

Convergent Algorithms for Statistical Image Reconstruction in Emission Tomography

by

Sangtae Ahn

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical Engineering: Systems)
in The University of Michigan
2004

Doctoral Committee:

Professor Jeffrey A. Fessler, Chair
Professor Alfred O. Hero III
Professor Robert A. Koeppe
Professor David L. Neuhoff
Assistant Professor Thomas E. Nichols

© Sangtae Ahn 2004
All Rights Reserved

This dissertation is dedicated to my parents and my wife.

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to my advisor Professor Jeff Fessler for his guidance, support, and kindness. It was my greatest luck that I had a chance to work with him. I would also like to thank Professors Alfred Hero, Robert Koeppe, David Neuhoff, and Thomas Nichols in my committee for their valuable input on my dissertation.

I extend my gratitude to my past and current colleagues including Idris Elbakri, Matt Jacobson, Jeongtae Kim, Sangwoo Lee, Nan Sotthivirat, Somesh Srivastava, Web Stayman, Anastasia Yendiki, Rongping Zeng, and Yingying Zhang. I owe my thanks to many other friends including Hyoseok Lee.

Finally, I would like to thank my parents for their love and encouragement. And I would like to thank my wife, Hwajeong Lee, for her loving support and patience.

This work was supported by National Science Foundation grant BES-9982349, National Institutes of Health grants CA-60711 and CA-87634, Department of Energy grant DE-FG02-87ER60561, and by a Rackham Predoctoral Fellowship.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	v
LIST OF TABLES	vi
LIST OF APPENDICES	vii
ABSTRACT	viii
CHAPTER	
1. Introduction	1
2. Background	5
2.1 Emission Tomography	5
2.2 Transmission Tomography	7
2.3 Background Noise	8
2.3.1 Accidental Coincidence (or Random)	8
2.3.2 Scatter	10
2.4 System Model	10
2.4.1 Transmission Scans	10
2.4.2 Emission Scans	11
2.5 Tracer Kinetics	12
2.5.1 Dynamic Model	12
2.5.2 Compartmental Model	13
3. Statistical Image Reconstruction Methods for Tomography	15
3.1 Emission Tomography	15
3.1.1 Statistical Model	15
3.1.2 Penalized-Likelihood Reconstruction	16
3.1.3 Penalty Function	18
3.1.4 Existence and Uniqueness	19
3.1.5 Boundedness	21
3.1.6 Modified Objective Function	21
3.2 Transmission Tomography	23
3.2.1 Statistical Model	23
3.2.2 Penalized-Likelihood Reconstruction	24
3.3 Optimization Transfer Algorithms	25

3.3.1	Optimization Transfer Principles	25
3.3.2	EM Surrogate Algorithms	29
3.3.3	Quadratic Surrogate Algorithms	33
4.	Convergent Relaxed Ordered Subsets Algorithms	36
4.1	Introduction	36
4.2	Ordered Subsets Algorithms	39
4.2.1	PL Emission Reconstruction Problem	40
4.2.2	Ordered Subsets (OS) Algorithms	41
4.3	Relaxed OS Algorithms	44
4.3.1	Modified BSREM	45
4.3.2	Diagonally-Scaled Incremental Gradient Method	48
4.3.3	Regularization into OS Algorithms	50
4.3.4	Subiteration-Independent Scaling Matrices Are Essential	51
4.4	Results	52
4.5	Conclusion	56
5.	Convergent Incremental Optimization Transfer Algorithms	63
5.1	Introduction	63
5.2	Incremental Optimization Transfer	66
5.2.1	Incremental Optimization Transfer Algorithms	66
5.2.2	Special Case: Incremental EM Algorithms	69
5.2.3	Convergence Properties	70
5.3	Application to Transmission Tomography	72
5.3.1	PL Attenuation Map Reconstruction Problem	72
5.3.2	Transmission Incremental Optimization Transfer (TRIOT)	73
5.3.3	OS-SPS	76
5.3.4	Acceleration	77
5.4	Results	80
5.4.1	Simulation	80
5.4.2	Real Data	82
5.5	Conclusion	82
6.	Statistical Image Reconstruction for Randoms-Precorrected PET Scans	91
6.1	Introduction	91
6.2	Measurement Model and Exact Log-likelihood	94
6.3	Approximations to Exact Log-likelihood	96
6.3.1	Ordinary Poisson (OP) Approximation	96
6.3.2	Shifted Poisson (SP) Approximation	97
6.3.3	Saddle-Point (SD) Approximation	97
6.3.4	Log-likelihood for Prompt Data	98
6.4	Effects of Zero-thresholding on Bias	98
6.5	Asymptotic Analysis	101
6.5.1	Asymptotic Bias and Asymptotic Covariance	101
6.5.2	Comparison with Weighted Least Squares Methods	105
6.5.3	Comparison of Covariances	108
6.6	Reconstruction Algorithms	108
6.6.1	SPS for OP ⁻ , SP ⁻ , and SD	110
6.6.2	Variation of ML-EM for OP ⁻ and SP ⁻	112
6.7	Simulations	114
6.7.1	Methods	114

6.7.2	Results	116
6.8	Conclusion	119
7.	Covariance of Kinetic Parameter Estimators Based on Time Activity Curve Reconstructions: 1D Study	123
7.1	Introduction	123
7.2	Problem	124
7.2.1	TAC Reconstruction	125
7.2.2	Kinetic Parameter Estimation	127
7.3	Covariance of kinetic parameter estimators	128
7.3.1	Derivation	128
7.3.2	Information Matrix for List-Mode Data	129
7.4	Results	129
7.5	Conclusion	131
8.	Summary and Future Work	132
8.1	Summary	132
8.2	Future Work	133
APPENDICES		136
BIBLIOGRAPHY		159

LIST OF FIGURES

<u>Figure</u>		
2.1	Diagram of a PET system with a cylindrical ring of detectors (left) and a SPECT system (right).	6
2.2	Background events in PET: accidental coincidence (AC) event (left) and scattered event (right).	10
2.3	One tissue compartment model.	13
3.1	Illustration of optimization transfer in 1D. In this figure, Φ and ϕ denote an objective function and a surrogate function, respectively.	26
3.2	This is a 1D example, whose constraint is $x \geq 0$, showing that the “matched gradient” condition (3.31) does not hold at $x^n = 0$ while the “matched Φ value” condition (3.29) and “lies below” condition (3.30) hold.	28
4.1	Toy example of OS algorithms. (a) Trajectory of iterates of a nonOS gradient method with a constant stepsize and its OS version with 3 subsets. The optimal point is $\hat{x} = (0.5, 0.5)$ and the initial point is $x^0 = (5, 5)$. (b) $(\Phi(\hat{x}) - \Phi(x^{n,m})) / (\Phi(\hat{x}) - \Phi(x^0))$ versus iteration number. For the OS method, each subiterate is denoted.	58
4.2	Comparison of normalized Φ difference $(\Phi(\hat{\lambda}) - \Phi(\lambda^n)) / (\Phi(\hat{\lambda}) - \Phi(\lambda^0))$ versus iteration number for non-OS algorithms including SPS and DPEM; and unrelaxed OS algorithms including unrelaxed OS-SPS and unrelaxed modified BSREM (with 8 and 40 subsets).	59
4.3	Comparison of normalized Φ difference $(\Phi(\hat{\lambda}) - \Phi(\lambda^{n,m})) / (\Phi(\hat{\lambda}) - \Phi(\lambda^0))$ versus iteration number for unrelaxed modified BSREM and relaxed modified BSREM with 8 and 40 subsets. For relaxed modified BSREM-8 (top) and relaxed modified BSREM-40 (bottom), $\alpha_n = 1/(\frac{1}{15}n + 1)$ and $\alpha_n = 1/(n + 1)$ are used, respectively. This figure shows every subiterate.	59
4.4	Comparison of normalized Φ difference $(\Phi(\hat{\lambda}) - \Phi(\lambda^{n,m})) / (\Phi(\hat{\lambda}) - \Phi(\lambda^0))$ versus iteration number for unrelaxed OS-SPS and relaxed OS-SPS with 8 and 40 subsets. For relaxed OS-SPS-8 (top) and relaxed OS-SPS-40 (bottom), $\alpha_n = 1/(\frac{1}{5}n + 1)$ and $\alpha_n = 1/(n + 1)$ are used, respectively. This figure shows every subiterate.	60
4.5	Comparison of normalized Φ difference $(\Phi(\hat{\lambda}) - \Phi(\lambda^n)) / (\Phi(\hat{\lambda}) - \Phi(\lambda^0))$ versus iteration number for unrelaxed OS algorithms and relaxed ones. (a) Unrelaxed modified BSREM and relaxed modified BSREM. This figure is the same as Fig. 4.3 except that it shows only each iterate. (b) Unrelaxed OS-SPS and relaxed OS-SPS. This figure is the same as Fig. 4.4 except that it shows only each iterate.	61

