. An MR-scanner exploit that a nucleus in a magnetic field precesses around an axis. This precession can be manipulated with radio-waves, allowing different characteristics of the nucleus to be determined. In medical MR it is primarily protons (i.e. hydrogen) that are in play, and characteristics such as proton density can give very detailed anatomical information.

The importance of MR in this thesis is that its high resolution provide anatomical information that is useful for improving the resolution in PET and SPECT.

## Image gallery

The best way to get acquainted with the different scanner types is probably to see some images, as they usually come out of the scanners. Fig. 2.2 to 2.5 display sample images from all the scanners discussed above. This shows the relative strength and weaknesses of the scanners, and can safely be skipped by those who are already familiar with such images.

### 2.2 Resolution

Before discussing the scanners, we need a precise definition of resolution. As shown in the image gallery, the scanners produce images that are blurred versions of the function we would really like to measure.

In medicine, the resolution is usually quantified by measuring the Full-Width at Half Maximum (FWHM). This concept is associated with a Point Spread Function (PSF), which is the function measured of a point-source. Ideally, this would be a sharp spike, but in practice it is a function with a peak at the point-source that gradually tapers off.

The importance of the FWHM illustrated in fig. 2.1. Since the PSF are, by definition, $1 / 2$ of the peaks at the intersection point, the sum at this point is the average of the peaks. If the PSFs are monotonously decreasing, two peaks of equal intensity more than one FWHM apart will result in two crests separated by a cleft.

Figure 2.1:
If the PSF
decreases monotonously, then two equally intense peaks can be distinguished when they are separated by more than one FWHM.

This definition generalizes to any dimension. Generally, the FWHM is either measured as the longest distance or, if it as an ellipse, the major and minor axes are given separately.

Figure 2.2:
Example CT images. CT offer the best resolution of the ordinary scanners, but does not have a good contrast in brain tissue.

Figure 2.3:
Example MR images. $M R$ and $C T$ are the primary means for identifying anatomical structures in the brain.

Figure 2.4:
Example of images produced with PET on the same person as in fig. 2.3. PET is a difficult technique, and requires a facility with a large staff of physicists and chemists in addition to the medical staff.

(a) Sagittal CT-image of a normal person.

(b) Transaxial slice of the same person.

(a) Sagittal MR-image of a normal person.

(b) Transaxial slice of the same person.

(a) Transaxial cut of an altanserin scanning. (Altanserin images a serotonin receptor.)
(b) Another cut, corresponding roughly to fig. 2.3(b)

$\oplus$

Figure 2.5:
Transaxial images produced using SPECT. The poor resolution limits the utility of the technique.


### 2.3 The CT scanner

Transmission tomography investigates a semi-transparent object by sending radiation through it from a number of angles as sketched in fig. 2.6. The projections are then fed to a computer and the amount of absorption is reconstructed. In medical imaging the relevant radiation is x-rays or $\gamma$-radiation. In the case of $x$ rays the scanning is called a CT-scan as an initialism (i.e. an unpronounceable abbreviation) for Computer Tomography. The term transmission scan is broader, and includes $\gamma$-radiation.

The CT-scanner is not only important in itself but also in emission tomography. First, emission tomography resembles transmission tomography from a technical point of view. Second, a transmission scan is often performed in addition to an emission scan in order to correct for attenuation of the radioactive tracer. Finally, the higher resolution of a CT-scan can be used to identify anatomical structures which can be used to interpret the emission scan or even improve it.

### 2.4 Attenuation

As the x-ray traverses through an object, the ray is attenuated according to the tissue-type. The simple example in fig. 2.6 only has two types, but in general

Figure 2.6:
Principle of a modern $C T$-scanner An x-ray source in the periphery of the scanner ring is rotated around the center. Each position of the source gives a projection as illustrated.

(a) Transmission from one angle. Measurement

(c) The measured intensity profile.

## CT principle


(b) A slightly different angle. Measurement

(d) Profile from the other angle.
it can have many different values. Physics tells us that the intensity through a homogeneous absorbing substance can be written as follows:

$$
f(x)=I_{0} e^{-\mu x}
$$

Here $x$ is the length through the medium, $\mu$ depends on the material, and $I_{0}$ is the intensity of the incoming beam.

The coefficient $\mu$ is known as the linear attenuation coefficient, and it depends on the substance as well as the energy of the photons (i.e. the wavelength) and has dimension of $1 /$ length ( $1 / \mathrm{m}$ in SI-units). Most tables do not list $\mu$ directly but rather the mass attenuation coefficient $\mu / \rho$ where $\rho$ is the density (usually given in $\mathrm{g} / \mathrm{cm}^{3}$ ).

Some typical values for relevant energies are found in table 2.3. It can be seen that the main difference between the absorption is due to differences in the tissue density. In neurological studies bone is the only tissue with a density that differs significantly from water. Therefore CT is not usually the method of choice for retrieving anatomical information about brain tissue, except when a contrast agent ${ }^{1}$ is used. However, the access to MR scans are limited and CT provides valuable information in many cases, and is an excellent technique for identifying intracranial bleeding.

When the attenuation is a function of position, the line must be split up in small segments. For simplicity we let the segments be distributed equidistantly with spacing $\Delta x$. This gives $f(x)=I_{0} e^{-\Delta x \mu\left(x_{0}\right)} e^{-\Delta x \mu\left(x_{1}\right)} \cdots e^{-\Delta x \mu\left(x_{N}\right)}$ or simply $f(x)=I_{0} e^{-\left(\Delta x \mu\left(x_{0}\right)+\cdots+\Delta x \mu\left(x_{N}\right)\right.}$. In the limit of an infinite number of segments this becomes $f(x)=I_{0} e^{-\int_{0}^{x} \mu(x) d x}$.

This means that the measured profiles in fig. 2.6(c) and 2.6(d) contains information about the line-integrals of the attenuation-map since $\int_{0}^{x} \mu(x) d x=\ln (-$

[^2]Table 2.3:
Absorption in different tissues. The mass attenuation coefficient (MAC) and nominal densities (from [36]) are used to get the linear attenuation coefficient (LAC). The last column gives the absorption through one cm of tissue. Typical photon energies are 7080 keV or lower for CT, 100-150 keV for SPECT, and 511 keV for PET.

| Tissue | Density $(\rho)$ <br> $\mathrm{g} / \mathrm{cm}^{3}$ | Photon <br> keV | MAC $(\mu / \rho)$ <br> $\mathrm{cm}^{2} / \mathrm{g}$ | LAC $(\mu)$ <br> $\mathrm{cm}^{-1}$ | Absorption <br> $\%$ pr. cm |
| :--- | :---: | :---: | :---: | :---: | :---: |
| Bone | 1.92 | 500 | 0.090 | 0.173 | 16 |
|  |  | 100 | 0.186 | 0.356 | 30 |
| Water | 1.00 | 50 | 0.424 | 0.814 | 56 |
|  |  | 500 | 0.097 | 0.097 | 9 |
| Blood | 1.06 | 500 | 0.171 | 0.171 | 16 |
|  |  | 500 | 0.227 | 0.227 | 20 |
|  |  | 100 | 0.170 | 0.102 | 10 |
| Brain $^{\dagger}$ | 1.04 | 50 | 0.228 | 0.180 | 16 |
|  |  | 500 | 0.096 | 0.100 | 21 |
|  |  | 100 | 0.170 | 0.177 | 10 |
|  |  | 50 | 0.228 | 0.237 | 21 |

$\dagger$ Brain includes both gray- and white-matter as they are virtually indistinguishable.
$f(x) / I_{0}$ ). In the general case, $x$ represents the position on the line inside the volume.

In fig. 2.6 each point on the graph represents such line integral. In this simple example the $\mu$-function is either zero or a constant, say, $c$. Thus the logarithm of the measurement will be $-c$ times the length that the line intersects the object.

The $\mu$-function is sometimes called a $\mu$-map. The CT-images are slices computed from these projections, usually trans-axial e.g. planes parallel to the floor when standing. The procedure for doing this will be explained in chapter 4.

### 2.5 The SPECT scanner

The SPECT-scanner is an emission tomographic scanner, i.e. the scanned object must itself be radiating. This allows the scanner to track a radioactive substance within the patient.

The principle of a basic SPECT-scanner is depicted in fig. 2.7. Much as the CTscanner, the SPECT-scanner measures the radiation from a number of directions. However, since the object is radiating in all directions simultaneously, the $\gamma$-rays must be limited prior to the measurements. Also, most SPECT-scanners are multihead scanners, i.e there are two or more $\gamma$-cameras rotating simultaneously.

For completeness, we list a couple of tracers used in SPECT-studies in tab. 2.4. One noteworthy observation is that the half-lives of the tracers are sufficiently long that they can generally be created and transported from another site, e.g. a commercial vendor.

Most tracers used clinically are based on technetium $\left({ }^{99} \mathrm{Tc}\right)$, which is produced by decaying molybdenum ( ${ }^{99} \mathrm{Mo}$ ). The half-life of ${ }^{99} \mathrm{Mo}$ is about 66 h , and can typically produce ${ }^{99} \mathrm{Tc}$ for about a week. The tracer is labeled with the technetium using a standard kit, making SPECT relatively inexpensive.

## Collimators

To get directional information from the beams, they are limited by a device known as a collimator. A collimator consists of a number of pipes (usually of lead), and

## SPECT principle

Figure 2.7:
Principle of a basic SPECT-scanner with a single head and a parallel-hole collimator. The activity of a radiating object is measured by a $\gamma$-camera. A collimator, placed in front of the camera, limits the radiation to beams normal to the surface. The camera is rotated around to object to produce measurements from a number of directions.

Table 2.4:
Examples of isotopes and tracers used in SPECTstudies.

(a) Transmission from one angle. Measurement

(c) The measured intensity profile.

Table 2.5:
Examples of commercial collimators for the Philips (formerly Marconi) IRIX-scanner. (Private communications.)

SPECT principle

(b) A slightly different angle. Measurement

(d) Profile from the other angle.

| Isotope | Half-life | Tracer | Use |
| :---: | :---: | :--- | :--- |
| ${ }^{99} \mathrm{Tc}$ | 6.0 h | HMPAO $^{\dagger}$ | Blood flow |
| ${ }^{133} \mathrm{Xe}$ | 5.2 d | Xenon | Blood flow |
| ${ }^{123} \mathrm{I}$ | 13.2 h | PE2I | Receptor studies |
|  |  | ADAM | Receptor studies |

$\dagger$ HMPAO (Hexa-Methyl-Propylene Amine Oxime) is a lipophilic tracer used to diagnose e.g. strokes and dementia. PE2I, a cocaine derivative, is used to study dopamine transporters, while ADAM is used to study serotonine transporters.
a beam is only allowed to pass through when it is aligned with the tube. The scanner depicted in fig. 2.7 is a parallel-hole collimator, but many other geometries are possible.

A fan-beam collimator has all tubes directed towards a line parallel to the scanner z-axis. Thus, all trans-axial cuts through the collimator will have the same structure. A cone-beam collimator has tubes all directed at a single point in space. Finally, a pin-hole collimator consists of a single hole, much like a pin-hole camera.

The collimator acts as a filter and blocks the majority of $\gamma$-rays. Therefore, the collimator chosen for a task must be a trade-off between sensitivity (how many $\gamma$ rays are blocked) and resolution (how closely the beam is aligned with the pipes). Typical examples of the pipe dimensions of low-energy high-resolution and ultra-high-resolution collimators are given in tab. 2.5.

| Name | Geometry | Hole <br> size <br> mm | Hole <br> length <br> mm | Thickness | Focal length |
| :--- | :--- | :---: | :---: | :---: | :---: |
| LEHR |  | mm | mm |  |  |
| LEUHR-PAR | Parallel | 1.22 | 27.0 | 0.20 | $\infty$ |
| LEHR-FAN | Fan-beam | 1.78 | 58.4 | 0.15 | $\infty$ |
| LEUHR-FAN | Fan-beam | 1.40 | 25.4 | 0.18 | 500 |

Although SPECT does not have a theoretical lower resolution limit, the collimator introduces severe practical limits. First, to obtain a sensitivity which is useful in medical applications, it is necessary to allow beams that diverge from the ideal. Also, the thickness of the tubes are limited, allowing some photons to be reflected or penetrate the tube. In practice, the resolution of a SPECT-scanner is limited to about 8 mm in typical clinical studies.

These effects are minimal for a pin-hole collimator. However, to compensate for the poor sensitivity of such a collimator, it is usually necessary to inject a very high dose of radioactivity for a complete brain-scan. Normally, the high dose of radioactivity will kill the subject, which limits the utility to animal studies. The resolution, however, is excellent.

## SPECT imperfections

There are a number of imperfections that limits the resolution of SPECT. First, the attenuation discussed in sec. 2.4 is equally valid for SPECT. However, where the CT-scanner is used to measure the attenuation, it is a nuisance in SPECT we would rather have been without.

Typically, the energies in SPECT are higher than in CT. As can be seen in tab. 2.3, the attenuation becomes less problematic for the higher photon energies. Furthermore, the difference between the mass attenuation coefficient for bone and water decreases radpidly in the energy-range $100-120 \mathrm{keV}$.

Secondly, some of the "attenuated" photons are not absorbed, but scattered in new directions. This comes at a loss of energy, which depends on the deflection angle. This effect can, to some extent, be disregarded by only considering photons in a certain energy window.

There are many other imperfections in SPECT, but it is beyond the scope of this thesis to explain them in detail. We refer the interested reader to the references in sec. 2.9.

### 2.6 The PET scanner

Like SPECT, PET is an emission tomographic technique. However, the PETscanner is based on radioactive isotopes that has a $\beta^{+}$decay, i.e. emits a positron. Shortly after emission, the positron annihilates with an electron and they are both converted into pure energy in the form of $\gamma$-radiation. As it is impossible to conserve both energy and momentum with a single photon, the energy is converted into two $\gamma$-quanta of (nearly) opposite direction as sketched in fig. 2.8.

Figure 2.8:
The basic principle of a PET-scanner. The ligand emits a positron, which quickly annihilates into two 511 keV $\gamma$-photons of opposite directions.


Table 2.6:
Isotopes and example of tracers used in PET-studies.

The PET-scanner consists of rings of scintillation crystals, and the line of response (LOR) of the annihilation is determined by the pair of crystals that detects the emission nearly simultaneously. This eliminates the need for a collimator, which makes PET far more sensitive then SPECT.

Because of the $\beta^{+}$-decay, most isotopes has a relatively short life as shown in table 2.6. This is both an advantage and a problem. It is a problem because the isotopes has to be produced the place they are used. It is difficult to build the isotope into the tracer, test the purity and administering the drug during the short time available. One advantage is that the substance is quickly removed which reduces the dose absorbed by the subject and allows consecutive scans.

| Isotope | Half-life | Substance | Use |
| :---: | :---: | :--- | :--- |
| ${ }^{15} \mathrm{O}$ | 2 m | Oxygen | Oxygen metabolism |
|  |  | Carbon monoxide | Blood volume |
|  |  | Water | Blood flow |
| ${ }^{11} \mathrm{C}$ | 20 m | Flumazenil | Receptor studies |
|  |  | Raclopride | Receptor studies |
| ${ }^{18} \mathrm{~F}$ | 110 m | FDG $^{\dagger}$ | Glucose metabolism |
|  |  | Altanserin | Receptor studies |
| ${ }^{13} \mathrm{~N}$ | 10 m | Ammonia | Cardiac blood flow |
| FDG (Fluoro-Deoxy-Glucose) is a sugar used for activation studies and cancer <br> treatment (i.e. oncology). |  |  |  |

## PET imperfections

As in SPECT, there are a large number of non-ideal effects that ultimately limits the resolution in PET. It is beyond the scope of this thesis to go into details with most of them, but we briefly discuss a few with practical consequences.

## Physical limitations

There are several unavoidable physical effects that limits the resolution of a PETscanner. First, the positron usually contains a residual momentum when it annihilates with the electron. Since momentum is conserved, the result is that the two beams are not completely collinear. This effect depends on the speed of the emitted positron, which depends on the tracer used in the study.

Secondly, there is a distance from the emission of the positron to the anihilation. This is called the "positron range", and depends on the speed of the emitted positron and the nature of the surrounding tissue. This is further discussed in sec. 2.7.

Thirdly, attenuation and scatter have much the same effect on PET as it has on SPECT. However, if one of the photons is scattered, the measured line may not be inside the object. This makes scatter more severe for PET than it is for SPECT. As in SPECT, most (but not all) of the scattered photons can be identified by measuring the energy-loss by the photon.

## Equipment limitations

The scintillation crystals in the PET-scanner converts the $511 \mathrm{keV} \gamma$-rays to lower energies which can be detected using conventional electronics. Each crystal is
used to detect emissions that occurs on all the lines that intersects it. This is unlike SPECT where the crystal residing under a collimator tube only detects photons from (approximately) one direction.

This greatly increase the sensitivity of PET compared to SPECT. However, a crystal has a dead-time after measuring a photon where it returns to equillibrium and is unable to detect new events. This is a physical effect, and is not merely limited by the electronic circuits. The result is that the number of detected events is a non-linear function of the actual number of photons emitted on a line.

Most algorithms used to form images depends on the linearity of the response. The dead-time makes it necessary to correct the measurements before further processing. Fortunately, this is standard on modern PET-scanners.

## Randoms

Finally, we mention randoms which may be measured if two emissions occur at nearly the same time. If two of the four photons are absorbed, the measured LOR may be completely unrelated to the object.

The standard way to correct for randoms is to measure and subtract known randoms. These can be identified when two emissions are detected with a timeinterval that is too large to be emitted from a point inside the scanner. Thus, there is a window of "acceptance" where counts are included, followed by a window of "randoms" of the same duration. When a random is detected, it is subtracted instead of added. This technique is known as on-line subtraction.
This arrangement is special because it allows "negative counts". On the average, the measured number will be correct, but the existence of negative counts have implications for some algorithms. In principle, the "trues" and "randoms" can be measured independently, but not all PET-scanners support this.

### 2.7 Intrinsic resolution in PET

Table 2.7:
Isotopes and typical ranges [16].

Unlike SPECT, the positron range and non-collinearity puts a fundamental limit to the resolution of PET. The size of the scintillation crystals and other non-ideal properties of the PET-scanner further limits the resolution.

The positron range depends on the isotope, with some examples given in table 2.7. Here $r$ is the FWHM-equivalent of the standard deviation and can thus be combined with other factors.

| Isotope | ${ }^{18} \mathrm{~F}$ | ${ }^{11} \mathrm{C}$ | ${ }^{68} \mathrm{Ga}$ |
| :--- | :---: | :---: | :---: |
| FWHM [mm] | 0.13 | 0.13 | 0.31 |
| FW(0.1)M [mm] | 0.38 | 0.39 | 1.60 |
| $r[\mathrm{~mm}]$ | 0.54 | 0.92 | 2.80 |

The effect arising from non-collinearity depends on the diameter of the detectorring, $D$, and at the center it is given roughly by $.0022 D$. Even if technology could allow perfect detection, the intrinsic resolution of a PET-scanner would still be limited to $R=1.25 \sqrt{r^{2}+(.0022 D)^{2}}$ [16]. For a dedicated brain-scanner this could be less than 1 mm for ${ }^{18} \mathrm{~F}$-images.

Practical scanners must have crystals of finite width, $d$, and there may be a significant coupling between crystals which also limits the resolution. In this case
the range is approximately given by $R=1.25 \sqrt{(d / 2)^{2}+r^{2}+(.0022 D)^{2}+b^{2}}$, where $d$ is the crystal width and $b$ is the coupling. Tomographs with a photomultiplier tube (PMT) for each crystal typical have $b<.3 \mathrm{~mm}$, but for economic reasons crystals ${ }^{2}$ often share PMTs. This usually limits coupling to $b \approx 2 \mathrm{~mm}$.

Thus the achievable resolution of modern scanners useful for brain-scanning is around $4-6 \mathrm{~mm}$, and 2 mm is a reasonable estimate of the best resolution that can be achieved with current technology.

### 2.82 D and 3D acquisitions

So far we only discussed how a transaxial slice is acquired. In order to produce a 3D volume, a number of slices are measured simultaneously. In SPECT, the $\gamma$-camera acquires a number of plane-projections, and a line from each plane is used as the data for a transaxial slice. This may only be approximately true if the object moves during the acquisition, but this is usually ignored.

Modern PET-scanners can be operated in 2D and 3D mode. The PET-scanner consists of a number of rings, and in 2D-mode each ring collects the data for a slice. The number of slices can be nearly doubled by treating adjacent rings as a single ring, but we shall not distinguish between a "virtual ring" formed by a ring-pair and a real pair.

Although the object radiates in all directions, only beams that are not perpendicular to the scanner $z$-axis are desirable in 2D-mode. Much like the SPECT-scanner, other beams are limited using a metal shield. In PET, the shields that blocks this radiation are known as septa. The most common material for the septa is tungsten.

In 3D-mode the septa are removed during the acquisition. This drastically improves the sensitivity of the scanner, but also exacerbates the imperfections discussed in sec. 2.6.1. More importantly to the work in this thesis, the amount of data acquired during a scan increases enormously.

While the software developed in this thesis was intended to allow full 3D reconstruction, the memory requirements is at least 4 Gb for a practical problem. This can not be accommodated by the 32 -bit computers available at NRU during the project. We shall therefore, with some regret, disregard 3D reconstructions from the thesis. We look forward to working on 3D-problems as 64 -bit computers become mainstream.

### 2.9 References

Most of the material in this chapter is standard. The tome by Bushberg et. al. [12] includes almost everything described in this chapter, and can be recommended for both practioners and technical oriented people. It is particularly strong on CT-scanners, but also includes much material on gamma-cameras.

For a hands-on guide to SPECT, the primer by English [19] is a good starting point. However, it mostly leaves out technical details.

A good source for PET, for those who are technological oriented is [77]. Most of the descriptions of PET-instrumentation in this chapter was derived from this

[^3]book. A shorter, but informational, paper for PET is [74], which should be readable for most involved with PET.

Positron range is treated in [16], while analytical approximations for several isotopes can be found in [28].

For 3D-acquisition, algorithms, and scanner limitations, [5] is a thorough and complete reference. It is, however, primarily oriented towards people with a strong mathematical background.
Many issues regarding imperfect hardware in PET is treated statistically in [84]. Randoms, in particular, is treated this way in [85], while a hardware-oriented and experimental treatment can be found in [35].


This chapter develops mathematical models of the emission tomographic scanners. We define test-images, phantoms, which will be used throughout the thesis. We also show typical responses measured by the scanners in the form of sinograms.

The bulk of the chapter is mathematically oriented. We define the Radon transform, which is followed by specific models of the SPECT- and PET-scanner, followed by a more general linear model. We also describe the Poisson distribution, which is the fundamental statistical model underlying emission tomographic measurements.

### 3.1 Phantoms

As a prelude modeling the scanners, we shall define some test-images which will be used to illustrate the procedures. In tomography, a test-image is called a phantom. Physical phantoms are usually made of plastic, and may contain a number of chambers. These chambers may in turn be filled with radioactive liquid, which can be measured separately. This allows a quantitative measurement, which can be used to evaluate the accuracy of a real reconstruction.

We shall use the following synthetic phantoms in this thesis. In fig. 3.1 we have a test-phantom, which is suitable for estimating the resolution. It simulates a number of rods, arranged in a hexagonal grid. The smallest rods which can be distinguished individually represents the resolution of the scanner.

Figure 3.1:
This phantom is useful in testing the resolution of a reconstruction image.

## Phantom 1



The second phantom, depicted in fig. 3.2, is a classical phantom used in neurological studies. It consists of a sum of ellipses, which can be treated analytically. Although we shall not use this property, we use it because it is familiar to many. We shall use a variant due to Peter Aundal Toft [72], which has a better contrast than the original.

Figure 3.2:
The Shepp-Logan phantom is a classical phantom used to in neurological tomography [72].

Phantom 2


All phantoms and reconstructed images are represented as $256 \times 256$ pixels. There is nothing magic about this size, except that images are traditionally rendered as a $2^{n} \times 2^{n}$ for some $n$. For the purpose of this thesis, a size such as $257 \times 257$ (i.e. a prime) would serve just as well.

For precision, we shall adopt the convention that all images represents a function, $f(x, y)$, defined in the range $(x, y) \in[-1 ; 1] \times[-1 ; 1]$, where $\times$ is the cartesian product. By definition we set $f(x, y)=0$ for any value outside this area, i.e. it has compact support.

### 3.2 The Radon transform

## SJS

normal form

The scanners described in chapter 2, i.e. CT, SPECT, and PET, all measure the response along a large number of Lines Of Response (LORs). In the case of emission tomography, the measurements are the sum of events occuring on each LOR during the scan.

The number of events is proportional to the activity of the substance at each point. Thus, if $\ell$ represent a LOR and $f(x, y)$ is the activity at the point $(x, y)$, the measured activity is proportional to the quantity:

$$
g(\ell)=\int_{\ell} f(x, y) d x
$$

For this to be useful, we need to parameterize the line $\ell$. While several possibilities exist, the most convenient representation in tomography is the normal form, which is specified by the distance from origo, $\rho$, and the angle, $\theta$, as shown in fig. 3.3. It is worth noting that for a parallel-hole collimator in SPECT, the angle of the LOR, $\theta$, is the inclination of the collimator.

Figure 3.3:
A line represented in normal form.


Using this convention, the function $g(\theta, \rho)$ is now a function of $f(x, y)$ as specified by eq. (3.1). This is known as the Radon-transform, after Johann Karl August Radon [81] who showed that the functions were linked by a differential equation that has a unique solution under mild restrictions.

It follows directly from eq. (3.1) that the Radon transform is symmetric, i.e. it satisfies $g(\theta, \rho)=g\left(\theta+k 360^{\circ}, \rho\right), k \in \mathbb{Z}$, and $g(\theta, \rho)=g\left(\theta+180^{\circ},-\rho\right)$. This means we only need to consider the interval $0 \leq \theta<180^{\circ}$, or alternatively, $0 \leq \theta<360^{\circ}$ and $\rho \geq 0$.

The first form is by far the most convenient in tomographic problems, so we shall adopt this convention throughout. Furthermore, we only consider images with compact support. Specifically we have $f(x, y)=0$ for $\max \{|x|,|y|\}>1$. Therefore, we have $g(\theta, \rho)=0$ for $\rho>\sqrt{2}$.

However, this range is often limited further by a field of view (FOV). In this case we define $f(x, y)=0$ for $\sqrt{x^{2}+y^{2}}>R$, where $0<R \leq \sqrt{2}$ is the largest distance from the origin. In this case we have $g(\theta, \rho)=0$ for $|\rho|>R$.

### 3.3 The discrete Radon transform

The Radon-transform of an image is called a sinogram. This terms stems from the fact that the Radon transform of a single point is a sine-function.

In practice, the Radon transform is discretized. First off, only a finite number of projections are made. Considering SPECT as an example, the simplest mode of operation is "step-and-shoot" where the heads moves to a location (a single value of $\theta$ ), and remains there for some time. The angles are discretized to the positions of the head (or heads for multi-head scanners).

Even in "continuous" mode, where the heads moves slowly along a prescribed orbit, the data is collected in planes at discrete values of $\theta$. Here, however, each value represents an integral over the range of angles which result in a slight blurring of the images.

Secondly, the image is discretized. As described in sec. 3.1, the images in this thesis will always be $256 \times 256$ pixels.

Finally, the position (i.e. $\rho$ ) is discretized in a finite number of bins. When a scanner detects an event, it determines a bin corresponding to the LOR, and increases it. The number of bins depends on the scanner type and will be discussed below. For synthetic experiments, it is reasonable to sample the sinogram at the same density as the image. In our case, this would mean $\rho$ is sampled at $256 \sqrt{2} \approx 362$ locations.
With these caveats, we can now show the sinograms of the phantoms in fig. 3.4 and 3.5.

## Implementation issues

The first issue that arise concerning an implementation of the Radon transform is to decide how the image should be represented. There are several options, but a simple approach that caries over to more complex models, is to represent the image as a set of point-masses centered at pixels of equal size.

Figure 3.4: The Radontransform of the resolutionphantom in fig. 3.1, sampled every $3^{\circ}$. The sinogram is scaled so the bins sum to 100,000 , which will be used to illustrate noise.

Figure 3.5:
The Radontransform of the Shepp-Logan phantom in fig. 3.2 with bins scaled to a sum of 20,000 . The sinogram appears smoother than fig. 3.4 as the image has fewer sharp edges.


A problem with this approach is that the sampling density of $\rho$ may conflict with the density of the point-masses for certain angles. This problem can be nearly eliminated by interpolating the sinogram using a squared-cosine, i.e. if a point of mass $w$ is projected at sample 3.25 for some angle $\theta$, the mass is distributed to sample 3 and 4 as $w \cos (.25 \pi / 2)^{2}$ and $w \cos (.75 \pi / 2)^{2}$.

Another option is to ensure the image is sampled densely compared to $\rho$.

### 3.4 Basic SPECT model

A collimator is a complex device, and we shall use a simple idealized model. Consider a single point emitting photons and a collimator of lead with a single cylindrical hole. In the ideal case, photons will only pass the collimator if they pass through empty space.

Our goal is to derive the fraction of the photons that pass the collimator as a function of position. Since the geometry is symmetric, we only need to consider the distance from the collimator, $y$, and the distance from the center of the hole, $x$. We use $L$ to denote length of the hole, and $D$ for the diameter. The geometry is shown in fig. 3.6.

Figure 3.6:
A collimator with a single hole with diameter $D$ and length $L$. The figure is the plane of the hole axis and emission point.

## Single-hole collimator



An unshielded point-source radiates with equal intensity in all directions. The intensity of the beam is therefore propertional to the angle allowed to pass the collimator. We shall simplify the analysis by assuming that $y \gg D+L$, i.e. the hole has infinitesimal dimensions. This implies that $\alpha$ is small, and the intensity is negligible. We shall therefore normalize the intensity to 1 for a source that is placed on the hole axis, i.e. $x=0$.

With these conventions in mind, we first note the following function points:

$$
I(x, y)= \begin{cases}1 & \text { for }|x| \leq D / 2 \\ 0 & \text { for }|x| \geq D / 2+y(D / L)\end{cases}
$$

For intermediate points, the desired intensity is the ratio $\alpha / \alpha_{0}$, where $\alpha_{0}$ is the angle for a point at the center:

$$
I(x, y)=\alpha \alpha_{0}^{-1}=\left(\tan ^{-1}\left(\frac{|x|+D / 2}{y+L}\right)-\tan ^{-1}\left(\frac{|x|-D / 2}{y}\right)\right) \alpha_{0}^{-1}
$$

Using the identity $\tan ^{-1}(x)-\tan ^{-1}(y)=\tan ^{-1}((x-y) /(1+x y))$ [86]:

$$
I(x, y)=\tan ^{-1}\left(\left[\frac{|x|+D / 2}{y+L}-\frac{|x|-D / 2}{y}\right] /\left[1+\left(\frac{|x|+D / 2}{y+L}\right)\left(\frac{|x|-D / 2}{y}\right)\right]\right) \alpha_{0}^{-1}
$$

Simplifying and ignoring second order terms:

$$
I(x, y) \approx \tan ^{-1}\left(\frac{D y-L|x|}{x^{2}+y^{2}}\right) \alpha_{0}^{-1}
$$

Since the angle $\alpha$ is infinitesimal, we may further use the relation $\tan ^{-1}(x) \approx x$. The normalizing term, $\alpha_{0}$, is defined by the relation $I(0, y)=1$, giving the pleasing result:

$$
I(x, y)=\frac{1-(L / D)(|x| / y)}{1+(x / y)^{2}}
$$

It is reasonable to simplify this even further. A practical collimator must limit the photons to small angles, which implies that $L / D$ is much larger than one (see e.g. tab. 2.5 on p. 15). Therefore, the term $x / y$ does not change appreciably from 0 in the valid range, i.e. $|x| \leq y(D / L)$.

We may combine this approximation with eq. (3.2), which simplifies under the assumption that $L$ and $D$ are infinitesimal, whereas the ratio $L / D$ is non-negligible:

$$
I(x, y)= \begin{cases}1-(L / D)(|x| / y) & \text { for }|x| \leq y(D / L) \\ 0 & \text { otherwise }\end{cases}
$$

## Multi-hole collimator

A single-hole collimator is not very useful, except for pin-hole collimators which require an entirely different analysis. However, when the response of a single hole is known, it is easy to determine the response of a collimator consisting almost entirely of holes.

Consider fig. 3.7, which shows the reponse of a number of holes. The measured response from a single point consists of the sum of the response of all the holes. The figure shows both the discrete values of the angles and the continuous model given by eq. (3.3). Clearly, the continuous approximation is quite good, even though the infinitesimal requirement is violated to make the drawing legible.

Figure 3.7:
A multi-hole collimator. The continuous model, from eq. (3.3), is close to the ratio of angles, which is shown as discrete points.


We still need to take normalization into account. Real collimators will collect nearly the same number of counts regardless of the distance of the point. Therefore, the area under the curve is nearly constant, and the intensity under the point should be scaled accordingly.

If depth-dependent sensitivity needs to be modeled, it is probably best to measure this dependence directly from the scanner. Note that only the collimator response should be included in this. We shall treat attenuation separately.

## Model flaws

It should be noted that eq. (3.3) was derived from 2D-considerations only. A more detailed analysis involves the solid angle, and scaling becomes a somewhat complicated matter. Nevertheless, the result does not change appreciably when this is taken into account, and the assumptions used to derive the simple model is too far from reality to justify such an analysis.

In reality, collimator holes are not cylindrical as this would waste area. Neither are the holes infinitesimal, photons are reflected, scattered, and may penetrate the walls between the holes. The accuracy of the $\gamma$-camera is another limiting factor. In short, an exhaustive list of significant issues neglected by the model would be very long.

When these factors are combined, the measured point-spread function is much closer to a Gaussian function than the simple triangular function in eq. (3.3). Nevertheless, it captures an essential property of the collimator: beam divergence.

An important consequence of beam-divergence is that the resolution of a point decreases with the distance from the collimator. Therefore it is critical to get the
head of the scanner as close to the object as possible. It also means that the measurements can not adequately be represented by the limiting range $0 \leq \theta<180^{\circ}$, but must be extended to the whole range $0 \leq \theta<360^{\circ}$.

### 3.5 Basic PET model

Figure 3.8:
Some lines connected by crystalpairs. Each line is a point in the sinogram, and the number of detected counts are considered to be the line-integral.

Fortunately, it is somewhat easier to model PET than SPECT, at least in 2D-mode. For simplicity, we shall consider a scanner with only $S=60$ crystals, but the observations caries over to any number of crystals. A real scanner, such as the GE Advance at the PET and Cyclotron Unit at Copenhagen university Hospital, has $S=672$ crystals.

The simplest PET-model assumes that a detected event occured on the line that connects the center of two crystals. These lines are shown in fig. 3.8, and it should be clear that even a few crystals provides a large number of LORs.


Using the convention in fig. 3.8, it is easy to determine the parameter of LOR in normal form. If an event is detected by the crystals $s_{1}$ and $s_{2}$, and the scanner ring has unit radius, then we have:

$$
\begin{align*}
\theta & =\bmod \left(s_{1}+s_{2}+S / 2, S\right) 180^{\circ} / S  \tag{3.4}\\
\rho & = \begin{cases}-\cos \left(180^{\circ}\left(s_{2}-s_{1}\right) / S\right) & \text { for } S / 2 \leq s_{1}+s_{2}<S 3 / 2, \\
+\cos \left(180^{\circ}\left(s_{2}-s_{1}\right) / S\right) & \text { otherwise. }\end{cases} \tag{3.5}
\end{align*}
$$

Here $\bmod (a, b)$ is the modulus function, i.e. $0 \leq \bmod (a, b)=a+k b<|b|$ for some $k \in \mathbb{Z}$.

In practice, the field-of-view (FOV) limits $\rho$ so only LORs near the center are considered. Converting to crystals using eq. (3.5), this limit may be specified by a constant, $B$, so we only consider LORs that satisfy:

$$
\begin{equation*}
S / 2-B \leq \bmod \left(s_{2}-s_{1}, S\right) \leq S / 2+B \tag{3.6}
\end{equation*}
$$

## Sinogram bins

Just as for SPECT, the probability that an event is measured by a crystal pair is proportional to the solid angle subtended by the LORs that intersects the crystal pair and point of emission. Unlike SPECT, disregarding the 3D geometry is justified. If the scanner ring is thin, i.e. the distance from the crystals to the emission

Figure 3.9:
The probability of an ideal scanner detecting an anihilation is proportional to the angle. The probability that crystal 0 detects the event is proportional to $\alpha$, while the pair $p(0,23) \propto \beta_{1}$.

Figure 3.10:
Measured sinogram bins for an idealized PETscanner, where all events are binned by the crystal pair intersected by the LOR. The dots represents the center lines used for a reaslistic FOV.
point is large compared to the width of the ring, the angle in the $z$-axis is nearly constant.

The geometry of a transaxial slice is illustrated in fig. 3.9. The probability that an emission intersects a given pair is proportional to the angle. In the current example, there are three pairs where the emission intersects crystal 0 .

Point emission ( $S=60, B=13$ )


In PET, the measured data is the number of counts detected by each pair of crystals (within the FOV). Thus, each event is put in a bin corresponding to an area of the sinogram. These bins are shown in fig. 3.10, together with the point corresponding to the LOR of the crystal centers.


One thing evident in fig. 3.10 is that the bins do not correspond to equal areas. This is a consequence of the circular geometry, and can not be correctly classified as a flaw. Those who are mathematically inclined may view this as a consequence of eq. (3.5).

This fact is not too important when the FOV is restricted to small values of $\rho$. However, since most processing algorithms are designed for a uniformly sampled sinogram, it is customary to apply a geometric correction to the sinogram, which resamples the sinogram to take this into account.

## Model flaws

A number of considerations have been ignored. As briefly mentioned in sec. 2.7, p. 18, the scintillation crystals in most PET-scanners are not single entities. Instead the ring consists of blocks with slit of different depths carved into it. The part of the block detecting an event is then inferred by interpolating the response of several photo-multiplier tubes on the back of each block.

This system is not perfect, and it is usually necessary to correct for the sensitivities of each block by calibration. Furthermore, many of the limitations mentioned in sec. 2.6 .1 p. 17 are difficult to model accurately. Positron range can be taken into account, but dead-time can only be partially corrected.

### 3.6 Attenuation model

As described in sec. 2.4 p. 12, some photons may be absorbed inside an object without reaching the scanner. The probability if this is an exponential function of the length traversed inside the object and the attenuation constant of the medium.

For brain scans, it is adequate to model a medium with a constant attenuation constant. Other scans that includes, say, airways and lungs, may not have this property. A further simplification in neurology is that transaxial cuts are generally convex (see e.g. p. 10). That is, if two points are inside the object, then any point on the connecting line-segment is also inside the object.

This implies that a line intersects the object boundary at exactly 0,1 , or 2 points. Now assume a LOR intersects the object as distances $d_{1} \leq d_{2}$ from the scanner, and the emission occurs at distance $d_{o}$. Then the length of the LOR that is inside the object is given by:

$$
l= \begin{cases}0 & \text { for } d_{o} \leq d_{1}  \tag{3.7}\\ d_{o}-d_{1} & \text { for } d_{1}<d_{o}<d_{2} \\ d_{2}-d_{1} & \text { otherwise }\end{cases}
$$

The distances $d_{1}$ and $d_{2}$ may be precomputed at discrete points for each value of $\theta$. If the object is a polygon, these distances can be interpolated linearly from these samples. Once $l$ is determined, the attenuation can be applied by applying eq. (2.1) p. 13.

In PET, the position of the object point is not needed. A point inside the object, $d_{1} \leq d_{o} \leq d_{2}$, is detected when both photons reach the scanner which has the probability $p=e^{-\left(d_{o}-d_{1}\right) \mu} e^{-\left(d_{2}-d_{o}\right) \mu}=e^{-\left(d_{2}-d_{1}\right) \mu}$. Points outside the object also obey this law.

### 3.7 General linear models

All the models above has one crucial property: the measured response is a linear function of the number of emissions at each point in the scanner. Each emission has a certain probability of being detected by a particular bin in the sinogram, but the number of detected events is linear in the number of emissions (ignoring non-linear properties such as dead-time).
Mathematically we may express this as follows:

$$
\begin{equation*}
\boldsymbol{\lambda} \triangleq \mathscr{E}\{\mathbf{b}\}=\mathbf{A} \mathbf{x} \tag{3.8}
\end{equation*}
$$

hypercube

Here $\mathscr{E}\{\cdot\}$ is the expectation operator, so $\boldsymbol{\lambda}$ is defined as a vector that contains the expected number of counts in each bin of the measured sinogram, $\mathbf{b}$. The number of events emitted from each point of the scanned image is represented by the vector $\mathbf{x}$, and they are related by the system matrix $\mathbf{A}$.

Generally, each element in $\mathbf{x}$ represents a distributed area. In that case, each emission is assumed to occur from a point with a prescribed distribution. This distribution can also be modeled by $\mathbf{A}$.

All the discrete models in this chapter can be expressed as a system matrix. However, we stress that the vector x represents an image in terms of a set of basis functions. This notion may be unfamiliar to programmers, who are accustomed to representing an image as a 2D-array (and a volume as a 3D array).

Thus, even if the continuous basis of the image is chosen to represent pixels (or voxels in 3D), it does not change the fact that $\mathbf{x}$ is a vector. For simplicity, all vectors are column-vectors in this thesis. Thus, even if $\mathbf{x}$ represents a rectangular image, it does not represent a matrix. Please refer to chapter 10 for a full discussion of this point.

To distinguish between vectors representing rectangular data and matrices, we shall use the term hyper-cube for a vector that represents an array. This has the advantage of being unambigous (at least for the purposes of this thesis), and is not easily mistaken. It also conveys the important notion that the number of elements of the underlying data-array is irrelevant. In tomography, for instance, the FOV is typically defined by a circle or ellipse, and restrincting an using a rectangular array to represent an image is inappropriate.

With this caveat out of the way, we shall nevertheless refer to elements of x as pixels. While the underlying basis is mostly unimportant in this thesis, it remains helpful to use familiar terms. We refer to elements of $\mathbf{b}$ and $\boldsymbol{\lambda}$ as bins for the same reason.

### 3.8 Noise model

The linear model in eq. (3.8) only describes the average measurement $\boldsymbol{\lambda}$. However, only a finite number of counts are measured, and this needs to be taken into account.

Generally, a very large number of radioactive nuclei are injected into the person being scanned. These are distributed spatially by a probability distribution; determining this distribution is the true objective of the scanning.

The probability that a nucleus emits a photon in a small volume is very small. Therefore, the number of photons emitted in such a volume is a stochastic variable, $X$, governed by the Poisson distribution discovered by Simeon Denis Poisson [60]:

$$
P\{X=k\}=e^{-\lambda} \lambda^{k} / k!, \quad k \in \mathbb{N}
$$

Here the parameter, $\lambda$, turns out to be both the mean and variance of $X$.
If each nucleus are considered to be distributed independently (i.e. a tracer dose of the substance is administered), then the number of photons emitted from the volumes are independent.

Figure 3.11:
The sinogram in fig. 3.4 sampled from a Poisson distribution with approximately 100,000 counts.

Each photon emitted from the volume has a fixed probability distribution of being counted by the bins of the sinogram. This distribution is exactly the models derived in this chapter. Thus, if we include an extra bin for undetected photons, the distribution of the bins is governed by a multinomial distribution.

It turns out that if the bins, $b_{1}, b_{2}, \ldots$, are multinomial distributed with a parameter drawn from a Poisson distribution, then each of the bins are independent Poissondistributed variables [45].

Thus, the number of counts detected by a bin, $b_{i}$, is a sum of independent Poissondistributed variables. This, in turn, is merely another Poisson-distributed variable [60]:

$$
\begin{equation*}
P\left\{b_{i}=k\right\}=e^{-\lambda_{i}} \lambda_{i}^{k} / k!, \quad k \in \mathbb{N} \tag{3.9}
\end{equation*}
$$

Note that the derivation showed that the measured counts are independent. This is a very powerful result, which will be used extensively.

In terms of linear models, we may express this as an error-term:

$$
\begin{equation*}
\mathbf{b}=\mathbf{A} \mathbf{x}+\mathbf{e}=\boldsymbol{\lambda}+\mathbf{e} \tag{3.10}
\end{equation*}
$$

## Sinogram examples

We are now ready to show what the sinograms of the phantoms look like when noise is considered. For variety, the phantoms in fig. 3.4 and fig. 3.5 were scaled to incoorporate about 100, 000 (fig. 3.11) and 20, 000 counts (fig. 3.12) respectively. Clearly, there is a substantial amount of information lost in the process.


## Model flaws

The Poisson-distribution is only valid when no amount of restoration has been performed on the sinogram. Thus, geometric correction, callibration of photomultiplers, linearization of dead-time, and other factors result in correlated measurements.

In principle, the Poisson-statistic could be maintained for linear operations by randomized sampling. However, the idea of applying even more noise to the measurement is not appealing.

Figure 3.12:
The sinogram of the Shepplogan phantom in fig. 3.5 sampled with approximately 20, 000 counts.


At a more fundamental level, a PET-scan corrected for randoms using on-line subtraction can result in negative counts as described in sec. 2.6.1. This breaks fundamentally with the Poisson-distribution.

To model this more accurately, a shifted Poisson-distribution can be used [85]. This use a Poisson-distribution of the combined variance, but shifted to match the average.

### 3.9 Software

001010
100001 111100

All phantoms in this thesis were created by bbphantom found in BBTools. The system-matrix corresponding to a discrete Radon-transform, using the algorithm in sec. 3.3.1, can be created as a black-box operator by bbradon (see chapter 10 for information about black-box operators).

A more elaborate model is available as bbtomoemit, which includes beamdivergence (sec. 3.4), and uniform attenuation of a convex body (sec. 3.6). Geometric correction is not currently implemented, but may be included by a interpolation.

The Matlab ${ }^{\circledR}$ image processing toolbox from Mathworks includes phantom, which can create a Shepp-Logan phantom. Unlike bbphantom, however, the average value of the pixels depends on the image size. The function radon computes the Radon transform of an image, but the corresponding back-projection is not included.

More elaborate system-matrices can be generated by the package ASPIRE from Fessler [23]. He also distributes an image reconstruction toolbox for MATLAB ${ }^{\circledR}$ that contains simple implementations of the Radon transform.

Toft showed in [72] how to implement a high-speed Radon-transform using slantstacking. He also distributes implementations of the algorithms along with his thesis.

A different approach is to use Monte-Carlo simulations, although these are mostly useful for validation purposes. One of the most popular packages is SimSET (Simulation System for Emission Tomography) [75], but a number of packages are freely available.

### 3.10 References

The Radon transform is treated analytically in a tomographic setting by a number of authors. One of the most complete is the thesis by Peter Aundal Toft [72], which derives the Radon transform of a number of analytical functions.

Both the SPECT- and PET-models were derived from first principles, and more elaborate models is tantamount to instrumentation models. We refer to sec. 2.9 p. 19 for references.

Linear models and the Poisson statistic are used throughout the reconstruction litterature. In particular we highlight [45], as it is both clear and includes a list of properties of the Poisson distribution.

Fessler argued that samples of a Gaussian distribution could be a better noisemodel than the Poisson-distribution on PET-systems with on-line subtraction [20]. A more comprehensive analysis of non-Poisson distributed counts and corresponding likelihood-functions is the thesis of Yavuz [84], which advocates the use of the shifted Poisson-distribution.
$\oplus$
$\oplus$

## Classical reconstruction

This chapter goes through the classical reconstruction algorithms used in emission tomography. We start with a description of back-projection, which is a basic building-block also used in other part of the thesis.

This is followed by a detailed description of the two methods implemented on most scanners: FBP (Filtered Back-Projection) and EM (Expectation Maximization). We give examples of images reconstructed with each of these methods.

For reference, we also mention a few other methods found in the litterature.

### 4.1 Methods and abbreviations

As usual, there are a plethora of abbreviations found in the litterature. The abbreviations most used in the literature are shown in table 4.1. The first part are common, whereas the lower part as encountered more rarely.

Table 4.1:
Terms used in classical reconstruction methods.

| Abbrevation | Meaning |
| :--- | :--- |
| FBP | Filtered Back Projection |
| ML | Maximum Likelihood |
| EM | Expectation Maximization |
| OSEM | Ordered Subset EM |
| ART | Algebraic Reconstruction Technique |
| WLS | Weighted least-squares |
| PWLS | Penalized weighted least-squares |

### 4.2 Back-projection

The back-projection is used in some form in most reconstruction algorithms. Although it is rarely used by itself, we shall study it in some detail before proceeding to more sophisticated algorithms.

Back-projection is a one-liner: the counts detected in each bin are distributed evenly on the corresponding LOR. If we consider a point-source, all LORs intersect the point and the total length of LORs in a small area decreases with the distance from the point.

## $\iiint$ <br> Intensity function

Consider an annulus around the emission-point as shown in fig. 4.1. If the radii of the annulus are given by $r_{1}<r_{2}$, then the area of the annulus is $A=\pi\left(r_{2}^{2}-r_{1}^{2}\right)$. If the sinogram contains $P$ samples of $\theta$ (i.e. $P$ LORs intersects the point of emission), then the total length of LORs in the annulus is $2 P\left(r_{2}-r_{1}\right)$.

This means that the intensity of the annulus is given by:

$$
\begin{equation*}
\frac{2 P\left(r_{2}-r_{1}\right) c}{\pi\left(r_{2}^{2}-r_{1}^{2}\right)}=\frac{2 P c}{\pi\left(r_{1}+r_{2}\right)} \rightarrow \frac{P c}{\pi r_{1}} \text { for } r_{2} \rightarrow r_{1} \tag{4.1}
\end{equation*}
$$

Figure 4.1:
An annulus around a point-source and the intersecting LORs

## Back-projection



Here $c$ is a scale-factor that determines the number of counts deposited on a LOR of unit length

For completeness, the analogously result for a spherical shell for a 3D backprojection is:

$$
\frac{3 P c}{2 \pi\left(r_{1}^{2}-r_{1} r_{2}+r_{2}^{2}\right)} \rightarrow \frac{3 P c}{2 \pi r_{1}^{2}} \text { for } r_{2} \rightarrow r_{1} .
$$

If the LORs are distributed evenly (i.e. $\theta$ is sampled equidistantly), then any area (or volume) inside the annulus (or shell) will have approximately the same intensity. Therefore, the reconstructed point will result in a symmetric point-spread function in this case.

## $\iiint$

## Discrete back-projection

For a discrete system, the forward projection of bin $i$ is given by the linear equation $\boldsymbol{\lambda}=\mathbf{A x}$ :

$$
\begin{equation*}
\boldsymbol{\lambda}_{i}=\sum_{j} \mathbf{A}_{(i, j)} \mathbf{x}_{j} \tag{4.3}
\end{equation*}
$$

To back-project $\boldsymbol{\lambda}_{i}$ or $\mathbf{b}_{i}$, we must distribute the counts to $\mathbf{x}$ in proportion to the amount ("LOR-length" for a Radon transform) each element contributed to the measurement. Specifically, the back-projection is:

$$
\begin{equation*}
\hat{\mathbf{x}}_{j}=\sum_{i} \mathbf{A}_{(i, j)} \mathbf{b}_{i} \tag{4.4}
\end{equation*}
$$

More simply, this can be written as $\hat{\mathbf{x}}=\mathbf{A}^{\mathrm{T}} \mathbf{b}$.

## Implementation

Scaling can not be ignored in a real implementation. The simplest way to deal with scaling is to distribute each bin evenly along the LOR which is inside the field of view. Note, however, that a point-source will no longer be reconstructed as a symmetric function.