4.6	(a) Shepp-Logan digital phantom (true image). (b) FBP reconstruction (starting image). (c) PL reconstruction using 20 iterations of relaxed modified BSREM with 8 subsets. (d) PL reconstruction using 20 iterations of relaxed OS-SPS with 8 subsets.	62
5.1	Illustration of a geometric view of OS algorithms. In this example, there are two subobjective functions f_1 and f_2 such that $\Phi = f_1 + f_2$. When \mathbf{x}^n is far from the optimal point \mathbf{x}^* , all $\nabla\Phi(\mathbf{x}^n)$, $\nabla f_1(\mathbf{x}^n)$, and $\nabla f_2(\mathbf{x}^n)$ point roughly towards \mathbf{x}^* . However, then \mathbf{x}^k is near \mathbf{x}^* , they do not.	77
5.2	Comparison of non-OS algorithms (SPS-MC/OC), an OS algorithm (OS-SPS), incremental optimization transfer algorithms (TRIOT-MC/OC/PC), and enhanced TRIOT (E-TRIOT-MC/OC/PC) for 2D attenuation map reconstruction using simulation data. The normalized Φ difference is defined by $(\Phi(\hat{\boldsymbol{\mu}}^{\text{PL}}) - \Phi(\hat{\boldsymbol{\mu}}^n))/(\Phi(\hat{\boldsymbol{\mu}}^{\text{PL}}) - \Phi(\hat{\boldsymbol{\mu}}^0))$ where $\hat{\boldsymbol{\mu}}^{\text{PL}}$ is a PL estimate image. OS-SPS and TRIOT algorithms used 16 subsets, and TRIOT algorithms include one iteration of OS-SPS initially. The starting image was a FBP image for all cases.	85
5.3	Same as Fig. ??, but TRIOT algorithms include <i>three</i> iterations of OS-SPS initially.	86
5.4	Same as Fig. ??, but 64 subsets are used for OS-SPS and TRIOT algorithms.	86
5.5	(a) A slice of the (modified) Zubal phantom (true image). (b) FBP reconstruction (starting image). (c) PL estimate image $\hat{\boldsymbol{\mu}}^{\text{PL}}$ obtained using 30 iterations of OS-SPS-PC with 16 subsets followed by 800 iterations of SPS-OC. (d) PL reconstruction using 1 iteration of OS-SPS-PC and 19 iterations of TRIOT-PC with 64 subsets.	87
5.6	Same as Fig. ??, but real PET data are used.	88
5.7	Same as Fig. ??, but real PET data are used.	88
5.8	Same as Fig. ??, but real PET data are used.	89
5.9	Reconstructed attenuation maps from real PET data. (a) FBP reconstruction (starting image). (b) PL estimate image $\hat{\boldsymbol{\mu}}^{\text{PL}}$ obtained using 30 iterations of OS-SPS-PC with 16 subsets followed by 800 iterations of SPS-OC. (c) PL reconstruction using 20 iterations of OS-SPS-PC with 64 subsets (an image that is one point of a limit cycle). (d) PL reconstruction using 1 iteration of OS-SPS-PC and 19 iterations of TRIOT-PC with 64 subsets.	90
6.1	Effective means of trues increased by zero-thresholding for OP^+ and SP^+ . In this figure, randoms fractions in % denote (mean of randoms)/(mean of trues).	100
6.2	Illustration of a linear surrogate q^{OP^-} [see (6.41)] at $l^n = 1$ for an OP^- log-likelihood h^{OP^-} for a negative value $y_i < 0$. The concave surrogate q^{OP^-} lies below the objective h^{OP^-} that is convex. One can see that $q^{\text{OP}^-}(l; l^n) \geq q^{\text{OP}^-}(l^n; l^n)$ implies that $h^{\text{OP}^-}(l) \geq h^{\text{OP}^-}(l^n)$ for $l \geq 0$	110
6.3	Digital phantom used in simulations. The background, left cold disc, and right hot disc have relative emission activities of 2, 0.5, and 4, respectively.	115
6.4	(a) Horizontal profile through sample mean of estimators for 2K counts. (b) Horizontal profile through sample standard deviation of estimators for 2K counts.	120
6.5	(a) Horizontal profile through sample mean of estimators for 2M counts. (b) Horizontal profile through sample standard deviation of estimators for 2M counts.	121

6.6	Histogram of the ratio of the standard deviation of different reconstruction methods to the standard deviation of PR method for $2M$ counts.	122
F.1	Comparison of mean root-convergence factors of incremental SPS with different numbers M of subobjective functions for a one-parameter transmission problem. Nonincremental ordinary SPS corresponds to the case $M = 1$	158

LIST OF TABLES

Table

4.1	Outline for the algorithms presented in this chapter	52
4.2	Comparison of algorithms	53
5.1	Outline for incremental optimization transfer algorithms. The right side of (5.9) is due to (5.4) and (5.5).	68
5.2	Outline for transmission incremental optimization transfer (TRIOT) algorithm using maximum curvature (MC).	75
7.1	Comparison of predicted and empirical sample means of kinetic parameter estimators . . .	131
7.2	Comparison of predicted and empirical standard deviations of kinetic parameter estimators	131

LIST OF APPENDICES

Appendix

A.	Computable Upper Bounds for PL Solutions in Emission Tomography	137
A.1	Computable Upper Bound: Example 1	137
A.2	Computable Upper Bound: Example 2	138
B.	Determination of ϵ for Modified Objective Function in Emission Tomography	140
C.	Convergence Proofs for modified BSREM	142
D.	Convergence Proofs for Diagonally-Scaled Incremental Gradient Method	147
E.	Convergence Proofs for Incremental Optimization Transfer Algorithms	150
F.	Local Convergence Rate Analysis for Incremental Optimiazaiton Transfer Algorithms	154
F.1	Asymptotic Convergence Rate	154
F.2	One-Parameter Example	157

ABSTRACT

Convergent Algorithms for Statistical Image Reconstruction in Emission Tomography

by

Sangtae Ahn

Chair: Jeffrey A. Fessler

Emission computed tomography (ECT), including positron emission tomography (PET) and single photon emission computed tomography (SPECT), is a medical imaging technique that provides functional information about physiological processes. The goal of ECT is to reconstruct the distribution of the radioisotopes in the body by measuring the emitted photons. Statistical image reconstruction methods have shown improved image quality over conventional nonstatistical methods by using accurate physical models and appropriate noise models. However, statistical methods require huge computation and complex modeling. So computationally efficient algorithms and simple yet accurate statistical models are essential. First, we develop fast and convergent algorithms for statistical image reconstruction. Ordered subsets or incremental gradient type algorithms have been popular due to their fast initial convergence rates. However, they do not converge to a solution in general. We achieve global convergence by two methods: relaxation and incremental optimization transfer principles. Those two families of algorithms are provably convergent yet converge fast. We apply the algorithms to emission and transmission to-

mography and to simulation and real PET data. Secondly, we develop statistical models for randoms-precorrected PET. Accidental coincidence (AC) events, or randoms, are one of primary sources of background noise in PET. Most PET scanners are corrected for AC events by real-time subtraction of delayed window coincidences. Although the randoms-precorrection compensates in mean for AC events but destroys the usual Poisson statistics, complicating statistical reconstruction. We propose new likelihood approximations that allow negative sinogram values without requiring zero-thresholding. Analysis and simulation results show that the new statistical model is nearly free of systematic bias yet keeps low variance. Finally, we analyze the parameterization of time activity curves (TACs) in dynamic imaging. We provide approximate expressions for the covariance matrix of kinetic parameter estimators based on TAC reconstructions when TACs are modeled as a linear combination of temporal basis functions such as B-splines. The approximations are useful tools for assessing and optimizing the temporal basis functions for TACs.

CHAPTER 1

Introduction

Emission computed tomography (ECT) is a medical imaging modality that provides functional information about physiological processes, as opposed to X-ray CT providing anatomical information. In ECT, a small amount of a radioactive compound labeled with a radioisotope, called a *radiotracer*, is introduced into a patient's body by injection or inhalation, and then the spatial and sometimes also the temporal distribution of the radioisotope, the decay of which generates photons, is reconstructed from the photon measurements. The radioactivity image reconstruction gives important information on a specific biochemical process or blood flows which is difficult to obtain through other imaging modalities. Although it is a challenging task to reconstruct high quality images from intrinsically low SNR data mainly due to the dosage limitation and low sensitivities, statistical methodology has successfully shown desirable performance.

A statistical approach, maximum likelihood (ML) estimation, for ECT was first proposed by Rockmore and Macovski in 1976 [109], and the expectation maximization (EM) algorithm [32] was applied to ML reconstruction by Shepp and Vardi in 1982 [112]. Since then, statistical image reconstruction methods like ML as well as penalized-likelihood (PL) or maximum *a posteriori* (MAP) estimation have enjoyed great popularity to date. The advantages of statistical methods over nonstatistical analytical methods like filtered

backprojection (FBP) [79, p. 125] are addressed as follows. Statistical methods can model physical effects and nonstandard scan geometries accurately, for example, nonuniform attenuation in SPECT, and PET scans with missing data [62]. Also, statistical methods take into account appropriate statistical noise models to lower the variance of reconstructed images. As a result, improved image quality [113] and better lesion detectability [101] can be acquired.

However, statistical methods require a huge computational load, and complex physics and noise modeling. Despite increasing computing power, we still face challenges in computation particularly when moving from 2D to fully 3D imaging, and from static to dynamic imaging. Additionally, establishing accurate yet simple models is the key to success in statistical image reconstruction. The objective of this dissertation is to develop improved statistical methods for ECT in three aspects: fast and reliable algorithms for image reconstruction, accurate statistical models for randoms-precorrected PET emission scans, and efficient parameterization of time activity curves (TACs) for dynamic imaging.

In the first part of our work (Chapters 4 and 5), we develop fast and convergent iterative algorithms for PL image reconstruction. There are several properties that we wish an ideal algorithm to have: fast convergence rate, global convergence, simplicity, parallelizability, flexibility to accommodate any type of system model, and capability to enforce object constraints like nonnegativity. Among these, we focus on convergence rate and global convergence. Since the ordered-subsets EM (OS-EM) algorithm was proposed by Hudson and Larkin in 1994 [55], OS type or block iterative algorithms, also known as incremental gradient methods in the optimization literature, have been very popular for tomographic image reconstruction because of their remarkably fast initial convergence rates. However, OS type algorithms generally do not converge to a PL optimal image, but become eventually stuck at a limit cycle that is different from the optimal image.

We provide two families of *convergent* OS type algorithms: relaxed OS algorithms and incremental optimization transfer algorithms. Firstly, the relaxed OS algorithms achieve global convergence by introducing relaxation parameters or diminishing stepsizes. Particularly, we provide two types of relaxed OS algorithms for emission tomography in Chapter 4: relaxed OS-SPS and modified BSREM. Secondly, we generalize the incremental EM approach proposed by Neal and Hinton [87], so we obtain a very broad family of incremental optimization transfer algorithms in Chapter 5. We apply the incremental optimization transfer algorithms to transmission tomography in which a challenging nonconcave maximization problem arises. All the algorithms above are convergent yet converge rapidly, and they can be applied to other image reconstruction problems.

In the second part of this dissertation (Chapter 6), we focus on statistical models specific to randoms-precorrected PET emission scans. Accidental coincidence (AC) events, also known as randoms, are one of primary sources of background noise in PET. The AC events affect a quantitative PET study unfavorably, in particular, for septaless fully 3D PET where AC rates are high, and for newer scanners with small crystals where photon counts per ray can be low. Many PET scans are corrected for AC events by real-time subtraction of delayed window coincidences which represent AC events. The randoms-precorrection reduces the amount of data by a factor of 2, which is desirable for large systems. However, the precorrection destroys the usual Poisson statistics and complicates ML or PL reconstruction. We provide a novel estimator allowing negative sinogram values, which is free of systematic bias and leads to low variance despite its simple implementation. We also develop algorithms tailored to the new estimators.

In the third part of our work (Chapter 7), we focus on time activity curve (TAC) parameterization for dynamic imaging. Usually, an object to be scanned is parameterized spatially using rectangular voxels, that is, rect basis functions (or sometimes blobs are

used [80, 134]). For dynamic imaging where the goal is to reconstruct TACs for each voxel, a conventional method is to reconstruct a series of images independently, which is equivalent to using rect temporal basis functions. Recently, Nichols *et al.* [90] have reconstructed continuous TACs using overlapping cubic B-splines and temporal regularization. Temporal regularization is carried out by including a temporal penalty function which encourages temporal smoothness in reconstructed TACs. To assess and optimize temporal basis functions (and temporal regularization), we analyze the covariance of kinetic parameter estimators from TACs estimated using those temporal basis functions. We derive the approximate formula for the covariance of kinetic parameter estimators as a function of temporal basis functions so that we can optimize the choice of basis functions without performing exhaustive simulation studies.

The original contributions made by this dissertation are as follows.

- New convergent and fast relaxed OS algorithms for emission tomography are presented: relaxed OS-SPS and modified BSREM [3, 4].
- A new family of incremental optimization transfer algorithms are presented, and a particular transmission incremental optimization transfer (TRIOT) algorithm is developed for transmission tomography [7, 8].
- A new shifted Poisson (SP) model allowing negative sinogram values is presented for randoms-precorrected PET emission scans [5, 6].
- New reconstruction algorithms (SPS and ML-EM based) for the new SP model are developed [5, 6].
- Approximate formulas for the covariance of kinetic parameter estimators from TACs reconstructed using temporal basis functions are developed for 1D [9].