To this end, consider eq. (4.4). If we distribute the counts in $\mathbf{b}_{i}$ according to the weights given by $\mathbf{A}_{(i,:)}$, then we need to scale by the sum of the weights:

$$
\begin{equation*}
\hat{\mathbf{x}}_{j}=\sum_{i} \mathbf{b}_{i} \mathbf{A}_{(i, j)} / \sum_{k} \mathbf{A}_{(i, k)} \tag{4.5}
\end{equation*}
$$

This assumes that $\mathbf{A}$ is non-negative, i.e. there are no negative elements of $\mathbf{A}$. All models described in chapter 3 satisfy this requirement.

With this scaling, the number of counts in $\hat{\mathbf{x}}$ matches the number of counts in b. However, we prefer images which corresponds to image originally put in the scanner. The Radon-transform and, by extension, all other models in this thesis, project a point of unit intensity to one count for each projection. Thus, if we sample $\theta$ at $P$ discrete positions, the reconstructed point will contain $P$ counts. Thus, if we divide by $P$, we get a scale corresponding to the original image.

This leads directly to alg. 4.1. Since this is the first algorithm in this thesis, a few comments are in place:

- The result is an operator which performs the back-projection. To backproject a specific sinogram, $\mathbf{b}$, the matrix-vector product $\hat{\mathbf{x}}=\mathbf{B b}$ must be computed, i.e. $x h=B * b$.
- The algorithm use BBTools described in chapter 10. The function bbdiag is a black-box implementation of the MATLAB-function diag and functions in much the same way. We will assume that functions in BBTools similar to MATLAB ${ }^{\circledR}$ can be used without explanation.
- Generally, the algorithms in the thesis are meant to be illustrative. However, since the system-matrix A can be a general black-box implemented as compiled mex-files, the operator produced by alg. 4.1 is competitive with any other algorithm based on compiled code.
Those who are theoretically inclined may find that the syntax $x h=B * b$ is merely cosmetic.


### 4.3 Examples of back-projection

With this implementation in mind, we may show the result of the back-projection algorithm. For convenience, the original phantom, fig. 3.1, is reproduced in fig. 4.2.

It should be evident that pure back-projection is inadequate as a reconstruction algorithm. It is difficult to discern even the largest hot-spots in the phantom and edges are not preserved well.

## Algorithm 4.1: Back-projection

Create a back-projection operator given a non-negative system-matrix, A, and the number of projections.

```
function \(\mathbf{B}=\) backproject \((\mathbf{A}, P)\)
    \(\mathbf{p}=\mathbf{A} * \operatorname{ones}(\operatorname{size}(\mathbf{A}, 2), 1) ; \quad\) Project to determine weights
    \(\mathbf{c}=\operatorname{double}(\mathbf{p} \neq 0)\); Zero out null entries
    \(\mathbf{p}(\mathbf{p} \equiv 0)=1 ; \quad\) Avoid division-by-zero
    \(\mathbf{B}=\mathbf{A}^{\mathrm{H}} * \operatorname{bbdiag}(\mathbf{c} . /(P * \mathbf{p})) ; \quad\) Construct operator
```

Figure 4.2:
A reproduction
of the resolutionphantom in fig. 3.1.


Back-projection


### 4.4 Filtered back-projection

The Filtered Back Projection algorithm (FBP) the dominant algorithm used in medical reconstruction. The images produces are linear in the inputs, and the algorithm is sometimes known simply as "linear reconstruction". As the name suggest, it is based on back-projection, but use a filter to compensate for the weakness of pure back-projection.

Again, the underlying idea is simple: solve the linear system $\mathbf{A x}=\boldsymbol{\lambda}$ for $\mathbf{x}$, where $\mathbf{A}$ is the discrete Radon-transform. Unlike back-projection, however, finding an efficient way to implement the scheme is not intuitive.

## $\iiint \quad$ Derivation of FBP

Filtered back-projection is usually derived in the continuous domain using the Fourier Slice Theorem (also known as the Central Slice Theorem). The details are messy, and require careful attention at singularities. Since most of this thesis is based on linear algebra, we shall give an informal derivation that lends itself to an implementation in this setting.

The starting-point for FBP is a direct solution of the linear system $\boldsymbol{\lambda}=\mathbf{A x}$ which, in the presence of noise, becomes $\mathbf{A} \hat{\mathbf{x}} \approx \mathbf{b}$. As mentioned previously, the solution

Figure 4.4:
Back-projection followed by a forward projection, representing the operator $\mathbf{A A}^{\mathrm{T}}$.
is given by a filter followed by a back-projection:

$$
\hat{\mathbf{x}}=\mathbf{A}^{\mathrm{T}} \mathbf{F b}
$$

We may now plug $\hat{x}$ into the linear equation:

$$
\mathbf{b} \approx \mathbf{A} \hat{\mathbf{x}}=\mathbf{A} \mathbf{A}^{\mathrm{T}} \mathbf{F} \mathbf{b}
$$

Thus, a good filter approximates the inverse (or pseudo-inverse in the general case) of the operator $\mathbf{A} \mathbf{A}^{\mathrm{T}}$. In other words, if a sinogram with a single non-zero bin is filtered and $\mathbf{A} \mathbf{A}^{\mathrm{T}}$ is applied, then the new sinogram should be approximately the same as the starting point.

Consider $\mathbf{A A}^{\mathrm{T}}$, which represents a back-projection followed by a forward projection. Assume the system-matrix $\mathbf{A}$ is the Radon-transform, and let us momentarily assume that only $\theta$ is discrete.

The situation for a circular FOV is sketched in fig. 4.4. A bin is first distributed on the LOR, which is in turn forward projected to the same point. On the other hand, when the LOR is projected on a plane with a different angle, the result is a rectangular function.


Since the projection may be placed anywhere, we need a position-invariant filter that blocks rectangular functions. Such a filter is most easily designed in the frequency domain. If $r(x)$ is a rectangular function of width $2 a$, symmetric around $x=0$, then the Fourier-transform, $R(\omega)=\mathscr{F}\{r(x)\}=\int_{-\infty}^{\infty} r(x) e^{-j \omega x} d x$, is given by [86]:

$$
R(\omega)= \begin{cases}2 a & \text { for } \omega=0 \\ 2 \sin (a \omega) / \omega & \text { otherwise }\end{cases}
$$

This can not be completely suppressed, but it can be rendered mostly harmless by damping low frequencies. In fact it turns out that it suffice to neutralize the denominator $\omega$ by applying a ramp-filter:

$$
F(\omega)= \begin{cases}|\omega| & \text { for }|\omega|<\omega_{c} \\ |\omega| / 2 & \text { for }|\omega|=\omega_{c} \\ 0 & \text { otherwise }\end{cases}
$$

The cut-off frequency, $\omega_{c}$, is a necessary for the inverse Fourier-transform of $F(\omega)$ to exist [44].

The inverse Fourier-transform of $R(\omega) F(\omega)$ has sharp peaks near $x= \pm a$, but the integral of an area around a peak is small and diminishes as $\omega_{c}$ becomes large.

In the discrete case these peaks may be visible in the form of "streak" or "star" artifacts as we shall show in the examples.

The case $\omega=0$ is difficult to analyze properly. The ordinary analysis requires Stieltjes-intergrals, a generalization of Riemann integrals. It also requires that back-projection occurs on an infinitly large image, so $a$ can become infinite in a limit. The result is that the ramp-filter also suffices in this case under mild restrictions.

Curiously, this result remains approximately valid, to within a factor of 2 , in the discrete case [78]. This is surprising, since the filter removes the average, i.e. the total number of counts in the filtered sinogram is (nearly) zero.

001010
100001 111100

## Implementation

While it is difficult to get the theoretic details of FBP right, it is relatively easy to implement once the filter has been established. A basic implementation is given in alg. 4.2, which constructs an operator. It is necessary to scale by the number of projections, for the same reason as the back-projection case. The factor of 2 was mentioned in sec. 4.4.1.

Again a few comments are in place:

- The filter is designed in the frequency domain and inverted to the spacial domain before being applied. It is common to find implementations which skips this step, which leads to unnecessary aliasing.
- The filter is applied to each column of the sinogram, which is a hypercube with samples of $\rho$ and $\theta$. In linear algebra this is known as a Kronecker product, usually written as $\otimes$. The operator that applies $\mathbf{A}_{i}$ along all data along dimension $i$ on an $N$-dimensional hyper-cube is given by $\mathbf{A}=\mathbf{A}_{N} \otimes \cdots \otimes \mathbf{A}_{1}$.
Since no filtering takes place between angles, we put in an identity operator on dimension 2.


## Algorithm 4.2: Filtered back-projection

Basic implementation of filtered back-projection. The input is a Radon-transform, A (created by bbradon), and the sinogram is an $S \times P$ hyper-cube with $S$ equidistant samples of $\rho$ and $P$ equidistant samples of $0 \leq \theta<180^{\circ}$. The cut-off frequency is normalized to $0 \leq \omega_{c} \leq 1$.

```
function \(\mathbf{B}=f b p\left(\mathbf{A}, S, P, \omega_{c}\right)\)
    \(L=2^{\wedge}\) nextpow2 \((\max (3 * S, 128))\);
    \(\boldsymbol{\omega}=(0: L)^{\mathrm{T}} / L\);
    \(\mathbf{f}=|\pi * \boldsymbol{\omega}| \cdot *\left(\boldsymbol{\omega} \leq \omega_{c}\right) ; \quad \quad\) Use radians and apply ramp with cut-off
    \(\mathbf{p}=\operatorname{real}(i f f t([\mathbf{f} ; \mathbf{f}(\mathbf{e n d}-1:-1: 2)])) ; \quad\) Compute impulse-response
    \(\mathbf{p}=[\mathbf{p}(\) end \(+2-S:\) end \() ; \mathbf{p}(1: S)] ; \quad\) Extract the useful part of the impulse
    \(\mathbf{p}=\mathbf{p} /(2 * P)\);
    \(\mathbf{C}=\operatorname{bbconvn}(S, \mathbf{p}\), ' same'); Create convolution operator
    \(\mathbf{F}=\operatorname{kron}(\operatorname{bbeye}(P), \mathbf{C}) ; \quad \quad\) Apply convolution on every value of \(\theta\)
    \(\mathbf{B}=\mathbf{A}^{\mathrm{T}} * \mathbf{F} ; \quad\) Create FBP-operator
```

- The operator produced by alg. 4.2 is an ideal pre-conditioner for an iterative algorithm that solves a system using a more complex model than the Radon transform.

For example, assume that A is a model that includes beam-divergence and attenuation, B is the FBP-operator constructed from the Radon transform, and $b$ is the measured sinogram. Then a standard iterative solver, e.g. $x h=1$ sqr $(B * A, B * b)$, converges in a few steps.

### 4.5 Examples of FBP

The basic operator created by alg. 4.2 was used to reconstruct the resolutionphantom using the noisy singram with 20,000 counts. The result is shown in fig. 4.5, which is clearly a useless approximation of the original phantom in fig. 4.2.

Figure 4.5:
The resolutionphantom reconstructed using FBP with cut-off $\omega_{c}=1$.

Figure 4.6:
A new reconstruction, but low-pass filtered with a cutoff $\omega_{c}=0.175$.

To get a more useful image, we need to filter out the noise by lowering the cut-off frequency. Eactly what value to use depends on the circumstances and intended use of the image, but as we only need to illustrate the technique, we will use the arbitrary value of $\omega_{c}=0.175$. The result is shown in fig. 4.6.


Note that smoothing comes with a loss of resolution. The ultimately smooth image (cut-off $\omega_{c}=0$ ) is blank and contains no information.

An artifact often seen in reconstructed images is "ringing". In this context is usually manifest itself as streaks radiating from the center and outwards. This is sometimes called "star artifacts" [19].

The term "ringing" comes from electronics and is meant in the sense of a bellringing. Here, it comes from the abrubt change at $\omega_{c}$ in eq. (4.9), which is converted to a highly oscillating function. As touched briefly in sec. 4.4.1, the average is negligible, but the peaks will always be high. This is a well-known result known as Gibbs phenomenon [54].

To reduce this kind of artifacts, it is common to modify the filter by multiplying by a window-function. It may be confusing, because there are plenty to choose from in the litterature. In practical terms this is not too important, and the author ${ }^{1}$ is inclined to agree with the following quote [58, p.534]:

There is a lot of unnecessary lore about choice of a window function, and practically every function that rises from zero to a peak and then falls again has been named after someone.

However, at the level of this book, there is effectively no difference between any of these (or similar) window functions.

We shall stick to the so-called Hamming window below. One thing to notice about windows is that it is often possible to increase the cut-off frequency. Fig. 4.7 below was reconstructed with $\omega_{c}=0.3$ to yield approximately the same resolution as fig. 4.6.

Figure 4.7:
Another reconstruction, low-pass filtered with cut-off $\omega_{c}=0.3$ but including a Hamming window to reduce star-artifacts.


### 4.6 EM reconstruction

After back-projection and, possibly, attempts at direct inversion of the systemoperator, the EM-algorithm is likely to be one of the first algorithms a novice would think of. To back up this claim, the author knows of 4 persons who came up with the idea independently ${ }^{2}$.

[^4]Let us rehash the idea of back-projection: the counts measured in a bin is distributed evenly on the LOR from which they could have come. Now suppose we already have a rough idea about what the image looks like. Instead if distributing the counts evenly, we may distribute them on the LOR in proportion to the concentrations in the approximate image.

This is the basic step of the EM-algorithm, and it produces an estimate which should be better than the image we started with. Thus, we may improve the new image by applying a new step. Thus we have an iterative procedure that, hopefully, converges to a fix-point.

The EM-algorithm is sometimes known as "iterative reconstruction" (as opposed to "linear reconstruction"), because it is the second major algorithm implemented on many commercial scanners.

## Implementation

As explained, the counts in bin $\mathbf{b}_{i}$ should be distributed in the LOR in the proportion given by the pixels in the previous estimate, $\hat{\mathbf{x}}_{j}$, in the LOR. Since the length of the LORs are represented by a system-matrix, the counts must also be weighted with the elements of $\mathbf{A}$. We may therefore express the EM-step by an equation similar to eq. (4.5):

$$
\hat{\mathbf{x}}_{j}^{\prime}=\sum_{i} \mathbf{b}_{i} \mathbf{A}_{(i, j)} \hat{\mathbf{x}}_{j} / \sum_{k} \mathbf{A}_{(i, k)} \hat{\mathbf{x}}_{k}
$$

Like the back-projection algorithm, this produces an image with a total number of counts that matches the number of acquired counts, and we must divide by the number of projections to get an image comparable with the original.

There is, however, a more general way to handle this. Instead of dividing by the number of projections, we can divide each pixel by the amplification factor:

$$
\begin{aligned}
\hat{\mathbf{x}}_{j}^{\prime} & =\frac{1}{\sum_{i} \mathbf{A}_{(i, j)}} \sum_{i} \mathbf{b}_{i} \mathbf{A}_{(i, j)} \hat{\mathbf{x}}_{j} / \sum_{i} \mathbf{A}_{(i, j)} \hat{\mathbf{x}}_{j} \\
& =\left(\frac{\hat{\mathbf{x}}_{j}}{\sum_{i} \mathbf{A}_{(i, j)}}\right) \sum_{i} \mathbf{A}_{(i, j)}\left(\frac{\mathbf{b}_{i}}{\sum_{k} \mathbf{A}_{(i, k)} \hat{\mathbf{x}}_{k}}\right)
\end{aligned}
$$

This is the final iteration of the EM-algorithm, and the bulk of the work can be performed efficiently as matrix-vector products as shown in alg. 4.3.

## Ordered subset

The EM-algorithm converges rather slowly, and since it is used in practical work it is desirable to speed it up. The ordered subset algorithm, invented by Harold Malcolm Hudson and Richard S. Larkin [37], is one way implemented on many scanners, often under the name OSEM.

The idea is to update only a subset of the projections at each iteration. We shall not discuss this algorithm further, as it is merely intended to be a faster version of the EM-algorithm.

## $\iiint \quad$ Justification of EM

The idea described above is simple to understand, but there was little to motivate it and even a cursory analysis reveals problems. Suppose we draw the LORs corresponding to all bins with non-zero counts. If we create an iterate by placing a single point on each LOR at random, except where it intersects another LOR, then it converges to a fix-point in one step.

While this result appears alarming, we note that it may be impossible to construct such a solution in a discretized image. It does, however, suggest that algorithms that are not backed by solid theory should be regarded with suspicion.

Happily, the simple iteration described above happens to coincide with an Expectation Maximization (EM) algorithm, which computes a maximum-likelihood solution for a likelihood defined by the Poisson-distribution.

It is beyond the scope of this text to delve in the theoretical foundation of this algorithm and refer to the references in sec. 4.9 for details. We will, however, state the basic result about convergence.
Theorem 4.1: (EM-Convergence) Let $\mathbf{b}$ be a sample of random vectors with independent Poisson-distributed elements with parameters $\boldsymbol{\lambda}=\mathbf{A x}$. Furthermore, let $\hat{\mathbf{x}}_{0}$ be an image with entirely positive pixels.

Assume that $\mathbf{A}$ is non-negative, and that no element in $\mathbf{A}^{\mathrm{T}} \mathbf{b}$ is zero, i.e. all pixels in the image are intersected by a LOR corresponding to a bin in $\mathbf{b}$ with non-zero counts. Finally, assume that null $\{\mathbf{A}\}=\{0\}$.

Then the iterates of the EM-algorithm, $\hat{\mathbf{x}}_{0}, \hat{\mathbf{x}}_{1}, \ldots$, converges to the maximumlikelihood solution:

$$
\lim _{i \rightarrow \infty} \hat{\mathbf{x}}_{i}=\underset{\mathbf{x}}{\operatorname{argmax}} \operatorname{Prob}(\mathbf{b} \mid \boldsymbol{\lambda}=\mathbf{A x})
$$

## Proof

## See [45]

The most revealing about this theorem is the conditions required to make it true. First, the number of pixels in the image must necessarily be smaller than the number of sinogram bins. Secondly, sampling the sinogram denser to accomodate

## Algorithm 4.3: EM-Reconstruction

Reconstruct an image using the EM-algorithm, given a non-negative systemmatrix, A, a measured sinogram of independent Poisson-distributed counts, b, and a limit on the number of iterations.

```
function \(\hat{\mathbf{x}}=\operatorname{em}\left(\mathbf{A}, \mathbf{b}, i_{\max }\right)\)
    \(\mathbf{w}=\mathbf{A}^{\mathrm{T}} * \operatorname{ones}(\operatorname{size}(\mathbf{A}, 1), 1) ;\)
    \(\mathbf{w}(\mathbf{w} \equiv 0)=1 ;\)
    \(\hat{\mathbf{x}}=\operatorname{ones}(\operatorname{size}(\mathbf{A}, 2), 1)\);
    for \(i=1: i_{\max }\)
        \(\hat{\mathbf{x}}_{\mathbf{p}}=\hat{\mathbf{x}}\); Save iterate for convergence test
        \(\mathbf{p}=\mathbf{A} * \hat{\mathbf{x}} ; \quad\) Forward-project current estimate
        \(\mathbf{p}\left(\mathbf{p}<\epsilon_{M}\right)=\epsilon_{M} ; \quad\) Avoid division by zero
        \(\hat{\mathbf{x}}=(\hat{\mathbf{x}} . / \mathbf{w}) \cdot *\left(\mathbf{A}^{\mathrm{T}} *(\mathbf{b} \cdot / \mathbf{p})\right) ; \quad \quad\) Evaluate eq. (4.11)
        f \(\left\|\hat{\mathbf{x}}_{\mathbf{p}}-\hat{\mathbf{x}}\right\|_{2}<\|\hat{\mathbf{x}}\|_{2} * 1 e-3, \quad\) break Stop if we reached a fix-point
```

more pixels in the image is only helpful if the number of acquired counts is large enough to make most of them non-zero.

In other words, the EM-algorithm relies on pixelation of the image to function. The EM-algorithm must be able to "see" the pixels, i.e. the images produced by the method are not an approximation of an underlying continuous function. This is a distinct disadvantage.

### 4.7 Examples of EM

If the EM-algorithm is used naively it often produces useless results. Fig. 4.8 shows the resolution-phantom reconstructed using the EM-algorithm. Although better than an unfiltered version of a FBP-reconstruction, it is not useful when used in this form.

## EM Reconstruction



One way to produce a useful image, is simply to apply a post-filter on the reconstructed image. The image in fig. 4.9 is the same as fig. 4.8 after a Gaussian filter with a FWHM of about 3.52 was applied. It is clear that such a filter is necessary.

Figure 4.9:
Fig. 4.8 postfiltered with a Gaussian kernel with $F W H M \approx$ 3.52 pixels.


Another popular way to avoid excessive noise, is early stopping. This simply put a limit ( $i_{\max }$ in alg. 4.3) on the number of iterations taken and prevents the EMalgorithm from converging to the maximum-likelihood solution. The result after 8

Figure 4.10:
EM-reconstruction using early stopping.

Figure 4.11:
EM-reconstruction using a Gaussian convolutive sieve with $F W H M \approx 1.2$ pixels.
iterations is shown in fig. 4.10. Early stopping is another way to speed up the EM-

algorithm, since the number of iterations are limited to a fixed number. However, each iteration still require approximately twice the work of the FBP-algorithm, and can not be recommended in general.

A better way to incorporate a filter is to use a sieve [65]. In the elementary form, each iteration in the EM-algorithm is followed by a Gaussian filter. This filter should have a substantially smaller kernel than a post-filter, and significantly reduces the number of iterations required for convergence.

The resolution-phantom reconstructed using a sieve is shown in fig. 4.11. Given the speed of this algorithm, together with the solid underlying theory, this should method can generally be recommended.


### 4.8 Other algorithms

Historically, the first algorithm was Algebraic Reconstruction Technique (ART), which is a direct solution of the linear system using a Gauss-Seidel iteration. There are several similar algorithms, e.g. SIRT (simultaneous iterative reconstruction technique), SART (simultaneous algebraic reconstruction technique), and Kaczmarz method. All of these are obsolete, and should not be used any more.

An alternative to filtered back-projection is "filtering after back-projection". The method applies a filter after back-projection to compensate for the spread given by eq. 4.1. Unlike FBP, the problem with singularities is easily circumvented in the discrete case by integrating over a pixel or putting the singularity between pixel-centers.

### 4.9 References

The original paper by Radon [59] is in the domain of people who cope with Stieltjes-, Cauchy-, and ordinary Riemann-integrals. A fair knowledge of German is recommended, but it is a worthy read for anyone demanding mathematical rigor. Many equations found in the litterature can be found here together with the assumptions necessary for them to be valid; a point often glossed over elsewhere.

Traditional derivations of filtered back-projection, based on the Fourier slice theorem, can be found in numerous places. A terse derivation is given in [5], while a more verbose version can be found in [72]. While such derivations may appear to motivate the ramp-filter, they do, in the oppinion of this author, fall short in rigour. If rigor is desirable, the paper by Radon should be consulted.
The treatment of FBP given here is new, but the intepretation as an inverse of the operator $\mathbf{A A}^{\mathrm{T}}$ was noted by Kawata and Nalciogli [40]. However, they made no attempt at analyzing the operator other than stating that the inverse "corresponds to the two-dimensional rho-filter".

EM-algorithms are effective at finding maximum-likelihood solutions for a broad class of problems, and several excellent texts, such as [8], are available. The original papers in the context of reconstruction is by Shepp and Vardi [64], and Lange and Carson [45]. Of these, the latter should be prefered as it includes scaling and the convergence-proof in theorem. 4.1.
The method of sieves was first used for tomographic reconstruction by Snyder and Miller [65]. The paper includes a clear explanation of the problem with the basic EM-algorithm and can be recommended for this reason alone. It also argues that pixelation, in isolation, does not help the EM and establish that Gaussian sieves are backed by solid theory.

Eugene Veklerov and Jorge Llacer took the sieve method a step further in [76], where they gave a rigorous method for choosing the sieve in the EM-algorithm.

A classical reference for the early methods described in sec. 4.8 is [39], while "filtering after back-projection" is described in a continuous setting in [72].


## Part II

## Improved Reconstruction

$\oplus$
$\oplus$

## Chapter

## Reconstruction goals and strategies

This chapter contains a discussion about the practical requirements imposed on a reconstruction algorithm. The majority of this chapter is collected from personal experience, and is therefore less rigorous than other parts of the thesis. Many of the opinions expressed in this chapter are personal opinions of the author, and should not be construed as being universal.

When other requirements than accuracy are considered, linear methods have many desirable properties that are important from a user-point of view. We discuss previous research in this context, and go forward with a description of regularized reconstruction.

### 5.1 User requirements

One truth holds in engineering as well as other aspects of life: if you can not formulate what you want, then the chance of getting it is remote. The usually way for engineers to express "what they want" is through a cost-function. When a cost-function has been decided upon, a solution can be quantitatively evaluated and therefore optimized.

It follows that it is important to consider the properties of the cost-function. Usually it is necessary to balance several conflicting goals, and this is especially true when it comes to reconstruction. Since the work in this thesis aims to provide useful solutions, the demands of the users should be considered.

Many issues are important from a clinical point of view. Although some requirements may be relaxed in research studies, it still remains that methods with little clinical prospect are unlikely to be used much. The following list is a compilation of issues drawn from the author's experience from working with potential users.

1. Low variability: If a person is scanned twice the results should be alike. Scans between persons must also look the same, as this is the basis that clinicians use for diagnosing illness by inspecting the images.
2. Minimal clinical work: The system must be useful without demanding much extra clinical work. For instance it is not reasonable to assume generally that several MR-images are available.
3. Minimal manual processing: The level of required (i.e. mandatory) manual work on the collected data must be low.
4. Consistent with old methods: If at all possible, the reconstruction should allow a gradual introduction of new features. This allows clinicians to start with something familiar and introduce more advanced methods as they expand their understanding of the system.
5. Homogeneous: It greatly facilitates interpretation of the images if they are homogeneous, i.e. the point spread function (PSF) should be the same over all parts of the image. For specificity, we shall use the term "homogeneous" if the PSFs are spatial invariant and isotropic, i.e. independent of position and direction.
6. Intuitive: The reconstructed image and the true concentration should be related in an "obvious" way. If the reconstruction is inhomogeneous, it should at least be possible to characterize it accurately in a form which can be conveyed to the user in a comprehensible form.
7. Free of artifacts: While noise is unavoidable and images will never be perfect, a good reconstruction algorithm should suppress artifical artifacts. In other words, noise must not fool a reconstruction algorithm to produce an image with brain-like features.
8. Few assumptions: Clinicians who uses the images for diagnosing patients can be uneasy with methods based on elaborate models of the brain. Their interests are mainly parts of the brain that are abnormal or lesioned, and it is hard to build a prior that is acceptable to them.
9. Coverage: It should be possible to use the method for most scannings seen by an individual doctor. As an example, an algorithm should not build on assumptions proven only for a single tracer. The training necessary for interpreting images reconstructed with a new method is not justified unless it can cover a reasonable amount of the total number of scannings.
10. Comparable: In dynamic studies the different time frames must be reconstructed similarly, to make it possible to interpret the sequence. For this reason the FBP-algorithm is currently preferred to EM in such studies.
11. Numerically robust: The method should always work. The users are trained to understand human physiology, not numerical algorithms. The user will generally be unable to intervene if the reconstruction misbehaves.
12. Rigorous and simple: Engineers, physicists, and mathematicians are generally easier to convince than clinicians when it comes to declaring one method better than another. Nevertheless it remains helpful to use well known mathematical theories. It is important that this group "likes the method", since the clinicians are unlikely to even consider a method unless the other experts agree.
13. Quantitative and accurate: Reconstructed images should reflect the concentration distribution as closely as possible. The ability to compute ratios between concentrations in varios regions of the brain has a big impact on both diagnostics and research. In particular the ability to model tracerkinetics depends on this.

## Discussion

A project that focuses on accuracy alone is not likely to have much success; it was no mistake that this item is last on the list. Even though accuracy is perhaps the easiest requirement to put into a cost-function, we should not forget that the
people who will use the images are trained at interpreting them. This means that consistency between images are usually a more important parameter.

Looking at the list as a whole, one may get the impression that clinical users are overly conservative. Generally, new methods are introduced only if the diagnostic power is demonstrably improved. To illustrate this point, we mention that it is still subject to debate whether attenuation-correction should be used in clinical SPECT. If users could be convinced by, say, a root-mean-square metric, this question would have been settled a decade ago. In 2002, only 9 of 19 Danish SPECT-centers used it routinely ${ }^{1}$.

The ability to produce results similar to existing methods can not be underestimated. In engineering, this goes under the heading "backwards-compatibility" and is by no means limited to reconstruction algorithms. Consider the likely success of the following products, regardless of technical merits:

- A text-processor that can not read any existing document.
- A CPU unable to run any previously purchased software.
- A car running on an exotic fuel, but neither gasoline nor diesel fuel.

It should be noted that lives are potentially at stake. A method should reflect this in the sense that if it delivers an answer it better be a good one.

A method should be viewed in light of the discussion above, although a research project can choose to ignore some aspects in order to set the stage for progress. However, principles that inherently fails to balance the issues are unlikely to be adopted for medical use.

### 5.2 Technical requirements

There are also a number of technicalities that should be considered. Unlike the user requirements, these are relatively easy to address.

## Computational power

Computational power represents a one-time (possibly large) investment but not clinical work. It is reasonable to expect that such investments will be made if they offer significant advantages over old methods.

Another factor is the continued increase in computational power. Any method that runs on existing hardware is likely to be useful for routine use on computers developed in 5 years. According to the industry's own predicitions [63], based on factories under construction and laboratory technologies, the current trend is expected to continue to at least 2016.

[^5]
## Model mismatch

It should be evident from chapter 3 that scanning models used in reconstruction are somewhat crude approximations. This is a result of the complicated interactions in the scanning process, uncertainties in calibrations, and approximations necessary to make the reconstructions practically feasible.

However, there is another source of model mismatch. The user may deliberately turn off some levels of realism for a number of reasons. For example, a model of scatter or attenuation could be turned off in order to approximate the usual results, to increase the reconstruction speed, or simple distrust.

There are situations where each of these reasons are perfectly valid. It is important to use methods that do not assume that the models are the actual truth, and robustness with respect to the model is very desirable.

## Non-negativity

Since concentrations are non-negative it is a question whether the reconstructed image should be allowed to contain negative values. This may also be posed as a question about whether it is acceptable to introduce bias to reduce variance. No reasonable ${ }^{2}$ method can guarantee that a region, which is truly zero, will be reconstructed as zero. Thus, a non-negativity constraint implies that such regions will be biased.

From a technological point of view one may ask whether it is reasonable to expect non-negative concentrations when corrections result in negative counts, as discussed in sec. 2.6.1 p. 18.

Neither of these arguments are likely to convince a practioneer of the plausability of negative concentrations. However, negative values occur mainly in low-count regions, which is mostly areas outside the head and low-affinity regions. While the former region is without much interest, the latter is often used as a reference. In this context a large part (e.g. cerebellum) is chosen and averaged to give a reference value. In this kind of studies it is crucial to have unbiased estimates of the reference region.

Nevertheless, the negative counts represents an inaccuracy it would be best to avoid. If areas are known to be zero (e.g. because they are outside the head), then the reconstruction should be restricted in this way. We shall return to this in sec. 5.6.

## Statistical analysis

Not all uses of the reconstructed images involves human interpretation. Currently, most statistical analysis of of emission tomographic imaging is based on the theory of Gaussian Random Fields.

The theory was first developed by Keith John Worsley et al. in [82], although this should be regarded as a precursor for [83] where the theory was solidified. Those equipped with the mathematical skills necessary to read this paper are encouraged to read it.

[^6]While random field theory is well beyond the scope of this thesis, we may nevertheless note that the noise of the reconstructed image is modeled as white noise filtered by a Gaussian kernel. Thus, to produce images useful in statistical analysis, the reconstruction algorithm should approximate this model.

## $\iiint \quad$ Image continuity

Discretization of the reconstructed image is a necessary prerequisite for a reconstruction algorithm. Theoretically, one could attempt to determine the position of the emission of each count acquired during the scanning process. However, since the counts are modeled as independent events, we are in reality approximating a probability distribution of the emission positions.

If we assume this distribution is in $L^{1}$ and non-zero values are confined to the interior of the scanner it follows from the Riemann-Lebesgue lemma [30] that it is band-limited ${ }^{3}$. The only relevant functions this author can imagine without this property is a finite sum of Dirac functions.

Band-limited functions can be discretized, as established by the sampling theorem [54]. For a sum of Dirac-functions, it is probably acceptable to bin the emission points in pixels of finite size, especially if we can make the pixels arbitrarily small.

If we reconstruct a series of images with an increasing number of pixels, it is desirable that these converges towards a stable solution which estimates this distribution. In a nutshell, a reconstructed image of size $1024 \times 1024$ should be no worse than an image of size $128 \times 128$, if the underlying functions satisfies the requirements above.

As explained in sec. 4.6.1 the unmodified EM-algorithm relies on pixelation and does not have this property. In light of the discussion above, the method should always be regularized, e.g. using a sieve.

### 5.3 Statistical reconstruction

After considering the requirements, we now turn to the different strategies. The reconstruction literature is vast, and it is beyond the scope of this text to even attempt summarizing the existing work. For early methods, we refer to the references in sec. 4.9 p. 47.

We note, however, that the majority of work in the field in the past 10 years has centered around statistical methods. Those interested in these methods should consult the excellent review by Fessler [22], which treats many of these methods from an algorithmic point of view in a general setting of surrogate functions. All EM-algorithms, for example, may be viewed as an optimization algorithm using a surrogate function derived from a Kullback-Leibler distance [8].

Generally, these methods use an iterative scheme to maximize a probability that involves the scanner (described in chapter 3) and a prior. This has led some authors to characterize these algorithms under headings such as SIRA (Statistical

[^7]Iterative Reconstruction Algorithms) or SIR. The following describes a few representative samples of this kind of information used by such methods.

Most algorithms use Bayesian inference at their core. A typical example of this approach was given by Bowsher et al. [11], who used a prior that encouraged regions of constant concentrations. They showed that this worked well on a phantom which matched this prior.

The approach was taken to an extreme by Philipsen et al. [57], who used adaptive snakes (an active contour model) in a weak-membranes model as a prior. The resulting objective function was optimized using a mean-field approximation in a simulated annealing algorithm. The author finds that papers like this should convince most people that the quest for improved images has sent researchers to areas very far from the algorithms currently used in medical imaging.

We also mention a paper by Lee et al. [48], who used a Bayesian method with a prior that encouraged piecewise linear areas. While the algorithms are quite typical, it differentiates itself by justifying the prior using autoradiography.

In a different approach, Ardekani et al. [3] used cross-entropy to reconstruct images using several MR-scans together with the emission scan. Most of the remaining literature builds a prior from the other modalities before the reconstructing begins.

A remarkable paper by Dougherty et al. [18] demonstrates how difficult it is to validate statistical reconstruction methods. The paper describes a phantom of a rat-brain, which was derived from tomographic scannings followed by autoradiography. This should result in a gold standard for a realistic image, but coregistering the emission data with the autoradiographic sections proved difficult.

## Discussion

Given the enormous body of work in statistical reconstruction, there has been remarkably little impact in medical usage. While this could be ascribed to inertia, there are factors that make statistical reconstruction unattractive in clinical use.

First, statistical methods rely on mathematical brain-models. This means that noisy measurements are guided towards meaningful brain-images, implying a tendency to turn noise into brain-like artifacts. This is inherent in the approach taken by such algorithms. We note that this is a common objection even for the EMalgorithm described in sec. 4.6 p. 42; arguably the simplest possible statistical reconstruction method.

Second, most statistical methods use statistical distributions in a hyper-space. It is very difficult to visualize this in a way that is meaningful to the user. Although this may be regarded as a "visualization problem", it must be solved if statistical methods are to be practical.

Thirdly, the most advanced statistical methods are difficult to compare across patients. However, in clinical use it is common to compare an image with a set of normals, i.e. images produced from healthy volunteers.

Finally, the images produced by the statistical methods usually looks qualitatively different from the classical methods. As already argued in sec. 5.1, this does not bode well for clinical adoption.

Nevertheless, statistical methods have the ability to use very detailed models of the scanning process. If an appropriate prior can be build from experimental knowledge, and quantitative accuracy is the only concern, then statistical methods may well be the proper choice.

### 5.4 Linear reconstruction

## j.j.

Figure 5.1:
Point-spread functions for the FBPreconstruction using the parameters in fig. 4.7 on p. 42.

Filtered back-projection, as described in sec. 4.4 p. 38, is a linear reconstruction method. This means that the reconstructed image can be described in terms of a matrix: $\hat{\mathbf{x}}=\mathbf{A}^{\dagger} \mathbf{b}$. For FBP, this matrix was given by $\mathbf{A}^{\dagger}=\mathbf{A}^{\mathrm{T}} \mathbf{F}$, where $\mathbf{A}$ is the discrete Radon transform and $\mathbf{F}$ is a filter, usually specified as a cut-off frequency and a window. Generally, the matrix $\mathbf{A}^{\dagger}$ may originate from a number of algorithms.

One crucial property common to all linear methods, is that they are easy to characterize. Since they are linear, a given image will, on the average, be reconstructed as $\mathscr{E}\left\{\mathbf{A}^{\dagger} \mathbf{b}\right\}=\mathbf{A}^{\dagger} \boldsymbol{\lambda}=\mathbf{A}^{\dagger} \mathbf{A} \mathbf{x}$. Thus, if we let $\mathbf{x}$ consists of a number of points, we can visualize the reconstructed PSFs. An example for the FBP is given in fig. 5.1.


This leads to the question: how can this be improved without violating the userrequirements from sec. 5.1? Limiting new reconstruction methods to the exact same result as FBP does not constitute progress.

In the opinion of this author there is room for improvement in at least the following areas:

- Make PSFs closer to homogeneous Gaussians: It is worth noting from fig. 5.1 that the PSFs are elongated near the edges. Furthermore, the space between them contains negative values, even in this noiseless case. However, the PSFs are close enough to Gaussian that this should not make a qualitative difference in the look of the reconstructed images.
- Increase resolution: The whole reason for the project is to improve the resolution, which in the case of linear reconstruction is equivalent to make the PSFs more narrow. There would be no objections to this from the users, provided the PSFs remain near-Gaussian.
- Incorporate anatomy: Anatomical information should optionally be employed to get improved images. However, this is only likely to be successful if the user is in absolute control. The user must understand exactly what is going on, and be able to turn off any modification. This is particularly important in areas that are suspected to contain lesions and other irregularities.
- Use more accurate system-matrices: As described in chapter 3, the Radon transform is a crude scanner model, and incorporating effects such as scatter, attenuation, and beam-divergence is appropiate. This, however, is only meaningful if the user is allowed to switch them off.
- Improve noise-models: Only poor noise-models can be incorporated in strictly linear methods. However, quasi-linear methods may be applicable, where the reconstruction-matrix $\mathbf{A}^{\dagger}$ is constructed using data. If necessary, the same matrix can be used to reconstruct any number of images, including normal material or different times in a dynamic series. Preferably, this model should still result in homogeneous Gaussian PSFs.


## Discussion

The most prominent example of a quasi-linear method is "Penalized Weighted Least Squares" by Fessler [20]. We note that this is one of the few methods that have been in routine use ${ }^{4}$.

Generally, linear methods satisfies the requirements in sec. 5.1 much better than statistical reconstruction methods. However, this comes at the prize of less accurate models:

- Concentrations may be negative.
- The discrete nature of the count-statistics can not be modeled.

This author believes that linear methods are easier to get adopted for clinical use than statistical methods, despite their shortcomings. Hence, the remaining part of the thesis will focus entirely on these methods. Chapter 6 deals with the issue of accuracy

### 5.5 Regularized reconstruction

It should be evident from chapter 4 that any reconstruction method must suppress noise to create useful images. In general, this field is known as regularization and is important in many disciplines, e.g. seismology, astronomy, and signal processing to name a few [30].

As described above, the goal is to compute an approximate solution resulting in homogeneous Gaussian PSFs. In the litterature, this is known as mollifier methods, and we can use an algorithm from chapter 4 to achieve this: first compute a solution using FBP ( $\omega_{c}=1$ and no window) and apply a post-filter with a Gaussian kernel. If a more accurate solution is required, FBP may be used as a pre-conditioner in an iterative method (see p. 41).

[^8]This offer only a small advantage over FBP, since the resolution is basically limited in the same way. However, the literature offers plenty of alternatives to postfiltering. Thus, we may use a post-filter (to imitate FBP), and increase the resolution but avoid excessive noise using regularization.

This appears like a nearly ideal way to deal with the requirements from the users. This is particularly true if the amount of post- and regularization-filter can be changed in real-time while the user is trying different settings. This is the method pursued in the remaining part of the thesis.

Table 5.1:
Example of filterfactors studied in the literature [30].

## Standard regularization

There are many standard regularization-schemes that can be expressed in terms of the singular value decomposition (SVD). The SVD of the system-matrix $\mathbf{A}$ is given by:

$$
\mathbf{A}=\mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^{\mathrm{H}}
$$

Here $\mathbf{U}$ and $\mathbf{V}$ have orthonormal columns, i.e. $\mathbf{U}^{\mathrm{H}} \mathbf{U}=\mathbf{V}^{\mathrm{H}} \mathbf{V}=\mathbf{I}$ (this variant is known as the thin $S V D$ ), and $\boldsymbol{\Sigma}$ is a square diagonal matrix with real, non-negative elements on the diagonal.

With this definition, the regularization schemes considered in this thesis may be written in terms of filter-factors [30]:

$$
\mathbf{A}^{\dagger}=\sum_{i: \sigma_{i}>0} \frac{f_{i}}{\sigma_{i}} \mathbf{v}_{i} \mathbf{u}_{i}^{\mathrm{H}}
$$

Here $f_{i}$ are the filter-factors, which can usually be specified as a function $f_{i}=$ $f\left(\sigma_{i}\right)$. A few examples are given in tab. 5.1, although we shall not go further into the specific merits of these choices. The pseudo-inverse is as close to an inverse of $\mathbf{A}$ as possible. All remaining filters share the property that $f_{i} \approx 1$ for $\sigma_{i} \gg \lambda$, while $f_{i} \approx 0$ for $\sigma_{i} \ll \lambda$. The parameter $\lambda$ plays the same role as the cut-off frequency, $\omega_{c}$, did in the FBP studied in sec. 4.4 p. 38.

| Name | Filter-function |
| :--- | :--- |
| Pseudo-inverse | $f(\sigma)=1$ |
| Tikhonov | $f(\sigma)=\frac{\sigma^{2}}{\sigma^{2}+\lambda^{2}}$ |
| Truncated SVD (TSVD) | $f(\sigma)= \begin{cases}1 & \text { for } \sigma \geq \lambda, \\ 0 & \text { otherwise. }\end{cases}$ |
| Rutishauser | $f(\sigma)=\frac{\sigma^{2}}{\sigma^{2}+\lambda^{2}+\lambda^{2} /\left(\sigma^{2}+\lambda^{2}\right)}$ |
| Damped SVD | $f(\sigma)=\frac{\sigma}{\sigma+\lambda}$ |

## $\iiint \quad$ Tikhonov regularization

Apart from the pseudo-inverse, which is not ordinarily regarded as a form of regularization, by far the most studied method is Tikhonov regularization. It can be shown that the the Tikhonov solution, $\hat{\mathbf{x}}=\mathbf{A}^{\dagger} \mathbf{b}$, of the linear system $\mathbf{A} \hat{\mathbf{x}} \approx \mathbf{b}$ is equivalent to the following optimization problem [30]:

$$
\begin{equation*}
\hat{\mathbf{x}}=\underset{\hat{\mathbf{x}}}{\operatorname{argmin}}\left(\|\mathbf{A} \hat{\mathbf{x}}-\mathbf{b}\|_{2}^{2}+\lambda^{2}\|\hat{\mathbf{x}}\|_{2}^{2}\right) \tag{5.3}
\end{equation*}
$$

Thus, the solution attempts to keep the residual error, $\mathbf{A} \hat{\mathbf{x}}-\mathbf{b}$ small in norm, but adds a penalty to discourage solutions with large norms.

In a statistical setting, Tikhonov regularization is known as ridge regression [52]. If $\mathbf{A x}=\mathbf{b}+\mathbf{e}$, where the elements in the noise-vector $\mathbf{e}$ are uncorrelated random variables with zero mean but the same variance, $s^{2}<\infty$, then there is a regularization parameter, $\lambda$, where the mean-square-error, $\mathscr{E}\left\{\|\hat{\mathbf{x}}-\mathbf{x}\|_{2}^{2}\right\}$, is less than the unregularized solution [43].

Tikhonov regularization may also be derived using a Bayesian approach. The regularized solution, $\hat{\mathbf{x}}$, is a maximum a posteriori estimate, assuming the elements of $\mathbf{x}$ has expected zero mean and variance $s^{2} / \lambda[43]$.

Finally, it can be shown that among the methods that can be defined in terms of filter-factors, the choice that minimizes $\mathscr{E}\left\{\|\hat{\mathbf{x}}-\mathbf{x}\|_{2}^{2}\right\}$ is given by [30]:

Here $\mathbf{W}$ is a diagonal matrix where the element $\mathbf{W}_{(i, i)}$ is the inverse of the approximated standard deviation of $\mathbf{e}_{i}$. This function is the subject of chapter 6 , and leads to a quasi-linear algorithm.

Another common variation of Tikhonov regularization is to change the penaltyterm $\|\hat{\mathbf{x}}\|_{2}$ to $\|\mathbf{L} \hat{\mathbf{x}}\|_{2}$, where $\mathbf{L}$ is typically a discrete approximation of, say, the Laplace operator. However, in the current context this approach has several disadvantages:

- It introduces spatially correlated artifacts.
- The PSFs becomes dissimilar to Gaussian kernels.
- The form of $\mathbf{L}$ must be justified from autoradiographic studies for each tracer, and such validation studies are extremely scarce.

For these reasons, we shall not use these methods, but instead rely on postfiltering.

## $\iiint \quad$ Model mismatch

One convincing reason to use regularization methods based on filter-factors is their ability to reject model-mismatch. This stems from the fact that the scannermodel can be regarded as a perturbation of the true scanner, and we may therefore use results from perturbation-theory of the SVD to establish bounds on this error.

Several key results are, for other reasons, included in appendix C: Ostrowski's theorem (theorem C3 p. 147) states that multiplicative factors will at most change the singular values by the factors themselves. Thus, an inaccurate calibration of the photo-multiplier tubes or the absence of an attenuation model, will not change the singular values significantly.

Weyl's theorem (theorem C1 p. 147) takes care of the remaining errors, although somewhat less convincingly.

The reconstructed image depends not only on the singular values, but also singular vectors. A rigorous treatment depends on angles between subspaces, but the main result is that these are not severely affected by small perturbations. Thus, if we use a reconstruction method where the filter-factors do not change too rapidly, then we can expect that the reconstructed image is not detrimentally affected by the model mismatch.

We note that considerations of model-mismatch is very uncommon in the reconstruction literature [22].

### 5.6 Incorporating anatomy

There are two principal ways to use anatomy in reconstruction:

1. Build a prior, using probability- or information-theory, and apply a statistical reconstruction algorithm.
2. Impose constraints on the reconstructed image during the reconstruction.

As discussed in sec. 5.3, statistical methods relies on information that is difficult to convey accurately to non-technical users. This issue is made even more difficult when we add the requirement that the user must be able to control the process, including turning it off in different parts of the image.

This leads us to consider the second approach. We note that the idea of restricting the images was proposed at least as early as 1984 in one of the most cited papers on the EM-algorithm [45, p. 309]:

Certain equality constraints also become trivial to implement in the above framework.

Neighboring pixels can be forced to have the same intensities.

The paper also mentions the simplicity of incorporating areas known to be zero. One way to accomplish this is to set the initial iterate to zero, although this scheme must be modified if a convolutive sieve is used.

Using restrictions to incorporate anatomical knowledge has the attractive property that it is easy to visualize: a simple sketch of the constant regions suffice. It also has the advantage that many workflows used clinically already depends on identifying regions. Therefore the approach can be used together with familiar tools.

## $\iiint$ <br> Anatomy and regularization

If we know that some pixels in $\mathbf{x}$ are zero, then we may create an operator, $\mathbf{P}$, that zero out pixels known to be zero. Similarly, we may let $\mathbf{P}$ be an operator that averages all pixels in one or more disjoint regions.

To be specific, let all pixels be categorized into $K$ disjoint regions, many possibly containing a single pixel. Let $\mathbf{p}_{k}$ be a vector with one element per pixel, and let an element be 1 if a pixel belonging to region $k$ and zero otherwise. We may then write $\mathbf{P}$ as follows:

$$
\begin{equation*}
\mathbf{P}=\sum_{k=1}^{K}\left(\mathbf{p}_{k} \mathbf{p}_{k}^{\mathrm{T}}\right) /\left(\mathbf{p}_{k}^{\mathrm{T}} \mathbf{p}_{k}\right) \tag{5.6}
\end{equation*}
$$

It follows immediately that $\mathbf{P}=\mathbf{P}^{\mathrm{T}}$. Since all regions are assumed to be disjoint we have $\mathbf{p}_{i}^{\mathrm{T}} \mathbf{p}_{j}=0$ for $i \neq j$, and we get:

$$
\begin{aligned}
\mathbf{P P} & =\left(\sum_{k 1=1}^{K}\left(\mathbf{p}_{k 1} \mathbf{p}_{k 1}^{\mathrm{T}}\right) /\left(\mathbf{p}_{k 1}^{\mathrm{T}} \mathbf{p}_{k 1}\right)\right)\left(\sum_{k 2=1}^{K}\left(\mathbf{p}_{k 2} \mathbf{p}_{k 2}^{\mathrm{T}}\right) /\left(\mathbf{p}_{k 2}^{\mathrm{T}} \mathbf{p}_{k 2}\right)\right) \\
& =\sum_{k 1=1}^{K} \sum_{k 2=1}^{K}\left(\mathbf{p}_{k 1} \mathbf{p}_{k 1}^{\mathrm{T}} \mathbf{p}_{k 2} \mathbf{p}_{k 2}^{\mathrm{T}}\right) /\left(\mathbf{p}_{k 1}^{\mathrm{T}} \mathbf{p}_{k 1} \mathbf{p}_{k 2}^{\mathrm{T}} \mathbf{p}_{k 2}\right) \\
& =\sum_{k=1}^{K}\left(\mathbf{p}_{k} \mathbf{p}_{k}^{\mathrm{T}} \mathbf{p}_{k} \mathbf{p}_{k}^{\mathrm{T}}\right) /\left(\mathbf{p}_{k}^{\mathrm{T}} \mathbf{p}_{k} \mathbf{p}_{k}^{\mathrm{T}} \mathbf{p}_{k}\right) \\
& =\sum_{k=1}^{K}\left(\mathbf{p}_{k} \mathbf{p}_{k}^{\mathrm{T}}\right) /\left(\mathbf{p}_{k}^{\mathrm{T}} \mathbf{p}_{k}\right) \\
& =\mathbf{P}
\end{aligned}
$$

Thus, $\mathbf{P}$ is an orthogonal projector [73]. In any case we have $\mathbf{x}=\mathbf{P x}$, since $\mathbf{P}$ does not change an image that satisfies the specified restrictions. Thus we may compute a regularized solution to the following system:

$$
\begin{equation*}
\mathbf{A P} \hat{\mathbf{x}} \approx \mathbf{b} \tag{5.7}
\end{equation*}
$$

If we let $\mathbf{A P}=\mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^{\mathrm{T}}$ be an SVD where only non-zero singular values are retained, then we have $\hat{\mathbf{x}} \in \operatorname{range}(\mathbf{V}) \subseteq \operatorname{range}\left(\mathbf{P}^{\mathrm{T}}\right)=\operatorname{range}(\mathbf{P})$. Thus, the solution produced by any regularized method described in sec. 5.5 will automatically satisfy the specified constraints.

### 5.7 References

The user-requirements in sec. 5.1 are rarely treated in the literature, except at commercial websites. As the author did not have access to their software, these websites can not be endorsed by giving references.

Anyone depending on the computational progress should be aware of the ITRS: International Technology Roadmap for Semiconductors [62, 63]. This organization produces high-quality reports with facts-based predictions.

Statistical reconstruction is a very active area of research. For a small sample we refer to the references given in sec. 5.3. Apart from pointing at a single review
[22], the best advice for those who wants to get into the field is to study the papers by Jeffrey A. Fessler.

Regularization of linear systems is standard material, which is covered in many places. However, in the setting of linear algebra (as opposed to more general Hilbert spaces), the place one should turn first is the doctoral thesis by Per Christian Hansen [30].

Projectors are described in many text-books on linear algebra, e.g. [71]. Another excellent reference is [73], which also use projectors in a numerical setting.
$\oplus$
$\oplus$

## Weighted least-squares

Chapter 5 described a number of advantages of linear and quasi-linear reconstruction methods. However, one area where it lacked behind statistical methods was in quantitative accuracy. This chapter analyzes this behavior, and establish bounds on the difference between a simple weighted least-squares objective and a proper model of Poisson noise.

This study considers a broad class of reconstruction algorithms, and analyzes the worst-case behavior when a least-squares objective replaces a Poisson loglikelihood. We derive a bound on the Hausdorff-distance, in sinogram space, and compare it with statistical noise.

We show that in many circumstances a broad class of reconstruction algorithms can be expected to produce similar images using either objective. For rigor, we give a bound in the important special case where the penalty-function (or logprior) is itself a weighted least-squares objective.

This chapter is mathematically oriented.

### 6.1 Chapter overview

Most reconstruction algorithms can be stated as an optimization-problem of a function that contains a term that models the scanning process and a term that models an a priori knowledge or, equivalently, penalizes undesired images.

We consider a scanner that produces Poisson-distributed counts, and analyze the solutions when the Poisson log-likelihood is replaced with two different leastsquares objectives. We bound the distance between solutions of equal likelihood and show that the reconstructed images should be similar in many circumstances. In the important special case, where both the log-likelihood and penalty-function (or log-prior) are weighted least-squares objectives, we give a rigorous bound.

Most studies evaluate a likelihood function as a part of a reconstruction process. The goal in this chapter is to isolate the effect of the likelihood function, and make as few assumptions as possible. To the best knowledge of the author, no previous work have established bounds that applies as generally as those presented here.

## Methods

First, the likelihood functions considered are introduced in sec. 6.2. The LSerror is seen as an approximation to the Poisson log-likelihood, and objectives for weighted LS (WLS) and model-weighted LS (MLS) schemes are derived in sec. 6.3 and 6.4.

Next, sec. 6.5 introduces a generalized view of regularized reconstruction algorithms. We consider iso-surfaces of the likelihood-functions, and define the Hausdorff distance between such surfaces. We show that small distances between such surfaces result in similar reconstructed images.

The bound obtained will in general depend on the level of regularization, and sec. 6.7 considers the statistical distribution of the Poisson log-likelihood. This leads to a range of practical threshold values for the log-likelihood.

Section 6.8 analyzes the WLS-scheme, and derives a bound on the distance between the likelihood surfaces. This bound is compared with the inherent statistical variation of a Poisson process.
Finally, sec. 6.9 analyzes the model-weighted LS scheme. Since the WLS-scheme is too simple to encompass regions with very few counts, the new relative bound is given. This demonstrates that the Poisson-likelihood and WLS-objective differ mainly by the quality of the variance estimates.

### 6.2 Likelihood functions

The acquisition process in emission tomography is inherently statistical, and is described in chapter 3 . We shall assume that the system matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$ maps $N$ pixels into $M$ sinogram bins.
It is assumed the components of $\mathbf{b}$ are independent and Poisson-distributed with parameters given by $\boldsymbol{\lambda}=\mathbf{A x}$, where $\mathbf{x}$ is the true object in the scanner. The likelihood of $\boldsymbol{\lambda}$ can therefore be derived from the Poisson-distribution as described in sec. 3.8 p. 30:

$$
\begin{equation*}
L_{\mathrm{P}}(\boldsymbol{\lambda}, \mathbf{b})=\prod_{i=1}^{M} e^{-\lambda_{i}} \lambda_{i}^{b_{i}} / b_{i}! \tag{6.1}
\end{equation*}
$$

We do not consider the likelihood functions directly. Since the projections in $\mathbf{b}$ are independent stochastic variables, it is easier to work with the log-likelihood. The functions considered peak at $\boldsymbol{\lambda}=\mathbf{b}$, and we subtract this peak to get non-positive values.

The Gaussian likelihood resulting in least-squares problems is given by [52]:

$$
L_{\mathrm{LS}}(\boldsymbol{\lambda}, \mathbf{b})=\frac{e^{-\frac{1}{2}(\mathbf{b}-\boldsymbol{\lambda})^{\mathrm{T}} \boldsymbol{\Sigma}^{-1}(\mathbf{b}-\boldsymbol{\lambda})}}{\sqrt{\operatorname{det}(2 \pi \boldsymbol{\Sigma})}}
$$

where $\boldsymbol{\Sigma}=\operatorname{diag}\left(\left[\begin{array}{lll}\sigma_{1}^{2} & \ldots & \sigma_{M}^{2}\end{array}\right]\right)$.
Briefly postponing the estimation of $\sigma_{i}^{2}$, we may now define the log-likelihood functions. For convenience define $0 \ln (0)=0$ as a compound symbol and get:

It will be necessary to distinguish between solutions associated with different likelihood functions. In this chapter, a hat decorates quantities that involve the Poisson-likelihood (e.g. $\hat{\lambda}$ and $\hat{\mathbf{x}}$ ).

### 6.3 Weighted least-squares

If the variances are estimated directly from the measured sinogram, maximizing $\ell_{\text {LS }}$ becomes a weighted least-squares (WLS) problem. Quantities will be decorated with an inverted hat, e.g. $\check{\boldsymbol{\lambda}}$.

In the simplest case, the approximation is restricted to estimate $\sigma_{i}^{2}$ from a single projection, i.e. $b_{i}$. This represents a worst-case scenario with no smoothing in the sinogram, but have the advantage that no assumptions are required from $\mathbf{A}$ or $\mathbf{x}$.

Expanding a single term of eq. (6.2) around $\lambda_{i}=b_{i}$ gives:

$$
\begin{equation*}
\ell_{\mathrm{P}}\left(\lambda_{i}, b_{i}\right) \approx-\frac{\left(\lambda_{i}-b_{i}\right)^{2}}{2 b_{i}}+\frac{\left(\lambda_{i}-b_{i}\right)^{3}}{3 b_{i}^{2}}-\frac{\left(\lambda_{i}-b_{i}\right)^{4}}{4 b_{i}^{3}}+\cdots \tag{6.4}
\end{equation*}
$$

Choosing $\sigma_{i}^{2}=b_{i}$ makes the first term equal to eq. (6.3), but can not be used for $b_{i}=0$. The analysis is based on the limit $\sigma_{i}^{2} \rightarrow 0^{+}$, which effectively enforces $\lambda_{i}=b_{i}=0$. Since $\ell_{\mathrm{P}}(0,0)=0$, we also define $\ell_{\mathrm{WLS}}(0,0)=0$ and may now write the WLS-objective in the following form:

$$
\begin{equation*}
\ell_{\mathrm{WLS}}(\boldsymbol{\lambda}, \mathbf{b})=\sum_{\substack{i=1 \\ b_{i} \neq 0}}^{M}-\frac{\left(\lambda_{i}-b_{i}\right)^{2}}{2 b_{i}}, \quad \lambda_{i}=0 \text { for } b_{i}=0 \tag{6.5}
\end{equation*}
$$

In practical methods, such as penalized weighted least-squares (PWLS), $\sigma_{i}^{2}$ is estimated by filtering over the sinogram and considering detector efficiencies, deadtime correction, and transmission errors among other factors [20]. We therefore expect the analysis based on this function to be much worse than that encountered in practice, in particular for regions with very few counts.

### 6.4 Model-weighted least-squares

If the variances are allowed to depend on the estimated parameters, i.e. $\sigma_{i}^{2}=$ $g\left(\lambda_{i}\right)$, we get a model-weighted least-squares (MLS) problem. Quantities are decorated with a tilde, e.g. $\tilde{\boldsymbol{\lambda}}$.

Although the MLS-problem is essentially as hard to optimize as the Poisson likelihood [22], it is of theoretical interest for several reasons. First, the analysis explains the behavior when a good approximation of the variance is available. Also, it helps to explain the behavior of the Poisson-likelihood in the region very near the measured sinogram where $\tilde{\boldsymbol{\lambda}} \approx \mathbf{b}$.

Inserting $\sigma_{i}^{2}=g\left(\lambda_{i}\right)$ in eq. (6.3) and expanding yields:

$$
\begin{aligned}
\ell_{\mathrm{MLS}}\left(\lambda_{i}, b_{i}\right) \approx & -\frac{\left(\lambda_{i}-b_{i}\right)^{2}}{2 g\left(b_{i}\right)}+\frac{\left(\lambda_{i}-b_{i}\right)^{3} g^{\prime}\left(b_{i}\right)}{2 g\left(b_{i}\right)^{2}} \\
& +\frac{\left(\lambda_{i}-b_{i}\right)^{4}}{4}\left(\frac{g^{\prime \prime}\left(b_{i}\right)}{g\left(b_{i}\right)^{2}}-2 \frac{g^{\prime}\left(b_{i}\right)^{2}}{g\left(b_{i}\right)^{3}}\right)+\cdots
\end{aligned}
$$

Setting $\sigma_{i}^{2}=g\left(\lambda_{i}\right)=\lambda_{i}$ reproduces the first term in eq. (6.4), while the error in the second term is $-\left(\lambda_{i}-b_{i}\right)^{3} /\left(6 b_{i}^{2}\right)$, i.e. half that of $\ell_{\text {WLS }}$ in absolute value. Thus, $\ell_{\text {MLS }}$ is much better than $\ell_{\text {WLS }}$ near the peak $\lambda_{i} \approx b_{i}$ for this choice of $\sigma_{i}^{2}$.

For our purpose, with emphasis on low-count regions, a more interesting situation is the case $\lambda_{i} \gg b_{i}$ and especially $b_{i}=0$. Therefore we want the ratio $\ell_{\mathrm{MLS}}\left(\lambda_{i}, 0\right) / \ell_{\mathrm{P}}\left(\lambda_{i}, 0\right)=\lambda_{i} /\left(2 g\left(\lambda_{i}\right)\right)$ to be 1 . It follows that $\sigma_{i}^{2}=g\left(\lambda_{i}\right)=\lambda_{i} / 2$.

The two equations differ only by a constant factor. From an algorithmic point of view, this corresponds to different choices of the regularization parameter discussed in sec. 6.5. All further analysis will be based on the latter:

$$
\begin{equation*}
\ell_{\mathrm{MLS}}(\boldsymbol{\lambda}, \mathbf{b})=\sum_{i=1}^{M}-\frac{\left(\lambda_{i}-b_{i}\right)^{2}}{\lambda_{i}} \tag{6.6}
\end{equation*}
$$

### 6.5 General reconstruction

As shown in chapter 4, the maximum-likelihood solution is rarely desirable, unless the spatial sampling is very crude. Instead, we search a smoother solution that strikes a balance between likelihood and other desirable features.

Most methods can be written in the form of an optimization problem.
Definition 6.1: (Optimized reconstruction) Let $\ell(\mathbf{A} \hat{\mathbf{x}}, \mathbf{b})$ be a log-likelihood function (e.g. $\ell_{P}$ or $\left.\ell_{W L S}\right)$ modeling the scanner acquisition, let $r(\hat{\mathbf{x}})$ be a non-negative convex penalty-function, and let $\alpha$ be a regularization constant.

Furthermore, define the following function:

$$
f(\hat{\mathbf{x}}, \mathbf{b})=\ell_{P}(\mathbf{A} \hat{\mathbf{x}}, \mathbf{b})-\alpha r(\hat{\mathbf{x}})
$$

Then an optimizing reconstruction algorithm computes the following image over a convex set $\mathcal{S}$ :

$$
\hat{\mathbf{x}}=\underset{\hat{\mathbf{x}} \in \mathcal{S}}{\operatorname{argmax}} f(\hat{\mathbf{x}}, \mathbf{b})
$$

The function $r(\hat{\mathbf{x}})$ is either a penalty or log-prior depending on the philosophy and derivation of the method. Often the regularization parameter, $\alpha$, is considered separately. Although $\alpha$ may be determined indirectly (e.g. by early stopping as described in sec. 4.7 p. 45), its function is to determine how large $\ell_{\mathrm{P}}(\mathbf{A} \hat{\mathbf{x}}, \mathbf{b})$ must be before a solution is acceptable.

Conversely, we may view the reconstruction as the process of maximizing the penalty-function over the iso-surface $\ell_{\mathrm{P}}(\mathbf{A} \hat{\mathbf{x}}, \mathbf{b})=k$. If $\ell_{\mathrm{LS}}$ approximates $\ell_{\mathrm{P}}$ well, then the solution $\hat{\mathbf{x}}$ should have $\ell_{\mathrm{LS}}(\mathbf{A} \hat{\mathbf{x}}, \mathbf{b}) \approx k$.

From an algorithmic point of view, fixing $k$ instead of $\alpha$ is often the goal. If $\ell_{\text {WLS }}$ is combined with a weighted least-squares penalty-function, this strategy is known as the discrepancy principle [30].

A similar view was advocated by Veklerov and Llacer [76], who used the notion of feasibility shell to determine an appropiate sieve for the EM-algorithm. The shell corresponds to iso-surfaces for a small range of $k$, and the appropiate sieve is essentially equivalent to a fixed value of $k$.

### 6.6 Hausdorff distance

To fix ideas we present a reduced example. Let the system matrix A project two voxels $(N=2)$ onto a two-element sinogram $(M=2)$. Define the penaltyfunction as $r(\mathbf{x})=\|\mathbf{x}\|_{2}^{2}$ and assume we have measured the sinogram $\mathbf{b}$ :

$$
\mathbf{A}=\left[\begin{array}{ll}
0.30 & 0.50 \\
0.28 & 0.70
\end{array}\right] \quad \mathbf{b}=\left[\begin{array}{l}
5 \\
2
\end{array}\right]
$$

Fig. 6.1 shows iso-curves of the likelihood- and penalty functions. The loglikelihood is set at $k=-0.58 \cdot 2$, a choice discussed in sec. 6.7. The enlargement on the right shows the distance from each solution to the nearest point on the opposite surface.

Figure 6.1:
Iso-surfaces of the likelihood- and the penalty-functions in sinogram-space, the maximum distance $D$, and solutions. The enlargement shows the solutions and their minimal displacement to the opposite surface.

Iso-surfaces of likelihood functions


Our approach will be to bound the distance between the surfaces in sinogram space. To make the statement precise, take $\|\cdot\|$ to be an unspecified norm, and define the distance $D$ as the maximum displacement from one surface to another:

$$
\begin{align*}
D_{1 \rightarrow 2} & =\max _{\boldsymbol{\lambda}_{1} \in S_{1}} \min _{\boldsymbol{\lambda}_{2} \in S_{2}}\left\|\boldsymbol{\lambda}_{1}-\boldsymbol{\lambda}_{2}\right\| \\
D_{2 \rightarrow 1} & =\max _{\boldsymbol{\lambda}_{2} \in S_{2} \boldsymbol{\lambda}_{1} \in S_{1}}\left\|\boldsymbol{\lambda}_{1}-\boldsymbol{\lambda}_{2}\right\| \\
D & =\max \left\{D_{1 \rightarrow 2}, D_{2 \rightarrow 1}\right\} \tag{6.7}
\end{align*}
$$

The surfaces are the iso-surfaces of the likelihood-functions and depend on the threshold-value $k$. Although the displacements $D_{1 \rightarrow 2}$ and $D_{2 \rightarrow 1}$ were identical in the example, it is generally necessary to consider them separately.

The distance $D$ is called the Hausdorff-distance between $S_{1}$ and $S_{2}$ with respect to the given norm [1]. It is named after Felix Hausdorff, who showed that $D$ is a distance, and $D_{1 \rightarrow 2}$ and $D_{2 \rightarrow 1}$ are metrics. $D$ is sometimes known as the "undirected Hausdorff distance".

## Implications in image-space

Bounds on $D$ are useful because they have consequences in image-space. The distance between the solutions in fig. 6.1 are larger than $D$ only because the penalty differ less between the solutions than the opposite points. This must be the case since, by definition, the solutions found on each surface have the lowest penalty. In other words, the LS-solution is better (measured by the penalty-function) than the Poisson-solution perturbed at most $D$ in the sinogram.

Consider the difference, in image-space, between the solutions above measured in the 2 -norm, i.e. the penalty function. The transformation (the inverse of $\mathbf{A}$ ) squeezes the penalty-curves into circles. Replacing the Poisson-surface with an LS-surface caused a smaller change in the penalty-function than any of the displacements.

To emphasize this point, fig. 6.2 shows LS-surfaces that arise from small perturbations of $\mathbf{b}$. All these surfaces must result in similar images in image-space, because they represent the same object perturbed by statistical noise. It should be evident that in general, any surface embedded in the "cloud" should result in similar images.

## LS-surfaces perturbed by noise



Cloud of LSsurfaces resulting from perturbing $\mathbf{b}$.

In the important special case where the penalty-function is on the form $r(\mathbf{x})=$ $\|\mathbf{L x}\|_{2}^{2}$, theorem A2 in appendix A shows that under mild assumptions there are images satisfying the relation $r\left(\hat{\mathbf{x}}-\check{\mathbf{x}}_{1}\right)+r\left(\hat{\mathbf{x}}-\check{\mathbf{x}}_{2}\right) \leq r\left(\check{\mathbf{x}}_{2}\right)-r\left(\check{\mathbf{x}}_{1}\right)$ for some WLS-reconstructed images, $\check{\mathrm{x}}_{1}$ and $\check{\mathbf{x}}_{2}$.

These solutions must, by definition, minimize the penalty-function on some surfaces in the cloud. If we can guarantee that all surfaces result in similar images, by bounding $D$, then it follows that the reconstructed image is near the image reconstructed using the Poisson surface.

In the general case this result must be weakened. Section A2 outlines a proof that any vector $\mathbf{d}_{\boldsymbol{\lambda}}$ have a $\mathbf{d}_{\mathbf{b}}$ so that $\ell_{\mathrm{WLS}}\left(\hat{\boldsymbol{\lambda}}+\mathbf{d}_{\boldsymbol{\lambda}}, \mathbf{b}\right)=\ell_{\mathrm{WLS}}\left(\hat{\boldsymbol{\lambda}}, \mathbf{b}-\mathbf{d}_{\mathbf{b}}\right)$ where $\left(d_{b}\right)_{i}=c_{i}\left(d_{\lambda}\right)_{i},\left|c_{i}\right|<1+\sqrt{2}<2.42$. This ensures that if $D$ is small, compared to the statistical noise of $\mathbf{b}$, the Poisson-surface will be embedded in the cloud. Thus, the images reconstructed using the Poisson-surface and the WLS-surface can only vary greatly, if the images reconstructed using the Poisson-surface is sensitive to statistical noise.

## Non-negativity

The distance bounds apply unmodified for algorithms that enforce non-negativity in the image-space. To see this, observe that such an algorithm must be able to choose a solution on the boundary of the non-negative set.

Since the non-negative vectors in image-space form a convex set, so do the set in sinogram-space after a linear transformation. The iso-surfaces treated in this section are therefore cut so they lie inside this convex set. However, the distance between two surfaces modified in this way can never exceed that of the original surfaces. If the original opposite point has been removed, then there must be a closer one lying on the surface of the non-negative set.

### 6.7 Log-likelihood range

Before comparing the likelihood functions we will start by determining the range of likelihood-thresholds that is useful in practice. To this end, consider the distribution of $\ell_{\mathrm{P}}(\boldsymbol{\lambda}, \mathbf{b})$ as a stochastic variable when $\boldsymbol{\lambda}$ is fixed.

Since the total log-likelihood $\ell_{\mathbf{P}}(\boldsymbol{\lambda}, \mathbf{b})=\sum_{i=1}^{M} \ell_{\mathrm{P}}\left(\lambda_{i}, b_{i}\right)$ is a sum of many independent terms, it should have a nearly Gaussian distribution. Fig. 6.3 shows the negated mean and variance for a single term.

Figure 6.3:
Mean and variance of the Poisson log-likelihood as a function of the true parameter $\lambda_{i}$. The mean is negated for illustrative convenience. Both are bounded, implying a limited range of useful values of the log-likelihood.


## Mean and variance

By definition, the mean log-likelihood of a projection is:

$$
\begin{aligned}
\mathscr{E}\left\{\ell_{\mathrm{P}}\left(\lambda_{i}, b_{i}\right)\right\} & =\sum_{b_{i}=0}^{\infty} \frac{e^{-\lambda_{i}} \lambda_{i}^{b_{i}}}{b_{i}!} \ell_{\mathrm{P}}\left(\lambda_{i}, b_{i}\right) \\
& =\sum_{b_{i}=0}^{\infty} \frac{e^{-\lambda_{i}} \lambda_{i}^{b_{i}}}{b_{i}!}\left(\lambda_{i} \ln \left(\lambda_{i}\right)-b_{i} \ln \left(b_{i}\right)\right)
\end{aligned}
$$

One important point is that $\mathscr{E}\left\{\ell_{\mathrm{P}}\left(\lambda_{i}, b_{i}\right)\right\} \geq-0.5802$. This gives a lower bound of the mean, irrespective of the unknown parameter $\lambda_{i}$. Similarly, the variance is bounded above by $\operatorname{var}\left(\ell_{\mathrm{P}}\left(\lambda_{i}, b_{i}\right)\right)<0.6015$ and asymptotically becomes $1 / 2$ for $\lambda_{i} \rightarrow \infty$.

## Likelihood threshold

As mentioned we expect that $\ell_{\mathrm{P}}(\boldsymbol{\lambda}, \mathbf{b})$ has a nearly Gaussian distribution. Combining this with the result above, it follows that the threshold $k=-0.5802 M$ should give $\ell_{\mathbf{P}}(\boldsymbol{\lambda}, \mathbf{b})>k$ at least $50 \%$ of the time for the true $\boldsymbol{\lambda}$.

Setting $k=-0.5802 M(1+\epsilon)$ adds a confidence interval, where $\epsilon$ represents a few standard deviations. Since the variance was bounded, it follows that $\epsilon$ decreases as $1 / \sqrt{M}$. Usually $M$ is sufficiently large that $\epsilon$ can safely be ignored.

For $k$ much lower than $-0.5802 M$, we accept solutions, $\hat{\mathbf{x}}$, significantly less likely to produce the measured sinogram, $\mathbf{b}$, than the true concentration $\mathbf{x}$. Usually $k$ should be higher, especially if there are many low-count projections ( $\lambda_{i}<$ $1)$ as is often the case for real scans.

### 6.8 Analysis of weighted least-squares

Let $\hat{\lambda}$ and $\check{\lambda}$ be regularized solutions obtained using the Poisson likelihood and the WLS-likelihood respectively. These are called primary estimates.

We now have to find a displacement vector from each primary estimate to the opposite surface (see fig. 6.1 in sec .6 .5 .) We want to bound a norm of these displacement vectors.

Consider replacing each component with a new value that matches the likelihood of the opposite surface. The displacement always overestimates the distance defined in eq. (6.7).
For instance, let $S_{1}=\left\{\hat{\boldsymbol{\lambda}} \mid \ell_{\mathrm{P}}(\hat{\boldsymbol{\lambda}}, \mathbf{b})=k\right\}$ and $S_{2}=\left\{\check{\boldsymbol{\lambda}} \mid \ell_{\mathrm{WLS}}(\check{\boldsymbol{\lambda}}, \mathbf{b})=k\right\}$. Now replace each component of $\hat{\boldsymbol{\lambda}}$ with a suitable entry so $\check{\boldsymbol{\lambda}}^{*}$ satisfies $\ell_{\mathrm{p}}(\hat{\boldsymbol{\lambda}}, \mathbf{b})=$ $\ell_{\text {WLS }}\left(\check{\boldsymbol{\lambda}}^{*}, \mathbf{b}\right)$. Finally, an upper bound on $\left\|\hat{\boldsymbol{\lambda}}-\check{\boldsymbol{\lambda}}^{*}\right\|$ must also be an upper bound on $D_{1 \rightarrow 2}$.

## Surface distance bound

First, assume that the primary Poisson-estimate is given. Equating terms in eq. (6.3) and eq. (6.5) and solving for the parameter yields:

Figure 6.4:
Relative displacement for equating a single term of the likelihoodfunctions. Solid: primary estimate on the Poissonsurface (undefined for $\lambda_{i} \leq 0$ ).
Dashed: relative displacement from a primary estimate on the $L S$-surface.

$$
\begin{equation*}
\check{\lambda}_{i}^{*}=b_{i}+\operatorname{sgn}\left(\hat{\lambda}_{i}-b_{i}\right) \sqrt{-2 b_{i} \ell_{\mathrm{P}}\left(\hat{\lambda}_{i}, b_{i}\right)} \tag{6.8}
\end{equation*}
$$

Plotting $\Delta_{i} / \ell_{\mathrm{P}}\left(\hat{\lambda}_{i}, b_{i}\right)$, where $\Delta_{i}=\check{\lambda}_{i}^{*}-\hat{\lambda}_{i}$, reveals an important property of the displacement: $\ell_{\mathrm{P}}\left(\hat{\lambda}_{i}, b_{i}\right) \leq \Delta_{i} \leq 0$. Fig. 6.4 shows only $b_{i}=0$ and $b_{i}=1$, since for $b_{i} \neq 1$ the displacement is a function of $\lambda_{i} / b_{i}$. Thus, for $b_{i} \geq 1$ one should read the $\lambda$-axis as $\lambda_{i} / b_{i}$ and use the curve for $b_{i}=1$.

Shift relative to the log-likelihood


Using the observation above and summing over all projections immediately gives the bound:

$$
\left\|\check{\lambda}^{*}-\hat{\lambda}\right\|_{1} \leq-\ell_{\mathrm{P}}(\hat{\lambda}, \mathbf{b})=-\ell_{\mathrm{WLS}}\left(\check{\lambda}^{*}, \mathbf{b}\right)
$$

Analogously, equating terms in eq. (6.3) and eq. (6.6) bounds the displacement from a point on the LS-surface. The dashed curve in fig. 6.4 shows the result, but it is necessary to solve for $\hat{\lambda}_{i}^{*}$ numerically. The same argument as before applies, and we obtain:

$$
\left\|\check{\lambda}-\hat{\lambda}^{*}\right\|_{1} \leq-\ell_{\mathrm{P}}\left(\hat{\lambda}^{*}, \mathbf{b}\right)=-\ell_{\mathrm{WLS}}(\check{\lambda}, \mathbf{b})
$$

Figure 6.5:
Ratio between displacementbound and statistical variation for the worst-case scenario $\lambda_{i}=\|\mathbf{b}\|_{1} / M$.

It is convenient to combine the equations into a single equation, without specifying the primary surface:

$$
\begin{equation*}
\|\check{\boldsymbol{\lambda}}-\hat{\boldsymbol{\lambda}}\|_{1} \leq D \leq-\ell_{\mathrm{P}}(\hat{\boldsymbol{\lambda}}, \mathbf{b})=-\ell_{\mathrm{WLS}}(\check{\boldsymbol{\lambda}}, \mathbf{b})=-k \tag{6.9}
\end{equation*}
$$

## Comparison with noise

Whether the bound given in eq. (6.9) is acceptable depends on the norm of $\mathbf{b}$ and the remaining errors. Since $\|\mathbf{b}\|_{1}$ is the total number of counts, the result in sec. 6.7 suggests that the error is negligible if $\|\mathbf{b}\|_{1} \gg 0.58 M$, where $M$ again represents the number of projections. However, to be quantitative we can compare the bound on $D$ with the statistical fluctuation $\mathscr{E}\left\{\|\mathbf{b}-\boldsymbol{\lambda}\|_{1}\right\}$, where $\boldsymbol{\lambda}=\mathbf{A x}$ is a true (unknown) fixed parameter vector.

The fluctuation can be arbitrarily small when $\boldsymbol{\lambda} \approx \mathbf{0}$, but then the log-likelihood threshold should also be close to zero. We assume that the reconstruction is consistent in the sense that $\ell_{\mathrm{P}}(\hat{\boldsymbol{\lambda}}, \mathbf{b}) \geq \mathscr{E}\left\{\ell_{\mathrm{P}}(\boldsymbol{\lambda}, \mathbf{b})\right\}(1+\epsilon)$, where the last expectation is over $\mathbf{b}$ and $\epsilon$ is small (see sec. 6.7.) It now follows from eq. (6.9):

$$
\begin{equation*}
q=\frac{\|\check{\boldsymbol{\lambda}}-\hat{\boldsymbol{\lambda}}\|_{1}}{\mathscr{E}\left\{\|\mathbf{b}-\boldsymbol{\lambda}\|_{1}\right\}} \leq \frac{-\ell_{\mathrm{P}}(\hat{\boldsymbol{\lambda}}, \mathbf{b})}{\mathscr{E}\left\{\|\mathbf{b}-\boldsymbol{\lambda}\|_{1}\right\}} \leq \frac{(1+\epsilon) \sum_{i=1}^{M} \mathscr{E}\left\{-\ell_{\mathbf{P}}\left(\lambda_{i}, b_{i}\right)\right\}}{\sum_{i=1}^{M} \mathscr{E}\left\{\left|b_{i}-\lambda_{i}\right|\right\}} \tag{6.10}
\end{equation*}
$$

The worst case occurs for $\lambda_{i}=\|\mathbf{b}\|_{1} / M$, shown in fig. 6.5 , which gives:

$$
\begin{equation*}
q \leq \mathscr{E}\left\{-\ell_{\mathrm{P}}\left(\lambda_{i}, b_{i}\right)\right\} / \mathscr{E}\left\{\left|b_{i}-\lambda_{i}\right|\right\} \tag{6.11}
\end{equation*}
$$



The displacement becomes larger than the fluctuation when $\lambda_{i}$ is very small. Nevertheless, if $\lambda_{i}$ is large for some $i$, the terms associated with very small $\lambda_{i}$ have little effect on the fraction in eq. (6.10). Break-even happens around $\lambda_{i} \approx 0.2464$. We conclude that unless the majority of projections have an expected value less than 0.25 , then the perturbation of the surface made by replacing the Poisson likelihood with the WLS-likelihood is no larger than what should be expected from statistical noise.

### 6.9 Analysis of model-weighted least-squares

The MLS-objective, defined by eq. (6.6), uses an improved estimate of the variance. This function equals the Poisson-function for $b_{i}=0$, and it suffices to study $b_{i}>0$.

Figure 6.6:
Relative displacement between surfaces for $b_{i}>0$ where the dashed curve is from LS to Poisson, i.e. $\tilde{s}_{i}$, and the solid curve is Poisson to LS, i.e. $\hat{s}_{i}$. The relative displacement is less than 0.42 .


Consider the relative displacements $\hat{s}_{i}=\left(\hat{\lambda}_{i}-\tilde{\lambda}_{i}\right) /\left(\hat{\lambda}_{i}-b_{i}\right)$ and $\tilde{s}_{i}=\left(\hat{\lambda}_{i}-\right.$ $\left.\tilde{\lambda}_{i}\right) /\left(\tilde{\lambda}_{i}-b_{i}\right)$. Fig. 6.6 shows both $\hat{s}_{i}$ and $\tilde{s}_{i}$ as functions of $\lambda_{i} / b_{i}$, and again the displacements are bounded.

The observation from fig. 6.6 leads to another bound. Writing $\tilde{\lambda}^{*}=\hat{\lambda}+\mathbf{e}$ or $\hat{\boldsymbol{\lambda}}^{*}=\tilde{\boldsymbol{\lambda}}-\mathbf{e}$ we have $0 \leq e_{i}<0.42 b_{i}$. Thus, $\|\mathbf{e}\|=\|\tilde{\boldsymbol{\lambda}}-\hat{\boldsymbol{\lambda}}\|<0.42\|\hat{\boldsymbol{\lambda}}-\mathbf{b}\|$ measured in any norm.

We conclude that the iso-surfaces are relatively close, in the sense that if the Poisson-surface is close to the measured sinogram, then the distance to the MLSsurface is not much further away. This result holds for any choice of likelihood threshold and any subset of projections.

### 6.10 Summary of the analysis

This chapter analyzed the consequences of replacing a Poisson-likelihood with a Gaussian likelihood. Two simple least-squares objectives were defined as approximations to the Poisson log-likelihood.

We bounded the Hausdorff-distance between the Poisson- and LS-likelihood surfaces in sinogram-space. It was shown that exchanging the Poisson-likelihood with a WLS-objective perturbed the surface less than the expected statistical noise in regions with more than about 0.25 average counts per projection. A modelweighted LS-objective always generated solutions with a bounded relatively distance between the Poisson-solution.

We argued that perturbing the likelihood-surfaces less than the expected statistical noise implies that the reconstructed images are similar in image-space. This was prooved rigorously in the case of Tikhonov-regularization.

The analysis derived loose upper bounds of a worst-case analysis of a greatly simplified scenario. No statistical cancellation was introduced to improve the least-squares errors, and the weighted LS-scheme was restricted to estimating the variance from a single projection. It is surprising that a scheme that simple can achieve any degree of accuracy in areas with fractional average counts.

Even though no proof can be given that applies generally to all penalty-functions, there is ample reason to believe that weighted LS-objectives form very good approximations to the Poisson log-likelihood under practical circumstances. The model-weighted least-squares objective demonstrated that the main reason the

Poisson-likelihood do better in regions with very few counts is the ability to estimate fractional variances more than the particular form of the likelihood.

An important consequence is that there is sound theoretical justification for applying the tools in linear algebra to the problem of reconstructing Poisson-sources.


## Part III

## Numerical Algorithms

$\oplus$
$\oplus$

## Chapter

## Iterative framework

This chapter develops a framework for computing the singular value decomposition. The algorithms are based on matrix-vector products, and operates on parts of the SVD, i.e. without requiring a full SVD to reside in memory at a single time.

The framework consists of a number of building-blocks, or modules. To put this in context, we show how several traditional algorithms can be expressed in terms of these modules. For generality, we allow any operator to be complex.

## All chapters in this part are mathematically or algorithmically oriented.

### 7.1 Partial SVD

The goal of this chapter is to develop modules for approximating the SVD of a matrix, $\mathbf{A}$, without forming $\mathbf{A}$ explicitly. We will only access the matrix through matrix-vector multiplications by $\mathbf{A}$ and $\mathbf{A}^{\mathrm{H}}$, since these operations are often the only practical operations in large-scale problems.

In large-scale problems the SVD must be defined without needing all vectors simultaneously. One way to do this is to define a singular triplet.

Definition 7.1: (Singular triplet) A singular triplet $\left(\mathbf{u}_{i}, \mathbf{v}_{i}, \sigma_{i}\right)$ of a matrix, $\mathbf{A} \in \mathbb{C}^{m \times n}$, consists of two unit vectors, $\mathbf{u}_{i} \in \mathbb{C}^{m \times 1}$ and $\mathbf{v}_{i} \in \mathbb{C}^{n \times 1}$, and a scalar, $\sigma_{i}$, which satisfies the equations:

$$
\begin{align*}
\mathbf{A} \mathbf{v}_{i} & =\sigma_{i} \mathbf{u}_{i}  \tag{7.1a}\\
\mathbf{A}^{\mathrm{H}} \mathbf{u}_{i} & =\sigma_{i}^{*} \mathbf{v}_{i} \tag{7.1b}
\end{align*}
$$

For brevity the term "singular" may be dropped when the context is clear. An exact (singular) triplet satisfies these relations exactly. In contrast a numerical (singular) triplet only satisfies $\mathbf{A} \mathbf{v}_{i} \approx \sigma_{i} \mathbf{u}_{i}$ and $\mathbf{A}^{\mathrm{H}} \mathbf{u}_{i} \approx \sigma_{i}^{*} \mathbf{v}_{i}$. We refer to sec. C6 p. 151 in the appendix for more details about singular triplets.

We will say that two singular triplets are mutually orthogonal if their singular vectors are orthogonal, i.e. $\mathbf{v}_{n}^{\mathrm{H}} \mathbf{v}_{m}=\mathbf{u}_{n}^{\mathrm{H}} \mathbf{u}_{m}=0$ for $n \neq m$. Again the term "mutual" may be dropped, and numerically orthogonal triplets only have approximately orthogonal singular vectors.

It is often advantageous to work on a number of singular triplets at a time. This is most conveniently done by collecting the singular triplets in matrices:

$$
\mathbf{U}=\left[\begin{array}{lll}
\mathbf{u}_{1}, & \mathbf{u}_{2}, & \cdots
\end{array}\right] \quad \mathbf{B}=\operatorname{diag}\left(\sigma_{1}, \sigma_{2}, \cdots\right) \quad \mathbf{V}=\left[\begin{array}{lll}
\mathbf{v}_{1}, & \mathbf{v}_{2}, & \cdots
\end{array}\right]
$$

We shall refer to a collection of orthogonal singular triplets as a partial SVD (PSVD). For reference, we give a formal defintion of the partial SVD.

Definition 7.2: (Partial SVD) A size- $k$ PSVD of a matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ is a triplet of matrices, $(\mathbf{U}, \mathbf{V}, \mathbf{B})$, where $\mathbf{U} \in \mathbb{C}^{m \times k}$ and $\mathbf{V} \in \mathbb{C}^{n \times k}$. Also $\mathbf{U}^{\mathrm{H}} \mathbf{U}=\mathbf{I}$ and $\mathbf{V}^{\mathrm{H}} \mathbf{V}=\mathbf{I}$ and finally the following equations must be satisfied:

$$
\begin{align*}
\mathbf{A V} & =\mathbf{U B}  \tag{7.2a}\\
\mathbf{A}^{\mathrm{H}} \mathbf{U} & =\mathbf{V B}^{\mathrm{H}} \tag{7.2b}
\end{align*}
$$

## Numerical Partial SVD (NPSVD)

A PSVD is either exact or numeric, depending on whether the triplets are exact or numeric. Exact triplets with distinct singular values are always orthogonal (see theorem C14 on p. 152).

Unlike exact triplets, numerical triplets are not automatically mutually orthogonal, and it is necessary to include an error-term in the form of a residual. This leads to the Numerical Partial SVD (NPSVD), which is defined in the box.

To keep the correspondence between individual triplets the matrix $\hat{\mathbf{B}}$ must be square. Note that some texts include part of the residual into $\hat{\mathbf{U}}$ or $\hat{\mathbf{V}}$ instead.

## Numerical Partial SVD (NPSVD)

A size- $k$ NPSVD of a matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ is a triplet of matrices, $(\hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}})$, and residuals, $\mathbf{E}_{\mathbf{U}} \in \mathbb{C}^{m \times k}$ and $\mathbf{E}_{\mathbf{V}} \in \mathbb{C}^{n \times k}$, where $\hat{\mathbf{U}} \in \mathbb{C}^{m \times k}, \hat{\mathbf{V}} \in$ $\mathbb{C}^{n \times k}$, and $\hat{\mathbf{B}} \in \mathbb{C}^{k \times k}$.
The matrices are related by the defining equations:

$$
\begin{gather*}
\mathbf{A} \hat{\mathbf{V}} \triangleq \hat{\mathbf{U}} \hat{\mathbf{B}}+\mathbf{E}_{\mathbf{U}}  \tag{7.3a}\\
\mathbf{A}^{\mathrm{H}} \hat{\mathbf{U}} \triangleq \hat{\mathbf{V}} \hat{\mathbf{B}}^{\mathrm{H}}+\mathbf{E}_{\mathbf{V}} \tag{7.3b}
\end{gather*}
$$

There must be a known constant, $\epsilon_{o}<1$, such that the columns of $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ satisfies $1-\epsilon_{o} \leq\left\|\hat{\mathbf{u}}_{i}\right\|_{2} \leq 1+\epsilon_{o}$ and $1-\epsilon_{o} \leq\left\|\hat{\mathbf{v}}_{i}\right\|_{2} \leq 1+\epsilon_{o}$.
If $\hat{\mathbf{B}}$ is diagonal and $\epsilon_{o} \ll 1$, then the columns of $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ approximates the left and right singular vectors, while the entries in $\hat{\mathbf{B}}$ approximates the singular values.

Ideally, $\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{U}}=\mathbf{I}, \hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}=\mathbf{I}$, and $\hat{\mathbf{B}}$ consist of a diagonal of real, non-negative entries. In this case the columns of $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ approximates left and right singular vectors and the diagonal of $\hat{\mathbf{B}}$ approximates the singular values. The quality of these approximations depends on the residual norms, $\left\|\mathbf{E}_{\mathbf{U}}\right\|$ and $\left\|\mathbf{E}_{\mathbf{V}}\right\|$. If the norms are zero the triplets are exact.

For the purpose of numerical software, we will not require that $\hat{\mathbf{B}}$ is diagonal (this is why we use $\hat{\mathbf{B}}$ instead of $\hat{\boldsymbol{\Sigma}}$ ). It will be shown in sec. 7.2 that we may diagonalize $\hat{\mathbf{B}}$ when needed.

As an aid to explain the theoretical background of the algorithms, it is helpful to name the quantities involved in an NPSVD. We shall use the terms shown in table 7.1. For brevity the residual matrices may be referred to as residuals if the context is clear.

## Numerical considerations

Table 7.1:
Terminology used in the descriptions.

| Quantity | Term |
| :---: | :--- |
| $\hat{\mathbf{B}}$ | Coupling matrix |
| $\hat{\mathbf{U}}, \hat{\mathbf{V}}$ | Left and right singular matrices |
| $\mathbf{E}_{\mathbf{U}}, \mathbf{E}_{\mathbf{V}}$ | Left and right residual matrices |

The goal is to develop numerically robust software that improves an NPSVD iteratively in a way that drives the residual norms to zero. Just as it is preferable to allow the coupling-matrix to be on a non-diagonal form, it is also preferable not to require that $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ have orthogonal columns.

In spite of these relaxations, it is convenient to think of the NPSVD as an approximation of a partial SVD. However, we will be careful to note when a derived quantity relies on a particular property of the singular matrices or the coupling matrix.

The NPSVD is given as a definition where $\hat{\mathbf{U}}, \hat{\mathbf{V}}$, and $\hat{\mathbf{B}}$ usually represents computed quantities (and are decorated with a hat). The residuals, $\mathbf{E}_{\mathbf{U}}$ and $\mathbf{E}_{\mathbf{V}}$, are known only through eq. 7.3 and absorbs the rounding errors.

The chapter derives a number of modules that transform an NPSVD. We use the term module, because "update" (and its analogue sometimes known as "downdate") have traditional meanings. Notationally, we will decorate quantities of a transformed NPSVD with tildes, i.e. $\tilde{\mathbf{U}}, \tilde{\mathbf{V}}, \tilde{\mathbf{B}}, \tilde{\mathbf{E}}_{\mathbf{U}}$, and $\tilde{\mathbf{E}}_{\mathbf{V}}$.

## Interpretation

It is helpful to have an intuitive understanding of the NPSVD, which may be attained by considering a specific problem. In scanner terminology range $(\hat{\mathbf{U}})$ represents a space of sinograms and range $(\hat{\mathbf{V}})$ represents a space of objects, i.e. brains in neurology.

When an object is put into the scanner, its emission concentration may be approximated by a vector in $\mathbf{v} \in \operatorname{range}(\hat{\mathbf{V}})$. This (approximated) object may then be projected to the sinogram as the matrix-vector product $\mathbf{u}=\mathbf{A v}$.

If the residuals in the NPSVD are known, then this multiplication which projects $\mathbf{v}$ into the sinogram $\mathbf{u}$, can be done using the NPSVD as illustrated ${ }^{1}$ in fig. 7.1.

Figure 7.1:
Forward projection of a vector. In general range( $(\hat{\mathbf{U}})$ and range $\left(\mathbf{E}_{\mathbf{U}}\right)$ may overlap. After the projection the component lying in range $\left(\mathbf{E}_{\mathbf{U}}\right)$ prevent a backprojection.


[^9]We can also take a sinogram $\mathbf{u} \in \operatorname{range}(\hat{\mathbf{U}})$ and compute $\mathbf{A}^{\mathrm{H}} \mathbf{u}$, i.e. back-project the sinogram into object space. However, after a projection, or back-projection, the residuals add a component that does not lie in the range of $\hat{\mathbf{U}}$ (or $\hat{\mathbf{V}}$ ), so we can not project again.

If the residual norms are negligible, we can also use the NPSVD to determine the object that projects to a given sinogram, i.e. solve the linear system. Although the NPSVD can be used for this purpose without decomposing it into singular triplets, it is nevertheless a useful way to think about the NPSVD.

### 7.2 Diagonalizing the NPSVD

Eventually we want the coupling-matrix $\hat{\mathbf{B}}$ to be diagonal. Fortunately, $\hat{\mathbf{B}}$ is generally of small dimension (compared to $\mathbf{A}$ ), and it is therefore possible to compute the SVD of $\hat{\mathbf{B}}=\mathbf{U}_{\mathbf{B}} \boldsymbol{\Sigma}_{\mathbf{B}} \mathbf{V}_{\mathbf{B}}^{\mathrm{H}}$. Inserting this into the NPSVD in eq. (7.3) gives:

$$
\begin{align*}
\mathbf{A}\left(\hat{\mathbf{V}} \mathbf{V}_{\mathbf{B}}\right) & =\left(\hat{\mathbf{U}} \mathbf{U}_{\mathbf{B}}\right) \boldsymbol{\Sigma}_{\mathbf{B}}+\mathbf{E}_{\mathbf{U}} \mathbf{V}_{\mathbf{B}}  \tag{7.4a}\\
\mathbf{A}^{\mathrm{H}}\left(\hat{\mathbf{U}} \mathbf{U}_{\mathbf{B}}\right) & =\left(\hat{\mathbf{V}} \mathbf{V}_{\mathbf{B}}\right) \boldsymbol{\Sigma}_{\mathbf{B}}^{\mathrm{H}}+\mathbf{E}_{\mathbf{V}} \mathbf{U}_{\mathbf{B}} \tag{7.4b}
\end{align*}
$$

This is an NPSVD where the coupling-matrix $\boldsymbol{\Sigma}_{\mathbf{B}}$ is diagonal by construction. Note that $\left\|\tilde{\mathbf{E}}_{\mathbf{U}}\right\|=\left\|\mathbf{E}_{\mathbf{U}}\right\|$ and $\left\|\tilde{\mathbf{E}}_{\mathbf{V}}\right\|=\left\|\mathbf{E}_{\mathbf{V}}\right\|$ in any unitarily invariant norm. As long as we can compute the SVD of $\hat{\mathbf{B}}$, there is no need to worry whether it is diagonal.

## NPSVD Diagonalization

Let $\hat{\mathbf{B}}=\mathbf{U}_{\mathbf{B}} \boldsymbol{\Sigma}_{\mathbf{B}} \mathbf{V}_{\mathbf{B}}^{\mathrm{H}}$ be an SVD of $\hat{\mathbf{B}}$. The NPSVD can then be transformed to one with a diagonal coupling-matrix using the following equations:

$$
\begin{align*}
\tilde{\mathbf{U}} & =\hat{\mathbf{U}} \mathbf{U}_{\mathbf{B}} & \tilde{\mathbf{B}}=\boldsymbol{\Sigma}_{\mathrm{B}} & \tilde{\mathbf{V}} \tag{7.5}
\end{align*}=\hat{\mathbf{V}} \mathbf{V}_{\mathrm{B}} .
$$

When $\hat{\mathbf{B}}$ is non-diagonal, the columns of $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ does not approximate singular vectors. Instead range $(\hat{\mathbf{U}})$ and range $(\hat{\mathbf{V}})$ approximate singular subspaces and the singular values of $\hat{\mathbf{B}}$ approximate singular values of $\mathbf{A}$.
One should be warned that even though the singular values of $\hat{\mathbf{B}}$ are accurate approximations to those of $\mathbf{A}$ when the residual norms are small, the dimension of the invariant subspace can only be assessed when $\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{U}} \approx \mathbf{I}$ and $\hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}} \approx \mathbf{I}$. Otherwise several singular values of $\hat{\mathbf{B}}$ might refer to the same singular value of A (see sec. C7 on p. 153).

### 7.3 Updating the NPSVD

We must devise methods for reducing the residual norms in the NPSVD. There are two options at this point: either augment the singular matrices, to better approximate the residuals, or remove a part of the residuals. This is a rather crude explanation, since the operations are connected.

The latter is known as restarting and is discussed in sec. 7.7, while this section pursue the former strategy. Preliminary work showed that little is gained from
augmenting range $(\hat{\mathbf{U}})$ and range $(\hat{\mathbf{V}})$ separately, and we will only consider the double-update:

$$
\tilde{\mathbf{U}}=\left[\begin{array}{ll}
\hat{\mathbf{U}}, & \hat{\mathbf{s}}_{\mathbf{u}}
\end{array}\right] \quad \tilde{\mathbf{V}}=\left[\begin{array}{ll}
\hat{\mathbf{V}}, & \hat{\mathrm{s}}_{\mathbf{v}}
\end{array}\right]
$$

The vectors $\hat{\mathbf{s}}_{\mathbf{u}}$ and $\hat{\mathbf{s}}_{\mathbf{v}}$ represent directions that will become available for approximating the subspaces. With a terminology partly borrowed from [32], we shall refer to these vectors as search vectors.

## Updating the coupling matrix

We motivate the update by considering the case where $\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{U}}=\mathbf{I}$ and $\hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}=\mathbf{I}$. In this case we may determine the matrix $\hat{\mathbf{B}}$ that minimizes the residual norms. The least-squares approximation in theorem $\mathbf{C 8}$ states that $\hat{\mathbf{B}}=\hat{\mathbf{U}}^{+} \mathbf{A} \hat{\mathbf{V}}$ minimizes $\left\|\mathbf{E}_{\mathbf{U}}\right\|$, while $\hat{\mathbf{B}}^{\mathrm{H}}=\hat{\mathbf{V}}^{+} \mathbf{A}^{\mathrm{H}} \hat{\mathbf{U}}$ minimizes $\left\|\mathbf{E}_{\mathbf{V}}\right\|$. However, in this particular case the pseudo-inverses are simply the adjoint matrices, and the two equations for $\hat{\mathbf{B}}$ are really the same:

$$
\hat{\mathbf{B}}=\hat{\mathbf{U}}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{V}}
$$

The updated NPSVD should also satisfy $\tilde{\mathbf{B}}=\tilde{\mathbf{U}}^{\mathrm{H}} \mathbf{A} \tilde{\mathbf{V}}$. Inserting eq. (7.7) leads to:

$$
\tilde{\mathbf{B}}=\left[\begin{array}{cc}
\hat{\mathbf{B}} & \hat{\mathbf{f}}_{v} \\
\hat{\mathbf{f}}_{u}^{\mathrm{H}} & \beta
\end{array}\right]
$$

where $\hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}=\hat{\mathbf{s}}_{\mathbf{u}}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{V}}, \hat{\mathbf{f}}_{\mathbf{v}}=\hat{\mathbf{U}}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}$, and $\beta=\hat{\mathbf{s}}_{\mathbf{u}}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}$.
Even when the NPSVD is non-orthogonal, it remains helpful to keep the update of $\hat{\mathbf{B}}$ in this form. Therefore we use eq. (7.9) as it stands, but revise the updates of $\hat{f}_{\mathrm{u}}, \hat{\mathrm{f}}_{\mathrm{v}}$, and $\beta$ in the general case.