CHAPTER 2

Background

2.1 Emission Tomography

Positron emission tomography (PET) and single photon emission computed tomography (SPECT) are two common types of emission computed tomography (ECT) [95, 98, 117]. They provide physiological information by reconstructing the spatial (and sometimes temporal) distribution of positron-emitting radioisotopes in a patient's body for PET, and, of γ -ray-emitting radioisotopes for SPECT. There are a variety of radionuclides that can be bound to organic molecules without changing the biological properties of the molecules, and the estimated concentrations of the radioisotopes enable us to investigate and study specific biochemical processes or blood flows.

For PET, commonly used positron-emitting radioisotopes include ^{11}C , ^{15}O , ^{13}N , and ^{18}F ; and the examples of a radiopharmaceutical labeled with them are [^{11}C]deoxyglucose (DG), [^{15}O]water, [^{13}N]ammonia, and [^{18}F]fluorodeoxyglucose (FDG). The radioisotope, introduced into the body, decays and emits a positron (β^+) which travels a short range of a few millimeters, called *positron range* [12, 75], before annihilating with an electron. The annihilation generates a pair of 511 keV gamma photons which travel in nearly opposite directions (Fig. 2.1). If both of these photons are detected within a coincidence timing window (about 5–10 ns) by some pair of detectors, then we determine that an annihila-

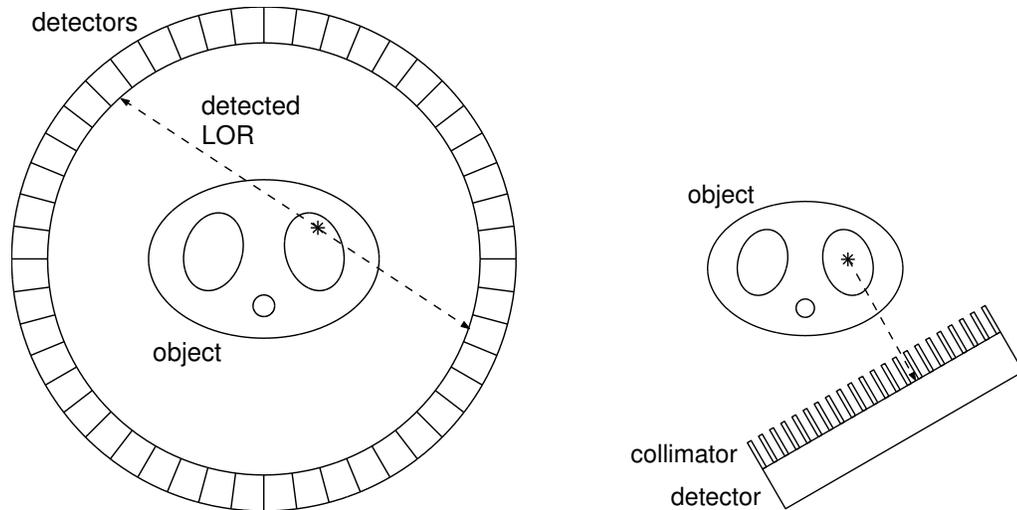


Figure 2.1: Diagram of a PET system with a cylindrical ring of detectors (left) and a SPECT system (right).

tion event has happened somewhere along the line of response (LOR), which is the line connecting the two detectors, and increments the number of events for the LOR.

For SPECT, a radioisotope that decays and emits a single gamma photon is used; for example, ^{99m}Tc emits a single 140 keV photon. The emitted photons are detected by a rotating array of detectors. To determine the corresponding LOR, collimators are attached to the detectors (Fig. 2.1). However, we focus mostly on PET in this dissertation.

Each LOR is characterized by a projection angle and a radial distance (for 2D imaging). The measurement data, a collection of the number of detected events for each LOR, is called *sinogram* or *projection data*. The process of acquiring the sinogram data is called *emission scan*. The sinogram data is said to lie in a *sinogram* or *projection* domain, and the object to be scanned is said to lie in an *image domain*.

The aim of emission tomography is to reconstruct the spatial distribution of the radioisotope from the sinogram data by considering geometrical factors, physical effects, and noise properties. To estimate the radioactivity distribution accurately, the effects of attenuation should be taken into account. One can estimate attenuation by performing a “transmission scan” as described in the next section.

2.2 Transmission Tomography

Photons traveling through the body experience two types of interactions: photoelectric absorption and Compton scatter. Consequently, some photons that get scattered or absorbed by the patient's body are undetected. This effect is called *attenuation*. The survival probability of a photon is determined by the length and the type of organ (*e.g.*, bone or tissue) that the photon traverses. By Beer's law, the survival probability of a photon traveling through a uniform medium is given by

$$\alpha = e^{-\mu L}$$

where μ is the linear attenuation coefficient of the medium and L is the length of the medium. The linear attenuation coefficient means the fraction of a beam of gamma rays that is absorbed or scattered per unit thickness, and it is typically reported in units of cm^{-1} . In other words, if we start with N_{in} photons, then after a thickness L , we have N_{out} photons such that

$$N_{\text{out}} = \alpha N_{\text{in}} = e^{-\mu L} N_{\text{in}}. \quad (2.1)$$

For example, suppose that a photon passes through two-layer media: the thickness and the attenuation coefficient of one layer are L_1 and μ_1 respectively, and those of the other layer are L_2 and μ_2 respectively. Then its survival probability is given by

$$\alpha = \alpha_1 \alpha_2 = e^{-\mu_1 L_1} e^{-\mu_2 L_2} = e^{-(\mu_1 L_1 + \mu_2 L_2)} \quad (2.2)$$

where one can see the attenuation coefficients play a role as “linear” weights.

The linear attenuation coefficient image (attenuation map), $\mu(\mathbf{x})$, can be reconstructed by using an outside radioactive source. Ring or rod sources with long half life positron emitters are used for PET, and the process of acquiring the sinogram data with those sources is called *transmission scan*. A transmission scan can be performed before or after

the emission scan. The mean number of photons recorded for the i th LOR for transmission scan would ideally be

$$b_i e^{-\int_{\text{LOR}_i} \mu(\mathbf{x}) d\mathbf{x}} \quad (2.3)$$

where LOR_i is the path of the corresponding LOR, and b_i is the mean number of photons recorded for the LOR when the object were absent. [Compare (2.3) with (2.1).] One can view (2.3) as a generalization of (2.2). An object (a patient's body) being present in the scanner reduces the number of detected photons by a factor of the survival probability for the corresponding LOR. The goal of transmission tomography is to estimate the attenuation map $\mu(\mathbf{x})$ from the sinogram data obtained in transmission scan. The reconstructed attenuation map is supposed to be used for radioactivity reconstruction in emission tomography.

The problem described here in transmission tomography is basically the same as that of X-ray computed tomography (CT). A major difference is that in PET or SPECT transmission scans, the source is usually a monoenergetic radiosotope whereas X-ray sources can be polyenergetic. In the polyenergetic case, one should consider effects such as “beam hardening” [59].

2.3 Background Noise

In PET, background noise consists mainly of accidental coincidences (ACs), also known as randoms, and scatters. For accurate quantification, the noise should be taken into account.

2.3.1 Accidental Coincidence (or Random)

For many of annihilation events, only one of the photon pair is detected because the other one is absorbed or scattered out of plane by the body, and the detection probability is

less than one even when the photon arrives at the scintillation crystal. Detections of those single photons are called *singles*. If two singles from separate annihilations are detected within the coincidence timing window, then they are mistakenly registered as having arisen from the same annihilation (Fig. 2.2). This is called an *accidental coincidence* (AC) event, also known as a *random* event.

Let S_{i_1} and S_{i_2} be the singles rate (photons per time) at detector i_1 and i_2 respectively, and let τ be the duration of a coincidence timing window. For a single event at detector i_1 that happens at $t = t_1$, on average, $2\tau S_{i_2}$ events are detected by detector i_2 within the coincidence timing window, that is, for $t_1 - \tau < t < t_1 + \tau$. Therefore, the rate of AC events between detector i_1 and i_2 is given by

$$R_{i_1 i_2} = 2\tau S_{i_1} S_{i_2}.$$

In other words, an AC event occurs when two unrelated singles are detected within a coincidence timing window.

The “prompt” sinogram data, obtained using an ordinary coincidence timing window, are contaminated by AC events. To remove the contribution of the AC events, most PET centers collect “delayed” sinogram data by using a delayed coincidence timing window that is offset in time by a predetermined time delay (larger than 2τ). The delayed sinogram data cannot contain true coincidences. Since the AC events are uniformly distributed locally in time, the delayed coincidences have the same mean as the AC events in a prompt coincidence timing window. Usually, sinogram data are precorrected for ACs by real-time subtraction of the delayed sinogram data. In other words, for each coincidence event in the prompt coincidence timing window, the corresponding sinogram bin is incremented. However, for coincidence events in the delayed coincidence timing window, the corresponding sinogram bin is decremented. By the real-time precorrection, the AC events are corrected for in the mean.

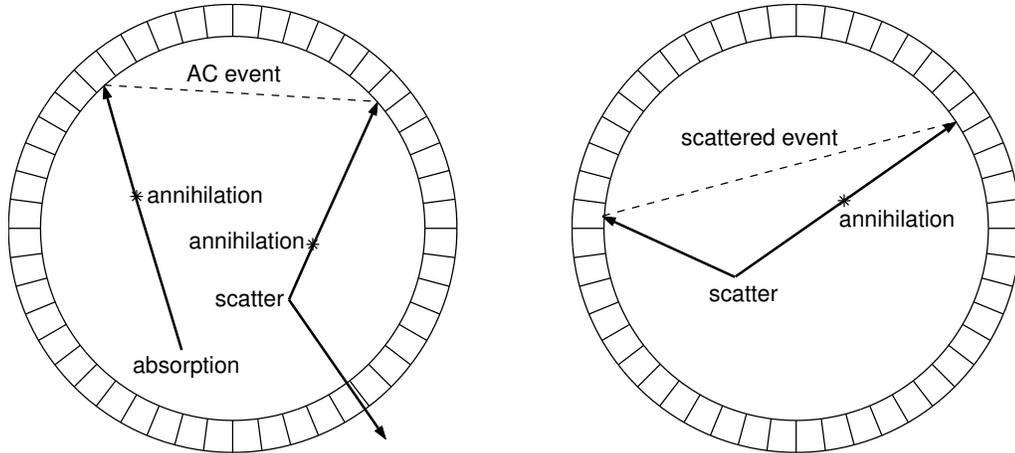


Figure 2.2: Background events in PET: accidental coincidence (AC) event (left) and scattered event (right).

2.3.2 Scatter

Although most scattered photons in the body leave the detector plane undetected, some scattered photons are still detected and registered as having an incorrect LOR (Fig. 2.2). A significant portion of scattered events can be removed by “energy discrimination” since photons lose a fraction of their energy during Compton interactions. However, scatter is a primary source of background noise particularly for septaless 3D PET, and scatter estimation is very important for accurate quantification. In this dissertation, we assume that the contribution of scattered events is known for simplicity.

2.4 System Model

2.4.1 Transmission Scans

By performing a transmission scan with exterior radioactive sources, one can reconstruct the attenuation map of the patient’s body. (The estimated attenuation map is then used for emission image reconstruction.) The attenuation map $\mu(\boldsymbol{x})$ is a real-valued non-negative function whose domain is \mathbb{R}^2 or \mathbb{R}^3 . It is usually parameterized through pixels or

voxels as follows:

$$\mu(\mathbf{x}) \cong \sum_{j=1}^p \mu_j I_j(\mathbf{x})$$

where $I_j(\cdot)$ is the indicator function representing the j th pixel or voxel, and p is the number of pixels or voxels. The integral in (2.3) is approximated by

$$\begin{aligned} \int_{\text{LOR}_i} \mu(\mathbf{x}) d\mathbf{x} &\cong \int_{\text{LOR}_i} \sum_{j=1}^p \mu_j I_j(\mathbf{x}) d\mathbf{x} = \sum_{j=1}^p \left(\int_{\text{LOR}_i} I_j(\mathbf{x}) d\mathbf{x} \right) \mu_j \\ &\triangleq \sum_{j=1}^p a_{ij} \mu_j \end{aligned}$$

where $\{a_{ij}\}$ defines a system matrix \mathbf{A} . For 2D, the element a_{ij} means the length of the intersection of LOR_i with the j th pixel, or the area of the intersection of a strip (with a finite width) representing the i th ray with the j th pixel. The mean of the i th sinogram data corresponding to the i th LOR in a transmission scan is modeled as

$$E[Y_i] = b_i e^{-\sum_{j=1}^p a_{ij} \mu_j} + r_i$$

where b_i is the i th measurement data from the blank scan, and r_i denotes backgrounds including scatter and AC events. The measurement data Y_i are modeled as independent Poisson random variables. The goal is to estimate the attenuation coefficients $\boldsymbol{\mu} = [\mu_1, \dots, \mu_p]'$ from the data $\mathbf{y} = [y_1, \dots, y_N]'$ where $'$ denotes transpose. Note that y_i are integer-valued while $E[Y_i]$ are not.

2.4.2 Emission Scans

The radioactivity $\lambda(\mathbf{x})$ is also parameterized using pixels or voxels by a discrete vector $\boldsymbol{\lambda} = [\lambda_1, \dots, \lambda_p]'$. For emission scans, the mean of the i th sinogram data is modeled as

$$E[Y_i] = \sum_{j=1}^p a_{ij} \lambda_j + r_i.$$

The a_{ij} accounts for geometry, attenuation (estimated from a transmission scan), detector efficiencies, positron range, scan time and so on. Note the a_{ij} here is different from that

in the previous subsection. The r_i denotes backgrounds including randoms, scatters, and background radiation. The measurement data Y_i are modeled as independent Poisson random variables. The goal is to estimate the emission activity $\boldsymbol{\lambda} = [\lambda_1, \dots, \lambda_p]'$ from the data $\mathbf{y} = [y_1, \dots, y_N]'$. Again, y_i are integer-valued while $E[Y_i]$ are not.

2.5 Tracer Kinetics

In the previous sections, we assumed the distribution of the radioisotopes in the body is temporally stationary (for emission tomography). We focus on static imaging of steady states in Chapters 3–6. However, in many cases, dynamic behavior of radioactivity in the body provides valuable information on specific physiological processes. In this section, we briefly review the dynamic model of radioactivity, on which Chapter 7 is based.