## Updating the residual matrices

The residual matrices can be evaluated directly from their definition in eq. (7.3):

$$
\begin{aligned}
& \tilde{\mathbf{E}}_{\mathbf{U}}=\mathbf{A} \tilde{\mathbf{V}}-\tilde{\mathbf{U}} \tilde{\mathrm{B}} \\
& =\left[\begin{array}{ll}
\mathbf{A} \hat{\mathbf{V}}, & \mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}
\end{array}\right]-\left[\begin{array}{ll}
\hat{\mathbf{U}}, & \hat{\mathbf{s}}_{\mathbf{v}}
\end{array}\right] \tilde{\mathbf{B}} \\
& =\left[(\mathbf{A} \hat{\mathbf{V}}-\hat{\mathbf{U}} \hat{\mathbf{B}})-\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}, \quad \mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}-\hat{\mathbf{U}} \hat{\mathrm{f}}_{\mathbf{v}}-\hat{\mathbf{s}}_{\mathbf{u}} \beta\right]
\end{aligned}
$$

Carrying out the same steps for $\tilde{\mathbf{E}}_{\mathbf{V}}$ gives similarly:

$$
\tilde{\mathbf{E}}_{\mathbf{V}}=\left(\mathbf{A}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}-\hat{\mathbf{V}} \hat{\mathbf{f}}_{\mathbf{u}}-\hat{\mathbf{s}}_{\mathbf{v}} \beta^{*}\right)\left[\begin{array}{llll}
0, & \cdots & 0, & 1
\end{array}\right]+\left[\begin{array}{ll}
\mathbf{E}_{\mathbf{V}}-\hat{\mathbf{s}}_{\mathbf{v}} & \hat{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}},
\end{array} \mathbf{0}\right]
$$

Thus, if we interpret $\hat{s}_{u} \hat{\mathbf{f}}_{\mathrm{u}}^{\mathrm{H}} \approx \mathrm{E}_{\mathrm{U}}$ as a rank one approximation, then the updated residual consists of the difference, $\mathbf{E}_{\mathbf{U}}-\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathrm{f}}_{u}^{\mathrm{H}}$, with an extra column appended. We may write this as an update of the search-vectors, which we scale so the updated search-vector is unit length.

## NPSVD Update

Given rank-one approximation of the residual matrices, $\mathbf{E}_{\mathbf{U}} \approx \hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}$ and $\mathbf{E}_{\mathbf{V}} \approx \hat{\mathbf{s}}_{\mathbf{v}} \hat{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}}$, and a constant, $\beta$, the following is an updated NPSVD:

$$
\begin{array}{rlrl}
\tilde{\mathbf{U}} & =\left[\begin{array}{ll}
\hat{\mathbf{U}}, & \hat{\mathbf{s}}_{\mathbf{u}}
\end{array}\right] & \tilde{\mathbf{V}}=\left[\begin{array}{ll}
\hat{\mathbf{V}}, & \hat{\mathbf{s}}_{\mathbf{v}}
\end{array}\right] \\
\tilde{\mathbf{B}} & =\left[\begin{array}{cc}
\mathbf{B} & \hat{\mathbf{f}}_{\mathbf{v}} \\
\hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}} & \beta
\end{array}\right] & & \\
\tilde{\mathbf{s}}_{\mathbf{u}} & =\left(\begin{array}{ll}
\left.\mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}-\hat{\mathbf{U}} \hat{\mathbf{f}}_{\mathbf{v}}-\hat{\mathbf{s}}_{\mathbf{u}} \beta\right) / \tilde{f}_{u} & \tilde{\mathbf{s}}_{\mathbf{v}}=\left(\mathbf{A}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}-\hat{\mathbf{V}} \hat{\mathbf{f}}_{\mathbf{u}}-\hat{\mathbf{s}}_{\mathbf{v}} \beta^{*}\right) / \tilde{f}_{v} \\
\tilde{\mathbf{f}}_{\mathbf{u}} & =\left[\begin{array}{c}
\mathbf{0} \\
\tilde{f}_{u}
\end{array}\right] \\
\tilde{\mathbf{f}}_{\mathbf{v}}=\left[\begin{array}{c}
\mathbf{0} \\
\tilde{f}_{v}
\end{array}\right] \\
\tilde{\mathbf{E}}_{\mathbf{U}} & =\tilde{\mathbf{s}}_{\mathbf{u}} \tilde{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}+\left[\begin{array}{ll}
\mathbf{E}_{\mathbf{U}}-\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}, & \mathbf{0}]
\end{array}\right] \tilde{\mathbf{E}}_{\mathbf{V}}=\tilde{\mathbf{s}}_{\mathbf{v}} \tilde{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}}+\left[\mathbf{E}_{\mathbf{V}}-\hat{\mathbf{s}}_{\mathbf{v}} \hat{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}},\right. \\
\mathbf{0}
\end{array}\right]
\end{array}
$$

## Choosing search-vectors

The update in eq. (7.15) implies that $\left\|\tilde{\mathbf{E}}_{\mathbf{U}}\right\|_{\mathrm{F}}^{2}=\left\|\mathbf{E}_{\mathbf{U}}-\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}\right\|_{\mathrm{F}}^{2}+\tilde{f}_{u}^{2}\left\|\tilde{\mathbf{s}}_{\mathbf{u}}\right\|_{\mathrm{F}}^{2}$. The goal is to minimize this quantity.

One way to proceed is the Jacobi-Davidson SVD [32]. This approach use a search-vector that keeps the norm small using an iterative procedure based on a preconditioner for A. However, it may not always be possible to find a suitable preconditioner, which makes the scheme unsuitable for a canned routine.

A more pragmatic approach is to minimize the first term by setting $\hat{\mathbf{s}}_{\mathbf{u}}$ as the first left singular vector of $\mathbf{E}_{\mathbf{U}}$ (from the low-rank approximation in theorem C9). Likewise, $\hat{\mathbf{s}}_{\mathbf{v}}$ should be the first left singular vector of $\mathbf{E}_{\mathbf{V}}$. In practice the searchvectors are generated iteratively.

## Rounding errors

It is worthwhile to consider the effect of rounding errors in some detail. If the computed matrix-vector product satisfies $\left\|f\left(\mathbf{A}, \hat{\mathbf{s}}_{\mathbf{v}}\right)-\mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}\right\|_{2} \leq c_{1}\|\mathbf{A}\|_{2} \cdot\left\|\hat{\mathbf{s}}_{\mathbf{v}}\right\|_{2}$ evaluating eq. (7.15) in floating-point leads to:

$$
\left\|\tilde{\mathbf{E}}_{\mathbf{U}}-\tilde{\mathbf{s}}_{\mathbf{u}} \tilde{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}\right\|_{\mathrm{F}} \leq\left\|\mathbf{E}_{\mathbf{U}}-\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}\right\|_{\mathbf{F}}+c_{1}\|\mathbf{A}\|_{2}+\left(c_{2}\|\hat{\mathbf{U}}\|_{2}\right)\left\|\hat{\mathbf{f}}_{\mathbf{v}}\right\|_{2}+\epsilon|\beta|
$$

While $c_{1}$ may differ between applications, $c_{2}$ is usually near machine-epsilon. Now consider running $k$ update steps after an initialization. The term $\|\hat{\mathbf{U}}\|_{2}$ can be dropped since only the last element of $\hat{\mathbf{f}}_{\mathbf{v}}$ is non-zero. Noting that $\hat{\mathbf{f}}_{\mathrm{v}}$ and $\beta$ are embedded in the updated coupling matrix leads to:

$$
\left\|\tilde{\mathbf{E}}_{\mathbf{U}}-\tilde{\mathbf{s}}_{\mathbf{u}} \tilde{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}\right\|_{\mathrm{F}} \leq k c_{1}\|\mathbf{A}\|_{2}+c_{2}\|\tilde{\mathbf{B}}\|_{\mathrm{F}}
$$

In practical algorithms the singular values of $\hat{\mathbf{B}}$ approximates those of $\mathbf{A}$, and we may consider $\left\|\tilde{\mathbf{E}}_{\mathbf{U}}-\tilde{\mathbf{s}}_{\mathbf{u}} \tilde{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}\right\|_{\mathrm{F}} \leq k\left(c_{1}+c_{2}\right)\|\mathbf{A}\|_{2}$ to be an upper bound. This implies that the quality of estimated singular triplets is limited by the matrix-vector product, not accumulation of rounding errors.

## Algorithms

To illustrate the development, several example implementations will be sprinkled throughout this chapter. It must be emphasized that they are meant for demonstration only (see sec. 1.7 p. 6).

Algorithm 7.1 creates an NPSVD from two starting-vectors, and uses eq. (7.8) for the coupling-matrix.

The update in alg. 7.3 is more complicated. Inspecting eq. (7.13) we see that only the last component is non-zero. Since the majority of work in the update is matrix multiplications with $\mathbf{A}, \hat{\mathbf{U}}$, and $\hat{\mathbf{V}}$ it is important to reduce this work. Therefore the algorithm allows $\hat{f}_{u}$ and $\hat{f}_{v}$ to be scalars, and also returns scalar values.

### 7.4 Updates with reorthogonalization

Most software packages keep the columns of the singular matrices $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ (nearly) orthogonal during the update. This material would be incomplete without

## Algorithm 7.1: Initialize NPSVD

Create an NPSVD from two start-vectors.

```
function \(\left[\tilde{\mathbf{U}}, \tilde{\mathbf{V}}, \tilde{\mathbf{B}}, \tilde{\mathbf{s}}_{\mathbf{u}}, \tilde{\mathbf{s}}_{\mathbf{v}}, \tilde{f}_{u}, \tilde{f}_{v}\right]=\operatorname{init}(\mathbf{A}, \hat{\mathbf{u}}, \hat{\mathbf{v}})\)
    \(\tilde{\mathbf{U}}=\hat{\mathbf{u}} /\|\hat{\mathbf{u}}\|_{2} ; \quad \tilde{\mathbf{V}}=\hat{\mathbf{v}} /\|\hat{\mathbf{v}}\|_{2} ; \quad \quad\) Normalize starting vectors
    \(\mathbf{x}_{\mathbf{u}}=\mathbf{A} * \tilde{\mathbf{V}} ; \quad \mathbf{x}_{\mathbf{v}}=\mathbf{A}^{\mathrm{H}} * \tilde{\mathbf{U}} ; \quad \quad\) Form products
    \(\tilde{\mathbf{B}}=\tilde{\mathbf{U}}^{\mathrm{H}} * \mathbf{x}_{\mathbf{u}} ; \quad \quad\) Coupling-matrix by eq. (7.8)
    \(\left[\tilde{\mathbf{s}}_{\mathbf{u}}, \tilde{f}_{u}\right]=\operatorname{unit}\left(\mathbf{x}_{\mathbf{u}}-\tilde{\mathbf{U}} * \tilde{\mathbf{B}}\right) ; \quad \quad\) Left residual by eq. (7.3a)
    \(\left[\tilde{\mathbf{s}}_{\mathbf{v}}, \tilde{f}_{v}\right]=\operatorname{unit}\left(\mathbf{x}_{\mathbf{v}}-\tilde{\mathbf{V}} * \tilde{\mathbf{B}}^{\mathrm{H}}\right) ; \quad \quad\) Right residual by eq. (7.3b)
```

Algorithm 7.2: $\quad$ Normalize a vector
Auxiliary routine that factorizes a vector, $\mathbf{x}=\alpha \mathbf{y}$, where $\mathbf{y}$ has unit length. The case $\mathbf{x}=\mathbf{0}$ occurs surprisingly often in applications, even in floating-point arithmetics.

| function $[\mathbf{y}, \alpha]=\operatorname{unit}(\mathbf{x})$ |  |
| :--- | :--- | :--- |
| $\alpha=\\|\mathbf{x}\\|_{2} ;$ | Compute norm |
| $\left[\begin{array}{rl}\mathbf{i f} \alpha \equiv 0 & \\ \mathbf{y}=\operatorname{randn}(\operatorname{size}(\mathbf{x})) ; \quad \mathbf{y}=\mathbf{y} /\\|\mathbf{y}\\|_{2} ; & \text { Create random unit vector } \\ \text { else } & \\ \mathbf{y}=\mathbf{x} / \alpha ; & \text { Normalize input vector } \\ \hline\end{array}\right.$ |  |

## Algorithm 7.3: Update NPSVD

Update a NPSVD as described on p. 84. If necessary, the search-vectors are prepended with zeros to avoid excessive work.

[^10]discussing this, even if the mechanism in sec. 7.6 is employed.
Assume that $\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{U}}=\mathbf{I}$ and $\hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}=\mathbf{I}$. It was shown in sec. 7.3.1 the update that minimizes any unitarily invariant norm of the residuals is $\hat{\mathbf{f}}_{\mathbf{u}}=\hat{\mathbf{V}}^{\mathrm{H}} \mathbf{A}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}$, $\hat{\mathbf{f}}_{\mathbf{v}}=\hat{\mathbf{U}}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}$, and $\beta=\hat{\mathbf{s}}_{\mathbf{u}}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}$. Inserting this into the formula for the updated search-vector, eq. (7.13), gives:
\[

$$
\begin{aligned}
\tilde{\mathbf{s}}_{\mathbf{u}} \tilde{f}_{u} & =\mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}-\hat{\mathbf{U}} \hat{\mathbf{f}}_{\mathbf{v}}-\hat{\mathbf{s}}_{\mathbf{u}} \beta \\
& =\mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}-\hat{\mathbf{U}} \hat{\mathbf{U}}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}-\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{s}}_{\mathbf{u}}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}} \\
& =\left(\mathbf{I}-\tilde{\mathbf{U}} \tilde{\mathbf{U}}^{\mathrm{H}}\right) \mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}} \\
\tilde{\mathbf{s}}_{\mathbf{v}} \tilde{f}_{v} & =\left(\mathbf{I}-\tilde{\mathbf{V}} \tilde{\mathbf{V}}^{\mathrm{H}}\right) \mathbf{A}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}
\end{aligned}
$$
\]

If we assume that the search-vectors are orthogonal to the singular matrices, $\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{S}}_{\mathbf{u}}=\hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{v}}=\mathbf{0}$, then it follows that the updated search-vectors are orthogonal to the updated singular matrices, i.e. $\tilde{\mathbf{U}}^{\mathrm{H}} \tilde{\mathbf{S}}_{\mathbf{u}}=\tilde{\mathbf{V}}^{\mathrm{H}} \tilde{\mathbf{S}}_{\mathbf{v}}=\mathbf{0}$. Unfortunately, this holds only in exact arithmetic; a major cause of algorithmic inefficiency.
Instead an algorithm can enforce the property by orthogonalizing the updated search-vectors against the singular matrices. Informally, orthogonalization removes a component of the search-vector, say $\hat{\mathbf{s}}_{\mathbf{u}}$, in the range of $\hat{\mathbf{U}}$ which is a result of rounding errors.

An algorithm that orthogonalizes the search-vectors every time they are updated is said perform "full reorthogonalization". This is sometimes used to approximate the behavior of an algorithm that runs in exact arithmetic [15].

It becomes expensive to use full reorthogonalization if the iteration is allowed to run for many steps. Instead it is possible to reorthogonalize either occasionally with all vectors, or with some previous search-vectors at each step.

This strategy is known as selective reorthogonalization. Usually the vectors are kept semi-orthogonal so that $\left|\hat{\mathbf{s}}_{\mathbf{u}}^{\mathrm{H}} \hat{\mathbf{u}}_{i}\right| \leq \sqrt{\epsilon_{M}}$ for $i=1, \cdots, k$ [15].

### 7.5 Krylov algorithms

The Lanczos- and the Arnoldi-iterations are classical algorithms for large-scale SVD- and eigenvalue-problems. Each algorithm result from a specific choice of the search-vectors $\hat{\mathbf{s}}_{\mathbf{u}}$ and $\hat{\mathrm{s}}_{\mathbf{v}}$.
We will develop the algorithms by keeping the residual matrices, $\mathbf{E}_{\mathbf{U}}$ and $\mathbf{E}_{\mathbf{V}}$, at rank one during (almost) all steps. As discussed in sec. 7.3.4 an update (i.e. alg. 7.3) returns a useful approximations of the residual matrices, even in the presence of rounding errors. For completeness we mention that there exist blockalgorithms using residuals with higher ranks [26].

Usually the classical algorithms are developed in terms of Krylov subspaces, e.g. [27, 73]. In this case the Krylov subspace range $\left(\mathbf{K}_{n}\right)$, where the Krylov matrix is given by $\mathbf{K}_{n}=\left[\begin{array}{lllll}\mathbf{x}, & \mathbf{A x}, & \mathbf{A}^{2} \mathbf{x}, & \cdots, & \mathbf{A}^{n-1} \mathbf{x}\end{array}\right]$, is taken as an "obvious place" to search.
The idea of minimizing the residual norms is rarely made explicit, although some derivations resembles those in the previous sections superficially, e.g. [67, p.8]. Instead the algorithms are derived by reducing the Krylov-matrix to a prescribed form of the coupling-matrix $\hat{\mathbf{B}}$, as shown in fig. 7.2.

Figure 7.2:
Structure of $\hat{\mathbf{B}}$ generated by the Krylov-algorithms. Only "•"-elements can be nonzero.

 (Lanczos bidiag).

(Arnoldi)

## Lanczos tridiagonalization

The Lanczos iteration attempts to find a partial SVD. It starts with two vectors and attempts to reduce the residual norms by repeatedly updating the NPSVD with the updated search-vectors.

In principle $\beta$ can be chosen freely, and the first idea is probably to set $\beta=0$. However, since the goal is to minimize the residual norms it is better to follow sec. 7.3.1 and use $\beta=\hat{\mathbf{s}}_{\mathbf{u}}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}$. This leads directly to algorithm 7.4.

Note that $\hat{\mathbf{f}}_{\mathbf{u}}$ and $\hat{\mathrm{f}}_{\mathrm{v}}$ only contains a single non-zero entry each, so $\hat{\mathbf{B}}$ is tridiagonal. However, rather than being the starting-point, this structure of the coupling-matrix is a byproduct of the desire to compute singular subspaces.

## Lanczos bidiagonalization

One can do better than keeping both errors at rank one. If we are given $\hat{\mathbf{U}}$ then we can choose $\hat{\mathbf{V}}$ such that range $(\hat{\mathbf{V}})=\operatorname{range}\left(\mathbf{A}^{\mathrm{H}} \hat{\mathbf{U}}\right)$, e.g. using a QR-factorization. This allows us to set $\mathbf{E}_{\mathbf{V}} \approx \mathbf{0}$.
If $\mathbf{E}_{\mathbf{V}} \approx \hat{\mathbf{s}}_{\mathbf{v}} \hat{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}}=\mathbf{0}$, then we may choose $\hat{\mathbf{s}}_{\mathbf{v}}$ freely. Since we want a recurrence where $\tilde{\mathbf{E}}_{\mathbf{V}} \approx \mathbf{0}$, it follows from eq. (7.13) that we should choose $\hat{\mathbf{s}}_{\mathbf{v}} \beta^{*}=$ $\mathbf{A}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}-\hat{\mathbf{V}} \hat{\mathbf{f}}_{\mathbf{u}}$. This immediately leads to algorithm 7.5.
The bidiagonal structure of the coupling-matrix follows from the fact that $\hat{f}_{v}=\mathbf{0}$ and $\hat{\mathbf{f}}_{\mathbf{u}}$ contains a single non-zero element. Keeping $\mathbf{E}_{\mathbf{U}} \approx \hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}=\mathbf{0}$ instead would similarly lead to an upper bidiagonal coupling-matrix.

## Algorithm 7.4: Lanczos tridiagonalization

The Lanczos tridiagonalization algorithm in terms of NPSVD updates.

```
function \(\left[\hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{\mathbf{s}}_{\mathbf{v}}, \hat{f}_{u}, \hat{f}_{v}\right]=\operatorname{tridiag}(\mathbf{A}, \hat{\mathbf{u}}, \hat{\mathbf{v}}, N)\)
    \(\left[\hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{\mathbf{s}}_{\mathbf{v}}, \hat{f}_{u}, \hat{f}_{v}\right]=\operatorname{init}(\mathbf{A}, \hat{\mathbf{u}}, \hat{\mathbf{v}}) ;\)
    for \(n=1: N-1 \quad\) Expand to \(N\) columns
        \(\beta=\hat{\mathbf{s}}_{\mathbf{u}}{ }^{\mathrm{H}} * \mathbf{A} * \hat{\mathbf{s}}_{\mathbf{v}} ; \quad\) Good choice, but could be anything
        \(\left[\hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{\mathbf{s}}_{\mathbf{v}}, \hat{f}_{u}, \hat{f}_{v}\right]=\operatorname{update}\left(\mathbf{A}, \hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{\mathbf{s}}_{\mathbf{v}}, \hat{f}_{u}, \hat{f}_{v}, \beta\right) ;\)
```

Algorithm 7.5: Lanczos bidiagonalization
The Lanczos lower bidiagonalization algorithm in terms of NPSVD updates.

```
function \(\left[\hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{\mathbf{s}}_{\mathbf{v}}, \hat{f}_{u}, \hat{f}_{v}\right]=\operatorname{bidiag}(\mathbf{A}, \hat{\mathbf{u}}, N)\)
    \(\left[\hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{\mathbf{s}}_{\mathbf{v}}, \hat{f}_{u}, \hat{f}_{v}\right]=\operatorname{init}\left(\mathbf{A}, \hat{\mathbf{u}}, \mathbf{A}^{\mathrm{H}} * \hat{\mathbf{u}}\right) ;\)
    for \(n=1: N-1 \quad\) Expand to \(N\) columns
        \(\hat{\mathbf{s}}_{\mathbf{v}}=\mathbf{A}^{\mathrm{H}} * \hat{\mathbf{s}}_{\mathbf{u}}-\hat{\mathbf{V}}(:, n) * \hat{f}_{u} ; \quad\) Search-vector that keep \(\tilde{\mathbf{E}}_{\mathbf{V}}=\mathbf{0}\)
        \(\left[\hat{\mathbf{s}}_{\mathbf{v}}, \alpha\right]=\operatorname{unit}\left(\hat{\mathbf{s}}_{\mathbf{v}}\right) ; \quad \beta=(\alpha)^{*} ; \quad\) Scale
        \(\left[\hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{\mathbf{s}}_{\mathbf{v}}, \hat{f}_{u}, \hat{f}_{v}\right]=\operatorname{update}\left(\mathbf{A}, \hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{\mathbf{s}}_{\mathbf{v}}, \hat{f}_{u}, 0, \beta\right) ;\)
```


## The Arnoldi iteration

The Arnoldi iteration is used in eigenvalue-problems, which is a digression from the topic of this chapter. We include it for completeness and because it is important in practice.

The goal is to find a unitary matrix, $\hat{\mathbf{Q}}$, that (nearly) spans an invariant subspace, i.e. $\mathbf{A} \hat{\mathbf{Q}}=\hat{\mathbf{Q}} \hat{\mathbf{B}}+\mathbf{E}, \mathbf{E} \approx \mathbf{0}$. Inserting a Schur-decomposition of $\hat{\mathbf{B}}$ leads to an approximate Schur-decomposition of $\mathbf{A}$. The hope is to find an acceptable approximation while the dimension of $\hat{\mathbf{B}}$ is small enough to allow this computation.

Comparing with the NPSVD in eq. (7.3) we see that the desired relation is obtained for $\hat{\mathbf{Q}}=\hat{\mathbf{U}}=\hat{\mathbf{V}}$ and $\mathbf{E}_{\mathbf{U}}=\mathbf{E}$. To keep $\tilde{\mathbf{U}}=\tilde{\mathbf{V}}$ it is necessary to set $\hat{\mathbf{s}}_{\mathbf{v}}=\hat{\mathbf{s}}_{\mathbf{u}}$. We obtain a recurrence that produces orthogonal search-vectors by setting $\hat{\mathbf{f}}_{\mathbf{v}}$ and $\beta$ so $\tilde{\mathbf{s}}_{\mathbf{u}} \tilde{f}_{u}=\mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}-\hat{\mathbf{U}} \hat{\mathbf{f}}_{\mathbf{v}}-\hat{\mathbf{s}}_{\mathbf{u}} \beta$ is orthogonal to $\tilde{\mathbf{U}}$.

Algorithm 7.6 implements the idea using Gram-Schmidt orthogonalization. Since $\hat{\mathbf{f}}_{\mathbf{u}}$ only contains a single non-zero entry, it follows that $\hat{\mathbf{B}}$ is upper Hessenberg.

This formulation reveals that $\left\|\mathbf{E}_{\mathbf{V}}\right\|$ is uncontrolled. Since $\left\|\mathbf{E}_{\mathbf{V}}-\hat{\mathbf{s}}_{\mathbf{v}} \hat{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}}\right\|$ can be arbitrarily large, so can $\left\|\tilde{\mathbf{E}}_{\mathbf{V}}\right\|$. As the precision of the eigenvalue problem depends on both left and right eigenvectors [27], this might be a problem. Since the Arnoldi iteration is of no use in SVD computations, this issue will not be pursued further.

001010
100001 111100

## Algorithms

The implementations in alg. 7.4-7.6 are unusual, and at first sight alg. 7.5 and its traditional counterpart, e.g. bidiag 1 in [55], might appear to be different. It is nevertheless trivial to optimize the NPSVD-based algorithms to their traditional version.

Lanczos tridiagonalization, as given by alg. 7.4, is very general. Normally it is applied to Hermitian matrices with a single starting-vector used twice. In this case alg. 7.4 can be considered as a special variant of the Arnoldi iteration.

One should not mistake alg. 7.4 with "unsymmetric Lanczos tridiagonalization" [26]. Although the coupling matrix is updated in a similar way, the search-vectors

## Algorithm 7.6: Arnoldi reduction to upper Hessenberg form.

The Arnoldi iteration written in terms of a NPSVD.

```
function \(\left[\hat{\mathbf{Q}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{f}_{u}\right]=\operatorname{arnoldi}(\mathbf{A}, \mathbf{x}, N)\)
    \(\left[\hat{\mathbf{Q}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{\mathbf{s}}_{\mathbf{v}}, \hat{f}_{u}, \hat{\mathbf{f}}_{\mathbf{v}}\right]=\operatorname{init}(\mathbf{A}, \mathbf{x}, \mathbf{x}) ;\)
    for \(n=1: N-1 \quad\) Iterate
            \(\hat{\mathbf{s}}_{\mathbf{v}}=\hat{\mathbf{s}}_{\mathbf{u}} ; \quad \tilde{\mathbf{s}}_{\mathbf{u}}=\mathbf{A} * \hat{\mathbf{s}}_{\mathbf{v}} ; \quad \quad\) Ensure \(\mathbf{s}_{\mathbf{u}}=\mathbf{s}_{\mathbf{v}}\) and form the product
            for \(j=1: n \quad\) Make \(\tilde{\mathbf{s}}_{\mathbf{u}}\) orthogonal to \(\hat{\mathbf{U}}\)
            \(\hat{\mathbf{f}}_{\mathbf{v}}(j, 1)=\hat{\mathbf{Q}}(:, j)^{\mathrm{H}} * \tilde{\mathbf{s}}_{\mathbf{u}} ;\)
            \(\tilde{\mathbf{s}}_{\mathbf{u}}=\tilde{\mathbf{s}}_{\mathbf{u}}-\hat{\mathbf{Q}}(:, j) * \hat{\mathbf{f}}_{\mathbf{v}}(j, 1) ;\)
            \(\beta=\hat{\mathbf{s}}_{\mathbf{u}}{ }^{\mathrm{H}} * \tilde{\mathbf{s}}_{\mathbf{u}} ; \quad\) Make \(\tilde{\mathbf{s}}_{\mathbf{u}}\) orthogonal to \(\tilde{\mathbf{U}}\)
    \(\tilde{\mathbf{s}}_{\mathbf{u}}=\tilde{\mathbf{s}}_{\mathbf{u}}-\hat{\mathbf{s}}_{\mathbf{u}} * \beta ; \quad\) Needlessly recomputed in update
    \(\left[\hat{\mathbf{Q}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{\mathbf{s}}_{\mathbf{v}}, \hat{f}_{u}, \hat{\mathbf{f}}_{\mathbf{v}}\right]=\operatorname{update}\left(\mathbf{A}, \hat{\mathbf{Q}}, \hat{\mathbf{Q}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{f}_{u}, \hat{\mathbf{f}}_{\mathbf{v}}, \beta\right) ;\)
```

are quite different. While unsymmetric Lanczos is an unstable method for eigenvalue problems, alg. 7.4 is perfectly stable for SVD-problems.

Starting from MATLAB ${ }^{\circledR} 6.0$, the function eigs is based on the Arnoldi iteration implemented in ARPACK [50]. In turn, svds for computing a PSVD, is based on eigs. However, the matrix involved is Hermitian and ARPACK uses Lanczos tridiagonalization, not Arnoldi.

### 7.6 Reorthogonalization

Just as we could restructure $\hat{\mathbf{B}}$ by diagonalizing it, it is also possible to reorthogonalize an NPSVD. Instead of keeping the columns of $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ nearly orthogonal, it will suffice to keep them well-conditioned.
Let $\hat{\mathbf{U}}=\tilde{\mathbf{U}} \mathbf{R}_{\mathbf{U}}$ and $\hat{\mathbf{V}}=\tilde{\mathbf{V}} \mathbf{R}_{\mathbf{V}}$ be QR-decompositions. The coupling-matrix that minimizes the residual norms is given in sec. 7.3.1 on p. 83:

$$
\begin{aligned}
\tilde{\mathbf{B}} & =\tilde{\mathbf{U}}^{\mathrm{H}} \mathbf{A} \tilde{\mathbf{V}} \\
& =\left(\hat{\mathbf{U}} \mathbf{R}_{\mathbf{U}}^{-1}\right)^{\mathrm{H}} \mathbf{A} \hat{\mathbf{V}} \mathbf{R}_{\mathbf{V}}^{-1} \\
& =\mathbf{R}_{\mathbf{U}}^{-\mathrm{H}}\left(\hat{\mathbf{U}}^{\mathrm{H}} \mathbf{A} \mathbf{\mathbf { V }}\right) \mathbf{R}_{\mathbf{V}}^{-1}
\end{aligned}
$$

Inserting $\mathbf{A} \hat{\mathbf{V}}=\hat{\mathbf{U}} \hat{\mathbf{B}}+\mathbf{E}_{\mathbf{U}}$ from eq. (7.3a) yields

$$
\begin{aligned}
\tilde{\mathbf{B}} & =\mathbf{R}_{\mathbf{U}}^{-\mathrm{H}}\left(\hat{\mathbf{U}}^{\mathrm{H}}\left(\hat{\mathbf{U}} \hat{\mathbf{B}}+\mathbf{E}_{\mathbf{U}}\right)\right) \mathbf{R}_{\mathbf{V}}^{-1} \\
& \left.=\mathbf{R}_{\mathbf{U}}^{-\mathrm{H}}\left(\mathbf{R}_{\mathbf{U}}^{\mathrm{H}} \mathbf{R}_{\mathbf{U}} \hat{\mathbf{B}}+\hat{\mathbf{U}}^{\mathrm{H}} \mathbf{E}_{\mathbf{U}}\right)\right) \mathbf{R}_{\mathbf{V}}^{-1} \\
& =\mathbf{R}_{\mathbf{U}} \hat{\mathbf{B}} \mathbf{R}_{\mathbf{V}}^{-1}+\tilde{\mathbf{U}}^{\mathrm{H}} \mathbf{E}_{\mathbf{U}} \mathbf{R}_{\mathbf{V}}^{-1}
\end{aligned}
$$

or, if we insert $\hat{\mathbf{U}}^{\mathrm{H}} \mathbf{A}=\hat{\mathbf{B}} \hat{\mathbf{V}}^{\mathrm{H}}+\mathbf{E}_{\mathbf{V}}^{\mathrm{H}}$ from eq. (7.3b) instead,

$$
\tilde{\mathbf{B}}=\mathbf{R}_{\mathbf{U}}^{-\mathrm{H}} \hat{\mathbf{B}} \mathbf{R}_{\mathbf{V}}^{\mathrm{H}}+\mathbf{R}_{\mathbf{U}}^{-\mathrm{H}} \mathbf{E}_{\mathbf{V}}^{\mathrm{H}} \tilde{\mathbf{V}}
$$

The new residuals follows from their definition in eq. (7.3) by inserting eq. (7.17b) in eq. (7.17a) and eq. (7.18b):

$$
\begin{align*}
& \tilde{\mathbf{E}}_{\mathbf{U}}=\mathbf{A} \tilde{\mathbf{V}}-\tilde{\mathbf{U}} \tilde{\mathbf{B}}=\left(\mathbf{I}-\tilde{\mathbf{U}} \tilde{\mathbf{U}}^{\mathrm{H}}\right) \mathbf{E}_{\mathbf{U}} \mathbf{R}_{\mathbf{V}}^{-1}  \tag{7.18a}\\
& \tilde{\mathbf{E}}_{\mathbf{V}}=\mathbf{A}^{\mathrm{H}} \tilde{\mathbf{U}}-\tilde{\mathbf{V}} \tilde{\mathbf{B}}^{\mathrm{H}}=\left(\mathbf{I}-\tilde{\mathbf{V}} \tilde{\mathbf{V}}^{\mathrm{H}}\right) \mathbf{E}_{\mathbf{V}} \mathbf{R}_{\mathbf{U}}^{-1} \tag{7.18b}
\end{align*}
$$

Assume we have residual approximations $\mathbf{E}_{\mathbf{U}} \approx \hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}$ and $\mathbf{E}_{\mathbf{V}} \approx \hat{\mathbf{s}}_{\mathbf{v}} \hat{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}}$. Adding and subtracting this to the residuals gives:

$$
\begin{aligned}
& \tilde{\mathbf{E}}_{\mathbf{U}}=\left(\mathbf{I}-\tilde{\mathbf{U}} \tilde{\mathbf{U}}^{\mathrm{H}}\right)\left(\mathbf{E}_{\mathbf{U}}-\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}\right) \mathbf{R}_{\mathbf{V}}^{-1}+\left(\mathbf{I}-\tilde{\mathbf{U}}^{\mathrm{H}}\right) \hat{\mathbf{s}}_{\mathbf{u}}\left(\hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}} \mathbf{R}_{\mathbf{V}}^{-1}\right) \\
& \tilde{\mathbf{E}}_{\mathbf{V}}=\left(\mathbf{I}-\tilde{\mathbf{V}} \tilde{\mathbf{V}}^{\mathrm{H}}\right)\left(\mathbf{E}_{\mathbf{V}}-\hat{\mathbf{s}}_{\mathbf{v}} \hat{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}}\right) \mathbf{R}_{\mathbf{U}}^{-1}+\left(\mathbf{I}-\tilde{\mathbf{V}} \tilde{\mathbf{V}}^{\mathrm{H}}\right) \hat{\mathbf{s}}_{\mathbf{v}}\left(\hat{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}} \mathbf{R}_{\mathbf{U}}^{-1}\right)
\end{aligned}
$$

The last term is the transform equations for the search-vectors:

$$
\begin{array}{ll}
\tilde{\mathbf{s}}_{\mathbf{u}}=\left(\mathbf{I}-\tilde{\mathbf{U}} \tilde{\mathbf{U}}^{\mathrm{H}}\right) \hat{\mathbf{s}}_{\mathbf{u}} & \tilde{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}=\hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}} \mathbf{R}_{\mathbf{V}}^{-1} \\
\tilde{\mathbf{s}}_{\mathbf{v}}=\left(\mathbf{I}-\tilde{\mathbf{V}} \tilde{\mathbf{V}}^{\mathrm{H}}\right) \hat{\mathbf{s}}_{\mathbf{v}} & \tilde{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}}=\hat{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}} \mathbf{R}_{\mathbf{U}}^{-1}
\end{array}
$$

Since $\tilde{\mathbf{U}}$ and $\tilde{\mathbf{V}}$ have orthonormal columns, we have the satisfying result that the updated search-vectors are orthogonal to the updated singular matrices.

## Algorithm

A straightforward implementation is given in alg. 7.7. As the update described in sec. 7.3, it also simplifies in the case where $\hat{f}_{u}$ and $\hat{f}_{v}$ have a single non-zero entry. Since $\mathbf{R}_{\mathbf{U}}^{-\mathrm{H}}$ is lower triangular it follows that $\mathbf{R}_{\mathbf{U}}^{-\mathrm{H}} \hat{\mathbf{f}}_{\mathbf{u}}$ also consists a single non-zero entry. In fact we have $\tilde{f}_{u}=\hat{f}_{u} /\left(\mathbf{R}_{\mathbf{U}}\right)_{(k, k)}^{*}$.

### 7.7 Restarting Lanczos

When dealing with large-scale problems the strategies developed so far are insufficient. The main problem is that there is no way to know in advance how many steps are required before the residual norms are sufficiently small.
In fact we are supposed to deal with problems so large that not even the requested singular triplets can be held in memory at a single time. We also want to avoid disk storage, if at all possible.

The solution is to restart the procedure. When we have exhausted the memory it is time to clean up and remove triplets that are of no interest. Conceptually we
form linear combinations of the vectors, $\tilde{\mathbf{u}}=\hat{\mathbf{U}} \mathbf{k}_{\mathbf{u}}$ and $\tilde{\mathbf{v}}=\hat{\mathbf{V}} \mathbf{k}_{\mathbf{v}}$ and then start from scratch with these vectors as starting-vectors.

In principle it should be most flexible to allow $\mathbf{k}_{\mathbf{u}}$ and $\mathbf{k}_{\mathbf{v}}$ to be different. However, in bidiagonalization only one vector is used, and the theory is almost invariably developed by rewriting eq. (7.3) in the following form:

$$
\mathbf{C Q}=\left[\begin{array}{ll}
\mathbf{A}^{\mathrm{H}} & \mathbf{A}
\end{array}\right]\left[\begin{array}{c}
\hat{\mathbf{U}} \\
\hat{\mathbf{V}}
\end{array}\right]=\left[\begin{array}{c}
\hat{\mathbf{U}} \\
\hat{\mathbf{V}}
\end{array}\right] \hat{\mathbf{B}}+\left[\begin{array}{c}
\mathbf{E}_{\mathbf{U}} \\
\mathbf{E}_{\mathbf{V}}
\end{array}\right]
$$

In exact arithmetic range $(\mathbf{Q})$ would form a Krylov-subspace of $\mathbf{C}$ with $\mathbf{q}=\left[\begin{array}{l}\hat{\mathbf{u}} \\ \hat{\mathbf{v}}\end{array}\right]$ as starting-vector. We could therefore write the new starting-vector as $\tilde{\mathbf{q}}=c_{o} \mathbf{q}+$ $c_{1} \mathbf{C q}+\cdots+c_{k} \mathbf{C}^{k} \mathbf{q}$. In other words, the restarting vector is found by multiplying a polynomial in $\mathbf{C}$ with the original starting-vector. By factoring this polynomial we may therefore write:

$$
\tilde{\mathbf{q}}=\left(\prod_{i}\left(\mathbf{C}-\mu_{i} \mathbf{I}\right)\right) \mathbf{q}
$$

In 1994 Åke Björck, Eric Grimme and Pail Van Dooren published a method for restarting the Lanczos bidiagonalization algorithm implicitly [9]. Their method was an adaption of a scheme by Danny C. Sorensen for restarting the Arnoldi method.

The general idea in implicit restarting is that all matrices are transformed to the state they would have had if the new starting vector had been used, but without going through all the steps again. The zeros of the polynomial are used, not the coefficients, in a series of shifted QR-decomposition. We refer to [50] for details.

## The exact shift strategy

A good introduction to different restarting strategies found in real software is [49]. Here we merely recognize that the most successful restarting strategy is the exact

## Algorithm 7.7: Reorthogonalization

Given an NPSVD with well-conditioned singular matrices, this function restore column-orthogonality to full machine precision. The resulting coupling-matrix is full in general.

```
function \(\left[\tilde{\mathbf{U}}, \tilde{\mathbf{V}}, \tilde{\mathbf{B}}, \tilde{\mathbf{s}}_{\mathbf{u}}, \tilde{\mathbf{s}}_{\mathbf{v}}, \tilde{\mathbf{f}}_{\mathbf{u}}, \tilde{\mathbf{f}}_{\mathbf{v}}\right]=\operatorname{orthogonalize}\left(\hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \mathbf{s}_{\mathbf{u}}, \mathbf{s}_{\mathbf{v}}, \mathbf{f}_{\mathbf{u}}, \mathbf{f}_{\mathbf{v}}\right)\)
    \(\left[\tilde{\mathbf{U}}, \mathbf{R}_{\mathbf{U}}\right]=q r(\hat{\mathbf{U}}, 0) ; \quad\left[\tilde{\mathbf{V}}, \mathbf{R}_{\mathbf{V}}\right]=q r(\hat{\mathbf{V}}, 0) ;\)
    \(\mathbf{x}_{\mathbf{u}}=\tilde{\mathbf{U}}^{\mathrm{H}} * \mathbf{s}_{\mathbf{u}} ; \quad \tilde{\mathbf{s}}_{\mathbf{u}}=\mathbf{s}_{\mathbf{u}}-\tilde{\mathbf{U}} * \mathbf{x}_{\mathbf{u}} ;\)
    \(\mathbf{x}_{\mathbf{v}}=\tilde{\mathbf{V}}^{\mathrm{H}} * \mathrm{~s}_{\mathrm{v}} ; \quad \tilde{\mathbf{s}}_{\mathbf{v}}=\mathbf{s}_{\mathbf{v}}-\tilde{\mathbf{V}} * \mathbf{x}_{\mathrm{v}} ;\)
    \(l_{u}=\operatorname{size}(\hat{\mathbf{B}}, 1)-\operatorname{length}\left(\mathbf{f}_{\mathbf{u}}\right) ; \quad \mathbf{z}_{\mathbf{u}}=\operatorname{zeros}\left(l_{u}, 1\right) ; \quad\) Form implicit zeros in \(\mathbf{f}_{\mathbf{u}}\)
    \(l_{v}=\operatorname{size}(\hat{\mathbf{B}}, 2)-\operatorname{length}\left(\mathbf{f}_{\mathbf{v}}\right) ; \quad \mathbf{z}_{\mathbf{v}}=\operatorname{zeros}\left(l_{v}, 1\right) ;\)
    \(\tilde{\mathbf{f}}_{\mathbf{u}}=\left(\left(\mathbf{R}_{\mathbf{V}}\right)^{-1}\right)^{\mathrm{H}} *\left[\mathbf{z}_{\mathbf{u}} ; \mathbf{f}_{\mathbf{u}}\right] ;\)
    \(\tilde{\mathbf{f}}_{\mathbf{v}}=\left(\left(\mathbf{R}_{\mathbf{U}}\right)^{-1}\right)^{\mathrm{H}} *\left[\mathbf{z}_{\mathbf{v}} ; \mathbf{f}_{\mathbf{v}}\right] ;\)
    if \(1, \quad \tilde{\mathbf{B}}=\mathbf{R}_{\mathbf{U}} * \hat{\mathbf{B}} *\left(\mathbf{R}_{\mathbf{V}}\right)^{-1}+\mathbf{x}_{\mathbf{u}} * \tilde{\mathbf{f}}_{\mathbf{u}}{ }^{H}\);
Update \(\mathrm{f}_{\mathrm{v}}\)
    Update \(\hat{B}\) using eq. (7.17a)
    else, \(\quad \tilde{\mathbf{B}}=\left(\left(\mathbf{R}_{\mathbf{U}}\right)^{-1}\right)^{\mathrm{H}} * \hat{\mathbf{B}} * \mathbf{R}_{\mathbf{V}}{ }^{\mathrm{H}}+\tilde{\mathbf{f}}_{\mathbf{v}} * \mathbf{x}_{\mathbf{v}}{ }^{\mathrm{H}}\)
    \(\left[\tilde{\mathbf{s}}_{\mathbf{u}}, \alpha\right]=\operatorname{unit}\left(\tilde{\mathbf{s}}_{\mathbf{u}}\right) ; \quad \tilde{\mathbf{f}}_{\mathbf{u}}=\tilde{\mathbf{f}}_{\mathbf{u}} * \alpha ; \quad \quad\) Normalize
    \(\left[\tilde{\mathbf{s}}_{\mathbf{v}}, \alpha\right]=\operatorname{unit}\left(\tilde{\mathbf{s}}_{\mathbf{v}}\right) ; \quad \tilde{\mathbf{f}}_{\mathbf{v}}=\tilde{\mathbf{f}}_{\mathbf{v}} * \alpha ; \quad \quad\) Normalize
```

shift strategy used in ARPACK. Other strategies, such as shifting at Leja-points or roots of a Chebyshev polynomial, may also be useful [13].

The exact shift strategy has the following appealing interpretation when run in exact arithmetic [50, p. 61-62]:

1. Diagonalize the coupling-matrix, $\hat{\mathbf{B}}$, as described in sec. 7.2.
2. Sort the diagonal into disjoint sets of "wanted" and "unwanted" values.
3. Form a linear combination, $\tilde{\mathbf{q}}=\tilde{\mathbf{Q}} \mathbf{k}$, of the columns in the wanted set.
4. Begin all over with $\tilde{\mathbf{q}}$ as the starting-vector, and update once for each entry in the wanted set.

The spectrum of the restarted coupling-matrix now consists of the wanted set.

### 7.8 Restarting the NPSVD

If $\hat{\mathbf{B}}$ is diagonal, then deleting a column from $\hat{\mathbf{U}}, \hat{\mathbf{V}}, \mathbf{E}_{\mathbf{U}}, \mathbf{E}_{\mathbf{V}}$, and both column and row from $\hat{\mathbf{B}}$ is still a valid NPSVD. Comparing with the description in the previous section, we see that deleting from a diagonal NPSVD is essentially equivalent to implicit restarting with the exact shift strategy.

The diagonalization itself was derived in sec. 7.2. As in the previous sections we extend this to the case where we are given residual approximations $\mathbf{E}_{\mathbf{U}} \approx \hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}$ and $\mathbf{E}_{\mathbf{V}} \approx \hat{\mathbf{s}}_{\mathbf{v}} \hat{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}}$. Adding and subtracting these approximations in eq. (7.6) leads to:

$$
\begin{aligned}
& \tilde{\mathbf{E}}_{\mathbf{U}}=\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}} \mathbf{V}_{\mathbf{B}}+\left(\mathbf{E}_{\mathbf{U}}-\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}\right) \mathbf{V}_{\mathbf{B}} \\
& \tilde{\mathbf{E}}_{\mathbf{V}}=\hat{\mathbf{s}}_{\mathbf{v}} \hat{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}} \mathbf{U}_{\mathbf{B}}+\left(\mathbf{E}_{\mathbf{V}}-\hat{\mathbf{s}}_{\mathbf{v}} \hat{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}}\right) \mathbf{U}_{\mathbf{B}}
\end{aligned}
$$

Apart from giving equations for the transformed search-vectors we also see that no amplification of previous rounding errors occurs (measured in any unitarily invariant norm).

## Algorithm

Instead of deleting columns of $\tilde{\mathbf{U}}$ we may delete columns of $\mathbf{U}_{\mathbf{B}}$ before applying the product, as shown in algorithm 7.8. This is a significantly improvement in large-scale applications.

One significant attribute is that the majority of work is expressed in the form of matrix-multiplications, which is available at peak-performance on modern computers. In fact, orthogonalization, alg. 7.7, and restart can be fused into a single efficient operation. This will be pursued in the applications in the comming chapters.

### 7.9 References

Krylov methods are very general, and the subject is treated in numerous places. Two excellent examples are books written by Trefethen and Bau [73] and Demmel [15].

## NPSVD Restart

Let $\mathbf{E}_{\mathbf{U}} \approx \hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}$ and $\mathbf{E}_{\mathbf{V}} \approx \hat{\mathbf{s}}_{\mathbf{v}} \hat{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}}$ be rank-1 approximations, and $\hat{\mathbf{B}}=$ $\mathbf{U}_{\mathbf{B}} \boldsymbol{\Sigma}_{\mathbf{B}} \mathbf{V}_{\mathbf{B}}^{\mathrm{H}}$ an SVD. Then the restarted NPSVD is given by:

$$
\begin{array}{rlr}
\tilde{\mathbf{U}}=\hat{\mathbf{U}} \mathbf{U}_{\mathbf{B}} \quad \tilde{\mathbf{B}}=\boldsymbol{\Sigma}_{\mathbf{B}} & \tilde{\mathbf{V}}=\hat{\mathbf{V}} \mathbf{V}_{\mathbf{B}} \\
\tilde{\mathbf{s}}_{\mathbf{u}}=\hat{\mathbf{s}}_{\mathbf{u}} & \tilde{\mathbf{s}}_{\mathbf{v}}=\hat{\mathbf{s}}_{\mathbf{v}} \\
\tilde{\mathbf{f}}_{\mathbf{u}}=\mathbf{V}_{\mathbf{B}}^{\mathrm{H}} \hat{\mathbf{f}}_{\mathbf{u}} & \tilde{\mathbf{f}}_{\mathbf{v}}=\mathbf{U}_{\mathbf{B}}^{H} \hat{\mathbf{f}}_{\mathbf{v}} \\
& \tilde{\mathbf{E}}_{\mathbf{U}}=\tilde{\mathbf{s}}_{\mathbf{u}} \tilde{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}+\left(\mathbf{E}_{\mathbf{U}}-\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}\right) \mathbf{V}_{\mathbf{B}} \\
& \tilde{\mathbf{E}}_{\mathbf{V}}=\tilde{\mathbf{s}}_{\mathbf{v}} \tilde{\mathbf{f}}_{\mathbf{v}}^{\mathrm{H}}+\left(\mathbf{E}_{\mathbf{V}}-\hat{\mathbf{s}}_{\mathbf{v}} \hat{\mathbf{f}}_{\mathbf{v}}^{H}\right) \mathbf{U}_{\mathbf{B}}
\end{array}
$$

Now split the singular values of $\hat{\mathbf{B}}$ into a "wanted" and "unwanted" set, and delete the following entries that corresponds to the unwanted set:

- Entries in $\tilde{\mathbf{s}}_{\mathbf{u}}$ and $\tilde{\mathbf{s}}_{\mathbf{v}}$,
- Columns in $\tilde{\mathbf{U}}$ and $\tilde{\mathbf{V}}$,
- Columns and rows in $\tilde{\mathbf{B}}$,

Note that the columns in $\tilde{\mathbf{E}}_{\mathbf{U}}$ and $\tilde{\mathbf{E}}_{\mathbf{V}}$ are deleted implicitly.

For the bidiagonalization, one of the most complete references is the paper where Paige and Saunders develops the iterative solver known as LSQR [55]. A thorough analysis of the effect of rounding errors in the case of partial reorthogonalization can be found in the PhD thesis by Rasmus Munk Larsen [46].

Implicit restarting was originally developed for the Arnoldi-iteration by Danny C. Sorensen [66]. The paper was expanded into a more explanatory technical report [67].

The restarting scheme is implemented in the highly successful software-package ARPACK written by Richard Bruno Lehoucq et. al. [50]. Those who are interested in restarting strategies in real software should also consult his survey in [49].

Exact restarting is the most successful strategy used in current software. However,

## Algorithm 7.8: Restart NPSVD

Example of a simple NPSVD restart, assuming the singular matrices have orthonormal columns (e.g. returned by alg. 7.7). The $N$ approximate triplets with singular vectors nearest $\alpha$ is chosen for the restart.

```
function \(\left[\tilde{\mathbf{U}}, \tilde{\mathbf{V}}, \tilde{\mathbf{B}}, \tilde{\mathbf{s}}_{\mathbf{u}}, \tilde{\mathbf{s}}_{\mathbf{v}}, \tilde{f}_{u}, \tilde{f}_{v}\right]=\operatorname{restart}\left(\hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \mathbf{s}_{\mathbf{u}}, \mathbf{s}_{\mathbf{v}}, \mathbf{f}_{\mathbf{u}}, \mathbf{f}_{\mathbf{v}}, \alpha, N\right)\)
    \(\mathbf{f}_{\mathbf{u}}=\left[\operatorname{zeros}\left(\operatorname{size}(\hat{\mathbf{B}}, 1)-\operatorname{length}\left(\mathbf{f}_{\mathbf{u}}\right), 1\right) ; \mathbf{f}_{\mathbf{u}}\right] ; \quad \quad\) Prepend implicit zeros in \(\mathbf{f}_{\mathbf{u}}\)
    \(\mathbf{f}_{\mathbf{v}}=\left[\operatorname{zeros}\left(\operatorname{size}(\hat{\mathbf{B}}, 2)-\operatorname{length}\left(\mathbf{f}_{\mathrm{v}}\right), 1\right) ; \mathbf{f}_{\mathbf{v}}\right]\);
    \(\left[\mathbf{U}_{\mathbf{B}}, \boldsymbol{\Sigma}_{\mathbf{B}}, \mathbf{V}_{\mathbf{B}}\right]=\operatorname{svd}(\hat{\mathbf{B}}, 0) ; \quad \boldsymbol{\sigma}=\operatorname{diag}\left(\boldsymbol{\Sigma}_{\mathbf{B}}\right)\);
    \([\varnothing, i n x]=\operatorname{sort}(|\boldsymbol{\sigma}-\alpha|) ; \quad\) inx \(=\operatorname{inx}(1: N) ;\)
    \(\tilde{f}_{u}=\mathbf{V}_{\mathbf{B}}{ }^{\mathrm{H}} * \mathbf{f}_{\mathbf{u}} ; \quad \tilde{f}_{u}=\tilde{f}_{u}(i n x) ; \quad \tilde{\mathbf{s}}_{\mathbf{u}}=\mathbf{s}_{\mathbf{u}} ;\)
    \(\tilde{f}_{v}=\mathbf{U}_{\mathbf{B}}{ }^{\mathrm{H}} * \mathbf{f}_{\mathbf{v}} ; \quad \tilde{f}_{v}=\tilde{f}_{v}(i n x) ; \quad \tilde{\mathbf{s}}_{\mathbf{v}}=\mathbf{s}_{\mathbf{v}} ;\)
    \(\tilde{\mathbf{B}}=\boldsymbol{\Sigma}_{\mathbf{B}}(i n x, i n x) ;\)
    \(\tilde{\mathbf{U}}=\hat{\mathbf{U}} * \mathbf{U}_{\mathbf{B}}(:, i n x) ; \quad \tilde{\mathbf{V}}=\hat{\mathbf{V}} * \mathbf{V}_{\mathbf{B}}(:, i n x) ;\)
```

Prepend implicit zeros in $\mathbf{f}_{\mathbf{u}}$
And for $f_{v}$
Compute the SVD of $\hat{\mathbf{B}}$
Find wanted set as entries closest to $\alpha$
Update $\tilde{\mathbf{E}}_{\mathbf{U}}$, search-vector is unmodified
Same for $\tilde{\mathbf{E}}_{\mathbf{V}}$
The coupling-matrix was given by the SVD
Update remaining matrices
other schemes are possible, with Leja points being one of the most interesting [13].

Those who want to use Krylov methods without too much thinking should be warned that it is neccessary to pay extreme attention to details in floating point arithmetics. Therefore, it is best to use canned routines written by experts. Many excellent routines are available, e.g. in MATLAB ${ }^{\circledR}$ or ARPACK.

Those who wants to implement Krylov algorithms anyway may start with the excellent book by Claerbout [14], and the book of template algorithms by Barett et. al. [4].

## Computing the SVD

This chapter applies the framework laid out in chapter 7 to compute large-scale SVD problems. A practical implementation faces a number of issues that must be dealt with.

First, we give a general description of high-performance implementations and strategies. This is followed by the basic iteration that forms the core of the SVD computations. After dealing with several practical problems, such as triplet aliasing, we give several deflation strategies.

Finally, these results are combined into the sliding window strategy.

### 8.1 On the meaning of "large-scale"

The term "large-scale" is a somewhat fluid concept and often depends on context. Sometimes the term is used for any problem whose dimensions are so large that direct methods are unpractical.

We will use the term for systems that are so large that only a fraction of the singular triplets will fit in main memory of the computer. If this "fraction" is small enough, the computational cost of several of the modules in the previous chapter are severely affected.

Consider a problem represented by a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}, m \geq n$, where only $K$ triplets can be held in memory. Assume the SVD of the $k \times k$ coupling-matrix, $\hat{\mathbf{B}}$, is needed at some point.

A general SVD algorithm runs in $\mathcal{O}\left(k^{3}\right)$ time while the matrix-matrix multiplication $\hat{\mathbf{U}} \mathbf{Q}, \mathbf{Q} \in \mathbb{R}^{k \times k}$ runs in $\mathcal{O}\left(m k^{2}\right)$ time. Although the order-factors for the two problems are very different, the multiplication still dominates the time spent if $k / m \leq K / m$ is small enough.

If the problem is very large, the work spent by any $\mathcal{O}\left(k^{3}\right)$ algorithm is negligible compared to operations performed on the singular matrices. Two direct consequences of this should be kept in mind:

1. it is essential to optimize the work performed on the singular matrices,
2. the computational cost is not dominated by work on $k \times k$ matrices.

Generally, it is sound advice to exploit the structure of a matrix. In contrast, the reorthogonalization discussed in sec. 7.6, and implemented in alg. 7.7 p. 91, destroys the structure of $\hat{\mathbf{B}}$. However, the extra cost of working on $\hat{\mathbf{B}}$ is insignificant compared to the advantage gained by collecting the work on the singular matrices.

### 8.2 The memory hierarchy

The memory system of a modern computed is implemented in a number of layers as shown in fig. 8.1. In order to perform a computation rapidly, it is important to utilize the small amount of fast memory available.

Figure 8.1:
Memory hierarchy of a typical modern computer. The small amount of fast memory must be utilized to achieve high performance.


From a practical point of view, this is achieved by blocking the operations. In practice, this means we should perform as much work as possible in highly optimized kernels. In the current context, these kernels come from two sources.

- BLAS (Basic Linear Algebra Subprograms). This is a standardized interface to basic operations. To achieve high performance, the library must be tuned to the computer architecture. A high-quality implementation depends on intricate details of the cache-implementations. An apparently simply routine, such as matrix-matrix multiplication, may use different methods depending on whether a matrix has a prime number of rows or is a power of two.
- LAPACK (Linear Algebra Package). The de-facto software package for small dense problem in linear algebra. It includes routines for a number of factorizations and solving basic problems. The package is built on BLAS, and is optimized for computers with deep cache hierarchies.

The BLAS-library is divided into level $1-3$, which can roughly be characterized in the following way [15].

- Level 1 Performs $\mathcal{O}(k)$ operations. Examples are inner products, vectorscaling, and vector-norms.
- Level 2 Performs $\mathcal{O}\left(k^{2}\right)$ operations. Examples are matrix-vector products, solving a triangular system, and adding a rank-one update to a matrix.
- Level 3 Performs $\mathcal{O}\left(k^{3}\right)$ operations. Generally, these operations are various incarnations of matrix-matrix products.

The level-3 BLAS is important because the amount of computation is large compared to the data that must accessed trhough memory. This makes it possible to exploit the speed of current computers, which is impossible with level-1 and level-2 operations.

Generally, it is wise to use level-3 operations if possible. CPU vendors, such as Intel and IBM, provide highly tuned libraries and using them can easily increase
the speed of an algorithm by a factor of 5-10, depending on the age of the computer. More importantly, this figure is expected to increase at an exponential rate at least until 2016 [63].

Finally, a procedure based on level-3 operations has the ability to scale to several processors for very large problems. Again this is only advantageous if the amount of processing is large compared to the data. This makes it reasonable to transmit the data over a slow link, such as a network, an internal cross-bar switch, or a cache residing on a different CPU.

### 8.3 Left residual factorization

From now on we restrict our attention to algorithms that would keep the right residual $\mathbf{E}_{\mathbf{V}}$ matrix at zero in absence of rounding errors. In other words we set $\mathbf{E}_{\mathbf{V}} \approx \mathbf{0}$ and seek to minimize $\mathbf{E}_{\mathbf{U}} \approx \hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}$, hence the name.

It is convenient to write the residuals in the form $\mathbf{E}_{\mathbf{U}}=\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}+\mathbf{G}_{\mathbf{U}}$ and $\mathbf{E}_{\mathbf{V}}=$ $\mathbf{G}_{\mathbf{V}}$. We can then rewrite the NPSVD in the following form:

$$
\begin{aligned}
\mathbf{A} \hat{\mathbf{V}} & =\hat{\mathbf{U}} \hat{\mathbf{B}}+\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}+\mathbf{G}_{\mathbf{U}} \\
\mathbf{A}^{\mathrm{H}} \hat{\mathbf{U}} & =\hat{\mathbf{V}} \hat{\mathbf{B}}^{\mathrm{H}}+\mathbf{G}_{\mathbf{V}}
\end{aligned}
$$

All rounding errors are represented in $\mathbf{G}_{\mathbf{U}}$ and $\mathbf{G}_{\mathbf{V}}$. We shall refer to these matrices as rounding residual matrices, or just rounding residuals.

## Lanczos and residual factorization

By our definition Lanczos bidiagonalization, described in sec. 7.5 .2 p. 87 , is a left residual factorization. The new term is introduced because the coupling matrix $\hat{\mathbf{B}}$ will not generally be on bidiagonal form.

Ronald B. Morgan gave an elegant proof in [53] that an orthogonalized NPSVD where $\mathbf{E}_{\mathbf{U}}$ is rank one, $\mathbf{E}_{\mathbf{V}}$ is zero, and the singular values of $\hat{\mathbf{B}}$ are distinct corresponds to a bidiagonalization with a certain starting vector. Thus, it does not matter what form the coupling matrix has.

Since the algorithms to be presented in this chapter are similar to the Lanczos algorithms, it might be appropriate to highlight some differences between the framework from the previous chapter and the Lanczos theory as it is developed traditionally.

Accuracy Lanczos algorithms have a striking ability to determine singular values accurately, in spite of rounding errors. This is quite surprising, since the inner loop is a recurrence using only the previous vector. The result is nevertheless fully established in sec. 7.3, where all rounding errors are taken into account. To the best knowledge of the author, this has not been done previously.

In more practical terms, the systematic modularization of the routines is good engineering practice. It allows the same code to be used in several different applications, and makes it easier to construct reliable software. Since all the modules were shown to be stable, it is straightforward to combine them into more complex algorithms and prove stability.

Efficiency Computational efficiency is perhaps best illustrated by delayed orthogonalization. The bulk of work in this module is naturally formulated as matrix-matrix multiplications, which are level-3 operations. However, this advantage can only be gained by abandoning the requirement that the coupling-matrix remains bidiagonal.

Furthermore, the implicit restarting scheme has some advantages not found in previous methods. Apart from the fact that the work is mainly performed as level3 operations, it also has the advantage that computing a converged triplet and removing it is done in the same operation. This fact will be exploited extensively in the algorithms discussed in this chapter.

## Alternatives

We reiterate that the advantages only applies to large-scale problems as described in sec. 8.1. While some optimization could possibly make the algorithms competitive for "medium-scale" problems (in terms of computation time), they are not suitable for small-scale problems such as unstructured dense systems.
When this project was initiated, there was no codes available that combined implicit restarting and partial reorthogonalization. In the meantime, Rasmus Munk Larsen released a Fortran package [47] which extends the MATLAB-code from his PhD [46] with this ability.

It is possible that this approach can be used as an alternative to the full reorthogonalization described in chapter 7. For now, the numerical behavior of the approach has not yet been fully analyzed.

Finally, it is possible that an algorithm based on partial reorthogonalization without restart will suffice for a particular problem. However, although implicitly restarted methods may be slower in a particular case, their ability to scale may justify their use.

### 8.4 Accuracy of matrix-vector product

The algorithms are based on functions for computing the matrix-vector products $\mathbf{A v}$ and $\mathbf{A}^{\mathrm{H}} \mathbf{u}$, where $\mathbf{u}$ and $\mathbf{v}$ are unit vectors. The numerical properties of the algorithms are affected by the accuracy of these functions, and we need to quantify this.

We will assume there is a constant, $\tau$, so the functions satisfies $f_{A}(\mathbf{v})=\mathbf{A} \mathbf{v}+$ $r_{A}(\mathbf{v})$ and $f_{A}^{\mathrm{H}}(\mathbf{u})=\mathbf{A}^{\mathrm{H}} \mathbf{u}+r_{A}^{\mathrm{H}}(\mathbf{u})$ where $\left\|r_{A}(\mathbf{v})\right\|_{2} \leq \tau$ and $\left\|r_{A}^{\mathrm{H}}(\mathbf{u})\right\|_{2} \leq \tau$.

This model is well suited for applications where the rounding errors generated by the functions are insensitive to the input. It implies that the rounding errors are consider to be of the same order irrespective of the norm of the output.

In many cases $\tau=c \epsilon_{M}\|\mathbf{A}\|_{2}$, where $c$ is a modest factor. If it is impossible to analyze the accuracy of the routines, setting $c$ to be a comfortably large constant may be the only practical way to proceed.

## Sparse matrices

For concreteness, we consider the case where $\mathbf{A}$ is sparse, and the matrix-vector product is computed by inner products. It follows from sec. 1.8 that we can set
$\tau=\nu_{l}\|(|\mathbf{A}|)\|_{2}$, where $l$ is the maximum number of nonzero elements in any row or column of $\mathbf{A}$.

Since the 2-norm is not absolute, we can not simply use $\|\mathbf{A}\|_{2}$ in this equation. Instead we may estimate the norm separately, or we can use the bound $\|(|\mathbf{A}|)\|_{2} \leq \sqrt{\|\mathbf{A}\|_{1}\|\mathbf{A}\|_{\infty}}$ [27]. However, such a bound is usually pessimistic since the rounding errors tends to be normally distributed and reaching the bound becomes extremely improbable.

## Asymmetric accuracy

It is possible that one computational routine is more accurate than the adjoint. For instance, consider a sparse matrix with only a few non-zero elements in each row, but some relatively dense columns. We saw in sec. 1.8 that in this case the product Av will usually be computed more accurately than $\mathbf{A}^{\mathrm{H}} \mathbf{u}$.
Another example introducing asymmetry by structure is preconditioning. A preconditioner is often a diagonal matrix which can be applied to high relative accuracy. It has the effect of scaling elements to roughly the same order, and all else being equal, this commonly leads to smaller rounding errors. However, the routine implementing the adjoint operator applies the preconditioner at the end of the computation, which does not improve the accuracy.

Asymmetry also arises for algorithmic reasons. Consider a pair of algorithms that accesses a sparse matrix with non-zero elements stored in column-order. The product $\mathbf{A}^{\mathrm{H}} \mathbf{u}$ can be computed element by element in a tight inner loop. A computer with extended-precision registers can often compute each component very accurately. In contrast, the product $\mathbf{A u}$ is computed by accumulating a vector in memory, and does not benefit much from extended precision.

## Direct measurement

A different approach is to measure $\tau$ based on the functions themselves. If the routines are available in different precisions, it may be possible to estimate $\tau$ by comparing the different variants.
In practice, it is more useful to estimate $\tau$ from a single set of algorithms. This is also useful for checking that the functions $f_{A}$ and $f_{A}^{\mathrm{H}}$ actually implements the adjoint of each other.

Given that the accuracy of the matrix-vector products may be asymmetric, and often have a component whose size depends on the norm of the output, it is by no means a trivial task to estimate $\tau$ reliably. We shall not give a full analysis here, but mention that the routine bbtauest in BBTools works well in many cases.

### 8.5 Basic iteration

The modules derived chapter 7 are given in a general form. This chapter concerns how these modules may be pieced together in the case of left residual factorization.

The basic iteration, or base-iteration, is essentially started as alg. 7.5 p. 87. This algorithm lacks a stopping criterion; the details are worked out in appendix B and a simplified implementation is given in alg. B1 p. 146.

When the iteration completes, we may orthogonalize and restart (alg. 7.8 and 7.7). Although the routines are limited to an exact restarting strategy, it is still necessary to choose the "wanted" and "unwanted" sets described in sec. 7.7. This will be addressed in the following sections.

The restarted, i.e. reduced, NPSVD is then expanded again via the base-iteration, until the desired singular triplets have been extracted. However, when we get to reorthogonalize, only the updated part of the singular matrices needs to be reorthogonalized.

To implement this, it is convenient to include reorthogonalization in the base iteration. The QR-decompositions in alg. 7.7 can be computed efficiently by computing $\mathbf{R}_{\mathbf{U}}$ and $\mathbf{R}_{\mathbf{V}}$ as the Cholesky-factorizations $\mathbf{R}_{\mathbf{U}}^{\mathrm{H}} \mathbf{R}_{\mathbf{U}}=\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{U}}$ and $\mathbf{R}_{\mathbf{V}}^{\mathrm{H}} \mathbf{R}_{\mathbf{V}}=\hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}$. The singular matrices may then be updated by post-multiplying by $\mathbf{R}_{\mathbf{U}}^{-1}$ and $\mathbf{R}_{\mathbf{V}}^{-1}$.

This is not generally a stable approach, but is a viable option in this case where the condition numbers of $\mathbf{R}_{\mathbf{U}}$ and $\mathbf{R}_{\mathbf{V}}$ are kept under strict control. Apart from being a fast way to compute the QR -factorization, it also means that we can update $\hat{\mathbf{B}}$ before $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$.

However, when we know $\hat{\mathbf{B}}$ we may diagonalize it and update $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ as a single matrix-matrix multiplication. Combining the basic updates, reorthogonalization, and diagonalization into the base-iteration reduces the work on the singular matrices to the following operations:

- Update $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ from, say, $k_{0}$ to $k$ columns (alg. B1).
- Compute columns ${ }^{1} k_{0}+1$ to $k$ of the matrices $\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{U}}$ and $\hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}$.
- Fuse reorthogonalization and diagonalization to update of $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ with a single matrix-matrix multiplication for each.

As discussed in sec. 8.1, the bulk of work in large-scale problems is performed in work involving the large matrices. In this case these are the singular matrices, $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$, and the operator $\mathbf{A}$.

The fused procedure above updates the singular matrices mainly with level-3 operations. The expansion itself requires two matrix-vector products, with $\mathbf{A}$ and $\mathbf{A}^{\mathrm{H}}$, to add a column to each singular matrix. It is difficult to imagine this can be improved much further.

For details about the fused implementation we refer to [34], which also includes benchmarks of the process.

### 8.6 Accuracy and aliasing

When the base-iteration completes, we have an NPSVD on the form given by eq. (8.1) where the coupling matrix is diagonal, $\hat{\mathbf{B}}=\operatorname{diag}(\hat{\boldsymbol{\sigma}})$, and the singular matrices have columns that are numerically orthonormal.

The rounding residual matrices, $\mathbf{G}_{\mathbf{U}}$ and $\mathbf{G}_{\mathbf{V}}$, are unknown, but their norms are on the order of $\tau$. They may generally ignored in the restarting strategy; their role served in the analysis of the base iteration.

[^11]Once we have the ortho-diagonal NPSVD, we also have estimates of the singular triplets of A. Ignoring the residual matrices, the quality of the triplets are determined by $\hat{\mathbf{s}}_{\mathbf{u}}$. Specifically, corollary $\mathbf{C} 18 \mathrm{p} .155$ states that $\mathbf{A}$ has a singular value, $\sigma$, in the range $|\hat{\boldsymbol{\sigma}}|_{i}-\left|\hat{\mathbf{s}}_{\mathbf{u}}\right|_{i} \leq \sigma \leq|\hat{\boldsymbol{\sigma}}|_{i}+\left|\hat{\mathbf{s}}_{\mathbf{u}}\right|_{i}$.

Thus, if triplet $i$ satisfies the $\left|\hat{\mathbf{s}}_{\mathbf{u}}\right|_{i} \leq \tau$, then it has converged to the accuracy provided by the matrix-vector product routines. After convergence, a triplet may be removed to make space for new iterates.

## Triplet aliasing

The fallacy of this is that there are many other triplets that may contain useful information, and it would be wasteful at best to throw these away. An in depth discussion of this point is given in sec. C7 on p. 153, and the result is that we should also include all triplets that may alias triplets in the wanted set.

In short, an aliasing triplet is a triplet that may be rich in components corresponding to triplets already chosen for the restart. If we do not include these triplets for the restart, the convergence may be slow because the algorithm is required to "rediscover" these directions. Even worse, we may throw away a triplet with a large components corresponding to a singular value that is smaller than those found already. This is not acceptable, and leads to algorithms that may not converge at all.

Fortunately, it is not hard to determine which triplets that may alias those in the wanted set. All we need is a bound on the singular values corresponding to each triplet. This is supplied by theorem C17 and implemented in alg. 8.1.

## Combating rounding errors

Rounding errors are the bane of efficiency for iterative algorithms. While detailed analysis is necessary to get numerical algorithms stable, an element of art remains as defined by Donald E. Knuth [56]:

Science is what we understand well enough to explain to a computer.
Art is everything else we do.

After a restart, the norm of $\mathbf{E}_{\mathbf{U}}$ may be very small. This is particularly true if the restart discards some triplets with large errors. The result is that the norm of $\mathbf{E}_{\mathbf{U}} \approx \hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}$ may be reduced by a significant factor during the restart.

Algorithm 8.1: Expand a set of triplets to include aliasing triplets.
Given an NPSVD and error-bounds, this routine find intervals for singular values of each triplets. It returns the triplets that aliases those in the wanted set.

```
function inx \(=\operatorname{widen}\left(\operatorname{inx}, \boldsymbol{\sigma}, \hat{\mathbf{f}}_{\mathbf{u}}\right)\)
    \(\sigma_{\text {lo }}=\sigma+a b s\left(\hat{\mathbf{f}}_{\mathbf{u}}\right) ; \quad\) Maximum aliased singular value
    \(\boldsymbol{\sigma}_{\mathrm{hi}}=\max \left(\boldsymbol{\sigma}-a b s\left(\hat{\mathbf{f}}_{\mathbf{u}}\right), 0\right) ; \quad\) Minimum aliased singular value
    \(\operatorname{Max}=\max \left(\boldsymbol{\sigma}_{\mathrm{lo}}(i n x)\right) ; \quad \operatorname{Min}=\min \left(\boldsymbol{\sigma}_{\mathrm{hi}}(i n x)\right) ; \quad\) Find extreme bounds for requested triplets
    \(\operatorname{inx}=\operatorname{find}\left(\left(\sigma_{\mathrm{hi}} \leq \operatorname{Max}\right) \wedge\left(\sigma_{\mathrm{lo}} \geq \operatorname{Min}\right)\right) ; \quad\) Find aliasing triplets
```

Although it is known that $\left\|\mathbf{E}_{\mathbf{U}}-\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}\right\|$ is small at the point of the restart, any component in $\hat{\mathbf{s}}_{\mathbf{u}}$ corresponding to singular triplets of $\mathbf{A}$ with large singular values, will be picked up immediately.

One way to improve this is to improve the rank-one approximation by taking a power-step of $\mathbf{E}_{\mathbf{U}}=\mathbf{A} \hat{\mathbf{V}}-\hat{\mathbf{U}} \hat{\mathbf{B}}$. We note that this improvement is empirical, and does not change any error-bounds, and is redundant in the absence of rounding errors.

### 8.7 Singular value decomposition

We are now ready to attack the problem of computing the SVD of a large-scale matrix $\mathbf{A}$. For this we will assume that $\mathbf{A} \in \mathbb{C}^{m \times n}$ where $m \geq n$. Otherwise we may instead compute the SVD of $\mathbf{A}^{\mathrm{H}}$.

Our problem is to compute as many triplets, corresponding to the largest singular values, with a restricted amount of memory. We note that computing the eigendecomposition $\hat{\mathbf{V}} \hat{\boldsymbol{\Sigma}}^{2} \hat{\mathbf{V}}=\mathbf{A}^{\mathrm{H}} \mathbf{A}$ only needs to store vectors of length $n$, which may be a very substantial saving.

This is a viable approach for relatively well-conditioned systems. For problems that does not need singular values below about $\sqrt{\epsilon_{M}}$, and have modest requirements about accuracy, this can be done efficiently using ARPACK [50].

As an example, assume that $n=10000$, wherease $m \gg n$ is large. If the matrix $\hat{\mathbf{V}}$ is stored as IEEE double-precision, then it can be stored in roughly 1.6 GB of memory. We should therefore be able to handle such a system on an off-the-shelf desktop computer, regardless of $m$.

Our goal is to compute the SVD as accurately as possible, without using more resources than the scheme above. We also want to compute triplets beyond those that will fit in memory.

The first thing to notice is that the first few singular triplets can be computed efficiently by restarting. This requires storage of both left and right singular vectors, and after convergence they will be returned to the user for whatever problem the SVD is used to solve.

Since we do not have storage to keep these triplets in memory they must be removed from further processing. However, these triplet will eventually reappear because of rounding errors. Even in the absence of rounding errors they will reappear, unless the removed triplets were exact. This would rarely be the case in any iterative scheme.

### 8.8 Explicit deflation

One solution is to deflate the triplets. Basically, this amount to continue processing on the matrix $\mathbf{A}^{\prime}=\mathbf{A}-\mathbf{Q}_{\mathbf{U}} \mathbf{S Q}_{\mathbf{V}}^{\mathrm{H}}$, where $\mathbf{Q}_{\mathbf{U}}, \mathbf{S}$, and $\mathbf{Q}_{\mathbf{V}}$ represents converged triplets:

$$
\begin{aligned}
\mathbf{A} \mathbf{Q}_{\mathbf{V}} & =\mathbf{Q}_{\mathbf{U}} \mathbf{S}+\mathbf{C}_{\mathbf{U}} \\
\mathbf{A}^{\mathrm{H}} \mathbf{Q}_{\mathbf{U}} & =\mathbf{Q}_{\mathbf{V}} \mathbf{S}^{\mathrm{H}}+\mathbf{C}_{\mathbf{V}}
\end{aligned}
$$

According to the Schmidt-Mirsky theorem (see p. 150) this minimizes the norm of $\mathbf{A}^{\prime}$, if the singular triplets are accurate. We shall refer to this as two-sided deflation, since both the left and right singular triplets are explicitly kept orthogonal to converged triplets.
Unfortunately, two-sided deflation stores all the converged triplets and does not save memory. Nevertheless, the scheme remains useful as a component of a larger strategy where only a few components are stored.
As an alternative we may turn to one-sided deflation:

$$
\begin{aligned}
\mathbf{A}^{\prime} & =\mathbf{A}\left(\mathbf{I}-\mathbf{Q}_{\mathbf{V}} \mathbf{Q}_{\mathbf{V}}^{\mathrm{H}}\right) \\
& =\mathbf{A}-\left(\mathbf{Q}_{\mathbf{U}} \mathbf{S}+\mathbf{C}_{\mathbf{U}}\right) \mathbf{Q}_{\mathbf{V}}^{\mathrm{H}} \\
& =\mathbf{A}-\mathbf{Q}_{\mathbf{U}} \mathbf{S} \mathbf{Q}_{\mathbf{V}}^{\mathrm{H}}-\mathbf{C}_{\mathbf{U}} \mathbf{Q}_{\mathbf{V}}^{\mathrm{H}}
\end{aligned}
$$

This is almost as good, since $\left\|\mathbf{C}_{\mathbf{U}}\right\|$ is less than the tolerance we require for the singular triplets.
It is necessary to ensure that changing the matrix does not disturb the restart. We note that before restarting all triplets have been orthogonalized against each other. It follows that all non-discarded triplets could have been computed using $\mathbf{A}^{\prime}$ instead of $\mathbf{A}$. It is therefore stable to move converged vectors into $\mathbf{Q}_{\mathbf{V}}$ before restarting.

However, to keep the updates consistent we are required to compute all further matrix-vector products using $\mathbf{A}^{\prime}$ instead of $\mathbf{A}$. Otherwise the NPSVD will be inconsistent, and unpredictable results may occur. While this represents a significant amount of work in level- 2 operations, it is acceptable if $n \ll m$.

The second point is that we must ensure the accuracy of the remaining triplets. So if $\hat{\mathbf{U}}, \hat{\mathbf{B}}$, and $\hat{\mathbf{V}}$ is an NPSVD of $\mathbf{A}^{\prime}$ we get:

$$
\begin{aligned}
\mathbf{A}\left(\mathbf{I}-\mathbf{Q}_{\mathbf{V}} \mathbf{Q}_{\mathbf{V}}^{\mathrm{H}}\right) \hat{\mathbf{V}} & =\hat{\mathbf{U}} \hat{\mathbf{B}}+\mathbf{E}_{\mathbf{U}} \\
\left(\mathbf{I}-\mathbf{Q}_{\mathbf{V}} \mathbf{Q}_{\mathbf{V}}^{\mathrm{H}}\right) \mathbf{A}^{\mathrm{H}} \hat{\mathbf{U}} & =\hat{\mathbf{V}} \hat{\mathbf{B}}^{\mathrm{H}}+\mathbf{E}_{\mathbf{V}}
\end{aligned}
$$

This implies:

$$
\begin{align*}
\mathbf{A} \hat{\mathbf{V}} & =\hat{\mathbf{U}} \hat{\mathbf{B}}+\mathbf{E}_{\mathbf{U}}+\mathbf{A} \mathbf{Q}_{\mathbf{V}} \mathbf{Q}_{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}} \\
& =\hat{\mathbf{U}} \hat{\mathbf{B}}+\mathbf{E}_{\mathbf{U}}+\left(\mathbf{Q}_{\mathbf{U}} \mathbf{S}+\mathbf{C}_{\mathbf{U}}\right)\left(\mathbf{Q}_{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}\right)  \tag{8.3a}\\
\mathbf{A}^{\mathrm{H}} \hat{\mathbf{U}} & =\hat{\mathbf{V}} \hat{\mathbf{B}}^{\mathrm{H}}+\mathbf{E}_{\mathbf{V}}+\mathbf{Q}_{\mathbf{V}} \mathbf{Q}_{\mathbf{V}}^{\mathrm{H}} \mathbf{A}^{\mathrm{H}} \hat{\mathbf{U}} \\
& =\hat{\mathbf{V}} \hat{\mathbf{B}}^{\mathrm{H}}+\mathbf{E}_{\mathbf{V}}+\mathbf{Q}_{\mathbf{V}}\left(\mathbf{S}^{\mathrm{H}} \mathbf{Q}_{\mathbf{U}}^{\mathrm{H}}+\mathbf{C}_{\mathbf{U}}^{\mathrm{H}}\right) \hat{\mathbf{U}} \\
& =\hat{\mathbf{V}} \hat{\mathbf{B}}^{\mathrm{H}}+\mathbf{E}_{\mathbf{V}}+\mathbf{Q}_{\mathbf{V}} \mathbf{C}_{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{U}}+\mathbf{Q}_{\mathbf{V}} \mathbf{S}^{\mathrm{H}}\left(\mathbf{Q}_{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{U}}\right) \tag{8.3b}
\end{align*}
$$

Since we orthogonalize all vectors in $\hat{\mathbf{V}}$ explicitly, we know that $\left\|\mathbf{Q}_{V}^{\mathrm{H}} \hat{\mathbf{V}}\right\|$ is negligible. However, it takes quite a bit of work to get a rigorous analysis of the complete case. The theoretical underpinnings of one-sided deflation is developed in appendix C, and we refer in particular to the orthogonality theorem on p. 152.

From a practical point of view, the result is that one-sided deflation is suitable for the following purposes:

- Accurate singular values (i.e. no singular vectors).
- Low-rank approximations.
- Any problem where orthogonality of the vectors in the long-space is unimportant.


### 8.9 Implicit deflation

Another option is to implicit deflatation, which has far more potential for memory conservation. This is based on theorem C14 p. 152, which states that numerical singular triplets are numerical orthogonal, provided they are accurate and have different singular values.

In principle, we may keep a list of the singular values returned to the user, and if we find a new triplet that correspond to the same singular triplet, then we simply discard it without returning it again. This is attractive, because it keeps the storage requirement independent of both $m$ and $n$, and does not introduce any additional rounding errors. However, it suffers from two drawbacks: some triplet may be missed and a very large number of matrix-vector products with A may be required.

The first issue can be dealt with by observing that triplets converge roughly in decreasing order of singular values. However, if there are clustered triplets it takes many iterations before a single triplet will split and resolve these individually. An elaborate explanation of this phenomena can be found in [73].

We may deal with this by explicitly deflating a number of the most recently converged triplets. This makes it very unlikely that any triplet will be glossed over, but still avoid the need to keep all previous triplets in storage.

The other issue is more problematic, and arises from the fact that rounding errors quickly generate components in directions corresponding to large singular values as discussed in sec. 8.6.2. Unfortunately, there seems to be no other cure than either deflating explicitly or accept a large number of matrix-vector products.

However, there are important applications where implicit deflation is useful. In many applications the singular values of A reaches a plateau. This is typical for applications that depends on low-rank approximations, e.g. latent semantic indexing [6]. It is often interesting to know both whether the singular values eventually starts to drop, and if they do, where. Implicit deflation can answer this question in an inexpensive way.

We may then use implicit deflation as long as the largest implicitly deflated singular value is not too large compared to the remaining triplets. Theoretically, there is no bound but in practice the number of matrix-vector products becomes excessive if the range of implicit deflation is about a factor of 10 .

### 8.10 Sliding window

The procedures described above can be combined into an aggressive restarting strategy, which we shall refer to as "sliding window". The idea is sketeched in fig. 8.2: explicit deflation (one- or two-sided) is used to keep large singular triplets from disturbing the computation (the head). Another set (the tail) is used to create a gap of singular values. This allows a large number of triplets to be deflated implicitly: any triplet with a singular value that exceeds the smallest implicitly

Figure 8.2:
Sketch of a step where a sliding window is used to compute the SVD of a discrete Radon transform.

Singular values of a discrete Radon transform

singular value is removed from the restart (unless it aliases triplets in the sliding window).

The best number of triplets to allocate for the head and tail depends on the problem at hand. The default behavior for the implementation in BBTools is to use a tail for all requested singular triplets. Tuning these to the problem at hand may significantly improve performance, but the other approach is very safe provided the system has enough memory.

### 8.11 References

Most of the material on computer architecture can be found in text-books on numerical linear algebra, e.g. [15] and [69]. Those who wants to implement their own BLAS may appreciate the difficulties by looking at ATLAS (Automatically Tuned Linear Algebra Software) [80].

The de-facto software package for numerical linear algebra with dense systems is LAPACK [2]. There is also a more modern variant, ScaLAPACK, which is intended for parallel computations [10].

Iterative computation of the SVD is supported by the following packages:

- ARPACK [50], described on p. 93, is written in Fortran. It forms the basis of the Matlab-functions eigs and svds.
- SVDPACK [7], by Michael W. Berry, contains several SVD-routines for large-scale computations. The routines, described in [6], are written in Fortran.
- PROPACK [47] includes Lanczos bidiagonalization with partial reorthogonalization. The original report [46] did not contain implicit restarting, but the software have now been expanded. The package exists for both Fortran and Matlab.

It should be noted that PROPACK only includes the bidiagonalization procedure. As should be evident from the contents of this chapter, using it for computing SVDs may require some additional work.
$\oplus$
$\oplus$

## Chapter

## Solving linear systems

This chapter builds on the two previous chapters to produce algorithms that solves linear systems. After treating the unregularized case, we show how regularized solutions, based on filter-factors, can be estimated efficiently.

The chapter finish with a description of a practical hybrid solver that is suitable for solving tomographic problems.

### 9.1 Prelude to iterative solvers

The goal of a linear solver is to determine a vector, $\hat{\mathbf{x}}$, so that $\mathbf{A} \hat{\mathbf{x}} \approx \mathbf{b}$ for a given b. An iterative solver works by approximately solving a system for a residual:

$$
\begin{equation*}
\mathbf{A} \hat{\mathbf{x}}_{i} \approx \mathbf{r}_{i} \tag{9.1}
\end{equation*}
$$

This is initialized with $\mathbf{r}_{0}=\mathbf{b}$, i.e. the original problem. At step $i$ an approximate solution, $\mathbf{A} \hat{\mathbf{y}}_{i} \approx \mathbf{r}_{i}$ is determined. Generally, this estimate may be quite poor, so we update the equation system:

$$
\begin{align*}
\mathbf{r}_{i+1} & =\mathbf{r}_{i}-\mathbf{A} \hat{\mathbf{y}}_{i}  \tag{9.2a}\\
\hat{\mathbf{x}}_{i} & =\sum_{k=0}^{i} \hat{\mathbf{y}}_{k} \tag{9.2b}
\end{align*}
$$

Applying eq. (9.2) repeatedly leads to the expression:

$$
\begin{aligned}
\mathbf{r}_{0} & =\hat{\mathbf{r}}_{1}+\mathbf{A} \mathbf{y}_{0} \\
& =\hat{\mathbf{r}}_{2}+\mathbf{A} \mathbf{y}_{1}+\mathbf{A} \mathbf{y}_{0} \\
& =\hat{\mathbf{r}}_{N}+\sum_{k=0}^{N-1} \mathbf{A} \mathbf{y}_{k} \\
& =\hat{\mathbf{r}}_{N}+\mathbf{A} \hat{\mathbf{x}}_{N-1}
\end{aligned}
$$

If $\left\|\mathbf{r}_{N}\right\|$ is small, then we have determined an accurate solution to the system $\mathbf{A} \hat{\mathbf{x}}=\mathbf{r}_{0}=\mathbf{b}$ as desired. This technique is sometimes used to combat rounding errors for direct solvers, where it is known as iterative refinement [69].

The strength is that the approximate solution may be quite inaccurate, as long as the residual norm is decreasing. To emphasize this, alg. 9.1 implements this by using a random vector to approximate the solution. This is allowed as long as the residual norm decreases in each step, which can be guaranteed by scaling.

Although simple, alg. 9.1 has most of the components necessary for an iterative solver, including a stopping-criterion that suffice if $\mathbf{A x}=\mathbf{b}$ has a solution.

### 9.2 Bidiagonalizing solver

The Lanczos bidiagonalization procedure, alg. 7.5 p .87 , is an excellent base for an iterative solver. We shall use the left residual factorization described in sec. 8.3. More precisely, the routines described in this chapter are based on the basic iteration described in sec. 8.5. Therefore, the effect of rounding-errors will not be addressed here.

Ignoring the rounding residual matrices, the basic iteration result in an NPSVD on the following form:

$$
\begin{equation*}
\mathbf{A} \hat{\mathbf{y}}_{0}=\left(\hat{\mathbf{U}} \hat{\mathbf{B}}+\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}\right) \hat{\mathbf{z}}_{0} \tag{9.6}
\end{equation*}
$$

From this there are two main ways to select $\hat{\mathbf{y}}_{0}$.

## Optimal (MINRES)

The MINRES algorithm (minimal residual) [73] is an algorithm which use Lanczos tridiagonalization on a Hermitian matrix $\mathbf{A}=\mathbf{A}^{\mathrm{H}}$. It has been generalized to the non-Hermitian case, where it is known as GMRES (generalized minimal residual).

Algorithm 9.1: $\quad$ Solve $A x=b$ iteratively using random approximations.
This program is the extreme case of a bad approximate solution. If the system has a solution, the algorithm slowly converges to a solution with arbitrarily small residual norm (ignoring rounding errors).

```
function \([\hat{\mathbf{x}}, \mathbf{r}]=\operatorname{randsolve}(\mathbf{A}, \mathbf{b}\), tol \()\)
    \([m, n]=\operatorname{size}(\mathbf{A}) ; \quad \hat{\mathbf{x}}=\operatorname{zeros}(n, 1) ; \quad \mathbf{r}=\mathbf{b} ; \quad\) Setup system
    \(T=\|\mathbf{b}\|_{2} * t o l ;\)
    while \(\|\mathbf{r}\|_{2}>T\)
        \(\hat{\mathbf{y}}=\operatorname{randn}(n, 1) ; \quad \mathbf{p}=\mathbf{A} * \hat{\mathbf{y}} ; \quad\) Approximate solution by a random (!!!) vector
        \(\alpha=\left(\mathbf{p}^{\mathrm{H}} * \mathbf{r}\right) /\left(\mathbf{p}^{\mathrm{H}} * \mathbf{p}\right) ; \quad \hat{\mathbf{y}}=\hat{\mathbf{y}} * \alpha ; \quad \mathbf{p}=\mathbf{p} * \alpha ; \quad\) Rescale to minimize residual
        \(\mathbf{r}=\mathbf{r}-\mathbf{p} ; \quad \hat{\mathbf{x}}=\hat{\mathbf{x}}+\hat{\mathbf{y}} ; \quad\) Update residual and solution
```

It is based on the idea that $\hat{\mathbf{y}}_{0}$ should be chosen to minimize the residual norm $\left\|\mathbf{r}_{0}\right\|$. From a practical point of view, this can be accomplished by solving the normal-equations on eq. (9.6) using Woodbury's formula [69]. While this is generally unstable, it is stable and efficient if the NPSVD is in ortho-diagonal form.

In exact arithmetic, this procedure is equivalent to MINRES on the normal equations:

$$
\mathbf{A}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{y}}_{0}=\mathbf{A}^{\mathrm{H}} \mathbf{r}_{0}
$$

The main problem with MINRES is that it requires an explicit restart: once the new residual is computed, the NPSVD can not be reused since eq. 9.4 fails to hold. Instead, a new NPSVD must be created starting from the new residual.

## Galerkin (Conjugate Gradient)

The deficiency of a MINRES restart leads us to look for a solution that allows an implicit restart of the NPSVD. This can be accomplished via the Galerkin criteria: $\hat{\mathbf{y}}_{0}$ is chosen so $\hat{\mathbf{r}}_{1}$ is orthogonal to the triplets removed during the restart. This is equivalent to CGNR: the conjugate gradient on the normal-equations in eq. (9.7).

More specifically, consider the case where we want to remove triplet $i$ from the NPSVD. For this to work, the relation in eq. (9.4) must still hold after the restart. To this end consider the following quantity:

By eq. (9.2a)
By eq. (9.6)

$$
\begin{aligned}
\mathbf{u}_{i}^{\mathrm{H}} \mathbf{r}_{1} & =\mathbf{u}_{i}^{\mathrm{H}}\left(\mathbf{r}_{0}-\mathbf{A} \hat{\mathbf{y}}_{0}\right) \\
& =\mathbf{u}_{i}^{\mathrm{H}} \mathbf{r}_{0}-\mathbf{u}_{i}^{\mathrm{H}}\left(\hat{\mathbf{U}} \hat{\mathbf{B}}+\hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}}\right) \hat{\mathbf{z}}_{0} \\
& =\mathbf{u}_{i}^{\mathrm{H}} \mathbf{r}_{0}-\hat{\sigma}_{i}\left(\hat{\mathbf{z}}_{0}\right)_{(i)}
\end{aligned}
$$

The last relation follows from the orthonormality of the NPSVD. It follows that the updated residual is orthogonal to $\mathbf{u}_{i}$ if we set $\left(\hat{\mathbf{z}}_{0}\right)_{(i)}=\mathbf{u}_{i}^{\mathrm{H}} \mathbf{r}_{0} / \hat{\sigma}_{i}$.

The Galerkin condition is justified in exact arithmetics where all search-vectors are orthogonal as described in sec. 7.4 p .85 . Thus, if the residual is orthogonal to all search-vectors, it must eventually be in the null-space of the operator.

In the presence of rounding errors, the reality is very different. If the conjugate gradient fails to converge on a few (say, in 10-30 iterations), the search vector has essentially lost any relationship with the starting vector. For fast convergence it is therefore highly desirable to use the largest NPSVD that will fit in memory.

The Galerkin orthogonalization removes a component of the residual of norm $\left|\mathbf{u}_{i}^{\mathrm{H}} \mathbf{r}_{0}\right|$. This is substituted for a component in the search-vector of norm $\left|\mathbf{u}_{i}^{\mathrm{H}} \mathbf{r}_{0}\right|$. $\left|\left(\mathbf{f}_{\mathbf{u}}\right)_{(i)}\right| / \hat{\sigma}_{i}$. It follows that the residual norm can not grow if the triplets discarded during the restart satisfy $\left|\left(\mathbf{f}_{\mathbf{u}}\right)_{(i)}\right| \leq \hat{\sigma}_{i}$. From sec. 8.6 p. 100 this means that the singular vector has converged to a relative accuracy better than 1 .

### 9.3 Unregularized solver

It is possible to combine the two restarting strategies in a single iterative solver. Since it is numerically advantageous to keep an NPSVD as large as possible, we can use implicit restarts as long as some singular triplets converges. In the unregularized case, it is always possible to limit the size of the NPSVD, since any triplet may be discarded at any time.

In practice, the difference between MINRES and Galerkin solutions are limited for triplets that converges to a high relative accuracy, i.e. $\left|\left(\mathbf{f}_{\mathbf{u}}\right)_{(i)}\right| \ll \hat{\sigma}_{i}$. However, a large-scale problems that runs over many iterations may build up a number of triplets that converges very slowly. In this case, an explicit restart may give a significiantly smaller residual norm than an implicit restart.

A simple solution is to monitor the norm resulting from using the two methods. If an explicit restart result in a residual norm that is significantly better than the Galerkin method, then such a step should be taken.

As an example, the default criteria used in BBTools is to use explicit restarting if the residual norm can be improved by a factor of 2 . There is no rigorous argument for this, but it appears to work well in practice.

After a large number of iterations has been completed, the residual is stripped for components corresponding to large singular values. These triplets are easily identified from eq. (9.3b): the singular value $\hat{\sigma}_{i}$ is small compared to $\|\mathbf{A}\|_{2}$.

Even if $\left|\left(\mathbf{f}_{\mathbf{u}}\right)_{(i)}\right|$ is large, these components should be removed from the residual. Certainly, this should happen if $\hat{\sigma}_{i} \leq \tau$, i.e. the accuracy of the matrix-vector product, but applications may choose to use a much larger threshold.

### 9.4 Regularized solver

At this point we are equipped to deal with the regularized solver. From sec. 5.5 p. 58 we would like a solution to the system $\mathbf{A} \hat{\mathbf{x}} \approx \mathbf{b}$ written in terms of an SVD, $\mathbf{A}=\mathbf{U} \mathbf{\Sigma} \mathbf{V}^{\mathrm{H}}$ using a set of filter-factors:

$$
\begin{equation*}
\hat{\mathbf{x}}=\sum_{i: \sigma_{i}>0} \mathbf{v}_{i} f_{i} \frac{\mathbf{u}_{i}^{\mathrm{H}} \mathbf{b}}{\sigma_{i}} \tag{9.9}
\end{equation*}
$$

In the current setting, we are mostly concerned with Tikhonov regularization as shown in fig. 9.1. In general, the filter-factors are most convenient to deal with when the singular values are given on a logarithmic scale.

Figure 9.1:
Tikhonov filterfactors for various degrees of regularization.


There are two important things to note at this point:

- The regularized solution is linear in the filter-factors.
- The common filter-factors (see tab. 5.1 p. 59) are smooth if the singular values are plotted on a logarithmic scale.

This implies that the filter-factors can be interpolated linearly by sampling the filter-factors sparsely on a logarithmic scale. In other words, we may keep a few representative vectors instead of $\hat{\mathbf{V}}$ and choose filter-factors (e.g. the regularization constant $\lambda$ ) after the linear system has been solved.

This split the solving process into two separate steps:

- Pre-solve: solve the system using Galerkin orthogonalization, except only triplets that has converged to an acceptable tolerance is used. Also, instead of updating the solution, update the set of representative solutions.
- Post-solve: compute filter-factors, scale the representative solutions accordingly, and add.

The post-solving procedure is reduces to a matrix-vector product. On a modern computer this is, for any practical purposes, instantaneous. This allows the solver to be relatively expensive but still offers the user to adjust the regularization parameter interactively.

Generally, the singular triplets do not need to be very accurate, provided the filterfactors do not change too rapidly as a function of the singular value. Theorem C14 p. 152 guarantees that the numerical triplet is a sum of exact triplets and any components far from the approximated singular value must be small.

For Tikhonov regularization it is usually safe to orthogonalize and discard triplets that satisfy $\left|\left(\mathbf{f}_{\mathbf{u}}\right)_{(i)}\right|<c \hat{\sigma}_{i}$ where $c \approx 0.1$. However, the required accuracy is application dependent and should be tested empirically.
It should be noted that if $c$ is very small ${ }^{1}$, the solver is forced to converge all singular triplets. In this case it may be more efficient, and memory conserving, to compute the SVD using the algorithm described in chapter 8.

### 9.5 Hybrid solver

Unfortunately, the requirement for relative accuracy of the singular triplets may prevent convergence. Generally, using a MINRES step is not an option, since the singular triplets are inaccurate.

However, after a large number of iterations, the residual mainly contains contributions from noise and small singular triplets that can not be resolved. This property hinges on the Picard condition [30]:
Definition 9.1: (Discrete Picard condition) Let $\mathbf{A}=\mathbf{U} \boldsymbol{\Sigma V}^{\mathrm{H}}$ be an SVD of an operator, and let $\mathbf{A x}=\boldsymbol{\lambda}$ be the expected measurement (i.e. the result measured in absence of noise). Then the discrete Picard condition is satisfied if the elements $\left|\mathbf{U}^{\mathrm{H}} \boldsymbol{\lambda}\right|_{(i)}$ decays faster than $\sigma_{i}$ for sufficiently large $i$.

Roughly speaking, the discrete Picard condition is satisfied if $\mathbf{x}$ has limited bandwidth, as argued in sec. 5.2 p. 55.

Furthermore, the unconverged singular triplets tends to correspond to small singular values. This occurs because any component corresponding to large singular

[^12]values are picked up in a few iterations and is decoupled from the remaining component.

The last property implies that if a numerical triplet is expanded into components of exact singular triplets, then the singular values of the exact triplets tends to be much smaller than the numerical singular value of the triplet.

Therefore it is not dangerous to end the procedure with a MINRES step, and apply filter-factors according to the estimated singular values. This result in a hybrid method [30] where the exact algorithm used influence the result.

In practice, the number of steps required before the difference between the hybrid and true Tikhonov solution becomes negligible is problem dependent. Generally, the scheme works best if the norm of the measurement-noise is on the order of a few percent or more of the noiseless measurement, as is the case for most tomographic problems.

### 9.6 References

The literature on iterative solvers is vast because of the large number of practical applications. A pedagogical text-book that describes Krylov methods is [73], while a more complete references that includes algorithmic details is [4]. A good review with historical notes is [61].

Another text-book is [15], which also includes several Krylov methods and several other methods and references. It also raises the noteworthy point that Krylov methods are no panacea, and methods such as multi-grid may be more suitable for some applications.

One of the most successful general solvers is LSQR developed by Paige and Saunders [55]. The paper also includes an excellent discussion of bidiagonalization.

Hybrid methods are described in [30], which also include many references. The current knowledge is centered around methods where some results of early stopping is known.

The idea that, say, 1000 iterations is small compared to the size of the operator is not an ordinary theme in the literature. Hybrid methods are, at the moment, associated with some amount of handwaving. The author is unaware of any hybrid method that guarantees a solution that can be expressed independently of the algorithm in the presence of rounding errors.


BBTools is a toolbox for MATLAB implementing the algorithms described in the previous chapters and much more. It represents a major part of the work produced during the PhD.

This chapter is intended to be self-contained, and may be regarded as an introduction to BBTools.

### 10.1 Introduction to BBTools

BBTools, or black-box toolbox, is a collection of routines for MATLAB that implements almost everything discussed in the thesis in a way that should be accessible for most people familiar with Matlab ${ }^{\text {® }}$.

The main focus of BBTools is to handle inverse ill-posed problems such as tomographic problems. However, the routines are general and can be used for a number of purposes, including:

- Inverse problems, regularized solutions of ill-posed problems and SVD analysis.
- Data mining, dimensionality reduction of large data sets.
- Iterative methods, research in algorithms based on Krylov subspaces.

The main reason the toolbox exists is that it can cope with problems that are much larger then any software package available to the author. It has been designed to utilize the available memory as much as possible.

## The meaning of a black-box

In the context of BBTools, a "black-box operator" (or simply "black-box") is a set of functions that implements the matrix-vector products $\mathbf{A x}$ and $\mathbf{A}^{\mathrm{H}} \mathbf{x}$. Here $\mathbf{A}$ is the operator and $\mathbf{x}$ is a normal MATLAB ${ }^{\circledR}$ vector (or matrix).