2.5.1 Dynamic Model

Let $\boldsymbol{\tau} = [\tau_1, \dots, \tau_{K_T}]'$ denote the emission event occurrence times for $t \in [0, T]$ in a region of interest or a pixel. Then the counting process represented by $(\boldsymbol{\tau}; K_T)$ is modeled as an inhomogeneous Poisson process whose rate (or intensity) function is $\lambda(\cdot)$ [114]. In other words, the number of events that occurred for $t \in [t_1, t_2)$, where $t_1, t_2 \in [0, T]$ and $t_1 < t_2$, is modeled as a Poisson random variable whose mean is the definite integral of $\lambda(\cdot)$ on $[t_1, t_2)$:

$$|\{\tau_1, \dots, \tau_{K_T}\} \cap [t_1, t_2)| \sim \text{Poisson} \left\{ \int_{t_1}^{t_2} \lambda(t) dt \right\}$$

where $|A|$ is the number of the elements in the set A . The rate function $\lambda(\cdot)$ is called a *time activity curve (TAC)*. The goal of dynamic imaging is to reconstruct the TACs for each ROI or pixel from the measurement data, contaminated by background noise, in sinogram domain.

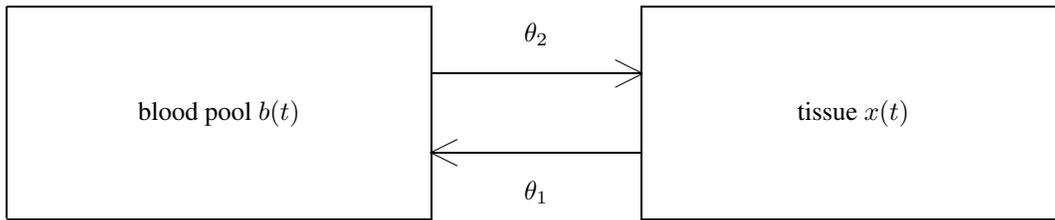


Figure 2.3: One tissue compartment model.

2.5.2 Compartmental Model

In tracer kinetics studies, TACs are often parameterized using a compartmental model based on knowledge of a physiological process of interest [27, 64]. A compartmental model lumps a system into discrete groups, called compartments, which are interconnected by pathways representing fluxes of material and/or biochemical conversions. A compartment is an amount of material that acts as though it is well-mixed and kinetically homogeneous. Compartments can represent different physical spaces, or different forms in the same physical space.

For example, consider a one tissue compartment model shown in Fig. 2.3. Let $b(t)$ denote a blood input function and let $x(t)$ denote radioactivity in the tissue. The blood input function can be measured through direct blood sampling. Assume $b(t)$ is known. Let θ_1 and θ_2 denote rate constants that are kinetic parameters to be estimated. From the governing differential equation

$$\frac{dx}{dt}(t) = \theta_2 b(t) - \theta_1 x(t),$$

the TAC can be expressed as

$$x(t) = b(t) \star h(t)$$

where \star denotes the convolution, and the impulse response of the given kinetic model, that

is, a one tissue compartment model, can be obtained as

$$h(t) = \theta_2 e^{-\theta_1 t} u(t).$$

The $u(t)$ is the unit step function. The goal is to estimate the physiologically meaningful kinetic parameters θ_1 and θ_2 from a realization, contaminated by background noise, of a Poisson process whose rate function is $h(t)$.

CHAPTER 3

Statistical Image Reconstruction Methods for Tomography

In this chapter we define the problems for statistical (static) image reconstruction in both emission and transmission tomography, which subsequent chapters, particularly Chapters 4 and 5, focus on. In transmission tomography, one estimates the attenuation map of a patient's body by using exterior radioactive sources. For emission tomography, we reconstruct the radioactivity distribution in the body by incorporating the attenuation map. In this dissertation, we consider the emission and transmission problems separately. For analysis of noise propagation from the attenuation map into emission reconstruction, see [51, 132].

We also review “optimization transfer algorithms” on which the new reconstruction algorithms we develop in this dissertation are based.

3.1 Emission Tomography

3.1.1 Statistical Model

We focus on the linear Poisson statistical model [73] that has been used extensively for emission computed tomography including positron emission tomography (PET) and single photon emission computed tomography (SPECT) as well as for photon-limited optical applications like fluorescence confocal microscopy [119]. The goal is to estimate an unknown parameter vector $\boldsymbol{\lambda} = [\lambda_1, \dots, \lambda_p]'$ from a realization $\boldsymbol{y} = [y_1, \dots, y_N]'$ of the

projection measurement random vector $\mathbf{Y} = [Y_1, \dots, Y_N]'$ where λ_j is the mean activity at the j th pixel, y_i is the number of events detected at the i th detector unit (detector pair for PET), and $'$ denotes matrix transpose. Assuming usual Poisson distributions, the measurement model¹ for emission scans is as follows:

$$Y_i \sim \text{Poisson}\{\bar{y}_i(\boldsymbol{\lambda}^{\text{true}})\}, \quad i = 1, \dots, N$$

where the measurement means are given by

$$\bar{y}_i(\boldsymbol{\lambda}) = \sum_{j=1}^p a_{ij}\lambda_j + r_i = [\mathbf{A}\boldsymbol{\lambda}]_i + r_i. \quad (3.1)$$

The vector $\mathbf{r} = [r_1, \dots, r_N]'$ denotes the mean number of background events such as scatters, random coincidences and background radiation, and $\mathbf{A} = \{a_{ij}\}$ is a system matrix incorporating scan geometry, attenuation, detector efficiencies, scanning time and so on.

We assume that r_i and a_{ij} are *known* nonnegative constants, and that the sensitivity factors, $a_j \triangleq \sum_{i=1}^N a_{ij}$, are nonzero (that is, positive) for all j , which is reasonable in practice.

3.1.2 Penalized-Likelihood Reconstruction

Assuming that the measurements Y_i are independent, the probability mass function (PMF) of \mathbf{Y} is given by

$$P(\mathbf{y}|\boldsymbol{\lambda}) = \prod_{i=1}^N \frac{1}{y_i!} e^{-\bar{y}_i(\boldsymbol{\lambda})} (\bar{y}_i(\boldsymbol{\lambda}))^{y_i}$$

where y_i are nonnegative and integer-valued. By taking the logarithm of the above PMF and ignoring constants independent of $\boldsymbol{\lambda}$, the log-likelihood of $\boldsymbol{\lambda}$ given \mathbf{y} is obtained as follows:

$$L(\boldsymbol{\lambda}) = \sum_{i=1}^N h_i([\mathbf{A}\boldsymbol{\lambda}]_i) \quad (3.2)$$

¹For randoms-precorrected PET scans, see Chapter 6.

where the marginal log-likelihood function h_i are defined by

$$h_i(l) \triangleq y_i \log(l + r_i) - (l + r_i). \quad (3.3)$$

The functions $h_i : [-r_i, \infty) \rightarrow \mathbb{R} \cup \{-\infty\}$ can be shown to have the following properties:²

- $h_i(l) \leq h_i(\hat{l}_i), \forall l \geq -r_i$ where $\hat{l}_i \triangleq y_i - r_i$. (3.4)

- h_i is monotone increasing on $[-r_i, \hat{l}_i]$ and monotone decreasing on $[\hat{l}_i, \infty)$. (3.5)

- h_i is concave. (3.6)

The property (3.6) implies that the log-likelihood L is concave.