It can be useful to think of a black-box as an extension of sparse matrices. However, instead of being a "matrix with a lot of zero elements", a black-box is "any matrix where matrix-vector products can be computed efficiently". This class of operators is much broader than sparse matrices.

The operator is a "black-box" because the algorithms used to compute the products are known only to the implementation. The internal representation of the operator may have little resemblance with the elements of the matrix with the same linear transform.

Example 10.1: The function bbconvn creates an operator whose internal representation consists of a Fourier transform of the convolution kernel. Fortunately, neither the user nor BBTools need to know this fact. (Except, of course, for bbconvn, which is part of BBTools).

## How BBTools helps

In BBTools, the set of functions is encapsulated in an object that can be treated like a MATLAB-matrix in many cases. In most other packages (including functions shipped with MATLAB ${ }^{\circledR}$ such as lsqr), a black-box operator must be given as a function (or set of functions).

Since a black-box can be manipulated like an ordinary MATLAB-matrix, it is usually much easier to create and/or change an operator. Instead of reinventing the wheel, things such as Kronecker products and concatenations are available using the normal MATLAB ${ }^{\oplus}$ syntax.

Furhermore, BBTools provides a set of standard black-boxes to use as a starting point. This includes convolutions, Hadamard matrices, and basic operators such as zeros, diagonal matrices, and identity operators. Anyone familiar with MATLAB ${ }^{\circledR}$ knows how even simple operators can simplify a complex task.

The more general question "where do black-box operators come from?" is really two different questions:

1. Where do nature create problems with exploitable structure?
2. Where do the black-box operators used in BBTools come from?

We answer the first question with a quotation from [73]:
Large matrices ... usually arise indirectly in the discretization of differential or integral equations. One might say, that if $m$ is large, it is probably an approximation to $\infty$. It follows that most large matrices of computational interest are simpler than their vast number of individual entries might suggest.

The second question is addressed in sec. 10.2.

## How black-box operators can be utilized

Black-box operators are only useful if they can be used to solve a problem. Although a matrix-vector product may be a useful operation in itself, there are many important iterative algorithms that are based solely on this operation.

From a practical point of view, the most useful procedures involves so called Krylov subspaces. BBTools is focused on the following problems:

- Solving large-scale SVD-problems.
- Computing regularized solutions to large-scale linear problems.

Theoretically, any operation performed on a matrix could be performed on a black-box operator. This is possible because the matrix corresponding to the linear transform can be computed using the matrix-vector product. From a pragmatic point of view, this would defeat the purpose of using black-box operators.

### 10.2 Central features

The fundamental unit in BBTools is the black-box operator. From this startingpoint, BBTools targets 4 key areas in linear algebra.

## Creating and manipulating operators

The fundamental idea in BBTools is to encapsulate linear operations in operators. BBTools sports a number of features for creating and working with such operators:

- Library of operators: Many common operators are included to provide a starting point. Although some of these may occasionally be sufficient to model real-world problems, most operators are meant to provide "buildingblocks, not solutions". Most functions available in BBTools is provided in appendix D.
- Operators can be combined: Creating operators by combining simpler ones is standard in MATLAB ${ }^{\text {® }}$, but BBTools supports huge operators by storing an evaluation tree rather than explicit elements. This is transparent to the user.
- Support for custom operators: Operators can be highly specialized and require large programs to implement efficiently. If such programs are embedded in a black-box, they can be manipulated which often leads to better modularization and reuse of code.
- Support for testing operators: When complex operators are implemented, it is useful to have a set of standard functions for testing the correctness of the algorithms. This is a natural part of BBTools.


## SVD Analysis

The singular value decomposition (SVD) is one of the most important tools for analyzing linear systems. It forms a cornerstone of BBTools, due to its strong connection with regularized solutions of linear systems.

In its simplest form, the function bbsvds can replace svds in the common case where some of the triplets corresponding to the largest singular values are desired. It is usually faster, partly because svds does not utilize ARPACK in the most efficient way, and partly because BBTools is designed to exploit modern hardware more efficiently.

A bit more under the hood, we find some more drastic differences:

1. bbsvds computes singular values with an amount of memory that scales as $\mathcal{O}(m+n+\min (m, n) k)$, where $k$ is the number of requested singular values. Although many packages can do this, BBTools do so without compromising the accuracy.
2. bbsvds can often compute more singular values/triplets than will fit in memory. This is accomplished without swapping data to and from disk; BBTools does not contain out-of-core routines. Instead, it works by trading memory for computational work. Rounding-errors may prevent BBTools from computing a full SVD of a large operator, but the extension can be very useful.

Computing singular values is easier than full singular triplets. If singular vectors are required, BBTools will normally use memory that scales as $\mathcal{O}((m+n) k)$. The last point remains true: it is still possible to compute more triplets than memory can hold.

## Solving ill-posed problems

The goal of a large linear problem is often to solve the system on the form $\mathbf{A x}=\mathbf{b}$. Unfortunately, the system-matrix $\mathbf{A}$ is ill-conditioned in many practical problems and $\mathbf{b}$ have measurement-, rounding-, and representation-errors.

Regularization deals with these problems. BBTools is concerned with solutions that can be expressed in terms of an SVD of $\mathbf{A}$ and a set of filter factors. This includes many standard techniques.

Unfortunately, the SVD of $\mathbf{A}$ is infeasible to compute in practice, and we must resort to approximate methods. BBTools accomplishes the task in two steps:

1. bbpresolve An iterative solver, except the result is not returned as a vector (bbsolve can do that). Rather, the solution is split into a number of solutions, each corresponding to a set of local filter-factors.
2. bbregsolve Form a regularized solution by approximating the wanted filter-factors as a sum of the local filter-factors. The solution, $\mathbf{x}$, is then a weighted sum of the solutions computed by bbpresolve.

It is typically necessary to compute a number of regularized solutions in order to balance noise and accuracy. The L-curve method (bblcurve), for instance, requires a number of of such evaluations. BBTools supports this efficiently: bbregsolve costs almost nothing when it is fed with the output from bbpresolve.

The solver in BBTools is primarily intended for ill-posed problems. When a "real" (i.e. unregularized) solution is meaningful, it may be better to use other routines.

Note that preconditioning destroys the singular values of the operator, and therefore regularization. As a consequence, the code was developed to cope with more iterations than other solvers normally expect is necessary.

## Iterative routines

Black-box operators generally require iterative routines, and both the SVD routines and solvers described above are iterative. However, some problems are already handled well by existing MATLAB ${ }^{\circledR}$ routines. For instance, a function like eigs is excellent.

There are several cases where one may want to use the iterative solvers shipped with Matlab:

1. An accurate, unregularized, solution is required.
2. A test-battery of algorithms is needed for comparison or benchmarking.
3. A known algorithm must be used to allow reproducible results.

BBTools supports most iterative routines shipped with MATLAB.

## Research in iterative routines

Since the 1960s we had libraries of canned routines for most dense problems in linear algebra [17]. The libraries have been rewritten to match the technology; we went from LINPACK and EISPACK to LAPACK [2], and later ScaLAPACK [10] emerged.

As computers get faster, we expect to handle larger problems, and this eventually leads to more use of iterative routines [73]. The literature is extensive, and many iterative solvers have been proposed; a good introduction can be found in [61].

Such a plethora exists because there is no single algorithm that consistenty beats the competition. No canned routine have emerged where we can leave its implementation to specialist and envoke with $\mathrm{A} \backslash \mathrm{b}$, trusting that it "just works".

BBTools should be helpful with research of such algorithms. For instance, it allows one to evaluate the effect of degraded accuracy and count products without changing the code of the solvers.

However, the most important aspect of BBTools is perhaps that it allows the enduser a familiar entrance to the iterative routines. The popularity of MATLAB ${ }^{\circledR}$ can be seen as an example of the importance; most of the functionality in MATLAB ${ }^{\circledR}$ is available as free libraries on the Internet. Nevertheless, MATLAB ${ }^{\circledR}$ is the de facto language for research in numerical linear algebra.

### 10.3 Working with black-box operators

BBTools is rooted in the idea that any linear transformation can be represented as an operator that acts on a column-vector ${ }^{1}$. While this may sound innocent, it implies a mental model that may be unfamiliar to many.

In most cases we do not think much about the meaning of a "matrix". To work efficiently with BBTools, it is advantageous to distinguish between operators (linear transforms) and hypercubes (data arrays). The most immediate consequence of the policy is that any matrix/array must be converted to a column-vector before an operation can be applied.

Example 10.2: Assume that IMG is an $100 \times 100$ image. An operation might traditionally do something like the following:

IMG2 $=5 *$ IMG ( $1: 50,20: 90$ );
IMG2=convn (IMG2,K);
Using BBTools this would instead look something like this:
[A1, dim1] =bbredim ([100,100], [1,20;50,90]);
[A2, dim2] =bbconvn (dim1,K) ;
A=A 2 * 5 *A1;
IMG2=reshape (A*IMG(:), dim2);

This may look unnecessarily complex, but it has the advantage that the steps are captured in A. This allows us to pass it as an argument to a function, which can be blissfully ignorant that IMG represents an image.

[^13]Example 10.3: Continuing the last example, apply the adjoint operator on IMG2: IMG3=reshape ( $\mathrm{A}^{\prime}$ *IMG2(:), dim1)

While the adjoint could be worked out in a simple example like this, BBTools made it trivial because the adjoints of the two sub-functions has been worked out once and for all.

## Matrices as operators

The word "operator" is usually defined (in the current context) in a way similar to the following [25]:
$\ldots$ any symbol that indicates an operation to be performed.

That is, an operator is not a passive table of elements; rather it is something active like a function. This view is sometimes helpful in understanding BBTools.

Let us go back to basics: linear operators are functions that map a vector-space to another linearly. That is, $T$ is a linear operator if it satisfies the following for all $\check{\mathbf{x}}, \check{\mathbf{y}}$ in the input domain and scalars $c$ :

$$
\begin{align*}
T(\check{\mathbf{x}}+\check{\mathbf{y}}) & =T(\check{\mathbf{x}})+T(\check{\mathbf{y}})  \tag{10.1}\\
T(\check{\mathbf{x}} c) & =c T(\check{\mathbf{x}})
\end{align*}
$$

When the input and output domains of a linear operator have finite dimensions, then any element, say $\check{x}$, in the vector-space can be written in terms of a basis:

$$
\begin{equation*}
\check{\mathbf{x}}=x_{1} \check{\mathbf{b}}_{1}+x_{2} \check{\mathbf{b}}_{2}+\cdots \tag{10.3}
\end{equation*}
$$

As long as the basis is given, we may refer to $\check{\mathrm{x}}$ by giving the coefficients (or coordinates) of this basis. These are conveniently represented by a vector of coordinates: $\mathbf{x}$.

Any linear operation that maps a column-vector to another may be described uniquely in the form of a matrix. The transformation $\check{\mathbf{y}}=T(\check{\mathbf{x}})$ may be computed as a matrix-vector product $\mathbf{y}=\mathbf{A x}$.

It is a good idea to keep in mind that an object is represented by a vector in terms of a basis. In many cases the vector is simply related to the basis, and is is easy to forget it exists. If an object is an image, for example, each element may refer to a pixel.

## Matrices as containers

Matrices are used for other things than describing linear operators. A matrix may hold any kind of data which can be stored in a 2 -dimensional table. Common examples are images and pairwise relations such as correlation coefficients, forces, distances, etc.

BBTools is primarily concerned with data when operators are acting on them. Consider an image where the some common editing is wanted, say cropping or sharpenening. Although BBTools can create operators that performs such operations, it does not understand the notion of an image.

To apply the operation the data must be in canonical form, i.e. a vector. That is, a data-matrix it must be expanded in terms of a basis before BBTools can do anything with it. We shall use the term hypercube to emphasize that data-matrices and other arrays are the same to BBTools. We could have used a number of different words:

## 1. Hypercubes

2. Multi-dimensional arrays
3. Data matrices

To confuse matters further, a matrix may also represent a collection of column vectors, e.g. $\mathbf{X}=\left[\begin{array}{lll}\mathbf{x}_{1}, & \mathbf{x}_{2}, & \cdots\end{array}\right]$. In this case an operator may be applied en masse: $\mathbf{A X}=\left[\begin{array}{lll}\mathbf{A x}_{1}, & \mathbf{A} \mathbf{x}_{2}, & \cdots\end{array}\right]$.

It may occasionally be difficult to determine what category a matrix belongs to. From a practical point of view it usually suffice to distinguish between operators and data-arrays.

## Advantages and disadvantages

BBTools assumes that linear transformations are applied by creating an operator and applying it to the data, which has the form of a vector. There are several advantages with this approach:

- The implementation of an operator is separated from its invokation. This allows one to use standard routines and concentrate on the modelling problem (i.e. creating the operator) rather than the algorithms required for their solution. item Hypercubes are treated as any kind of data. BBTools does not care if a vector represents a rectangular image, a round one, or contains holes.
- Algorithms can be closer to their "text-book" versions than what is normally seen in practical code. This speeds up development, makes it relatively easy to analyze an algorithm, and facilitates debugging. The algorithms shown in this thesis are examples of this concept.
- Vectors are stored in matrices that can be manipulated via high-performance BLAS (see sec. 8.2 p. 96).
- The creation of "the adjoint operation" (i.e. computing $\mathbf{A}^{H} \mathbf{x}$ ) only needs to be considered at the level of fundamental building-blocks. Operators build from these automatically supports this operation.
- BBTools keeps track of the accuracy of the matrix-vector product, $\tau$. Again, this only needs to be determined at the level of building-blocks. If necessary, the accuracy can be estimated from the operator directly.

The downside is that it is necessary to handle data manually. One of the more annying examples is the necessity to keeping track of the dimensions of a hypercube. Functions that create operators corresponding to actions on a hypercube (e.g. bbconvn) often require the dimension as an input argument and often return the output dimension as a result.

### 10.4 Compatibility

One of the major design-goals was to be as compatible with MATLAB ${ }^{\oplus}$ as possible. Once a black-box has been created, it should act as closely as possible as an ordinary matrix. BBTools was designed to blend in with ordinary MATLAB programs as unobtrusively as possible.

The main computations that black-boxe operators lacks is operations that works on individual elements. This implies that functions such as abs and A. ${ }^{\wedge}$ c) are not supported.

## Operator combinations

Many common manipulations of MATLAB-matrices are supported by BBTools, as shown in tab. 10.1. While the list may appear short, it covers the majority of the operations used in practical programs.

Table 10.1:
Basic operations compatible with black-box operators.

| Syntax | Function | Description |
| :--- | :--- | :--- |
| A+B | plus (A, B) | Addition |
| [A, B] | horzcat (A, B) | Horizontal concatenation |
| [A;B] | vertcat (A, B) | Vertical concatenation |
| A/c | mrdivide (A, C) | Division by scalar |
| A.' | transpose (A) | Transposition |
| A' $^{\prime}$ | ctranspose (A) | Complex conjugate transpose |
| +A | uplus (A) | Unary plus |
| -A | uminus (A) | Unary negatation |
| A*B | mtimes (A, B) | Matrix multiplication (see below) |
|  | feval (A, B) | Evaluate A*B (discouraged) |
|  | feval (A, B, transp') | Evaluate A' *B (discouraged) |
|  | $\operatorname{conj}(A)$ | Complex conjugate |
|  | real (A) | Complex real part |
|  | imag (A) | Complex imaginary part |
|  | kron (A, B) | Kronecker tensor product |
|  | repmat (A, .. ) | Replicate and tile an operator |

## Matrix multiplication

Matrix multiplication is supported via the normal syntax $A * B$ and the corresponding function mtimes ( $\mathrm{A}, \mathrm{B}$ ). However, to maximize compatibility there are special rules that applies to this product:

1. If either $A$ or $B$ is a Matlab ${ }^{\circledR}$-matrix, the product will be evaluated and returned as a normal MATLAB-matrix.
2. If either $A$ or $B$ is a scalar, which can be interpreted as a $1 \times 1$-matrix, then the product will also be evaluated.
3. In any other case it returns a black-box operator corresponding to the product $A * B$.

The convention allows a black-box operator to replace a MATLAB-matrix in many circumstances. If $A$ is a black-box operator and $x$ is a normal matrix, then $A * x$
evaluates the matrix-vector product, while A*blackbox (x) creates a black-box operator.

If a black-box is desired without evaluation, the function bbprod may be used instead of mtimes. This is guaranteed to return a black-box operator. Conversely, bbmult always return a MATLAB-matrix. However, this is not possible if both inputs are black-box operators.
The forms using feval exist for compatibility purposes only. While this allows black-box operators to be used with some existing programs, the use should be discouraged.

## Complex operations

BBTools supports complex operators everywhere. However, if A is a complex operator, then $B=r e a l(A)$ or $B=i m a g(A)$ creates an operator that evaluates two matrix-vector products with $A$ to evaluate a matrix-vector product with $B$. If these functions are nested inside each other, this may lead to exponential explosion.

## Operator information

The ordinary functions used to get basic information about an operator are also supported, as shown in tab. 10.2. However, disp, display and, by extension, the result printed by a computation that evaluates to a black-box operator, is not the elements of the matrix but the evaulation tree.

Table 10.2:
Basic information about a black-box operator.

| Function | Description |
| :--- | :--- |
| size (A) | Dimension of operator |
| isreal (A) | true if A is real |
| isa(A,str) | true if A is of type str |
| disp(A) | Display evaluation tree (as text) |
| display (A) | As disp, except the variable name is printed |

Example 10.4: One of the most complex functions in BBTools is bbrescale, which creates a black-box that rescales a hyper-cube to one with another dimension. For example, the command A=bbrescale([100,100], $[200,200], ' s p l i n e ')$ creates an operator, $\mathbf{A} \in \mathbb{R}^{(200 \cdot 200) \times(100 \cdot 100)}=\mathbb{R}^{40000 \times 10000}$, that scales a $100 \times 100$ image to $200 \times 200$ using cubic spline interpolation. The command disp (A) prints the following evaulation tree:

```
kron 40000x10000
    prod 200x100
    > matrix 200x200 (sparse)
    > collage 200x100
        > eye 100x100
        > prod 100x100
            > inverse 100x100 (LU substitution)
            > index 100x100
            > matrix 100x100 (sparse, central difference)
    prod 200x100
    > matrix 200x200 (sparse)
    > collage 200x100
        > eye 100x100
        > prod 100x100
            > inverse 100x100 (LU substitution)
            > index 100x100
            > matrix 100x100 (sparse, central difference)
```

A more traditional (and less informative) rendition is given in fig. 10.1.

Figure 10.1:
Evaluation-tree for an operator created by bbrescale.

Table 10.3:
Analysis support of black-box operators.

Table 10.4: Iterative routines supported by BBTools (parameters are omitted).


## Norms

Analysis of linear systems often start with norms. While BBTools has powerful functions for SVD-analysis, the ordinary functions in MATLAB ${ }^{\circledR}$ also works with black-boxes as shown in tab. 10.3.

| Function | Description |
| :--- | :--- |
| norm (A) | Compute norm of an operator (2-norm only) |
| normest (A) | Rough estimate of 2-norm |
| normest1 (A) | Good estimate of the 1-norm |

The exception is the Frobenius norm. While this norm is often used because it is easy to compute for matrices (sparse or full), black-box implementations rely on iterative procedures to accomplish the same task. Unfortunately, it is difficult to estimate the Frobenius-norm reliably, and the author is unaware of good algorithms that can accomplish this task in the general case.

Note: BBTools includes the function bbnormest, which should be used instead of normest. It supports both 1-, 2-, and $\infty$-norms, and usually provides a much better accuracy for substantially less work than normest.

The function norm reverts to bbnormest with a warning for the 1 - and $\infty$ norms. Although bbnormest (which is, in turn, based on normest 1 ) is usually very accurate, there are rare cases where it may fail.

## Iterative routines

Some care have been taken to allow the iterative routines shipped with Matlab ${ }^{\circledR}$ to work with black-box operators. MATLAB ${ }^{\circledR}$ usually assumes that these functions are intended for sparse matrices, but they work for black-box operators as well.

| Function | Description |
| :--- | :--- |
| eigs | Find a few eigenvalues and eigenvectors |
| svds | Find a few singular values and vectors |
| lsqr | Conjugate Gradients on the Normal Equations |
| bicg | BiConjugate Gradients |
| bicgstab | BiConjugate Gradients Stabilized |
| cgs | Conjugate Gradients Squared |
| gmres | Generalized Minimum Residual |
| minres | Minimum Residual |
| pcg | Preconditioned Conjugate Gradients |
| qmr | Quasi-Minimal Residual |
| symmlq | Symmetric LQ |

## Miscellaneous

Finally, tab. 10.5 lists a few miscellaneous functions that do not fit in any of the previous categories.

Table 10.5:
Miscellaneous
functions supported by BBTools.

| Function | Description |
| :--- | :--- |
| sum (A) | Column sum (or rows) |
| mean (A) | Column average (or columns) |

### 10.5 References

BBTools, and its manual, is freely available online [33]. Once installed, the documentation is available in the MATLAB ${ }^{\oplus}$ environment similar to the toolboxes available from Mathworks. A list of most functions provided by BBTools is given in appendix D p. 157.
The toolbox "Regularization Tools" [29], by Dr. Per Christian Hansen, is a fairly complete toolbox for MATLAB ${ }^{\circledR}$, and contains numerous tools and test problems. It is very useful for people who wants an introduction to the topic of regularization and ill-posed problems. The software and its manual are available online together with a tecnical report [31].

The toolbox is mainly concerned with techniques for small dense systems. There are, however, a few iterative routines and some functions which work independently of the solutions (e.g. computation of filter-factors).
The Matlab-toolbox MOORE Tools offers a different approach to implement operators [38]. Instead of considering an operator as a black-box, it uses object oriented programming. That is, an operator is implemented as a class, and is not limited to matrix-vector products. In fact, the operations supported by MOORE Tools increase with the capabilities a particular class implements.

Both BBTools and MOORE Tools can combine operators in an evaluation tree. However, the main emphasis of $\mathrm{MO} \mathcal{O}$ RE Tools is abstraction and infrastructure, whereas BBTools is efficiency, minimalism, and problem solving.
For example, MOORE Tools does not require an explicit conversion between hypercubes and vectors; it understands the difference once the class has been implemented. On the othert hand, it is not necessary to teach BBTools there is a difference.

Someone who mainly use BBTools to combine operators, or is looking for a framework for implementing general methods, may be served best by MOORE Tools. On the other hand, BBTools may be preferable to someone who wants to solve a specific problem.


## Part IV

## Results

$\oplus$
$\oplus$
 the phantoms from chapter 3, and compare the regularized reconstruction with the traditional methods.

Since the goals described in chapter 5 do not lend themselves to quantitative analysis, the reconstructed images must stand on their own. Hopefully, the power of regularized reconstruction is evident from these images.

### 11.1 Introduction

One thing might as well be admitted from the start: the method chosen in chapter 5 was not chosen to optimize quantitative accuracy. This brings up to the disturbing fact that an analysis of the quantitative accuracy is pointless; much better methods for this purpose can be found in the literature.

Therefore, the only results given in this chapter will be images produced by the algorithms. While this may be disappointing for those who are mathematically oriented, it is good news for the medical users targeted in this project: there will not be any mathematics in the chapter.

The starting point is the Shepp-Logan phantom introduced in chapter 3. For convenience we reproduce it in fig. 11.1

Figure 11.1:
The classical Shepp-Logan phantom reproduced from fig. 3.2.


### 11.2 Low-count reconstruction

Starting with the most difficult case, we return to the case where only 20,000 counts are acquired. This represent a typical scenario in SPECT, where it is typical to acquire about 500,000 counts in medical scanning. In 2D reconstruction, these counts are distributed on a number of slices, which is determined by the size of the acquisition matrix.

Figure 11.2:
The sinogram in fig. 3.12 reconstructed using FBP.

Figure 11.3:
An underregularized reconstruction of fig. 11.2.

In the current scenario, the acquisition matrix typically consists of 64 slices. However, the central (and most interesting) slices will usually get more counts that the outer slices. A typical reconstruction of a slice using filtered back-projection is shown in fig. 11.2.


Envoking BBTools, we may reconstruct the same image using Tikhonov regularization. Since we intend to apply a post-filter with a Gaussian kernel, the regularization parameter should deliberately be chosen to provide an under-regularized solution. For the current example ${ }^{1}$, the result may look like fig. 11.3.


As described in sec. 9.4 p. 110 this is a two-step procedure: first a computationally expensive presolving procedure is performed, and then the regularization-constant can be chosen interactively.

After applying a post-filter, which was chosen to match the resolution of the FBPreconstruction in fig. 11.2, the under-regularized image (fig. 11.3) is transformed into fig. 11.4. At this point it might be useful to refresh one of the goals from chapter 5: the reconstruction algorithm should be backwards compatible, i.e. it should be possible to reconstruct images that are similar to the traditional algorithm.

[^14]Figure 11.4:
As fig. 11.3, but with a Gaussian post-filter to imitate fig. 11.2.


To judge whether this goal has been achieved, we merely need to compare whether fig. 11.2 and fig. 11.4 are similar. Those who agree they are similar also agree that the algorithm is backwards compatible as desired.

## Comparison of point-spread functions

As we also mentioned ${ }^{2}$, in chapter 5, there would be no progress if backwards compatibility was the sole goal; in that case we might as well use the old algorithm in the first place.

However, a closer look on the two algorithms do reveal a difference. Recall that a point-spread function (PSF) is the reconstructed image of a point in absence of noise. A number of point-spread functions for the FBP-reconstruction that produced fig. 11.2 are shown in fig. 11.5, while the same PSFs corresponding to fig. 11.4 are shown in fig. 11.6.

Comparing these we may draw two conclusions:

1. The PSFs of post-filtered Tikhonov-reconstruction are more homogeneous compared to FBP. This is most expressed near the edges.
2. The negative values are reduced.

Both are very desirable goals.

## Comparison with EM

The comparison would be incomplete without including an EM-reconstructed image. Since the EM-algorithm is non-linear, and the PSFs would therefore be meaningless, this comparison can only be performed on images.

To make the comparison as fair as possible, the EM-image in fig. 11.7 was reconstructed using the same post-filter as the Tikhonov image in fig. 11.4. Although a sieved image might yield slightly different result, choosing a set of parameters would be a subjective choice.

[^15]Figure 11.5:
Point-spread functions for the reconstruction in fig. 3.12.

Figure 11.6:
As fig. 11.5, but using regularized reconstruction with post-filtering.


## PSFs for Tikhonov with post-filter



At first sight, the EM-image may appear "nicer" than the other images. However, it is questionable whether it actually contains more information. There is also a tendency to create false features that look natural.

Perhaps the one quality that makes the EM-image look better is that it does not contain negative counts. We have consistently shown the negative values also, but for the purpose of comparison, fig. 11.8 is clipped so negative values are black. It is now much more questionable whether the EM reconstruction, fig. 11.7, offers a significant advantage over FBP in fig. 11.8.

## Envoking Tikhonov

When we are assured that we can approximate the FBP-algorithm closely, it is interesting to see the effect of the Tikhonov regularization. At this point we have two options:

- Suppress more noise at the same resolution.
- Increase resolution at the same noise-level.

The image in fig. 11.9 shows an example of more noise-suppresion. In this case, the regularization-constant was set 10 times higher than the previous image. The

Figure 11.7:
As fig. 11.12, but reconstruction using the EMalgorithm with the same post-filter used in fig. 11.3.

Figure 11.8:
A reproduction of fig. 11.3, except negative values have been clipped to zero.

Figure 11.9:
As fig. 11.4, but with a regularizationconstant that is 10 times higher.


Clipped Tikhonov

noise has been suppressed as desired, without any obvious loss of resolution.


Tikhonov regularization suppress noise by making the images less extreme. While this avoids artifacts, it also tends to make real features hard to detect.

Another possibility is to increase the resolution once the noise has been suppressed. For specificity, we shall use a post-filter with twice the resolution, i.e.
the FWHM of the PSFs is halved. The result is shown in fig. 11.10, which shows that the resolution is improved as desired.

Figure 11.10:
As fig. 11.9, except the resolution has doubled.

Tikhonov with post-filter


The reconstruction parameters on these examples were chosen without any rigorous criteria. While this is a weakness, there are only two parameters: the resolution and noise-suppression. Both may be selected while the image is updated interactively.

### 11.3 Medium-count reconstruction

We now turn to a scenario that lies in between a low-count and high-count scenario. While we use the same phantom, for easy comparison, we will work with a measurement where about 100, 000 counts are acquired as shown in fig. 11.11.

Figure 11.11:
Sinogram of the
Shepp-Logan phantom with approximately 100, 000 counts.


While this number of counts is difficult to achieve in clinical SPECT, it is very low for a PET-scanning. However, it is easier to see the effects of reconstruction when a limited number of counts are used, and it is a more difficult problem; if the low-count problem can be handled, then the high-count problem is easy.

For comparison, fig. 11.12 shows a FBP-image, reconstructed to match the resolution in the Tikhonov-enhanced image in fig. 11.10. Again, this can be imitated by an appropriate filter (fig. 11.13) and the resolution can be increased by suppressing extremes (fig. 11.14).

Figure 11.12: A FBP-
reconstruction of the sinogram in fig. 11.11.


Figure 11.13:
Post-filtered regularized solution to imitate fig. 11.12.


Tikhonov with post-filter

Figure 11.14:
The same as
fig. 11.13, except more filtering is performed using regularization and less by postfiltering.

$\cdots$

### 11.4 Using anatomy in reconstruction

So far we only mentioned anatomy in chapter 5, where it was stated that it would be easy to incorporate this into the reconstruction. It is time to make good on this promise.

First, we need to define a set of anatomical regions that are known to have constant concentrations. In a real workflow, this would typically be accomplished either by drawing regions on another set of images, e.g. MR, or using a template. In this example we cheat, and use regions based on the original phantom as shown in fig. 11.15.

Figure 11.15:
Anatomical regions as they might have been defined by a user. Region 0: nothing is assumed, Region 1: known to be zero, Region 2,3,4: constant.

Figure 11.16:
Underregularized solution to determine the concentration of the anatomical regions.


Anatomical regions

The regions shown in fig. 11.15 are well within the region on the original phantom. That is, the regions are "safe" and need not be perfectly delineated. In this sense, the example is realistic.

As described in sec. 5.6 p. 61, we may now reconstruct the image under the constraints defined by the regions. An under-regularized solution is shown in fig. 11.16. Clearly, the constrained regions are large enough to be determined accurately.


Finally, the known regions may be subtracted out while the remaining image is reconstructed using a more regularized filter. This leads to fig. 11.17, which is
the culmination of the project: a tomographic image reconstructed using anatomy and regularization.

Figure 11.17:
The final image, reconstructued using the anatomy and regularization.


## Conclusion

## Conclusion

After a general introduction to the scanners used in emission tomography, we identified the main weakness of the current image formation methods: the resolution.

Currently, nearly all images in emission tomography are produced using either filtered back-projection (FBP) or expectation maximization (EM), both of which have been available for a long time. Much work has been done to improve these algorithms, but penetration into the medical fields has been limited.

Traditionally, the criteria used for success has been to reduce a residual 2-norm, energy, root-mean-square, or other equivalent measure. Instead, we identified a number of requirements posed by the intended users. This resulted in a list where the traditional criteria did not figure prominently.

Many, or even most, of these objectives are either unmeasurable or prohibitively expensive and longwinded to test. In order to make progress, the method to be developed was required to be backwards compatible. Since images obtained by a classical method can be imitated, this would ensure that the new method would be acceptable for the user community.

To this end, regularization via filter-factors was chosen as the method to be pursued. It was proved that methods in this class remained competitive with a much larger class of algorithms, even if only the residual is considered.

We continued by developing numerical algorithms to compute the solutions. Although regularization is textbook material, implementing a viable algorithm was challenging. Nevertheless, an algorithm was developed where regularized reconstructions of 2D images could be performed on consumer laptops.

Previous algorithms for computing regularized solutions were unable to cope with problems on the scale required for practical medical reconstruction. This may be surprising, since the technique is widely used in many diverse fields. Conversely, the new algorithms may be useful in applications reaching beyond medicine.

The regularized reconstruction algorithm was shown to produce images without creating random artifacts, i.e. brain-looking features, other than those created by the traditional method. This is important, since it allows the resolution to be enhanced at the cost of suppressing elevated or lowered regions.

To reproduce the traditional algorithm, a post-filter was applied to an underregularized solution. The result is visually indistinguishable from the traditional algorithm it imitates. Furthermore, the reconstruction algorithms allow the user to balance the amount of filtering (either regularization or post-filtering) in realtime, once the initial processing is complete. Therefore, the user can start with the classical image, and choose a different setting by inspection.

The new software made it possible to incorporate anatomical information by constraining certain regions to constant concentrations. The practical consequence of this is that no assumptions were required in regions where the user would be unsure of the result, e.g. near lesions.

## Scientific contributions

There are several novelties in the project:

1. Rigorous bounds were established on linear inversion of Poisson-distributed measurements.
2. Algorithms were developed to handle large-scale SVD computations and computing regularized solutions for large-scale linear systems using arbitrary filter-factors.
3. Software was produced to make the algorithms accessible in a familiar MATLAB environment.
4. Regularized reconstruction was demonstrated to be viable on consumer hardware.
5. Finally, we showed that anatomical constraints could be incorporated in the framework.

It should be noted that most of these results have applications outside tomographic reconstruction. Indeed, regularization is fundamental for a large number of inverse ill-posed problems in e.g. geology and helioseismology. An extension to the range of problems that could be handled by previously available software is interesting for all such applications.


## Part V

## Appendix

$\oplus$
$\oplus$

## Appendix

## Addendum to WLS

This appendix contains technical proofs used in the analysis of weighted leastsquares in chapter 6.

## A1 Image-space bound

Lemma A1: (Point-distances on convex surfaces) Let $O$ be the origin of a Cartesian coordinate system, $B \in \mathcal{S}_{1}$ the point on a convex surface $\mathcal{S}_{1}$ closest to $O$, and $A \in \mathcal{S}_{2}$ the point closest to $O$ on another convex surface $\mathcal{S}_{2}$.

Assume neither surfaces contains $O$ in their interior. Finally, let $D$ be a point on $\mathcal{S}_{1}$ and $C$ the orthogonal projection of $D$ into the plane $O A B$ as shown in fig. A.1. Then we have the inequality $|B D|^{2}+|A B|^{2} \leq|O D|^{2}-|O A|^{2}$.

## WLS-solutions in image-space



Figure A.1:
Geometry of the points in the distance lemma.

Proof $\quad$ Since $A$ and $B$ are part of the convex surface $\mathcal{S}_{2}$, where $A$ by definition is closest to $O$, it follows that $\angle O A B \geq 90^{\circ}$. By the same reasoning on $\mathcal{S}_{1}$, we see that $\angle O B C \geq 90^{\circ}$.

From this we have $|O C|^{2} \geq|O B|^{2}+|B C|^{2}$ and $|O B|^{2} \geq|O A|^{2}+|A B|^{2}$, which combines into $|A B|^{2}+|B C|^{2} \leq|O C|^{2}-|O A|^{2}$. Since $C D$ is orthogonal to the $O A B$-plane we further have $|B D|^{2}=|B C|^{2}+|C D|^{2}$ and $|O D|^{2}=|O C|^{2}+|C D|^{2}$. This finally gives the inequality $|B D|^{2}+|A B|^{2} \leq$ $|O D|^{2}-|O A|^{2}$.

Theorem A2: Let $\ell_{P}$ and $\ell_{W L S}$ be as defined in sec. 6.3 p. 67. Furthemore, let the penaltyfunction be given by $r(\hat{\mathbf{x}})=\|\mathbf{L} \hat{\mathbf{x}}\|_{2}^{2}$.

Let $\hat{\mathbf{x}}$ be the reconstructed image, given by def. 6.1 p. 68, using the Poissonlikelihood. Further assume there there are regularization-constants, $\alpha_{a}$ and $\alpha_{b}$, resulting in WLS-solutions, $\mathbf{x}_{a}$ and $\mathbf{x}_{b}$, which satisfies:

$$
\begin{aligned}
\ell_{P}\left(\mathbf{A} \mathbf{x}_{a}, \mathbf{b}\right) & \leq \ell_{P}(\mathbf{A} \hat{\mathbf{x}}, \mathbf{b})
\end{aligned} \leq \ell_{P}\left(\mathbf{A x}_{b}, \mathbf{b}\right), ~=\ell_{W L S}(\mathbf{A} \hat{\mathbf{x}}, \mathbf{b}) \leq \ell_{W L S}\left(\mathbf{A} \mathbf{x}_{b}, \mathbf{b}\right) .
$$

If the reconstruction algorithm is continuous in the regularization parameter, $\alpha$, then there exist $W L S$-solutions, $\mathbf{x}_{1}$ and $\mathbf{x}_{2}$, that satisfy:

$$
\begin{gathered}
r\left(\hat{\mathbf{x}}-\mathbf{x}_{1}\right)+r\left(\hat{\mathbf{x}}-\mathbf{x}_{2}\right) \leq r\left(\mathbf{x}_{2}\right)-r\left(\mathbf{x}_{1}\right) \\
r\left(\mathbf{x}_{1}\right) \leq r(\hat{\mathbf{x}}) \leq r\left(\mathbf{x}_{2}\right)
\end{gathered}
$$

Proof Since the reconstructed image is a continuous function of $\alpha$, there exists a regularization parameter, $\alpha_{a} \leq \alpha_{2} \leq \alpha_{b}$ resulting in a WLS-image, $\mathbf{x}_{2}$, on the the Poisson-surface: $\ell_{\mathrm{P}}\left(\mathbf{A x}_{2}, \mathbf{b}\right)=\ell_{\mathrm{P}}(\mathbf{A} \hat{\mathbf{x}}, \mathbf{b})$. Since $\hat{\mathbf{x}}$ by definition minimizes the penalty-function, we have the relation $r(\hat{\mathbf{x}}) \leq r\left(\mathbf{x}_{2}\right)$.

Similarly, there is a WLS-solution, $\mathbf{x}_{1}$, on a WLS-surface containing the Poissonsolution: $\ell_{\mathrm{WLS}}\left(\mathbf{A x}_{2}, \mathbf{b}\right)=\ell_{\mathrm{WLS}}(\mathbf{A} \hat{\mathbf{x}}, \mathbf{b})$. As before we have $r\left(\mathbf{x}_{1}\right) \leq r(\hat{\mathbf{x}})$ by definition.

Now let $r(\mathbf{x})=\|\mathbf{x}\|_{2}^{2}$, and apply lemma A1 with $\mathbf{x}_{2}$ as point $D, \hat{\mathbf{x}}$ as point $B$ and $\mathbf{x}_{1}$ as point $A$. Then $|O A|^{2}=r\left(\mathbf{x}_{1}\right),|O D|^{2}=r\left(\mathbf{x}_{2}\right),|A B|^{2}=r\left(\mathbf{x}_{1}-\hat{\mathbf{x}}\right)$ and $|B D|^{2}=r\left(\mathbf{x}_{2}-\hat{\mathbf{x}}\right)$ which is the desired result.

Since any convex surface remains convex after a linear transform, it follows that convex surfaces in the sinogram-space will also be convex in image-space. In fact this holds for any matrix, not merely A. It follows that the bound holds for any penalty-function on the form $r(\mathbf{x})=\|\mathbf{L x}\|_{2}^{2}$.

## A2 Noise-perturbation of WLS-surface

We may determine a $\mathbf{q}$ that satisfies $\ell_{\mathrm{WLS}}(\boldsymbol{\lambda}+\mathbf{p}, \mathbf{b})=\ell_{\mathrm{WLS}}(\boldsymbol{\lambda}, \mathbf{b}-\mathbf{q})$ by equating terms in eq. (6.5) and get $\left(\lambda_{i}+p_{i}-b_{i}\right)^{2} / b_{i}=\left(\lambda_{i}-b_{i}+q_{i}\right)^{2} /\left(b_{i}-q_{i}\right)$. Defining $q_{i}=p_{i} c_{i}$ and dropping the indices, it is straightforward to show that:

$$
0=c^{2} p b+c\left(2 b(\lambda-b)+(\lambda+p-b)^{2}\right)-b(2 \lambda-2 b+p)
$$

Here $c=0$ becomes a solution in the limit $b \rightarrow 0$. Also, inserting $b^{\prime}=b / b=1$, $\lambda^{\prime}=\lambda / b$ and $p^{\prime}=p / b$ do not change the roots of $c$, so we need only consider $b=1$.

There are always real solutions for $\lambda \geq 0$, which is guaranteed since $\lambda$ is a Poisson-estimate. We now seek the point that maximizes the root of least absolute value. This occur for $\lambda=0$ and $p=1-\sqrt{2}$, in which case the solutions are $c= \pm(1+\sqrt{2})$.

## Base iteration

This appendix details the basic iterative procedure used for computing large-scale SVD-problems (chapter 8) and solving linear systems (chapter 9).

## B1 Stopping criteria

As it is often the case with numerical algorithms, one of the most difficult things of an iterative algorithm is to determine when to stop.

In principle, the base iteration consist of a loop that updates the NPSVD using alg. 7.3 p. 85 . For the left residual factorization, given by eq. (7.13) with the constraint $\hat{f}_{v}=\mathbf{0}$, this gives:

$$
\begin{align*}
& \tilde{\mathbf{s}}_{\mathbf{v}} \beta^{*}=\mathbf{A}^{H} \hat{\mathbf{s}}_{\mathbf{u}}-\hat{\mathbf{V}} \hat{\mathbf{f}}_{\mathbf{u}}  \tag{B1}\\
& \tilde{\mathbf{s}}_{\mathbf{u}} \tilde{f}_{u}=\mathbf{A} \hat{\mathbf{s}}_{\mathbf{v}}-\hat{\mathbf{s}}_{\mathbf{u}} \beta
\end{align*}
$$

Since we intend to use reorthogonalization, as described in sec. 7.6, the singular matrices must remain well conditioned. Therefore there are two ways to stop the iteration: (1) memory is exhausted or (2) rounding errors cause the singular matrices to loose orthogonality.

## Detecting loss of orthogonality

If we let $\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{U}}=\mathbf{I}+\mathbf{M}_{\mathbf{U}}$ and $\hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}=\mathbf{I}+\mathbf{M}_{\mathbf{V}}$, then the orthogonality-level of $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ is a number, $T$, such that $\left\|\mathbf{M}_{\mathbf{U}}\right\|_{\max } \leq T$ and $\left\|\mathbf{M}_{\mathbf{V}}\right\|_{\max } \leq T$.

It follows from theorem C 12 p .151 that if $T \leq .5 / k$, then singular matrices are well-conditioned. This is less strict than the semi-orthogonality criteria discussed in sec. 7.4 p. 85 , which requires $T \leq \sqrt{\epsilon_{M}}$.

One way to bound the orthogonality levels is to maintain matrices with elementwise bounds on the orthogonality, i.e. $\left|\mathbf{M}_{\mathbf{U}}\right| \leq \mathbf{O}_{\mathbf{U}}$ and $\left|\mathbf{M}_{\mathbf{V}}\right| \leq \mathbf{O}_{\mathbf{V}}$. This is easiest to accomplish if only the last element of $\hat{\mathbf{f}}_{\mathbf{u}}$ is non-zero.

For this reason, the iteration should start with one step of alg. 7.3 with full reorthogonalization. This has the additional benefit that the NPSVD is guaranteed to be expanded by at least one triplet.

Orthogonality of left singular vectors

From eq. (8.1a)

From eq. (8.1b)
Consider the level of orthogonality between the search-vector $\hat{\mathbf{s}}_{\mathbf{u}}$ and previous left singular vectors:

$$
\begin{aligned}
\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}} & =\frac{1}{\hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}} \hat{\mathbf{f}}_{\mathbf{u}}} \hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}} \hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}} \hat{\mathbf{f}}_{\mathbf{u}} \\
& =\frac{1}{\hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}} \hat{\mathbf{f}}_{\mathbf{u}}} \hat{\mathbf{U}}^{\mathrm{H}}\left(\mathbf{A} \hat{\mathbf{V}}-\hat{\mathbf{U}} \hat{\mathbf{B}}-\mathbf{G}_{\mathbf{U}}\right) \hat{\mathbf{f}}_{\mathbf{u}} \\
& =\frac{1}{\hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}} \hat{\mathbf{f}}_{\mathbf{u}}}\left(\left(\hat{\mathbf{B}} \hat{\mathbf{V}}^{\mathrm{H}}+\mathbf{G}_{\mathbf{V}}^{\mathrm{H}}\right) \hat{\mathbf{V}}-\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{U}} \hat{\mathbf{B}}-\hat{\mathbf{U}}^{\mathrm{H}} \mathbf{G}_{\mathbf{U}}\right) \hat{\mathbf{f}}_{\mathbf{u}} \\
& =\frac{1}{\hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}} \hat{\mathbf{f}}_{\mathbf{u}}}\left(\hat{\mathbf{B}} \mathbf{M}_{\mathbf{V}}-\mathbf{M}_{\mathbf{U}} \hat{\mathbf{B}}+\mathbf{G}_{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}-\hat{\mathbf{U}}^{\mathrm{H}} \mathbf{G}_{\mathbf{U}}\right) \hat{\mathbf{f}}_{\mathbf{u}}
\end{aligned}
$$

Taking absolute values of the above gives:

Inserting eq. (8.1a)

$$
\begin{align*}
\left|\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}\right| & \leq \frac{1}{\hat{\mathbf{f}}_{\mathbf{u}}^{\mathrm{H}} \hat{\mathbf{f}}_{\mathbf{u}}}\left(|\hat{\mathbf{B}}| \mathbf{O}_{\mathbf{V}}\left|\hat{\mathbf{f}}_{\mathbf{u}}\right|+\mathbf{O}_{\mathbf{U}}\left|\hat{\mathbf{B}} \hat{\mathbf{f}}_{\mathbf{u}}\right|+\left|\mathbf{G}_{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}} \hat{\mathbf{f}}_{\mathbf{u}}\right|+\left|\hat{\mathbf{U}}^{\mathrm{H}} \mathbf{G}_{\mathbf{U}} \hat{\mathbf{f}}_{\mathbf{u}}\right|\right)  \tag{B3}\\
& =\frac{1}{\left|\hat{f}_{u}\right|}\left(|\hat{\mathbf{B}}|\left(\mathbf{O}_{\mathbf{V}}\right)_{[k]}+\mathbf{O}_{\mathbf{U}}|\hat{\mathbf{B}}|_{[k]}+\left|\mathbf{G}_{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}_{[k]}\right|+\left|\hat{\mathbf{U}}^{\mathrm{H}}\left(\mathbf{G}_{\mathbf{U}}\right)_{[k]}\right|\right)
\end{align*}
$$

Section 7.3.4 showed that elements in each of the last two terms are bounded by a constant, $\tau_{c}$. Although single elements could be larger, it is usually safe to use $\tau_{c}=\tau$.

In the rare case where this fails, it is possible to retract the last one or two iteration in the following reorthogonalization. The most common cause for this is that an optimistic value of $\tau$ is used in the iteration.

## Orthogonality of right singular vectors

We need an analogous result for the right singular vectors. This is slightly complicated by the fact that the right search-vector is computed during the update.

$$
\hat{\mathbf{s}}_{\mathbf{v}} \beta^{*}=\mathbf{A}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}-\hat{\mathbf{V}} \hat{\mathbf{f}}_{\mathbf{u}}+\mathbf{r}
$$

where $\mathbf{r}$ represents rounding errors. Then we have:

$$
\begin{aligned}
\hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{v}} & =\frac{1}{\beta^{*}}\left(\hat{\mathbf{V}}^{\mathrm{H}} \mathbf{A}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}-\hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}} \hat{\mathbf{f}}_{\mathbf{u}}+\hat{\mathbf{V}}^{\mathrm{H}} \mathbf{r}\right) \\
& =\frac{1}{\beta^{*}}\left(\hat{\mathbf{B}}^{\mathrm{H}} \hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}+\hat{\mathbf{f}}_{\mathbf{u}} \hat{\mathbf{s}}_{\mathbf{u}}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}+\mathbf{G}_{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}-\left(\mathbf{I}+\mathbf{M}_{\mathbf{V}}\right) \hat{\mathbf{f}}_{\mathbf{u}}+\hat{\mathbf{V}}^{\mathrm{H}} \mathbf{r}\right)
\end{aligned}
$$

Using the fact that $\hat{\mathbf{s}}_{\mathbf{u}}$ is normalized:

$$
=\frac{1}{\beta^{*}}\left(\hat{\mathbf{B}}^{\mathrm{H}} \hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}+\mathbf{G}_{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}-\mathbf{M}_{\mathbf{V}} \hat{\mathbf{f}}_{\mathbf{u}}+\hat{\mathbf{V}}^{\mathrm{H}} \mathbf{r}\right)
$$

Taking absolute values gives the bound:

$$
\begin{equation*}
\left|\hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{v}}\right| \leq \frac{1}{|\beta|}\left(\left|\hat{\mathbf{B}}^{\mathrm{H}}\right|\left(\left|\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}\right|\right)+\mathbf{O}_{\mathbf{V}}\left|\hat{\mathbf{f}}_{\mathbf{u}}\right|+\left|\mathbf{G}_{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}\right|+\left|\hat{\mathbf{V}}^{\mathrm{H}} \mathbf{r}\right|\right) \tag{B4}
\end{equation*}
$$

## Notes on orthogonality

A few comments are in place. First of all, the bounds degrades quickly if $\left|\hat{f}_{u}\right|$ or $|\beta|$ becomes small compared to $\|\hat{\mathbf{B}}\|$. However, if $\left|\hat{f}_{u}\right|$ or $\beta$ is small, the NPSVD
has started to converge. In rough terms $\hat{\mathbf{B}}$ will decouple if one of these entries is small enough, and "almost" decouple if it is just small.

The phenomena is well-known and has a comprehensive theory [15]. It allows us to use the loss of orthogonality to double as an efficient way to detect convergence. There is no need to provide a separate criteria.

Finally, it is possible that the bound on the search-vector is smaller than the previous level of orthogonality if $\left|\hat{f}_{u}\right|$ (or $|\beta|$ ) is large compared to $\|\hat{\mathbf{B}}\|$. In this case the largest value must be used.

001010
100001
111100

## Algorithm

Once the stopping criteria has been determined, it is relatively easy to implement the base iteration as shown in alg. B1. The orthogonality-level of $\hat{\mathbf{U}}$ is not allowed to degrade to the same level as $\hat{\mathbf{V}}$. This is because $T_{v}$ can not be computed until $\beta$ is known, which requires a matrix-vector product. Thus, if the algorithm stops at this point the computation will be wasted.

The implementation used in BBTools does not even check $T_{v}$, since it can retraction the last iterations to guard against an optimistic value of $\tau$. In many cases the bounds are pessimistic, and the retraction is not needed at all.

## Algorithm B1: Base iteration

Basic implementation of the left residual factorization. The input is an orthogonalized NPSVD updated once with alg. 7.3 with full reorthogonalization. The iteration updates the NPSVD until the singular matrices start to loose orthogonality or memory is full.

```
function \(\left[\hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, f_{u}, k\right]=\operatorname{iter}\left(\mathbf{A}, \hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, f_{u}, k, K, \tau, f_{A}, f_{A}^{\mathrm{H}}\right.\), varargin)
    \(\mathbf{O}_{\mathbf{U}}=\operatorname{ones}(K, K) * \epsilon_{M} ; \quad T_{u}=\epsilon_{M} ; \quad\) Initial orthogonality levels
    \(\mathbf{O}_{\mathbf{V}}=\operatorname{ones}(K, K) * \epsilon_{M} ; \quad T_{v}=\epsilon_{M} ;\)
    while \(k<K\)
        \(\hat{\mathbf{B}}_{\mathbf{A}}=|\hat{\mathbf{B}}(1: k, 1: k)| ;\)
        \(\tilde{\mathbf{o}}_{\mathbf{u}}=\mathbf{O}_{\mathbf{U}}(1: k, k) * \hat{\mathbf{B}}_{\mathbf{A}}(k, k)+\hat{\mathbf{B}}_{\mathbf{A}} * \mathbf{O}_{\mathbf{V}}(1: k, k) ; \quad\) Evaluate eq. (B3)
        \(\tilde{\mathbf{o}}_{\mathbf{u}}=\left(\tilde{\mathbf{o}}_{\mathbf{u}}+\tau\right) /\left|f_{u}\right| ;\)
        \(\mathbf{O}_{\mathbf{U}}(1: k, k+1)=\tilde{\mathbf{o}}_{\mathbf{u}} ; \quad \mathbf{O}_{\mathbf{U}}(k+1,1: k)=\tilde{\mathbf{o}}_{\mathbf{u}}{ }^{\mathrm{H}} ; \quad\) Update orthogonality matrix
        \(\mathbf{O}_{\mathbf{U}}(k+1, k+1)=\epsilon_{M} ;\)
        \(T_{u}=\max \left(T_{u}, \max \left(\tilde{\mathbf{o}}_{\mathbf{u}}\right)\right) ; \quad\) Update orthogonality level
        if \(T_{u} * k>0.05, \quad\) break Stop if orthogonality-level is low
        \(\left[\hat{\mathbf{s}}_{\mathbf{v}}, \beta\right]=\operatorname{unit}\left(\mathbf{A}^{\mathrm{H}} \hat{\mathbf{s}}_{\mathbf{u}}-\hat{\mathbf{V}}(:, k) * f_{u}\right) ; \quad\) Search-vector by eq. (B1)
        \(\tilde{\mathbf{o}}_{\mathbf{v}}=\hat{\mathbf{B}}_{\mathbf{A}}{ }^{\mathrm{H}} * \tilde{\mathbf{o}}_{\mathbf{u}}+\mathbf{O}_{\mathbf{V}}(1: k, k) *\left|f_{u}\right| ; \quad \quad\) Evaluate eq. (B4)
        \(\tilde{\mathbf{o}}_{\mathbf{v}}=\left(\tilde{\mathbf{o}}_{\mathbf{v}}+\tau\right) /|\beta| ;\)
        \(\mathbf{O}_{\mathbf{V}}(1: k, k+1)=\tilde{\mathbf{o}}_{\mathbf{v}} ; \quad \mathbf{O}_{\mathbf{V}}(k+1,1: k)=\tilde{\mathbf{o}}_{\mathbf{v}}{ }^{\mathrm{H}} ; \quad\) Update orthogonality matrix
        \(\mathbf{O}_{\mathbf{V}}(k+1, k+1)=\epsilon_{M} ;\)
        \(T_{v}=\max \left(T_{v}, \max \left(\tilde{\mathbf{o}}_{\mathbf{v}}\right)\right) ; \quad\) Update orthogonality level
        if \(T_{v} * k>.5, \quad\) break \(\quad\) Stop if orthogonality-level is low
        \(k=k+1 ; \quad \hat{\mathbf{U}}(:, k)=\hat{\mathbf{s}}_{\mathbf{u}} ; \quad \hat{\mathbf{V}}(:, k)=\hat{\mathbf{s}}_{\mathbf{v}} ; \quad\) Update singular matrices
        \(\hat{\mathbf{B}}(k, 1: k)=0 ; \quad \hat{\mathbf{B}}(1: k, k)=0 ; \quad\) Update coupling matrix
        \(\hat{\mathbf{B}}(k, k-1)=\left(f_{u}\right)^{*} ; \quad \hat{\mathbf{B}}(k, k)=\beta ;\)
        \(\left[\hat{\mathbf{s}}_{\mathbf{u}}, f_{u}\right]=\operatorname{unit}\left(\mathbf{A} * \hat{\mathbf{s}}_{\mathbf{v}}-\hat{\mathbf{s}}_{\mathbf{u}} * \beta\right) ; \quad\) Search-vector from eq. (B2)
```

$\oplus$
$\oplus$

## Appendix

## Matrix analysis

This appendix contains useful results and definitions from matrix analysis that are used in the matrix algorithms in part III. Although some results are proved here we strive to give references where possible.

Many perturbation results found in the literature are given in terms of angles between subspaces. In real programs it is normally more convenient to use inner products and norms. This is the emphasis of the forms given here.

## C1 Perturbation theory

Numerical methods rarely give exact result, but perturbation theory gives bounds that we can use to guide the algorithms. The most important perturbation result about the SVD is given by Weyl's theorem.

Theorem C1: (Weyl) The matrices $\mathbf{A}, \mathbf{E} \in \mathbb{C}^{m \times n}$ satisfy $\left|\sigma_{k}(\mathbf{A}+\mathbf{E})-\sigma_{k}(\mathbf{E})\right| \leq\|\mathbf{E}\|_{2}$.
PROOF
See [69].
A useful consequence of Weyl's theorem is the following corollary:
Corollary C2: Let A be the following block matrix:

$$
\mathbf{A}=\left[\begin{array}{ll}
\mathbf{A}_{1,1} & \mathbf{A}_{1,2} \\
\mathbf{A}_{2,1} & \mathbf{A}_{2,2}
\end{array}\right]
$$

Then $\|\mathbf{A}\|_{2} \leq \max \left\{\left\|\mathbf{A}_{1,1}\right\|_{2},\left\|\mathbf{A}_{2,2}\right\|_{2}\right\}+\max \left\{\left\|\mathbf{A}_{1,2}\right\|_{2},\left\|\mathbf{A}_{2,1}\right\|_{2}\right\}$.
There are also relative versions of Weyl's theorem. We will use the most common bound, but others have been found recently [51].

Theorem C3: (Ostrowski) Assume $\tilde{\mathbf{A}}=\mathbf{D}_{1}^{\mathrm{H}} \mathbf{A D} \mathbf{D}_{2}$ and $\mathbf{A} \in \mathbb{C}^{m \times n}$, where $\mathbf{D}_{1}$ and $\mathbf{D}_{2}$ are square. Let $\sigma_{k}$ and $\tilde{\sigma}_{k}$ the $k$ 'th singular value of $\mathbf{A}$ and $\tilde{\mathbf{A}}$.

Further, define $c_{M}=\sigma_{1}\left(\mathbf{D}_{1}\right) \sigma_{1}\left(\mathbf{D}_{2}\right)$, the product of the largest singular values, and $c_{m}=\sigma_{m}\left(\mathbf{D}_{1}\right) \sigma_{n}\left(\mathbf{D}_{2}\right)$, the product of the smallest singular values. Then we have:

$$
c_{m} \sigma_{i} \leq \tilde{\sigma}_{i} \leq c_{M} \sigma_{i}
$$

## C2 Unitarily invariant norms

A norm, $\|\|\cdot\|$, is unitarily invariant if $\| \mathbf{A}\|=\| \mid \mathbf{U}^{\mathrm{H}} \mathbf{A V} \|$ for all unitary matrices $\mathbf{U}$ and V. Many stable algorithms are based on orthogonal transformations, and this puts unitarily invariant norms in a central position for measuring rounding errors. This is true even though they are not usually absolute ${ }^{1}$.

[^16]Since the SVD reduces any matrix to diagonal form, it is convenient to define $\boldsymbol{\sigma}(\mathbf{A})$ as a vector containing the singular values of $\mathbf{A}$. With this notation we may write the two most important norms as $\|\mathbf{A}\|_{2}=\|\boldsymbol{\sigma}(\mathbf{A})\|_{\infty}$ and $\|\mathbf{A}\|_{\mathrm{F}}=$ $\|\boldsymbol{\sigma}(\mathbf{A})\|_{2}$. However, the following theorem theorem, due to John Von Neumann, shows that the family is infinite.

Theorem C4: (Von Neumann's norm theorem) The function $\|\mathbf{A}\|=\|\boldsymbol{\sigma}(\mathbf{A})\|_{\mathrm{V}}$ is a unitarily invariant matrix-norm if, and only if, $\|\mathbf{x}\|_{\mathrm{V}}$ is an absolute permutation invariant vector-norm.

Proof
See [70].
This shows that a result that holds for any unitarily invariant norm is quite strong. One way to prove such results is to use the following theorem.

Theorem C5: (Ky Fan's dominance theorem) Let $\mathbf{A}$ and $\mathbf{B}$ have equal size. Then $\|\mathbf{A}\| \leq$ $\|\mathbf{B}\|$ for all unitarily invariant norms if, and only if, $\sum_{i=1}^{k} \sigma_{i}(\mathbf{A}) \leq \sum_{i=1}^{k} \sigma_{i}(\mathbf{B})$ for $k=1,2, \cdots$.

PROOF
See [70].
The dominance theorem is sometimes impractical to use directly, but a few theorems take some of the labor away.

Theorem C6: (Padded matrix dominance) Let A be any matrix. Then any unitarily invariant norm satisfies $\left\|\left\|\left[\begin{array}{ll}\mathbf{A}, & \mathbf{0}\end{array}\right]\right\|\right\| \leq\left\|\left[\begin{array}{ll}\mathbf{A}, & \mathbf{B}\end{array}\right]\right\| \|$.

Proof We need only show the case where $\mathbf{B}=\mathbf{b}$ is a single column, since the general result then follows by adding a single column at a time. Now set $\mathbf{X}=\left[\begin{array}{ll}\mathbf{A}, & \mathbf{b}\end{array}\right]$ and form

$$
\mathbf{X}^{\mathrm{H}} \mathbf{X}=\left[\begin{array}{ll}
\mathbf{A}^{\mathrm{H}} \mathbf{A} & \mathbf{A}^{\mathrm{H}} \mathbf{b} \\
\mathbf{b}^{\mathrm{H}} \mathbf{A} & \mathbf{b}^{\mathrm{H}} \mathbf{b}
\end{array}\right]
$$

By the Cauchy interlace theorem [27], the eigenvalues of $\mathbf{A}^{\mathrm{H}} \mathbf{A}$ interlaces those of $\mathbf{X}^{\mathrm{H}} \mathbf{X}$. For $\mathbf{b}=\mathbf{0}$ the eigenvalues of $\mathbf{X}^{\mathrm{H}} \mathbf{X}$ are the eigenvalues of $\mathbf{A}^{\mathrm{H}} \mathbf{A}$ and a zero. Since the eigenvalues of $\mathbf{X}^{\mathrm{H}} \mathbf{X}$ are the squared singular values of $\mathbf{X}$, we have $\sigma_{i}\left(\left[\begin{array}{ll}\mathbf{A}, & \mathbf{0}\end{array}\right]\right) \leq \sigma_{i}\left(\left[\begin{array}{ll}\mathbf{A}, & \mathbf{b}\end{array}\right]\right)$ and the dominance property is satisfied.

Theorem C7: (Projection dominance) Let $\mathbf{P}$ be an idempotent projection, i.e. $\mathbf{P}=\mathbf{P P}$. Then $\|\mathbf{P A}\| \leq\|\mathbf{A}\|$ in any unitarily invariant norm.

Proof The singular values of an idempotent matrix, $\mathbf{P}$, are 1 or 0 . The theorem now follows from Ostrowski's theorem C3 and C5.

## Normalized norms

We are interested in norms that bound the error of individual components so that $\|\mathbf{A}\| \geq \max _{i, j}\left|\mathbf{A}_{i, j}\right|$. To make this inequality sharp, we assume the norm satisfies $\|\mathbf{A}\|=|c|$ if $\mathbf{A}$ consist of a single non-zero entry $c$. This is equivalent to requiring $\|\mathbf{A}\|=\|\mathbf{A}\|_{2}$ if $\operatorname{rank}(\mathbf{A})=1$.

Ky Fan's dominance theorem implies that the 2-norm is minimal in the sense that $\|\mathbf{X}\|_{2} \leq\|\mathbf{X}\|$, which follows from inserting $\mathbf{A}=\operatorname{diag}\left(\|\mathbf{X}\|_{2}, 0, \cdots\right)$ and $\mathbf{B}=\mathbf{X}$.

The Ky Fan norm, defined as $\|\mathbf{A}\|_{\mathrm{KF}}=\|\boldsymbol{\sigma}(\mathbf{A})\|_{1}$, is the other extreme. If we set $\mathbf{A}=\mathbf{X}$ and $\mathbf{B}=\operatorname{diag}\left(\|\mathbf{X}\|_{\mathrm{KF}}, 0, \cdots\right)$ and apply the dominance theorem we get $\|\mathbf{X}\| \leq\|\mathbf{X}\|_{\text {KF }}$. If we combine these observations we get the following result.
(C1)

$$
\|\mathbf{A}\|_{2} \leq\|\mathbf{A}\| \leq\|\mathbf{A}\|_{\mathrm{KF}}
$$

## C3 Norm inequalities

All norms can be bounded in terms of each other, and it is handy to have a list of these bounds ready. In addition to the usual norms, $\|\cdot\|_{1},\|\cdot\|_{2},\|\cdot\|_{\infty}$, and $\|\cdot\|_{F}$, we also need the max-norm:

Table C.1:
Constants that bounds the norm on the left in terms of the norm on the top. The matrix $\mathbf{A}$ is $m \times n$, and $d=\min \{m, n\}$.

$$
\begin{equation*}
\|\mathbf{A}\|_{\max }=\max _{i, j}\left|\mathbf{A}_{(i, j)}\right| \tag{C2}
\end{equation*}
$$

The max-norm is not consistent, i.e. the inequality $\|\mathbf{A B}\| \leq\|\mathbf{A}\| \cdot\|\mathbf{B}\|$ may not be satisfied. However, the max-norm gives a component-wise error which makes it a natural choice in some situations. This thesis uses it mainly to measure the orthogonality level or near-orthogonal matrices.

Table C. 1 give constants for the norms of used in this thesis for a matrix $\mathbf{A} \in$ $\mathbb{C}^{m \times n}$. For convenience we also define $d=\min \{m, n\}$. The upper bound is found starting on the left, while the lower bound is found from the top. For instance we find $\|\mathbf{A}\|_{1} / \sqrt{m} \leq\|\mathbf{A}\|_{2} \leq \sqrt{n}\|\mathbf{A}\|_{1}$ in two lookups.

|  | $\\|\mathbf{A}\\|_{\max }$ | $\\|\mathbf{A}\\|_{1}$ | $\\|\mathbf{A}\\|_{\infty}$ | $\\|\mathbf{A}\\|_{2}$ | $\\|\mathbf{A}\\|_{\mathrm{F}}$ | $\\|\mathbf{A}\\|_{\mathrm{KF}}$ |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
| $\\|\mathbf{A}\\|_{\max }$ |  | 1 | 1 | 1 | 1 | 1 |
| $\\|\mathbf{A}\\|_{1}$ | $m$ |  | $m$ | $\sqrt{m}$ | $\sqrt{m}$ | $\sqrt{m}$ |
| $\\|\mathbf{A}\\|_{\infty}$ | $n$ | $n$ |  | $\sqrt{n}$ | $\sqrt{n}$ | $\sqrt{n}$ |
| $\\|\mathbf{A}\\|_{2}$ | $\sqrt{m n}$ | $\sqrt{n}$ | $\sqrt{m}$ |  | 1 | 1 |
| $\\|\mathbf{A}\\|_{\mathrm{F}}$ | $\sqrt{m n}$ | $\sqrt{n}$ | $\sqrt{m}$ | $\sqrt{d}$ |  | 1 |
| $\\|\mathbf{A}\\|_{\mathrm{KF}}$ | $\sqrt{m n d}$ | $\sqrt{n d}$ | $\sqrt{m d}$ | $d$ | $\sqrt{d}$ |  |

The constants in the table are attained for some matrices. For instance the matrix $\mathbf{A}_{(i, j)}=e^{2 \pi i j / \max \{m, n\}}$ satisfies $\|\mathbf{A}\|_{\mathrm{KF}}=\sqrt{m n d}\|\mathbf{A}\|_{\text {max }}$. It is not always possible to attain equality if $\mathbf{A}$ is restricted to be real (e.g. $3 \times 3$ ). It is still asymptotically sharp, and it can sometimes be reached. If, for example, $\mathbf{H}=\left[\begin{array}{ll}\mathbf{A}, & \mathbf{B}\end{array}\right]$ is a Hadamard matrix then equality is attained by $\mathbf{A}$.

Finally we also need the following bounds [27, 69]:

$$
\begin{align*}
\|\mathbf{A}\|_{2} & \leq\|(|\mathbf{A}|)\|_{2}  \tag{C3}\\
\|\mathbf{A}\|_{2} & \leq \sqrt{\|\mathbf{A}\|_{1}\|\mathbf{A}\|_{\infty}}  \tag{C4}\\
\left|\mathbf{x}^{\mathrm{H}} \mathbf{y}\right| & \leq\|\mathbf{x}\|_{2}\|\mathbf{y}\|_{2} \tag{C5}
\end{align*}
$$

We refer to eq. (C5) as the Cauchy inequality, although some texts refer to it as the Cauchy-Schwartz inequality.

## C4 Matrix approximations

Any good textbook on linear algebra discusses the least-squares problem where we find an approximate solution to $\mathbf{A x} \approx \mathbf{b}$ by minimizing $\|\mathbf{A x}-\mathbf{b}\|_{2}$. It
turns out that the notion of "least-squares" generalizes naturally from vectors to matrices. This result is established in the following theorem.

Theorem C8: (Least-squares approximation) Let $\mathbf{A} \in \mathbb{C}^{m \times n}$ and $\mathbf{B} \in \mathbb{C}^{m \times p}$ be arbitrary matrices. Then $\mathbf{X}=\mathbf{A}^{+} \mathbf{B}$ minimizes $\|\mathbf{A X}-\mathbf{B}\|$, and is the unique choice that minimizes $\|\|\mathbf{X}\|$, where $\| \cdot \|$ is any unitarily invariant norm.
Proof Let $\mathbf{A}=\mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^{\mathrm{H}}$ be a full SVD of $\mathbf{A}$, and define $\mathbf{X}=\mathbf{A}^{+} \mathbf{B}+\mathbf{D}$ where $\mathbf{D} \in \mathbb{C}^{n \times p}$. For $k=\operatorname{rank}(\mathbf{A})$ we have:
(unitary invariance)

$$
\begin{aligned}
\|\mathbf{A X}-\mathbf{B}\| & =\left\|\mathbf{A} \mathbf{A}^{+} \mathbf{B}-\mathbf{B}+\mathbf{A} \mathbf{D}\right\| \\
& =\left\|\left(\mathbf{A} \mathbf{A}^{+}-\mathbf{I}\right) \mathbf{B}+\mathbf{A D}\right\| \\
& =\left\|\left(\left(\mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^{\mathrm{H}}\right)\left(\mathbf{V} \boldsymbol{\Sigma}^{+} \mathbf{U}^{\mathrm{H}}\right)-\mathbf{I}\right) \mathbf{B}+\mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^{\mathrm{H}} \mathbf{D}\right\| \\
& =\left\|\left(\mathbf{U} \boldsymbol{\Sigma} \boldsymbol{\Sigma}^{+} \mathbf{U}^{\mathrm{H}}-\mathbf{U} \mathbf{U}^{\mathrm{H}}\right) \mathbf{B}+\mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^{\mathrm{H}} \mathbf{D}\right\| \\
& =\left\|\left(\boldsymbol{\Sigma} \boldsymbol{\Sigma}^{+}-\mathbf{I}\right) \mathbf{U}^{\mathrm{H}} \mathbf{B}+\boldsymbol{\Sigma} \mathbf{V}^{\mathrm{H}} \mathbf{D}\right\| \\
& =\left\|\left[\begin{array}{ll}
\mathbf{0} & \\
& -\mathbf{I}
\end{array}\right] \mathbf{U}^{\mathrm{H}} \mathbf{B}+\left[\begin{array}{ll}
\boldsymbol{\Sigma}_{k} & \\
& \mathbf{0}
\end{array}\right] \mathbf{V}^{\mathrm{H}} \mathbf{D}\right\|
\end{aligned}
$$

where $\boldsymbol{\Sigma}_{k}=\operatorname{diag}\left(\sigma_{1}, \cdots, \sigma_{k}\right)$.
Since the rows of the two terms do not overlap, it follows from the padded dominance theorem C6, that the first $k$ rows of $\mathbf{V}^{\mathrm{H}} \mathbf{D}$ must be zero. We also have:

$$
\|\mathbf{X}\|=\left\|\mathbf{V} \boldsymbol{\Sigma}^{+} \mathbf{U}^{\mathrm{H}} \mathbf{B}+\mathbf{D}\right\|=\left\|\boldsymbol{\Sigma}^{+} \mathbf{U}^{\mathrm{H}} \mathbf{B}+\mathbf{V}^{\mathrm{H}} \mathbf{D}\right\|
$$

Only the first $k$ rows of the first term can be non-zero, and we have just seen that these must be zero in the second term for $\mathbf{X}$ to minimize $\|\mathbf{A X}-\mathbf{B}\|$. As before the rows do not overlap, so the remaining rows of $\mathbf{V}^{\mathrm{H}} \mathbf{D}$ must be zero to minimize $\|\mathbf{X}\|$.
The following result is well-known, but is not usually stated in full generality.
Theorem C9: (Schmidt-Mirsky low-rank approximation) Let the SVD of $\mathbf{A} \in \mathbb{C}^{m \times n}$ be given by $\mathbf{A}=\mathbf{U} \mathbf{\Sigma} \mathbf{V}^{\mathrm{H}}$ and let $\|\cdot\|$ be any unitarily invariant norm. Then the matrix $\mathbf{B}=\sum_{i=1}^{k} \boldsymbol{\Sigma}_{(i, i)} \mathbf{U}_{(:, i)} \mathbf{V}_{(:, i)}^{\mathrm{H}}$ minimizes $\|\mathbf{A}-\mathbf{B}\|$ among all matrices with $\operatorname{rank}(\mathbf{B})=k \leq \operatorname{rank}(\mathbf{A})$.

Proof A number of texts include a proof for special cases. A particularly readable geometric argument that covers the 2-norm is given in [73]. A proof for the general case can be found in [70].

## C5 Near-orthogonal matrices

It is useful to collect a few facts about matrices on the form $\hat{\mathbf{Q}}^{\mathrm{H}} \hat{\mathbf{Q}} \approx \mathbf{I}$. To this end we consider the matrix $\mathbf{Q}^{\mathrm{H}} \mathbf{Q}=\mathbf{I}+\mathbf{M}_{\mathbf{Q}}$. The orthogonality level of the columns in $\mathbf{Q}$ is then defined as the scalar $T=\left\|\mathbf{M}_{\mathbf{Q}}\right\|_{\max }$. A key property of near-orthogonal matrices is that their singular values are near 1.
Theorem C10: (Singular values of a near-orthogonal matrix) Let the columns of $\mathbf{Q} \in \mathbb{C}^{m \times n}$ be orthogonal at level $T$, i.e. $\mathbf{Q}^{\mathrm{H}} \mathbf{Q}=\mathbf{I}+\mathbf{E}$ where $\|\mathbf{E}\|_{\max } \leq T$. If $n T \leq 1$, then $\sqrt{1-n T} \leq \sigma_{i}(\mathbf{Q}) \leq \sqrt{1+n T}$ for all $i$.
Proof $\quad$ From table C. 1 we have $\|\mathbf{E}\|_{2} \leq n\|\mathbf{E}\|_{\max } \leq n T$. Weyl's theorem then gives the bound $1-n T \leq \sigma_{i}(\mathbf{I}+\mathbf{E})=\sigma_{i}\left(\mathbf{Q}^{\mathrm{H}} \mathbf{Q}\right) \leq 1+n T$. But $\sigma_{i}(\mathbf{Q})=\sqrt{\sigma_{i}\left(\mathbf{Q}^{\mathrm{H}} \mathbf{Q}\right)}$, which establishes the theorem.

A near-orthogonal matrix is always well-conditioned. We start with a lemma.
Lemma C11: If $|x| \leq \kappa<1$ are real scalars, then $(1+x)^{2} \leq(1+x) /(1-x) \leq(1+x)^{2} /(1-$ $\left.\kappa^{2}\right)$.

Proof Multiplying through with $1+x$ we get:

$$
\frac{1+x}{1-x}=\frac{(1+x)^{2}}{1-x^{2}}
$$

The lemma follows since $|x| \leq \kappa<1$ implies that $1-\kappa^{2} \leq 1-x^{2} \leq 1$.
Theorem C12: (Condition of near-orthogonal matrix) Let the columns of $\mathbf{Q} \in \mathbb{C}^{m \times n}$ be orthogonal at level $T$, i.e. $\mathbf{Q}^{\mathrm{H}} \mathbf{Q}=\mathbf{I}+\mathbf{E}$ where $\|\mathbf{E}\|_{\max } \leq T$. If $n T \leq \kappa<1$, then $\operatorname{cond}(\mathbf{Q}) \leq(1+\kappa) / \sqrt{1-\kappa^{2}}$.

Proof $\quad$ Since the condition of $\mathbf{Q}$ is given by $\operatorname{cond}(\mathbf{Q})=\max _{i} \sigma_{i}(\mathbf{Q}) / \min _{i} \sigma_{i}(\mathbf{Q})$, theorem C10 states that $\operatorname{cond}(\mathbf{Q}) \leq \sqrt{(1+n T) /(1-n T)}$. The theorem follows from inserting $x=n T$ in lemma C11.

For instance, if $\kappa \leq 1 / 2$, then $\operatorname{cond}(\mathbf{Q}) \leq \sqrt{3}$.

## C6 Triplet orthogonality

One of the big hurdles in creating numerical software for the SVD (and many other problems), is the need to keep the triplets orthogonal. In order get meaningful results, many iterative methods must sacrifice efficiency. The time spent during orthogonalization may dominate the whole computation

Explicit reorthogonalization becomes intolerable for problems that are too large to contain a solution in memory. In this case reorthogonalization requires disk access which severely hampers the efficiency of the algorithm.

To circumvent this we will rely on implicit orthogonality. The theorems derived in this section tells us that numerical triplets will automatically be (numerically) orthogonal under certain conditions.

We consider numerical singular triplets $\left(\hat{\mathbf{u}}_{i}, \hat{\mathbf{v}}_{i}, \hat{\sigma}_{i}\right)$ of the form:

$$
\begin{equation*}
\mathbf{A} \hat{\mathbf{v}}_{i}=\hat{\sigma}_{i} \hat{\mathbf{u}}_{i}+\breve{\mathbf{u}}_{i} \quad \mathbf{A}^{\mathrm{H}} \hat{\mathbf{u}}_{i}=\hat{\sigma}_{i}^{*} \hat{\mathbf{v}}_{i}+\breve{\mathbf{v}}_{i} \tag{C6}
\end{equation*}
$$

It is assumed that all triplets have been normalized, i.e. $1-\epsilon_{o} \leq\left\|\hat{\mathbf{u}}_{i}\right\|,\left\|\hat{\mathbf{v}}_{i}\right\| \leq$ $1+\epsilon_{o}$ where $\epsilon_{o}$ is on the order of machine-precision.

The results are derived by expressing the inner products in terms of the errors.
Lemma C13: A pair of triplets, as given by eq. (C6), where $\left|\hat{\sigma}_{i}\right| \neq\left|\hat{\sigma}_{j}\right|$ satisfies the relations:

$$
\begin{aligned}
\hat{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i} & =\frac{\hat{\sigma}_{i}\left(\breve{\mathbf{u}}_{j}^{\mathrm{H}} \hat{\mathbf{u}}_{i}-\hat{\mathbf{v}}_{j}^{\mathrm{H}} \breve{\mathbf{v}}_{i}\right)+\hat{\sigma}_{j}^{*}\left(\breve{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}-\hat{\mathbf{u}}_{j}^{\mathrm{H}} \breve{\mathbf{u}}_{i}\right)}{\left|\hat{\sigma}_{i}\right|^{2}-\left|\hat{\sigma}_{j}\right|^{2}} \\
\hat{\mathbf{u}}_{j}^{\mathrm{H}} \hat{\mathbf{u}}_{i} & =\frac{\hat{\sigma}_{i}^{*}\left(\breve{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}-\hat{\mathbf{u}}_{j}^{\mathrm{H}} \breve{\mathbf{u}}_{i}\right)+\hat{\sigma}_{j}\left(\breve{\mathbf{u}}_{j}^{\mathrm{H}} \hat{\mathbf{u}}_{i}-\hat{\mathbf{v}}_{j}^{\mathrm{H}} \breve{\mathbf{v}}_{i}\right)}{\left|\hat{\sigma}_{i}\right|^{2}-\left|\hat{\sigma}_{j}\right|^{2}}
\end{aligned}
$$

Proof
Consider the entity $\hat{\mathbf{v}}_{j}^{\mathrm{H}} \mathbf{A}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{v}}_{i}$ :

$$
\begin{aligned}
\hat{\mathbf{v}}_{j}^{\mathrm{H}} \mathbf{A}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{v}}_{i} & =\hat{\mathbf{v}}_{j}^{\mathrm{H}} \mathbf{A}^{\mathrm{H}}\left(\hat{\sigma}_{i} \hat{\mathbf{u}}_{i}+\breve{\mathbf{u}}_{i}\right) \\
& =\hat{\sigma}_{i} \hat{\mathbf{v}}_{j}^{\mathrm{H}}\left(\mathbf{A}^{\mathrm{H}} \hat{\mathbf{u}}_{i}\right)+\left(\hat{\mathbf{v}}_{j}^{\mathrm{H}} \mathbf{A}^{\mathrm{H}}\right) \breve{\mathbf{u}}_{i} \\
& =\hat{\sigma}_{i} \hat{\mathbf{v}}_{j}^{\mathrm{H}}\left(\hat{\sigma}_{i}^{*} \hat{\mathbf{v}}_{i}+\breve{\mathbf{v}}_{i}\right)+\left(\hat{\sigma}_{j} \hat{\mathbf{u}}_{j}+\breve{\mathbf{u}}_{j}\right)^{\mathrm{H}} \breve{\mathbf{u}}_{i} \\
& =\left|\hat{\sigma}_{i}\right|^{2}\left(\hat{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}\right)+\hat{\sigma}_{i}\left(\hat{\mathbf{v}}_{j}^{\mathrm{H}} \breve{\mathbf{v}}_{i}\right)+\hat{\sigma}_{j}^{*}\left(\hat{\mathbf{u}}_{j}^{\mathrm{H}} \breve{\mathbf{u}}_{i}\right)+\left(\breve{\mathbf{u}}_{j}^{\mathrm{H}} \breve{\mathbf{u}}_{i}\right)
\end{aligned}
$$

Exchanging $\hat{\mathbf{v}}_{i}$ with $\hat{\mathbf{v}}_{j}$ and taking the Hermitian gives another expression:

$$
\hat{\mathbf{v}}_{j}^{\mathrm{H}} \mathbf{A}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{v}}_{i}=\left|\hat{\sigma}_{j}\right|^{2}\left(\hat{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}\right)+\hat{\sigma}_{j}^{*}\left(\breve{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}\right)+\hat{\sigma}_{i}\left(\breve{\mathbf{u}}_{j}^{\mathrm{H}} \hat{\mathbf{u}}_{i}\right)+\left(\breve{\mathbf{u}}_{j}^{\mathrm{H}} \breve{\mathbf{u}}_{i}\right)
$$

Isolating $\hat{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}$ from these equations gives the first equation in the lemma. An analogous argument on $\hat{\mathbf{u}}_{j}^{\mathrm{H}} \mathbf{A} \mathbf{A}^{\mathrm{H}} \hat{\mathbf{u}}_{i}$ gives the other equation.

The following theorem asserts that accurate triplets are orthogonal, as long as their singular values differ slightly.

Theorem C14: (Mutual triplet orthogonality) Assume a set of normalized numerical triplets $\left(\hat{\mathbf{u}}_{i}, \hat{\mathbf{v}}_{i}, \hat{\sigma}_{i}\right)$ satisfies $\left\|\mathbf{A} \hat{\mathbf{v}}_{i}-\hat{\sigma}_{i} \hat{\mathbf{u}}_{i}\right\| \leq t_{i}$ and $\left\|\mathbf{A}^{\mathrm{H}} \hat{\mathbf{u}}_{i}-\hat{\sigma}_{i}^{*} \hat{\mathbf{v}}_{i}\right\| \leq t_{i}$. Then the orthogonality of a pair of triplets, where $\left|\hat{\sigma}_{i}\right| \neq\left|\hat{\sigma}_{j}\right|$, is bounded by:

$$
\max \left\{\left|\hat{\mathbf{v}}_{i}^{\mathrm{H}} \hat{\mathbf{v}}_{j}\right|,\left|\hat{\mathbf{u}}_{i}^{\mathrm{H}} \hat{\mathbf{u}}_{j}\right|\right\} \leq \frac{t_{i}+t_{j}}{\| \hat{\sigma}_{i}\left|-\left|\hat{\sigma}_{j}\right|\right|}\left(1+\epsilon_{o}\right)
$$

Proof
First note that $\left|\left|\hat{\sigma}_{i}\right|^{2}-\left|\hat{\sigma}_{j}\right|^{2}\right|=\left|\left|\hat{\sigma}_{i}\right|-\left|\hat{\sigma}_{j}\right|\right| \cdot\left(\left|\hat{\sigma}_{i}\right|+\left|\hat{\sigma}_{j}\right|\right)$. This, the Cauchy inequality, and lemma C13 leads to:

$$
\begin{aligned}
& \left|\hat{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}\right| \leq \frac{\left|\hat{\sigma}_{i}\right| \cdot\left(\left\|\breve{\mathbf{u}}_{j}\right\|+\left\|\breve{\mathbf{v}}_{i}\right\|\right)+\left|\hat{\sigma}_{j}\right| \cdot\left(\left\|\breve{\mathbf{v}}_{j}\right\|+\left\|\breve{\mathbf{u}}_{i}\right\|\right)}{\left|\left|\hat{\sigma}_{i}\right|-\left|\hat{\sigma}_{j}\right|\right| \cdot\left(\left|\hat{\sigma}_{i}\right|+\left|\hat{\sigma}_{j}\right|\right)}\left(1+\epsilon_{o}\right) \\
& \left|\hat{\mathbf{u}}_{j}^{\mathrm{H}} \hat{\mathbf{u}}_{i}\right| \leq \frac{\left|\hat{\sigma}_{i}\right| \cdot\left(\left\|\breve{\mathbf{v}}_{j}\right\|+\left\|\breve{\mathbf{u}}_{i}\right\|\right)+\left|\hat{\sigma}_{j}\right| \cdot\left(\left\|\breve{\mathbf{u}}_{j}\right\|+\left\|\breve{\mathbf{v}}_{i}\right\|\right)}{\| \hat{\sigma}_{i}\left|-\left|\hat{\sigma}_{j}\right|\right| \cdot\left(\left|\hat{\sigma}_{i}\right|+\left|\hat{\sigma}_{j}\right|\right)}\left(1+\epsilon_{o}\right)
\end{aligned}
$$

Substituting $t_{i}$ and $t_{j}$ into these equations completes the proof.
If there is just a small relative difference in the singular values, then the orthogonality is determined by the accuracy relative to the largest singular value. To make this clear, assume that $\left|\hat{\sigma}_{i}\right|>\left|\hat{\sigma}_{j}\right|$ and rewrite the relation as follows:

$$
\max \left\{\left|\hat{\mathbf{v}}_{i}^{\mathrm{H}} \hat{\mathbf{v}}_{j}\right|,\left|\hat{\mathbf{u}}_{i}^{\mathrm{H}} \hat{\mathbf{u}}_{j}\right|\right\} \leq \frac{t_{i}+t_{j}}{\left|\hat{\sigma}_{i}\right|}\left(\frac{1+\epsilon_{o}}{1-\left|\hat{\sigma}_{j}\right| /\left|\hat{\sigma}_{i}\right|}\right)
$$

For instance, if $\left|\hat{\sigma}_{i}\right| \geq 1.01\left|\hat{\sigma}_{j}\right|$, then the last term is at most $100\left(1+\epsilon_{o}\right)$.
The following result treats the case where $\left|\hat{\sigma}_{i}\right| \approx\left|\hat{\sigma}_{j}\right|$. It shows that keeping one set of singular triplets orthogonal is good enough for the large singular values.

Theorem C15: (One-sided triplet orthogonality) Assume the normalized numerical triplets $\left(\hat{\mathbf{u}}_{i}, \hat{\mathbf{v}}_{i}, \hat{\sigma}_{i}\right)$ satisfies $\left\|\mathbf{A} \hat{\mathbf{v}}_{i}-\hat{\sigma}_{i} \hat{\mathbf{u}}_{i}\right\| \leq t_{i}$ and $\left\|\mathbf{A}^{\mathrm{H}} \hat{\mathbf{u}}_{i}-\hat{\sigma}_{i}^{*} \hat{\mathbf{v}}_{i}\right\| \leq t_{i}$. Then we have the bounds:

$$
\begin{aligned}
\left|\hat{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}\right| & \leq\left|\hat{\mathbf{u}}_{j}^{\mathrm{H}} \hat{\mathbf{u}}_{i}\right|+\frac{t_{i}+t_{j}}{\min \left\{\left|\hat{\sigma}_{i}\right|,\left|\hat{\sigma}_{j}\right|\right\}}\left(1+\epsilon_{o}\right) \\
\left|\hat{\mathbf{u}}_{j}^{\mathrm{H}} \hat{\mathbf{u}}_{i}\right| & \leq\left|\hat{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}\right|+\frac{t_{i}+t_{j}}{\min \left\{\left|\hat{\sigma}_{i}\right|,\left|\hat{\sigma}_{j}\right|\right\}}\left(1+\epsilon_{o}\right)
\end{aligned}
$$

Proof

We first assume $\left|\hat{\sigma}_{i}\right|<\left|\hat{\sigma}_{j}\right|$, which only leaves equality since the triplets can be exchanged. We then rewrite lemma C13 in the form:

$$
\begin{aligned}
\hat{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i} & =\hat{\sigma}_{i} c_{1}+\hat{\sigma}_{j}^{*} c_{2} \\
\hat{\mathbf{u}}_{j}^{\mathrm{H}} \hat{\mathbf{u}}_{i} & =\hat{\sigma}_{i}^{*} c_{2}+\hat{\sigma}_{j} c_{1}
\end{aligned}
$$

Here we have defined the values $c_{1}=\left(\breve{\mathbf{u}}_{j}^{\mathrm{H}} \hat{\mathbf{u}}_{i}-\hat{\mathbf{v}}_{j}^{\mathrm{H}} \breve{\mathbf{v}}_{i}\right) /\left(\left|\hat{\sigma}_{i}\right|^{2}-\left|\hat{\sigma}_{j}\right|^{2}\right)$ and $c_{2}=\left(\breve{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}-\hat{\mathbf{u}}_{j}^{\mathrm{H}} \breve{\mathbf{u}}_{i}\right) /\left(\left|\hat{\sigma}_{i}\right|^{2}-\left|\hat{\sigma}_{j}\right|^{2}\right)$.

Isolating $c_{1}$ from eq. (C8) and inserting into eq. (C9) gives:

$$
\begin{aligned}
\hat{\mathbf{u}}_{j}^{\mathrm{H}} \hat{\mathbf{u}}_{i} & =\hat{\sigma}_{i}^{*} c_{2}+\frac{\hat{\sigma}_{j}}{\hat{\sigma}_{i}}\left(\hat{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}-\hat{\sigma}_{j}^{*} c_{2}\right) \\
& =\frac{c_{2}}{\hat{\sigma}_{i}}\left(\left|\hat{\sigma}_{i}\right|^{2}-\left|\hat{\sigma}_{j}\right|^{2}\right)+\frac{\hat{\sigma}_{i}}{\hat{\sigma}_{j}}\left(\hat{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}\right) \\
& =\frac{\breve{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}-\hat{\mathbf{u}}_{j}^{\mathrm{H}} \breve{\mathbf{u}}_{i}}{\hat{\sigma}_{i}}+\frac{\hat{\sigma}_{i}}{\hat{\sigma}_{j}}\left(\hat{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}\right)
\end{aligned}
$$

Similarly, isolating $c_{2}$ from eq. (C9) and inserting in eq. (C8) gives:

$$
\hat{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}=\frac{\breve{\mathbf{u}}_{j}^{\mathrm{H}} \hat{\mathbf{u}}_{i}-\hat{\mathbf{v}}_{j}^{\mathrm{H}} \breve{\mathbf{v}}_{i}}{\hat{\sigma}_{i}^{*}}+\frac{\hat{\sigma}_{i}^{*}}{\hat{\sigma}_{j}^{*}}\left(\hat{\mathbf{u}}_{j}^{\mathrm{H}} \hat{\mathbf{u}}_{i}\right)
$$

Inserting the bounds required by the theorem completes the case where $\left|\hat{\sigma}_{i}\right| \neq$ $\left|\hat{\sigma}_{j}\right|$. In the case of equality we simply multiply $\sigma_{i}$ by a factor, $c$, which will increase $t_{i}$ and $t_{j}$ by an arbitrarily small amount. In the limit $c \rightarrow 1$ we obtain the same relation.

## C7 Accuracy of singular values

To control the accuracy of the computations we need to bound the accuracy of the singular values in terms of quantities that can either be computed or bounded. In practice this means that it should be expressed in terms of the norms of the residuals matrices, $\left\|\mathbf{E}_{\mathbf{U}}\right\|$ and $\left\|\mathbf{E}_{\mathbf{V}}\right\|$, and the orthogonality levels, $\left\|\mathbf{M}_{\mathbf{U}}\right\|$ and $\left\|\mathbf{M}_{\mathrm{V}}\right\|$.

To this end we start with the following lemma.
Lemma C16: Let $\hat{\mathbf{U}} \in \mathbb{C}^{m \times k}, \hat{\mathbf{V}} \in \mathbb{C}^{n \times k}, k \leq \min \{m, n\}$, where $\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{U}}=\mathbf{I}+\mathbf{M}_{\mathbf{U}}$ and $\hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}=\mathbf{I}+\mathbf{M}_{\mathbf{V}}$. Also, let $\hat{\boldsymbol{\Sigma}}$ be a diagonal matrix with positive real entries that satisfies:

$$
\mathbf{A} \hat{\mathbf{V}}=\hat{\mathbf{U}} \hat{\boldsymbol{\Sigma}}+\mathbf{E}_{\mathbf{U}} \quad \quad \mathbf{A}^{\mathrm{H}} \hat{\mathbf{U}}=\hat{\mathbf{V}} \hat{\boldsymbol{\Sigma}}^{\mathrm{H}}+\mathbf{E}_{\mathbf{V}}
$$

Then there exists matrices $\check{\mathbf{U}}=\left[\begin{array}{ll}\hat{\mathbf{U}}, & \hat{\mathbf{U}}^{\circ}\end{array}\right] \in \mathbb{C}^{m \times m}$ and $\check{\mathbf{V}}=\left[\begin{array}{ll}\hat{\mathbf{V}}, & \hat{\mathbf{V}}^{\circ}\end{array}\right] \in$ $\mathbb{C}^{n \times n}$ where $\left(\hat{\mathbf{U}}^{\circ}\right)^{\mathrm{H}} \hat{\mathbf{U}}^{\circ}=\mathbf{I}$ and $\left(\hat{\mathbf{V}}^{\circ}\right)^{\mathrm{H}} \hat{\mathbf{V}}^{\circ}=\mathbf{I}$ and satisfies:

$$
\begin{equation*}
\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{U}}^{\circ}=\mathbf{0} \quad \hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}^{\circ}=\mathbf{0} \quad\left(\hat{\mathbf{U}}^{\circ}\right)^{\mathrm{H}} \mathbf{A} \hat{\mathbf{V}}^{\circ}=\hat{\boldsymbol{\Sigma}}^{\circ} \tag{C10}
\end{equation*}
$$

Furthermore, the product $\mathbf{C}=\check{\mathbf{U}}^{\mathrm{H}} \mathbf{A} \check{\mathbf{V}}$ is given by:

$$
\begin{align*}
\mathbf{C} & =\left[\begin{array}{cc}
\hat{\mathbf{U}}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{V}} & \mathbf{E}_{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}^{\circ} \\
\left(\hat{\mathbf{U}}^{\circ}\right)^{\mathrm{H}} \mathbf{E}_{\mathbf{U}} & \hat{\boldsymbol{\Sigma}}^{\circ}
\end{array}\right]  \tag{C11}\\
& =\left[\begin{array}{cc}
\mathbf{I}+\mathbf{M}_{\mathbf{U}} & \\
& \mathbf{I}
\end{array}\right]\left[\begin{array}{cc}
\hat{\boldsymbol{\Sigma}} & \\
& \hat{\boldsymbol{\Sigma}}^{\circ}
\end{array}\right]+\mathbf{G}_{\mathbf{U}}, \quad \mathbf{G}_{\mathbf{U}}=\left[\begin{array}{cc}
\hat{\mathbf{U}}^{\mathrm{H}} \mathbf{E}_{\mathbf{U}} & \mathbf{E}_{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}^{\circ} \\
\left(\hat{\mathbf{U}}^{\circ}\right)^{\mathrm{H}} \mathbf{E}_{\mathbf{U}} & \mathbf{0}
\end{array}\right] \\
& =\left[\begin{array}{ll}
\hat{\boldsymbol{\Sigma}} & \\
& \hat{\boldsymbol{\Sigma}}^{\circ}
\end{array}\right]\left[\begin{array}{cc}
\mathbf{I}+\mathbf{M}_{\mathbf{V}} & \\
& \mathbf{I}
\end{array}\right]+\mathbf{G}_{\mathbf{V}}, \quad \mathbf{G}_{\mathbf{V}}=\left[\begin{array}{cc}
\mathbf{E}_{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}} & \mathbf{E}_{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}^{\circ} \\
\left(\hat{\mathbf{U}}^{\circ}\right)^{\mathrm{H}} \mathbf{E}_{\mathbf{U}} & \mathbf{0}
\end{array}\right]
\end{align*}
$$

Proof $\quad$ First choose matrices $\hat{\mathbf{U}}_{1}^{\mathrm{H}} \hat{\mathbf{U}}_{1}=\mathbf{I}$ and $\hat{\mathbf{V}}_{1}^{\mathrm{H}} \hat{\mathbf{V}}_{1}=\mathbf{I}$ so $\hat{\mathbf{U}}^{\mathrm{H}} \hat{\mathbf{U}}_{1}=\mathbf{0}$ and $\hat{\mathbf{V}}^{\mathrm{H}} \hat{\mathbf{V}}_{1}=$ $\mathbf{0}$. This can be done e.g. using a QR-decomposition. If $\hat{\mathbf{U}}_{2} \hat{\boldsymbol{\Sigma}}^{\circ} \hat{\mathbf{V}}_{2}^{\mathrm{H}}=\mathbf{X}$ is an SVD of $\mathbf{X}=\hat{\mathbf{U}}_{1}^{\mathrm{H}} \mathbf{A} \hat{\mathbf{V}}_{1}$, then $\hat{\mathbf{U}}^{\circ}=\hat{\mathbf{U}}_{1} \hat{\mathbf{U}}_{2}$ and $\hat{\mathbf{V}}^{\circ}=\hat{\mathbf{V}}_{1} \hat{\mathbf{V}}_{2}$ satisfies eq. (C10).
Equation (C11) can be verified by direct multiplication. Inserting A $\hat{\mathbf{V}}=\hat{\mathbf{U}} \hat{\boldsymbol{\Sigma}}+$ $\mathbf{E}_{\mathbf{U}}$ leads to eq. (C12). Similarly we get eq. (C13) by inserting the $\mathbf{A}^{\mathrm{H}} \hat{\mathbf{U}}=$ $\hat{\mathbf{V}} \hat{\boldsymbol{\Sigma}}^{\mathrm{H}}+\mathbf{E}_{\mathbf{V}}$.
Theorem C17: (NPSVD accuracy) Let the conditions be as in lemma C16, and assume that no element of $\mathbf{M}_{\mathbf{U}}$ or $\mathbf{M}_{\mathbf{V}}$ exceeds $T$ in magnitude, where $k T=\kappa<1$. Then there exists a permutation, $P(i)$, that maps the $k$ numerical triplets, $\hat{\sigma}_{i}$, to exact triplets, $\sigma_{P(i)}$, of $\mathbf{A}$ that satisfies:

$$
\left|\hat{\sigma}_{i}-\sigma_{P(i)}\right| \leq \hat{\sigma}_{i} \frac{2 \kappa}{1-\kappa}+\left(\left\|\mathbf{E}_{\mathbf{U}}\right\|_{2}+\left\|\mathbf{E}_{\mathbf{V}}\right\|_{2}\right) \frac{1+\kappa}{1-\kappa}
$$

Proof
Let $\check{\sigma}_{i}$ be the singular values of $\mathbf{C}=\check{\mathbf{U}}^{\mathrm{H}} \mathbf{A} \check{\mathbf{V}}$. Using Ostrowski's theorem on eq. (C11) gives $\check{\sigma}_{i}=\xi_{1} \sigma_{i}$, where $1-\kappa \leq\left|\xi_{1}\right| \leq 1+\kappa$ by theorem C10.

Applying Ostrowski's theorem on the first term in eq. (C12) followed by Weyl's theorem gives $\check{\sigma}_{i}=\xi_{2} \hat{\sigma}_{P^{-1}(i)}+\xi_{3}$, where $P(i)$ is a permutation, $1-\kappa \leq \xi_{2} \leq$ $1+\kappa$, and $\left|\xi_{3}\right| \leq\left\|\mathbf{G}_{\mathbf{U}}\right\|_{2}$.

Eliminating $\check{\sigma}_{i}$ from these expressions leads to:

$$
\hat{\sigma}_{P^{-1}(i)}-\sigma_{i}=\left(1-\xi_{2} / \xi_{1}\right) \hat{\sigma}_{P^{-1}(i)}-\xi_{3} / \xi_{1}
$$

or, if we invert the perturbation,

$$
\begin{equation*}
\hat{\sigma}_{i}-\sigma_{P(i)}=\left(1-\xi_{2} / \xi_{1}\right) \hat{\sigma}_{i}-\xi_{3} / \xi_{1} \tag{C14}
\end{equation*}
$$

Inserting the extreme values for $\xi_{1}$ and $\xi_{2}$ shows that $(1-\kappa) /(1+\kappa) \leq \xi_{2} / \xi_{1} \leq$ $(1+\kappa) /(1-\kappa)$, which leads to $\left|1-\xi_{2} / \xi_{1}\right| \leq 2 \kappa /(1-\kappa)$.

To bound $\left|\xi_{3} / \xi_{1}\right|$ we use corollary $\mathbf{C} 2$ on $\mathbf{G}_{\mathbf{U}}$ :

$$
\left|\xi_{3} / \xi_{1}\right| \leq \max \left\{\left\|\mathbf{E}_{\mathbf{U}}\right\|_{2},\left\|\mathbf{E}_{\mathbf{V}}\right\|_{2}\right\} \frac{1}{1-\kappa}+\frac{1+\kappa}{1-\kappa}\left\|\mathbf{E}_{\mathbf{U}}\right\|_{2}
$$

An analogous treatment of eq. (C13) gives:

$$
\left|\xi_{3} / \xi_{1}\right| \leq \max \left\{\left\|\mathbf{E}_{\mathbf{U}}\right\|_{2},\left\|\mathbf{E}_{\mathbf{V}}\right\|_{2}\right\} \frac{1}{1-\kappa}+\frac{1+\kappa}{1-\kappa}\left\|\mathbf{E}_{\mathbf{V}}\right\|_{2}
$$

Choosing the lowest of these bounds:

$$
\begin{aligned}
\left|\xi_{3} / \xi_{1}\right| & \leq \max \left\{\left\|\mathbf{E}_{\mathbf{U}}\right\|_{2},\left\|\mathbf{E}_{\mathbf{V}}\right\|_{2}\right\} \frac{1}{1-\kappa}+\min \left\{\left\|\mathbf{E}_{\mathbf{U}}\right\|_{2},\left\|\mathbf{E}_{\mathbf{V}}\right\|_{2}\right\} \frac{1+\kappa}{1-\kappa} \\
& \leq\left(\left\|\mathbf{E}_{\mathbf{U}}\right\|_{2}+\left\|\mathbf{E}_{\mathbf{V}}\right\|_{2}\right) \frac{1+\kappa}{1-\kappa}
\end{aligned}
$$

Inserting all these bound in eq. (C14) completes the proof.
Although theorem C17 may seem unwieldy, it bounds the difference between the exact singular values and those of an NPSVD in terms of quantities that are relatively easy to work with in a program.