For maximum likelihood (ML) reconstruction, one needs to find a maximizer of the log-likelihood L over \mathbb{R}_+^p where the nonnegativity constraint³ is given by

$$\mathbb{R}_+^p = \{\boldsymbol{\lambda} \in \mathbb{R}^p : \lambda_j \geq 0, \forall j\}.$$

However, the ML approach leads to unacceptably noisy images because tomographic image reconstruction is an ill-conditioned problem. So we focus on penalized-likelihood (PL) reconstruction where one must find a maximizer of the following PL objective function over a constraint set \mathcal{D} :

$$\Phi(\boldsymbol{\lambda}) = L(\boldsymbol{\lambda}) - R(\boldsymbol{\lambda}) \quad (3.7)$$

where R is a regularization term that encourages smoothness in the reconstructed images and that improves conditioning. If $R = 0$, the problem becomes ML reconstruction. We define the constraint set $\mathcal{D} \subset \mathbb{R}_+^p$ by

$$\begin{aligned} \mathcal{D} &\triangleq \{\boldsymbol{\lambda} \in \mathbb{R}_+^p : \Phi(\boldsymbol{\lambda}) \in \mathbb{R}\} \\ &= \{\boldsymbol{\lambda} \in \mathbb{R}_+^p : [\mathbf{A}\boldsymbol{\lambda}]_i > -r_i, \forall i \text{ such that } y_i > 0\} \end{aligned} \quad (3.8)$$

²For convenience, we adopt the convention that $\log 0 = -\infty$ and $0 \cdot \log 0 = 0$.

³The nonnegativity constraint is based on physical grounds.

In a physically realistic case⁴ where $r_i > 0$ for all i , we have $\mathcal{D} = \mathbb{R}_+^p$. However, if $r_i = 0$ for some i , then $\mathcal{D} \neq \mathbb{R}_+^p$. In this case, one can slightly modify Φ without changing its maximizer so that Φ and $\nabla\Phi$ are well-defined (that is, finite) over \mathbb{R}_+^p as described in Section 3.1.6.

3.1.3 Penalty Function

Although the methods described in this work can be easily generalized, for simplicity we assume that R in (3.7) is the following type of roughness penalty function:

$$R(\boldsymbol{\lambda}) = \frac{\beta}{2} \sum_{j=1}^p \sum_{k \in \mathcal{N}_j} w_{jk} \psi(\lambda_j - \lambda_k) \quad (3.9)$$

where $\beta \geq 0$ is a regularization parameter that controls the smoothness in the reconstructed image, \mathcal{N}_j denotes the neighborhood of the j th pixel, ψ is a potential function, w_{jk} are nonnegative and symmetric weights (ordinarily, $w_{jk} = 1$ for horizontal and vertical neighboring pixels, and $w_{jk} = 1/\sqrt{2}$ for diagonal neighboring pixels). Viewing the pixels of an image as nodes of a graph with neighboring pixels (say, \mathcal{N}_j for the j th pixel) connected by an edge, we assume that the graph is *connected* in the sense that it is always possible to find some sequence of edges leading from any pixel to any other pixel [67]. We

⁴Any PET scan will have nonzero randoms and any real SPECT scan will be contaminated by a scattered component and by a nonzero (but possibly quite small) component from background radiation.

also assume that the potential function ψ satisfies the following conditions (cf [1, 37, 54]):

- ψ is nonnegative. (3.10)

- ψ is symmetric. (3.11)

- ψ is continuously differentiable. (3.12)

- $\dot{\psi}(t) = \frac{d\psi}{dt}(t)$ is nondecreasing for $t \geq 0$. (3.13)

- $\omega_\psi(t) \triangleq \frac{\dot{\psi}(t)}{t}$ is nonincreasing for $t \geq 0$. (3.14)

- $\omega_\psi(0) = \lim_{t \rightarrow 0} \frac{\dot{\psi}(t)}{t}$ is finite and nonzero, that is, $0 < \omega_\psi(0) < \infty$. (3.15)

Some of the above conditions are used in deriving a surrogate function for the penalty function in Section 3.3.2. Condition (3.13) implies that the penalty function R is convex, and R is nonnegative due to (3.10).

The penalty function can be interpreted as representing *a priori* information in the following sense. Assuming that the *a priori* distribution of $\boldsymbol{\lambda}$ is $f(\boldsymbol{\lambda}) = \exp(-R(\boldsymbol{\lambda}))$, the *a posteriori* distribution is given as

$$f(\boldsymbol{\lambda}|\mathbf{Y}) = \frac{P(\mathbf{Y}|\boldsymbol{\lambda})f(\boldsymbol{\lambda})}{P(\mathbf{Y})}.$$

by Bayes' formula [121, p. 78]. One can see that the log posterior $\log f(\boldsymbol{\lambda}; \mathbf{Y})$ is equivalent to the PL objective in (3.7) (to within constants with respect to $\boldsymbol{\lambda}$), in other words, maximum *a posteriori* (MAP) reconstruction is computationally equivalent to PL reconstruction.

The penalty function in (3.9) can be expressed in a more general (and more easily implementable) form as follows:

$$R(\boldsymbol{\lambda}) = \beta \sum_{k=1}^K \psi_k([\mathbf{C}\boldsymbol{\lambda}]_k)$$

where \mathbf{C} is a $K \times p$ penalty matrix and

$$[\mathbf{C}\boldsymbol{\lambda}]_k = \sum_{j=1}^p c_{kj} \lambda_j.$$

See [37] for a specific penalty matrix \mathbf{C} to implement (3.9).

A variety of penalty functions with different noise-resolution trade-offs and edge-preserving properties have been proposed [17, 67]. See [123, 124] for methods of designing quadratic penalty functions to achieve spatially uniform resolution. In this dissertation, we assume the penalty function R is given, and we focus on maximizing the given objective function.

3.1.4 Existence and Uniqueness

One can verify that the level set $\{\boldsymbol{\lambda} \in \mathcal{D} : \Phi(\boldsymbol{\lambda}) \geq \Phi(\mathbf{1})\}$ is compact (that is, bounded and closed) where $\mathbf{1} \in \mathbb{R}^p$ is a column vector of ones,⁵ using the coerciveness⁶ of Φ , and the continuity of Φ on \mathcal{D} . Then, by the Weierstrass' Theorem [15, p. 654], there exists a (possibly nonunique) PL solution⁷ $\hat{\boldsymbol{\lambda}} = \arg \max_{\boldsymbol{\lambda} \in \mathcal{D}} \Phi(\boldsymbol{\lambda})$.

If the objective function Φ is strictly concave on \mathcal{D} , then there exists a unique PL solution [15, p. 685]. We assume strict concavity for proving convergence of some algorithms like modified BSREM in Section 4.3.1. However, we will allow a concave objective function (possibly having multiple solutions) for most cases. The following Lemma (cf Theorem 1 of [34] and Lemma 1 of [67]) provides a simple sufficient condition for the strict concavity of Φ .

Lemma 3.1. *If $\mathbf{y}'\mathbf{A}\mathbf{1} \neq 0$, then Φ in (3.7) [with (3.9) for $\beta > 0$] is strictly concave on \mathcal{D}*

⁵The vector $\mathbf{1}$ is arbitrary; any vector in \mathcal{D} can replace it.

⁶The function Φ is said to be *coercive* in the context of maximization if $\lim_{\|\boldsymbol{\lambda}\| \rightarrow \infty} \Phi(\boldsymbol{\lambda}) = -\infty$. This can be easily shown by the assumption of nonzero sensitivity factors: $\sum_{i=1}^N a_{ij} > 0, \forall j$.