## C8 Triplet accuracy and aliasing

The singular triplets are meant to represent the matrix $\mathbf{A}$, so we want to interpret the singular triplets as approximations to exact triplets. This leads to the following requirements, illustrated in fig. C.1:

1. Each numerical triplet should map to an exact triplet of $\mathbf{A}$.
2. The difference between the singular value of a numerical triplet and the exact one it maps to should be bounded by a small multiple of $\|\mathbf{A}\|_{2} \epsilon_{M}$.
3. The exact triplets in the map should be exactly orthogonal, i.e. the exact triplets should be unique.

Figure C.1:
All numerical triplets should map to a unique exact triplet.

Coupling between exact and numerical triplets
Exact singular values


Theorem C17 is the main vehicle we will use to prove that these requirements are met. However, by itself it leads to unnecessarily large bounds on the singular values. If we apply it to a single triplet instead, we get the following result.

Corollary C18: (Single triplet accuracy) Let $\left(\hat{\mathbf{u}}_{i}, \hat{\mathbf{v}}_{i}, \hat{\sigma}_{i}\right)$ be a numerical triplet where $\| \mathbf{A} \hat{\mathbf{v}}_{i}-$ $\hat{\sigma} \hat{\mathbf{u}}_{i} \| \leq t_{u}$ and $\left\|\mathbf{A}^{\mathrm{H}} \hat{\mathbf{u}}_{i}-\hat{\sigma} \hat{\mathbf{v}}_{i}\right\| \leq t_{v}$. Further assume that $\hat{\mathbf{u}}_{i}$ and $\hat{\mathbf{v}}_{i}$ are normalized, so that $1-\epsilon_{o} \leq\left\|\hat{\mathbf{u}}_{i}\right\|_{2},\left\|\hat{\mathbf{v}}_{i}\right\|_{2} \leq 1+\epsilon_{o}$. If $\epsilon_{o} \leq 0.05$ then there exist an exact singular triplet of $\mathbf{A},\left(\mathbf{u}_{j}, \mathbf{v}_{j}, \sigma_{j}\right)$, so that

$$
\left|\sigma_{j}-\hat{\sigma}_{i}\right| \leq 1.11\left(t_{u}+t_{v}\right)+2.11\left|\hat{\sigma}_{i}\right| \epsilon_{o}
$$

This asserts that the computed singular values are accurate as long as the numerical triplets maps to unique exact triplets. However, if we are given a set of numerical triplets and the intervals allowed by the corollary does not overlap, then the exact triplets must be unique.

If two singular triplets have overlapping ranges, they may both have a large component corresponding to the same exact triplet we wish to extract. We say that such triplets alias each other. In this case we must go back to theorem C17 and include all triplets that may potentially alias the same singular value.

## C9 One-sided reorthogonalization

One interesting observation is that theorem C15 and C17 are, in some sense, reciprocal. The first states that if either the right or left numerical vectors are kept orthogonal at level $\epsilon_{M}$, then the other vectors are roughly orthogonal at level $T \approx \epsilon_{M} / \hat{\sigma}$. On the other hand, theorem C17 says that to compute the singular values accurately (relative to the largest singular value), the level of orthogonality need only satisfy $2 k T \hat{\sigma} \approx \epsilon_{M}$.

This leads us to think that keeping one set of triplets orthogonal suffices to compute singular values to full accuracy. The following theorem is a precise statement of this argument.

Theorem C19: (One-sided orthogonality accuracy) Assume a set of sorted numerical triplets, $\left(\hat{\mathbf{u}}_{i}, \hat{\mathbf{v}}_{i}, \sigma_{i}\right), \hat{\sigma}_{1} \geq \hat{\sigma}_{2} \geq \cdots \geq \hat{\sigma}_{k}$, satisfies $\left\|\mathbf{A} \hat{\mathbf{v}}_{i}-\hat{\sigma}_{i} \hat{\mathbf{u}}_{i}\right\| \leq t_{u}$ and $\| \mathbf{A}^{\mathrm{H}} \hat{\mathbf{u}}_{i}-$ $\hat{\sigma}_{i} \hat{\mathbf{v}}_{i} \| \leq t_{v}$. Furthermore, assume that all pairs satisfies either $\left|\hat{\mathbf{u}}_{j}^{\mathrm{H}} \hat{\mathbf{u}}_{i}\right| \leq T$ or $\left|\hat{\mathbf{v}}_{j}^{\mathrm{H}} \hat{\mathbf{v}}_{i}\right| \leq T$. Finally assume that $\kappa_{k}=k T+k\left(t_{u}+t_{v}\right) / \hat{\sigma}_{k}<1$.

Then the numerical triplets correspond to unique exact triplets of $\mathbf{A}, \hat{\sigma}_{i} \approx \sigma_{P_{k}(i)}$, with the difference bounded by:

$$
\left|\hat{\sigma}_{i}-\sigma_{P_{k}(i)}\right| \leq \hat{\sigma}_{i} \frac{2 k T}{1-\kappa_{k}}+2 k\left(\frac{\hat{\sigma}_{i}}{\hat{\sigma}_{k}}\right)\left(t_{u}+t_{v}\right)+\left(\left\|\mathbf{E}_{\mathbf{U}}\right\|_{2}+\left\|\mathbf{E}_{\mathbf{V}}\right\|_{2}\right) \frac{1+\kappa_{k}}{1-\kappa_{k}}
$$

Proof In the case where $\left|\mathbf{M}_{\mathbf{U}}\right| \leq T$ we have, from theorem $\mathbf{C} 15$, that $\left|\mathbf{M}_{\mathbf{V}}\right| \leq T+$ $\left(t_{u}+t_{v}\right) / \min \left\{\hat{\sigma}_{i}, \hat{\sigma}_{j}\right\}$. Otherwise we have $\left|\mathbf{M}_{\mathbf{V}}\right| \leq T$, and we similarly get $\left|\mathbf{M}_{\mathbf{U}}\right| \leq T+\left(t_{u}+t_{v}\right) / \min \left\{\hat{\sigma}_{i}, \hat{\sigma}_{j}\right\}$. In either case theorem C17 applies with $\kappa=\kappa_{k}$, which is the stated result.

The analysis would be complete had it not been for the term $\hat{\sigma}_{i} / \hat{\sigma}_{k}$. We can eliminate this by setting $k=i$, but there is no guarantee that $P_{k}(i)$ is the same for all $k \geq i$.

The problem is illustrated in fig. C.2. At some point we have a number of triplets determined accurately, but when we introduce a new triplet it may not be possible to map the previous triplets without increasing their error-bars substantially. In the example the escalation occurs even though the triplets are added in order of strictly decreasing singular values. Also, the new triplet can be arbitrarily close to an exact singular value.

Escalating errors for aliasing triplets
Figure C.2:
Adding a numerical triplet may cause the errors of previously triplets to increase.


The cure is to give up the pure one-sided approach and orthogonalize each triplet against all potentially aliasing triplets. This enforces the condition that $\kappa_{i} \approx \kappa_{k}$ (by lowering $\kappa_{k}$ ) if the error in the $k$ 'th triplet could escalate to triplet $i$. With this modification the factor $\hat{\sigma}_{i} / \hat{\sigma}_{k}$ vanishes.

The key observation is that we only need to keep both left and right singular vectors for a few triplets at a time. This means that the singular values can be computed to full machine precision with memory that scales in the smallest dimension of $\mathbf{A}$.

## Appendix

## BBTools reference

This appendix contains a list of the functions provided with BBTools v0.2.4b, released 16-Sep-2005. Please note that function provided for compatibility is not addressed here. See chapter 10 for compatibility with MATLAB.

## List of categories and tables

Getting started with BBTools ..... D. 1
Elementary black-box operators ..... D. 2
Toeplitz and related operators ..... D. 3
Image- and hypercube operators ..... D. 4
Tomographic operators ..... D. 5
Miscellaneous operators ..... D. 6
Black-box creation, manipulation, and destruction ..... D. 7
Black-box information and verification ..... D. 8
Tests and probes for iterative algorithms ..... D. 9
Explicit black-box versions of standard functions ..... D. 10
Statistics ..... D. 11
Miscellaneous support functions ..... D. 12
Computing with black-box operators ..... D. 13
Equation solving and regularization ..... D. 14
SVD computations ..... D. 15

Table D.1:
Getting started with BBTools.

| Function | Description |
| :--- | :--- |
| bbdoc | Start online documentation for the toolbox |
| bbinstall | (GUI) Install, uninstall, download, and upgrade |

Table D.2:
Elementary blackbox operators.

Table D.3:
Toeplitz and related operators.

Table D.4:
Image- and hypercube operators.

Table D.5:
Tomographic operators.

Table D.6:
Miscellaneous operators.

Table D.7:
Black-box creation, manipulation, and destruction.

| Function | Description |
| :--- | :--- |
| bbmatrix | Create black-box from a Matlab matrixx |
| bbinv | Create black-box to invert a Matlab matrix |
| bbindex | Create indexing black-box |
| bbeye | Create identity black-box |
| bbdiag | Create diagonal black-box |
| bbconst | Create black-box with constant elements |
| bbzeros | Create black-box of zeros |
| bbmean | Create averaging black-box |


| Function | Description |
| :--- | :--- |
| bbtoeplitz | Create Toeplitz black-box |
| bbdbltoeplitz | Create doubly Toeplitz black-box |
| bbblktoep1s | Create sparse block-Toeplitz type-1 black-box |
| bbhankel | Create Hankel black-box |


| Function | Description |
| :--- | :--- |
| bbflipdim | Create black-box that flips a specified dimension |
| bbpermute | Create dimension-permuting black-box |
| bbredim | Create redimensioning black-box |
| bbrescale | Create rescaling black-box |
| bbconvn | Create black-box for $N$-dimensional convolution |


| Function | Description |
| :--- | :--- |
| bbtomoemit | Create black-box model of an emission tomograph |
| bbradon | Create black-box for a Radon transform |
| bbfbp | Create inverse of Radon using filtered backprojection |


| Function | Description |
| :--- | :--- |
| bbhadamard | Create Hadamard black-box |
| bblaplacian | Create a finite-difference Laplace operator |


| Function | Description |
| :--- | :--- |
| blackbox | Create black-box operator |
| bbclear | Remove a black-box without leaving garbage |
| bbqualify | Add informational qualifier to a black-box |
| bbsimplify | Simplify a black-box operator |
| bbinline | Inline and freeze a black-box operator |
| get | Get black-box parameter |
| set | Set black-box parameter |
| isa | Check type of matrix operator |

Table D.8:
Black-box information and verification.

Table D.9:
Tests and probes for iterative algorithms.

Table D.10:
Explicit black-box versions of standard functions.

Table D.11:
Statistics.

Table D.12:
Miscellaneous support functions.

Table D.13:
Computing with black-box operators.

Table D.14:
Equation solving and regularization.

| Function | Description |
| :---: | :---: |
| bbinspect | (GUI) Utility to explore an evaluation tree |
| bbshowtree | Show evaluation tree for a black-box |
| bbtauest | Estimate accuracy of a black-box statistically |
| bbcompare | Compare two operators for equality |
| Function | Description |
| bbnoisy | Create a noisy version of an operator |
| bbcount | Create black-box that counts invocations |
| bbcountstat | Get count statistics |
| bbcountreset | Reset counts statistics |
| Function | Description |
| bbadd | Add a number of operators |
| bbprod | Form a product of operators and scalars |
| bbkron | Form Kronecker-product of two operators |
| bbrepmat | Replicate and tile an operator |
| Function | Description |
| bbrand | Draw random samples from a given distribution |
| bbpca | Principal Component Analysis |
| bbempoisson | Solve linear system with Poisson noise |
| Function | Description |
| bbkernel | Make a standard convolutional kernel |
| bbphantom | Make a standard test-image |
| bbimagein | Display intensity-image with gamma-correction |
| Function | Description |
| bbmult | Multiply black-box with a matrix |
| bbmultq | Quick evaluation of product (no checks) |
| bbmultqh | Quick evaluation of adjoint product (no checks) |
| bbfunc | Get functions for matrix-vector products |
| bbfuncm | Get functions for matrix-matrix products |
| Function | Description |
| bbpresolve | Prepare iterative solution of a linear system |
| bbsolve | Compute a solution of a linear system |
| bbregsolve | Compute a regularized solution of a linear system |
| bbsolveopt | Create options for solving a linear system |
| bbsolvemem | Estimate memory usage for solving a linear system |
| bblcurve | Compute L-curve and estimate regularization parameter |
| Function | Description |
| bbsvds | Compute singular values/triplets of an operator |
| bbsvdf | Compute singular triplets of an operator with call-back |
| bbsvdopt | Create options for computing an SVD |
| bbsvdmem | Estimate memory usage for an SVD computation |

Table D.15:
SVD computations.

## Bibliography

[1] Helmut Alt, Bernd Behrends, and Johannes Blömer. Approximate matching of polygonal shapes. In Proceedings of the 7th Annual Symposium on Computational Geometry, pages 186-193. ACM Press, 1991. ISBN 0-89791-426-0.
[2] Edward Anderson, Zhaojun Bai, Christian Heinrich Bischof, L. Susan Blackford, James Demmel, Jack J. Dongarra, Jeremy Du Croz, Anne Greenbaum, Sven Hammarling, Alan McKenney, and Danny Chris Sorensen. LAPACK Users' Guide. SIAM, Philadelphia, PA, third edition, 1999. ISBN 0-89871-447-8.

- http://www.netlib.org/lapack/lug/index.html.
[3] Babak A. Ardekani, Michael Braun, Brian F. Hutton, Iwao Kanno, and Hidehiro Iida. Minimum crossentropy reconstruction of PET images using prior anatomical information. Physics in Medicine and Biology, 41(11):2497-2517, 1996.
- http://www.iop.org/EJ/abstract/0031-9155/41/11/018.
[4] Richard Barrett, Michael W. Berry, Tony F. Chan, James Demmel, June M. Donato, Jack Dongarra, Victor Eijkhout, Roldan Pozo, Charles Romine, and Henk A. van der Vorst. Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods. Web, 2nd edition, 1994. Note: The first edition is also published by SIAM, 1993, ISBN 0-89871-328-5.
- http://www.netlib.org/linalg/html_templates/Templates.html.
[5] Bernard Bendriem and David W. Townsend, editors. The Theory and Practice of $3 D$ PET. Kluwer Academic Publishers, 1998. ISBN 0-7923-5108-8.
[6] Michael W. Berry. Large-scale sparse singular value computations. International Journal of Supercomputer Applications, 6(1):13-49, 1992. ISSN 0890-2720.
- http://www.netlib.org/svdpack/ijsa.ps.
[7] Michael W. Berry. SVDPACK. University of Tennessee, 1992.
- http://www.netlib.org/svdpack/index.html.
[8] Christopher M. Bishop. Neural networks for pattern recognition. Oxford University Press, 1995. ISBN 0-19-853864-2.
[9] Åke Björck, Eric Grimme, and Paul Van Dooren. An implicit shift bidiagonalization algorithm for illposed systems. BIT, 34(4):510-534, 1994.
[10] L. Susan Blackford, Jaeyoung Choi, Andrew J. Cleary, Eduardo Francisco D’Azevedo, James Demmel, Inderjit Singh Dhillon, Sven Hammarling, Greg Henry, Antoine Petitet, Ken Stanley, David William Walker, and Richard Clint Whaley. ScaLAPACK Users' Guide. SIAM, Philadelphia, PA, 1997. ISBN 0-89871-397-8.
- http://www.netlib.org/scalapack/slug/.
[11] James E. Bowsher, Valen E. Johnson, Timothy G. Turkington, Ronald J. Jaszczak, Carey E. Floyd, Jr, and Edward R. Coleman. Bayesian reconstruction and use of anatomical a priori information for emission tomography. IEEE Transactions on Medical Imaging, 15(8):673-686, 1996.
[12] Jerrold T. Bushberg, J. Anthony Seibert, Edwin M. Jr. Leidholdt, and John M. Boone. The Essential Physics of Medical Imaging. Lippincott Williams \& Wilkins, 2nd edition, 2002. ISBN 0-683-30118-7.
[13] Daniela Calvetti, L. Reichel, and Danny Chris Sorensen. An implicitly restarted Lanczos method for large symmetric eigenvalue problems. Electronic Transactions on Numerical Analysis, 2:1-21, 1994.
[14] Jon F. Claerbout. Image Estimation by Example. Web, 2004.
http://sepwww.stanford.edu/sep/prof/toc_html/index.html.
[15] James Weldon Demmel. Applied Numerical Linear Algebra. SIAM, Philadelphia, PA, 1997. ISBN 0-89871-389-7.
[16] Stephen E. Derenzo, William W. Moses, Ronald H. Huesman, and Thomas F. Budinger. Critical instrumentation for resolution smaller than 2 mm , high sensitivity brain PET. In Kazuo Uemura, Terry Jones, Niels A. Lassen, and Iwao Kanno, editors, Quantification of Brain Function, pages 25-37. Elsevier Science Publishers B.V., 1993. ISBN 0-444-89859-x.
[17] Jack J. Dongarra and David William Walker. The design of linear algebra libraries for high performance computers. Technical Report UT-CS-93-188, University of Tennessee, 1993.
- http://www.netlib.org/lapack/lawnspdf/lawn58.pdf.
[18] D. A. Dougherty, J. Nissanov, and Gene R. Gindi. Construction details of an autoradiographic-based rat brain phantom for emission tomography. In Nuclear Science Symposium, 1998. Conference Record., volume 2, pages 1250-1254. IEEE, 1998. ISBN 0-7803-5021-9.
[19] Robert Joseph English. SPECT, Single-Photon Emission Computed Tomography: a Primer. Society of Nuclear Medicine, 3rd edition, 1995. ISBN 0-932004-43-1.
[20] Jeffrey A. Fessler. Penalized weighted least squares image reconstruction for PET. IEEE Transactions on Medical Imaging, 13(2):290-300, 1994.
[21] Jeffrey A. Fessler. TERSE - a transmission and emission reconstruction environment for SPECT. Web, 1996.
- http://www.eecs.umich.edu/~fessler/terse/.
[22] Jeffrey A. Fessler. Statistical image reconstruction methods for transmission tomography. In Milan Sonka and J. Michael Fitzpatrick, editors, Handbook of Medical Imaging: medical image processing and analysis, volume 2, pages 1-70. The International Society for Optical Engineering, 2000. ISBN 0-8194-3622-4.
[23] Jeffrey A. Fessler. ASPIRE - A Sparse Iterative Reconstruction Library. The University of Michigan, 2003.
- http://www.eecs.umich.edu/~fessler/aspire/index.html.
[24] Free Software Foundation, Inc. GNU general public license, version 2, 1991. In [68], pages 195-200.
- http://www.fsf.org/copyleft/gpl.html.
[25] Philip W. Goetz, editor. Encyclopadia Britannica. Encyclopædia Britannica, Inc., 15th edition, 1991. ISBN 0-85229-529-4.
[26] Gene Howard Golub and Charles Francis Van Loan. Matrix Computations. John Hopkins University Press, 2nd edition, 1989. ISBN 0-8018-3739-1.
[27] Gene Howard Golub and Charles Francis Van Loan. Matrix Computations. John Hopkins University Press, 3rd edition, 1996. ISBN 0-8018-5414-8.
[28] S.F. Haber, Stephen E. Derenzo, and D. Uber. Application of mathematical removal of positron range blurring in positron emission tomography. IEEE Transactions on Nuclear Science, 27(3):1293-1299, 1990.
- http://breast.lbl.gov/~wwwinstr/publications/Papers/range.pdf.
[29] Per Christian Hansen. Regularization tools: A matlab package for analysis and solution of discrete ill-posed problems. Numerical Algorithms, 6:1-35, 1994.
[30] Per Christian Hansen. Rank-Deficient and Discrete Ill-Posed Problems. SIAM, 1998. ISBN 0-89871-403-6.
[31] Per Christian Hansen. Regularization tools. Technical Report IMM-REP-98-6, Technical University of Denmark, 2001.
- http://www.imm.dtu.dk/~pch/Regutools/.
[32] Michiel E. Hochstenbach. A Jacobi-Davidson type SVD method. SIAM Journal on Scientific Computing, 23(2):606-628, 2001. ISSN 1064-8275.
[33] Esben Høgh-Rasmussen. BBTools - a Matlab Toolbox for Black-Box Computations. Neurobiology Research Unit, Copenhagen University Hospital, 2005.
- http://nru.dk/software/bbtools/.
[34] Esben Høgh-Rasmussen. A cache-friendly kin to bidiagonalization with orthogonalization and restart, 2005. Submitted to BIT.
[35] Søren Holm, Peter Aundal Toft, and Mikael Jensen. Estimation of the noise contributions from blank, transmission and emission scans in PET. In IEEE Medical Imaging Conference, pages 1470-1474, October 1995.
[36] John H. Hubbell and Stephen M. Seltzer. Tables of x-ray mass attenuation coefficients and mass energyabsorption coefficients. Technical report, National Institute of Standards and Technology, 1997. Originally published as NISTIR 5632, NIST, Gaithersburg, MD (1995).
- http://physics.nist.gov/xaamdi.
[37] Harold Malcolm Hudson and Richard S. Larkin. Accelerated image reconstruction using ordered subsets of projection data. IEEE Transactions on Medical Imaging, 13(4):601-609, 1994.
[38] Michael Jacobsen. Modular Regularization Algorithms. PhD thesis, 2004.
- http://www2.imm.dtu.dk/pubdb/p.php?id=3380.
[39] Avinash C. Kak and Malcolm Slaney. Principles of Computerized Tomographic Imaging. IEEE Press, 1988. ISBN 0-8794-2198-3.
- http://www.slaney.org/pct/pct-toc.html.
[40] Satoshi Kawata and Orhan Nalcioglu. Constrained iterative reconstruction by the conjugate gradient method. IEEE Transactions on Medical Imaging, 4(2):65-71, 1985.
[41] John Kirman. Good Style, Writing for Science and Technology. E \& FN Spon, 1994. ISBN 0419171908.
[42] Donald Ervin Knuth. Literate Programming. Center for the Study of Language and Inf., 1992. ISBN 0937073814.
[43] Samuel Kotz, Normal Lloyd Johnson, and Campbell B. Read, editors. Encyclopedia of Statistical Sciences, volume 8. John Wiley \& Sons, Inc., 1988. ISBN 0-471-05556-5.
[44] Erwin Kreyszig. Advanced Engineering Mathematics. John Wiley \& Sons, Inc., 7th edition, 1993. ISBN 0-47-155380-8.
[45] Kenneth Lange and Richard Carson. EM reconstruction algorithms for emission and transmission tomography. Journal of Computer Assisted Tomography, 8(2):306-316, 1982.
[46] Rasmus Munk Larsen. Lanczos bidiagonalization with partial reorthogonalization. Technical Report DAIMI PB-357, Department of Computer Science, Aarhus University, 1998.
- http://sun.stanford.edu/~rmunk/PROPACK/paper.ps.gz.
[47] Rasmus Munk Larsen. PROPACK. Stanford University, 2005.
- http://sun.stanford.edu/~rmunk/PROPACK/.
[48] Soo-Jin Lee, Anand Rangarajan, and Gene Gindi. Bayesian image reconstruction in SPECT using higher order mechanical models as priors. IEEE Transactions on Medical Imaging, 14(4):669-680, 1995.
[49] Richard Bruno Lehoucq and Jennifer A. Scott. An evaluation of Arnoldi based software for sparse nonsymmetric eigenproblems. Technical Report RAL-TR-96-023, Rutherford Appleton Laboratory, Didcot, England, 1996.
[50] Richard Bruno Lehoucq, Danny Chris Sorensen, and Chao Yang. ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods. Rice University, 1997.
- http://www.caam.rice.edu/software/ARPACK/.
[51] Ren-Cang Li. Relative perturbation theory: I. eigenvalue and singular value variations. SIAM Journal on Matrix Analysis and Applications, 19(4):956-982, 1998.
[52] Kanti V. Mardia, John T. Kent, and John M. Bibby. Multivariate Analysis. Academic Press, London, 1995. ISBN 0-12-471252-5.
[53] Ronald B. Morgan. GMRES with deflated restarting. SIAM Journal on Scientific Computing, 24(1): 20-37, 2002.
[54] Alan Victor Oppenheim and Ronald W. Schafer. Discrete-Time Signal Processing. Prentice Hall, Inc., 1989. ISBN 0-13-216771-9.
[55] Christopher Conway Paige and Michael Alan Saunders. LSQR: an algorithm for sparse linear equations and sparse least squares. ACM Transactions on Mathematical Software (TOMS), 8(1):43-71, 1982. ISSN 0098-3500.
[56] Marko Petkovsek, Herbert Wilf, and Doron Zeilberger. $A=B$. AK Peters, Ltd., 1996. ISBN 1-56881-063-6.
- http://www.cis.upenn.edu/~wilf/AeqB.html.
[57] Peter Alshede Philipsen, Lars Kai Hansen, and Peter Aundal Toft. Mean field reconstruction with snaky edge hints. In Proceedings of the Interdisciplinary Inversion Conference on Methodology, Computation and Integrated Applications IIC'95 1996, pages 312-319. Springer Verlag, Berlin, 1996.
[58] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. Numerical recipes in C: The art of scientific computing. Cambridge University Press, New York, 2nd edition, 1992. ISBN 0-521-43108-5.
[59] Johann Karl August Radon. Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten. Berichte Sachsische Akademie der Wissenschaften, 69:262-267, 1917. Note: The title roughly translates to: "On the determination of functions from the integrals on certain manifolds".
[60] Sheldon M. Ross. Introduction to probability and statistics for engineers and scientists. John Wiley \& Sons, Inc., 1987. ISBN 0-471-81752-X.
[61] Yousef Saad and Henk A. van der Vorst. Iterative solution of linear systems in the 20th century. Journal of Computational and Applied Mathematics, 123(1-2):1-33, 2000.
- http://www.sciencedirect.com/science/article/-B6TYH-41H9NJX-2/2/235bc014228315f21059108e5cef9940.
[62] Semiconductor Industry Association. The International Technology Roadmap for Semiconductors. ITRS, 2003 edition edition, 2003.
- http://public.itrs.net/.
[63] Semiconductor Industry Association. The International Technology Roadmap for Semiconductors. ITRS, 2004 update edition, 2004.
- http://public.itrs.net/.
[64] Lawrence A. Shepp and Yehuda Vardi. Maximum likelihood reconstruction for emission tomography. IEEE Transactions on Medical Imaging, 1:113-122, 1982.
[65] Donald Lee Snyder and Michael I. Miller. The use of sieves to stabilise images produces with the EM algorithm for emission tomography. IEEE Transactions on Nuclear Science, 32(5):3864-3872, 1985.
[66] Danny Chris Sorensen. Implicit application of polynomial filters in a k-step arnoldi method. SIAM Journal on Matrix Analysis and Applications, 13(1):357-385, 1992.
[67] Danny Chris Sorensen. Implicitly restarted Arnoldi/Lanczos methods for large scale eigenvalue calculations. Technical Report TR-96-40, Rice University, 1996.
- http://techreports.larc.nasa.gov/icase/1996/icase-1996-40.pdf.
[68] Richard M. Stallman. Free Software, Free Society: Selected Essays of Richard M. Stallman. Free Software Foundation, Inc., 2002. ISBN 1-882114-98-1.
- http://www.gnu.org/doc/book13.html.
[69] Gilbert Wright Stewart. Matrix Algorithms: Basic decompositions. SIAM, 1998. ISBN 0-89871-414-1.
[70] Gilbert Wright Stewart and Ji-guang Sun. Matrix Perturbation Theory. Academic Press, 1990. ISBN 0-12-670230-6.
[71] Gilbert Strang. Introduction to Linear Algebra. Wellesley-Cambridge Press, 2nd edition, 1998. ISBN 0-9614088-5-5.
[72] Peter Aundal Toft. The Radon Transform - Theory and Implementation. PhD thesis, Department of Mathematical Modelling, Technical University of Denmark, 1996.
-http://pto.linux.dk/PhD/.
[73] Lloyd Nicholas Trefethen and David Bau III. Numerical Linear Algebra. SIAM, Philadelphia, PA, 1997. ISBN 0-89871-487-7.
[74] Tomithy G. Turkington. Introduction to PET instrumentation. Journal of Nuclear Medicine Technology, 29(1):4-11, 2001.
[75] SimSET - Simulation System for Emission Tomography. University of Washington, 2004.
- http://depts.washington.edu/~simset/html/simset_home.html.
[76] Eugene Veklerov and Jorge Llacer. The feasibility of images reconstructed with the method of sieves. IEEE Transactions on Nuclear Science, 37(2):835-841, 1990.
[77] Gustav Konrad von Schulthess, editor. Clinical Positron Emission Tomography, Correlation with Morphological Cross-Sectional Imaging. Lipincott williams \& Wilkins, 2000. ISBN 0-7817-1756-6.
[78] Steve Webb, editor. The Physics of Medical Imaging. Institute of Physics Publishing, 1988. ISBN 0-85274-349-1.
[79] Wolfang Weinrich and Hans-Joachim Kretschmann. Cranial neuroimaging and clinical neuroanatomy: computed tomography and magnetic resonance imaging. Thieme Medical Publishers, Inc., 1992.
[80] Richard Clint Whaley, Antoine Petitet, and Jack J. Dongarra. Automated empirical optimization of software and the ATLAS project. Parallel Computing, 27(1-2):3-35, 2001.
- http://www.netlib.org/lapack/lawns/lawn147.ps.
[81] Peter Kjær Willendrup. Point-spread functions in tomography using filtered back-projection reconstruction. Master's thesis, University of Copenhagen, 2000.
- http://www.willendrup.org/thesis.pdf.
[82] Keith John Worsley, A.C. Evans, S. Marrett, and P. Neelin. A three-dimensional statistical analysis for cbf activation studies in human brain. Journal of Cerebral Blood Flow and Metabolism, 12:900-918, 1992.
- http://www.math.mcgill.ca/~keith/jcbf/jcbf.abstract.html.
[83] Keith John Worsley, S. Marrett, P. Neelin, A.C. Vandal, K.J. Friston, and A.C. Evans. A unified statistical approach for determining significant signals in images of cerebral activation. Human Brain Mapping, 4: 58-73, 1996.
http://www.math.mcgill.ca/~keith/unified/unified.abstract.html.
[84] Mehmet Yavuz. Statistical tomographic image reconstruction methods for randoms-precorrected PET measurements. Technical report, The University of Michigan, 2000. Ph.D.thesis.
[85] Mehmet Yavuz and Jeffrey A. Fessler. Objective functions for tomographic reconstruction from randoms-precorrected PET scans. In Proceedings of IEEE Nuclear Science Symposium and Medical Imaging Conference, volume 2, pages 1067-71. IEEE, 1996.
[86] Daniel Zwillinger, editor. Standard Mathematical Tables and Formulae. CRC Press, 30th edition, 1996. ISBN 0-8493-2479-3.


## Index

2D acquisition, 19
3D acquisition, 19
absolute norm, 147
accuracy
asymmetric, 99
acknowledgments, $v$
acquisition
continuous, 23
matrix, 127
step-and-shoot, 23
acquisitions mode, 19
active contour model, 56
adaptive snake, 56
algorithm
arnoldi, 88
back-project, 37
bidiag, 87
init, 85
iter, 146
randsolve, 108
reorthogonalize, 91
restart, 93
tridiag, 87
unit, 85
update, 85
aliasing, 40, 101, 155
Alzheimer's disease, iv
anatomy, 58
Arnoldi iteration, 87-88
ARPACK, 93
ART, 46
see also reconstruction.
artifact
star, 39, 41
streaks, 41
artifacts, 52
ASPIRE, 32
asymmetric accuracy, 99
attenuation, 12, 16
computing, 29
linear attenuation coefficient, 13
mass attenuation coefficient, 13
autoradiography, 56
back-projection, 35-37
backwards compatibility, 53, 128
basis, 30
Bayesian inference, 56

BBTools, 4, 6, 32, 37, 99, 113-123, 157160
compatibility, 120
beam divergence, 16,26
bidiag, 87 (see algorithm)
bidiagonalization, 87 (see Lanczos), 97, 108
bin, 23, 28, 30
black-box operator, 32
BLAS, 96
BLAS levels, 96
blocking, 96
CAT, 9
see also $C T$.
Cauchy inequality, 149
CGNR, 109
Chebyshev polynomia, 91
collimator, 14, 24
cone-beam, 15
fan-beam, 15
parallel hole, 14
pin-hole, 15, 26
sensitivity, 15
collinearity, 17
compact support, $22,23,55$
condition number
near-orthogonal matrix, 151
consistent norm, 149
continuous
acquisition, 23 (see acquisition)
convergence
EM algorithm, 44
Cormack, Allan MacLeod, 9
coronal plane, 5
coupling matrix, 80 (see NPSVD)
coverage, 52
CT, 12 (see scanner)
dead-time, 18
deflation
explicit, 102
implicit, 104
one-sided, 103
two-sided, 102
discrepancy principle, 68
discrete Picard condition, 111
distribution
Gaussian, 33
multinomial, 31
Poisson, 30
shifted Poisson, 32, 33
divergence, 16 (see beam divergence)
early stopping, 45
EM, 42-46
Kullback-Leibler distance, 55
sieve, 46
energy window, 16
epilsepsy, iv
EU Fifth Framework Programme, v exact
shift strategy, 91
singular triplet, 79
explicit
orthogonalization, 151
see also orthogonalization.
restart, 109
FBP, 38-42
feasibility shell, 68
Felix Hausdorff, 69
Fessler, Jeffrey A., 32
field of view, 23
filter-factors, 59, 110
filtered back-projection, 38 (see FBP)
Fourier transform, 39
FOV, 23 (see field of view), 36
full reorthogonalization, 86
full-width at half maximum, 10
funding, v
FWHM, 10
Galerkin, 109
gallery, 10
Gauss-Seidel, 46
Gaussian likelihood, 66
Gaussian random fields, 54
general public license, 6
geometric correction, 28, 31
GMRES, 108
GNU, 6
GPL, 6

H:S (Hovedstadens Sygehusfællesskab),

## v

Hausdorff
distance, 69
Felix, 69
homogeneity, 52, 57
Hounsfield, Sir Godfrey Newbold, 9
Hudson, Harold Malcolm, 43
hybrid solver, 111
hypercube, 30, 118
Høgh-Rasmussen, Nis, v
implicit
deflation, 104
orthogonality, 151
see also singular triplet.
restarting, 98, 109
implicit restart, 91
initialism, 12
isotropic, 52
iterative
reconstruction, 43
refinement, 107
solver, 107

Johann Karl August Radon, 22
John Von Neumann, 148
see also Von Neumann.
Jørgensens og Gunnar Hansens Fond, v

Kaczmarz method, 46
Keith John Worsley, 54
Kilmer, Misha, v
Kronecker product, 40
Krylov, 4, 86
see also Lanczos and Arnoldi.
matrix, 86
restart, 90
subspace, 86
Kullback-Leibler distance, 55
Ky Fan's dominance theorem, 148

Lanczos, 97, 108
bidiagonalization, 87
restart, 90
tridiagonalization, 86-87
unsymmetric tridiag., 88
LAPACK, 96
Laplace operator, 60
Larkin, Richard S., 43
latent semantic indexing, 104
least-squares
matrix approximation, 150
model-weighted, 67
weighted, 67
left residual factorization, 97, 108
legat
Jeppe J. og hustru Ovita J., v
Jørgensens og Gunnar H., v
Otto Mønsted, v
Leja points, 91, 93
level 1-3, 96
level of orthogonality, 150
license
GPL, 6
line of response, 22
literate programming, 6

LOR, 16, 22, 35
LSQR, 93, 112
machine epsilon, 7
machine precision, 7
matrix
least-squares approximation, 150
low-rank approximation, 150
see also Schmidt-Mirsky. near-orthogonal, 150
small, 6
sparse, 98
matrix structure, 87
MCI, iv (see mild cognitive impair.)
mean-field approximation, 56
median plane, 5
Mehmet Yavuz, 33
memory hierarchy, 96
method of sieves, 46
mild cognitive impairment, iv
MINRES, 108
Misha Kilmer, v
MLS, 67
model-weighted least-squares, 67
modularization, 97
mollifier methods, 58
Monte-Carlo simuations, 32
MOORE Tools, 123
Morgan, Ronald B., 97
$\mu$-map, 14
multi-head, 14
multinomial distribution, 31
near-orthogonal matrix, 150
Nis Høgh-Rasmussen, v
non-collinearity, 17
non-negative, 36
norm
absolute, 147
consistent, 149
Ky Fan, 148
max-norm, 149
padded matrix dominance, 148
unitarily invariant, 147
normal equations, 108
normal form, 22
NPSVD, 4, 80
see also SVD, Krylov, Arnoldi,
and Lanczos.
accuracy, 153
coupling matrix, 80
diagonalization, 82
orthogonalize, 90
restart, 92
search vectors, 83
singular error matrix, 80
singular values, 153
terminology, 80
update, 82
null-space, 44
numerical
notation, 7
partial SVD (NPSVD), 80
singular triplet, 79
on-line subtraction, 18, 32
one-sided
deflation, 103
reorthogonalization, 155
see also orthogonalization.
ordered subset, 43
orthogonal, 5
matrix, 150
orthogonal projector, 62
orthogonality
implicit, 151
see also singular triplet.
level, 143, 150
loss, 143
orthogonalization
explicit, 151
Gram-Schmidt, 88
one-sided, 155
orthonormal, 5
OSEM, 43
Ostrowski's theorem, 147
Otto Mønsteds Fond, v
partial SVD, 4, 79
see also $S V D$.
penalty function, 60
perturbation theory, 147
PET, 16 (see scanner)
model, 27
phantom, 21
photo multiplier tube, 18
Picard condition, 111
pixel, 30
pixelation, 45, 55
PMT, 18 (see photo multiplier tube)
point-spread function, 10, 129
Poisson
distribution, 30
Simeon Denis, 30
positron range, 17, 29
post-filter, 45,58
post-solve, 111
pre-solve, 111
preconditioner, 41, 58, 83, 99
prior, 55
programming
literate, 6
projector, 62
propaedeutic ( $\pi \rho o \pi \alpha \iota \delta \epsilon \dot{\epsilon} \epsilon \iota \nu$ ), 3
pseudo-inverse, 39
PSF, 10, 52, 129 (see point-spread function)
PSVD, 79, 79
see also $S V D$ and $N P S V D$.
quantitative, 52
quasi-linear reconstruction, 58

Radon
discrete transform, 23
Johann Karl August, 22
transform, 22
ramp filter, 39
Random fields
Gaussian, 54
randoms, 18
reconstruction, 3
ART, 46
EM, 42-46
iterative, 43
regularization, 58
reorthogonalization, 85, 151
see also orthogonalization.
residual, 107
residual factorization, 97
resolution, 10, 57
restart
exact shift strategy, 91
explicit, 109
implicit, 91, 98, 109
Lanczos, 90
NPSVD, 92
ridge regression, 60
Riemann-Lebesgue, 55
ringing, 42
Ronald B. Morgan, 97
rounding residual, $97,100,108$
sagittal plane, 5
sampling theorem, 55
Savværksejer Jeppe Juhl og hustru Ovita Juhls Mindelegat, v
scan
transmission, 12
scanner
abbreviations, 9
acquisitions mode, 19
CT, 12-14
gallery, 10
PET, 16-19
resolution, 10
septa, 19
SPECT, 14-16
multi-head, 14
tracer, 14
Schmidt-Mirsky's theorem, 150
Schwartz, 149 (see Cauchy inequality)
scintillation crystal, 16
search vectors, 83
selective reorthogonalization, 85
semi-orthogonal vectors, 86
semi-orthogonality, 90
septa, 19
shifted Poisson distribution, 32
sieve, 46
SimSET, 32
simulated annealing, 56
singular triplet, 79, 151
accuracy, 155
exact, 79
implicit orthogonality, 151
numerical, 79
orthogonal, 79
orthogonality, 151, 152
singular value decomposition (SVD), 79
singular values
near-orthogonal matrix, 150
SIR, 55
SIRA, 55
slant-stacking, 32
small, 6
solver, iterative, 107
sparse matrix, 98
spatial invariant, 52
SPECT, 14 (see scanner)
model, 24
star artifact, 39, 41 (see artifact)
step-and-shoot, 23 (see acquisition)
Stieltjes integrals, 40
stopping criteria, 143
streak artifact, 39, 41 (see artifact)
surrogate function, 55
SVD, 4, 79-105
NPSVD, 80
partial (PSVD), 79
perturbation, 147
singular triplet, 79
thin, 59
system matrix, 30, 32
theorem
Central Slice, 38
cond. of near-ortho. matrix, 151
EM convergence, 44
Fourier Slice, 38

Ky Fan dominance, 148
matrix approximation, 150
NPSVD accuracy, 154
Ostrowski, 147
padded matrix dominance, 148
projection dominance, 148
sampling, 55
Schmidt-Mirsky, 150
triplet accuracy, 155
triplet orthogonality, 152
Von Neumann, 148
Weyl, 147
thin SVD, 59
Tikhonov regularization, 59
Toft, Peter Aundal, 32
tomography, 3
tracer
dose, 30
SPECT, 14
transaxial plane, 5
transform
Fourier, 39
Radon, 22
transmission scan, 12
transverse plane, 5
tridiag, 87 (see algorithm)
tridiagonalization, 86 (see Lanczos)
triplet, 79 (see singular triplet)
triplet aliasing, 101, 155
two-sided deflation, 102
unitarily invariant norm, 147
normalized, 148
variability, 51
Von Neumann
unit. invariant norm theorem, 148
weighted least-squares, 60, 67
Weyl's theorem, 147
WLS, 67 (see weighted least-squares)
Woodbury's formula, 108
Worsley, Keith John, 54
Yavuz, Mehmet, 33

## Nomenclature

## General notation

| $x^{*}$ | Complex conjugate of $x$. |
| :--- | :--- |
| $\triangleq$ | Definition, e.g. $x \triangleq 2$ means that $x=2$ by definition. |
| $\epsilon_{M}$ | Machine epsilon, the smallest number such that $1+\epsilon_{M}$ rounds to 1. |

## Matrix notation

All vectors are column vectors. Row vectors are described as transposed (conjugated) column vectors. Matrices are written in uppercase boldface letters, like A, while vectors are written in lowercase, like $\mathbf{x}$.

Submatrices are indicated using MATLAB ${ }^{\oplus}$ notation. Thus the matrix $\mathbf{A}_{(:, 1: 3)}$ consists of the first three columns of $\mathbf{A}$.
$\mathbf{A}^{\mathrm{T}} \quad$ Transpose.
$\mathbf{A}^{\mathrm{H}} \quad$ "Hermitian", i.e. transpose-conjugated. Thus $\mathbf{A}^{\mathrm{H}}=\left(\mathbf{A}^{\mathrm{T}}\right)^{*}$.
$\mathbf{A}^{+}$
The pseudo-inverse, or Moore-Penrose inverse, of the matrix $\mathbf{A}$.
$\|\mathbf{A}\| \quad$ Unspecified norm of $\mathbf{A}$ (see p. 149 for other norms).
I Identity matrix. When used in an expression, it automatically have the necessary dimensions.
$0 \quad$ The zero-vector or -matrix of appropiate dimension.
range(A) The space spanned by the columns of $\mathbf{A}: \operatorname{range}(\mathbf{A})=\{\mathbf{y}: \mathbf{y}=\mathbf{A x}\}$, where $\mathbf{x}$ is taken over all possible vectors.
$\operatorname{null}(\mathbf{A}) \quad \operatorname{Null}-$ space of $\mathbf{A}: \operatorname{null}(\mathbf{A})=\{\mathbf{x}: \mathbf{A x}=\mathbf{0}\}$.
$\operatorname{diag}(\mathbf{x})$
Diagonal, i.e. $\left.\mathbf{A}=\operatorname{diag}(\mathbf{x}) \operatorname{implies} \mathbf{A}_{(i, i)}=\mathbf{x}_{( } i\right)$ and $\mathbf{A}_{(i, j)}=0$ for $i \neq 0$.

## Sets

$\mathbb{R} \quad$ Real numbers.
$\mathbb{C} \quad$ Complex numbers.
$\mathbb{Q} \quad$ Rational numbers.
$\mathbb{N} \quad$ Natural numbers (non-negative integers).
$\mathbb{Z} \quad$ Integers.

## Signal processing

The Fourier-transform. If $F(\omega)$ and $f(x)$ are Fourier-pairs, then we have:

$$
\begin{aligned}
F(\omega)=\mathscr{F}\{f(x)\} & =\int_{-\infty}^{\infty} f(x) e^{-j \omega x} d x \\
f(x)=\mathscr{F}^{-1}\{F(\omega)\} & =\int_{-\infty}^{\infty} F(\omega) e^{+j \omega x} d \omega /(2 \pi)
\end{aligned}
$$

## Statistics

Expectation. E.g. $\mathscr{E}\{x\}$ is the mean value of the stochastic variable $x$.
Variance. $\operatorname{var}\{x\}=\mathscr{E}\left\{(x-\mathscr{E}\{x\})^{2}\right\}$.

## Miscellaneous

Asymptotic upper bound. The function $g(x)$ is in the set $\mathcal{O}(f(x))$ if there are positive constants $c$ and $x_{0}$ such that $0 \leq f(x) \leq c g(x)$ for all $x>x_{0}$.
If an algorithm runs in $\mathcal{O}\left(n^{2}\right)$ time, it means that the time taken for the algorithm is less than $c n^{2}$ for some constant $c$, when $n$ is sufficiently large.


[^0]:    ${ }^{1}$ The paper is excluded from the thesis, as the copyrights have been assigned to the publisher.

[^1]:    ${ }^{1}$ To account for the multiplication, we need to multiply this by $\epsilon_{M}\left(1+\epsilon_{M}\right)$ in the real case.

[^2]:    ${ }^{1} \mathrm{~A}$ contrast agent is a substance with an attenuation that differs significantly from the area of interest. This property makes it visible in the CT scan.

[^3]:    ${ }^{2}$ Actually slits of different depth are carved into a single crystal block.

[^4]:    ${ }^{1}$ The author is an engineer formally trained in the "lore".
    ${ }^{2}$ This count includes the author.

[^5]:    ${ }^{1}$ From the seminar "New challenges in SPECT", held by The PET and Cyclotron Unit, Copenhagen University Hospital, 19-20 Sep. 2002.

[^6]:    ${ }^{2}$ One "unreasonable" method that do meet this requirement would be an algorithm which reconstructed any data as a zero image.

[^7]:    ${ }^{3}$ Roughly, $L^{1}$ means "absolutely integrable", but also includes a requirement from measure theory; a rather difficult branch of mathematics. From a practical point of view, the last property excludes the Dirac function but can be ignored otherwise. The confinement property is a rephrasing of compact support.

[^8]:    ${ }^{4}$ Adopted for routine use in cardiac scans at the University of Michigan Medical Center [21].

[^9]:    ${ }^{1}$ This figure was inspired by a similar figure in [71].

[^10]:    function $\left[\tilde{\mathbf{U}}, \tilde{\mathbf{V}}, \tilde{\mathbf{B}}, \tilde{\mathbf{s}}_{\mathbf{u}}, \tilde{\mathbf{s}}_{\mathbf{v}}, \tilde{f}_{u}, \tilde{f}_{v}\right]=\operatorname{update}\left(\mathbf{A}, \hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{B}}, \hat{\mathbf{s}}_{\mathbf{u}}, \hat{\mathbf{s}}_{\mathbf{v}}, \hat{\mathbf{f}}_{\mathbf{u}}, \hat{\mathbf{f}}_{\mathbf{v}}, \beta\right)$
    $\tilde{\mathbf{V}}=\left[\hat{\mathbf{V}}, \hat{\mathbf{s}}_{\mathbf{v}}\right] ; \quad \tilde{\mathbf{U}}=\left[\hat{\mathbf{U}}, \hat{\mathbf{s}}_{\mathbf{u}}\right] ; \quad \quad$ Update singular matrices by eq. (7.7)
    $l_{u}=\operatorname{size}(\hat{\mathbf{B}}, 1)-\operatorname{length}\left(\hat{\mathbf{f}}_{\mathbf{u}}\right) ; \quad \mathbf{z}_{\mathbf{u}}=\operatorname{zeros}\left(l_{u}, 1\right) ; \quad$ Form implicit zeros in $\mathbf{f}_{\mathbf{u}}$
    $l_{v}=\operatorname{size}(\hat{\mathbf{B}}, 2)-\operatorname{length}\left(\hat{\mathbf{f}}_{\mathbf{v}}\right) ; \quad \mathbf{z}_{\mathbf{v}}=\operatorname{zeros}\left(l_{v}, 1\right) ; \quad$ And for $\mathrm{f}_{\mathrm{v}}$
    $\tilde{\mathbf{B}}=\left[\hat{\mathbf{B}},\left[\mathbf{z}_{\mathbf{v}} ; \hat{\mathbf{f}}_{\mathbf{v}}\right] ;\left[\mathbf{z}_{\mathbf{u}} ; \hat{\mathbf{f}}_{\mathbf{u}}\right]^{\mathrm{H}}, \beta\right] ;$
    Update coupling-matrix using eq. (7.12)
    $\tilde{\mathbf{s}}_{\mathbf{u}}=\mathbf{A} * \hat{\mathbf{s}}_{\mathbf{v}}-\hat{\mathbf{U}}\left(:, l_{v}+1:\right.$ end $) * \hat{\mathbf{f}}_{\mathbf{v}}-\hat{\mathbf{s}}_{\mathbf{u}} * \beta ; \quad$ Unscaled search-vector from eq. (7.10a)
    $\tilde{\mathbf{s}}_{\mathbf{v}}={\underset{\sim}{A}}^{\mathbf{H}} * \hat{\mathbf{s}}_{\mathbf{u}}-\hat{\mathbf{V}}\left(:, l_{u}+1:\right.$ end $) * \hat{\mathbf{f}}_{\mathbf{u}}-\hat{\mathbf{s}}_{\mathbf{v}} *(\beta)^{*} ; \quad \quad$ Unscaled search-vector from eq. (7.10b)
    $\left[\tilde{\mathbf{s}}_{\mathbf{u}}, \tilde{f}_{u}\right]=\operatorname{unit}\left(\tilde{\mathbf{s}}_{\mathbf{u}}\right) ; \quad \quad$ Normalize search-vector
    $\left[\tilde{\mathbf{s}}_{\mathbf{v}}, \tilde{f}_{v}\right]=\operatorname{unit}\left(\tilde{\mathbf{S}}_{\mathbf{v}}\right) ; \quad \quad$ Normalize search-vector

[^11]:    ${ }^{1}$ This also gives the columns, since $\mathbf{X}^{\mathrm{H}} \mathbf{X}$ is Hermitian.

[^12]:    ${ }^{1}$ This is, for example, required for a truncated SVD.

[^13]:    ${ }^{1}$ One may also use row-vectors provided it is used consistently. This is often preferred by statisticians [52].

[^14]:    ${ }^{1}$ The regularization-constant was chosen according to the L-curve criteria [30], which occurs around $\lambda \approx 0.021\|\mathbf{A}\|_{2}$ in this case. However, we shall generally assume that the parameter is fixed or chosen interactively.

[^15]:    ${ }^{2}$ Section 5.4 to be precise.

[^16]:    ${ }^{1}$ A norm is absolute if $\|\mathbf{A}\|=\|(|\mathbf{A}|)\|$. The matrix 2-norm is not absolute.