⁷We use this “arg max” notation if and only if $\Phi(\hat{\boldsymbol{\lambda}}) \geq \Phi(\boldsymbol{\lambda})$ for all $\boldsymbol{\lambda} \in \mathcal{D}$.

for any ψ that is strictly convex and twice differentiable.⁸

Proof. The (negative) Hessian of Φ can be computed as follows:

$$-\nabla^2\Phi(\boldsymbol{\lambda}) = \mathbf{A}'\mathbf{W}(\boldsymbol{\lambda})\mathbf{A} + \nabla^2R(\boldsymbol{\lambda})$$

with

$$\mathbf{W}(\boldsymbol{\lambda}) = \text{diag}\left\{\frac{y_i}{\bar{y}_i^2(\boldsymbol{\lambda})}\right\} \quad (3.16)$$

for $\boldsymbol{\lambda} \in \mathcal{D}$, where we interpret $y_i/\bar{y}_i^2(\boldsymbol{\lambda})$ as 0 if $y_i = 0$. For any $\mathbf{x} \in \mathbb{R}^p$, using the symmetry of ψ and w_{jk} , we obtain:

$$\mathbf{x}'\nabla^2R(\boldsymbol{\lambda})\mathbf{x} = \frac{\beta}{2} \sum_{j=1}^p \sum_{k \in \mathcal{N}_j} w_{jk} \ddot{\psi}(\lambda_j - \lambda_k) (x_j - x_k)^2.$$

Since $\ddot{\psi} > 0$ and the neighborhood system is connected by assumption, for $\beta > 0$, $\mathbf{x}'\nabla^2R(\boldsymbol{\lambda})\mathbf{x} = 0$ only if $\mathbf{x} = \mathbf{0}$ or $\mathbf{x} = c\mathbf{1}$ for some $c \neq 0$. But $c\mathbf{1}'\mathbf{A}'\mathbf{W}(\boldsymbol{\lambda})\mathbf{A}c\mathbf{1} = c^2\|\mathbf{W}^{1/2}(\boldsymbol{\lambda})\mathbf{A}\mathbf{1}\|^2 \neq 0$ by assumption. Therefore, $\mathbf{x}'\nabla^2\Phi(\boldsymbol{\lambda})\mathbf{x} < 0, \forall \mathbf{x} \neq \mathbf{0}$. \square

Since y_i and a_{ij} are nonnegative, the assumption $\mathbf{y}'\mathbf{A}\mathbf{1} \neq 0$ is equivalent to $\mathbf{A}'\mathbf{y} \neq \mathbf{0}$. In other words, the backprojection of the data must be a nonzero image, which is reasonable in practice.

3.1.5 Boundedness

It is clear that a PL solution set

$$\Lambda^* \triangleq \{\boldsymbol{\lambda}^* \in \mathcal{D} : \Phi(\boldsymbol{\lambda}^*) \geq \Phi(\boldsymbol{\lambda}), \forall \boldsymbol{\lambda} \in \mathcal{D}\} \quad (3.17)$$

is bounded by the coerciveness of Φ . In fact, for given data \mathbf{y} , one can compute an upper bound $U = U(\mathbf{y}) \in (0, \infty)$ on the elements of Λ^* such that $\Lambda^* \subset \mathcal{B}$ where

$$\mathcal{B} \triangleq \{\boldsymbol{\lambda} \in \mathbb{R}^p : 0 \leq \lambda_j \leq U, \forall j\}. \quad (3.18)$$

⁸Such potential functions include the quadratic function $\psi(t) = t^2/2$ and many others suggested by Lange [67].

Thus, one can search for a solution over the *bounded* set $\mathcal{B} \cap \mathcal{D}$ instead of \mathcal{D} . This property is important because the boundedness of iterates generated by an algorithm plays an essential role in many convergence proofs.

We provide two types of such computable upper bounds in Appendix A. One can also compute a physically meaningful upper bound for emission images based on the dosage taken by the patient.

3.1.6 Modified Objective Function

For $\boldsymbol{\lambda} \in \mathcal{S} \triangleq \mathbb{R}_+^p \setminus \mathcal{D}$ where \mathcal{D} is defined in (3.8), we have $\Phi(\boldsymbol{\lambda}) = -\infty$ and $\|\nabla\Phi(\boldsymbol{\lambda})\| = \infty$, causing gradient-based methods to *collapse*. Although we have $\mathcal{S} = \emptyset$ in a physically realistic case that $r_i > 0$ for all i , many researchers have considered a case that $r_i = 0$. We provide a method of dealing with the problem related to the region \mathcal{S} .

The EM algorithm for ML reconstruction avoids the problem due to its intrinsic positivity, that is, the iterates never lie in \mathcal{S} as long as the starting point does not; however, in many other cases, positivity is not guaranteed. To circumvent the problem, we replace the log-likelihood near the problematic region \mathcal{S} with well-behaved functions, say, quadratic approximations. Define

$$\mathcal{I} \triangleq \{i = 1, \dots, N : r_i = 0 \text{ and } y_i > 0\}. \quad (3.19)$$

Consider the following modified objective function:

$$\tilde{\Phi}(\boldsymbol{\lambda}) = \sum_{i=1}^N \tilde{h}_i([\mathbf{A}\boldsymbol{\lambda}]_i) - R(\boldsymbol{\lambda})$$

where

$$\tilde{h}_i(l) \triangleq \begin{cases} \frac{\ddot{h}_i(\epsilon)}{2}(l - \epsilon)^2 + \dot{h}_i(\epsilon)(l - \epsilon) + h_i(\epsilon), & \text{for } l \leq \epsilon \text{ and } i \in \mathcal{I} \\ h_i(l), & \text{otherwise} \end{cases} \quad (3.20)$$

for some $\epsilon > 0$, and R is the same as in (3.9). The modified marginal log-likelihood \tilde{h}_i is a twice differentiable and strictly concave real-valued function defined on \mathbb{R} for $i \in \mathcal{I}$. Note that $\Phi(\boldsymbol{\lambda}) = \tilde{\Phi}(\boldsymbol{\lambda})$ for $\boldsymbol{\lambda} \in \mathcal{E}$ where

$$\mathcal{E} \triangleq \{\boldsymbol{\xi} \in \mathcal{D} : [\mathbf{A}\boldsymbol{\xi}]_i > \epsilon, \forall i \in \mathcal{I}\}, \quad (3.21)$$

and that $\tilde{\Phi}$ is well-defined on \mathbb{R}_+^p . The modified objective function $\tilde{\Phi}$ preserves the (strict) concavity⁹ of Φ . Remarkably, one *can compute* $\epsilon > 0$ such that

$$\Lambda^* = \tilde{\Lambda}^* \triangleq \{\boldsymbol{\lambda}^* \in \mathcal{B} : \tilde{\Phi}(\boldsymbol{\lambda}^*) \geq \tilde{\Phi}(\boldsymbol{\lambda}), \forall \boldsymbol{\lambda} \in \mathcal{B}\}, \quad (3.22)$$

where Λ^* is defined in (3.17), and \mathcal{B} is defined in (3.18) with such U as computed in the previous section. This means that the modified objective function with such ϵ has the same maximizer(s) as the original. Moreover, one can search for a solution over the *compact* set \mathcal{B} instead of \mathcal{D} or $\mathcal{B} \cap \mathcal{D}$. The modified objective function $\tilde{\Phi}$ is real-valued on \mathcal{B} , and it has a nice property that its gradient $\nabla \tilde{\Phi}$ is Lipschitz continuous¹⁰ on \mathcal{B} . We provide a method to determine such ϵ that ensures (3.22) in Appendix B.

One should be cautioned that ϵ given in (B.1) could be too small to be practical in finite precision computers; nevertheless, at least we can proceed to develop theory. Although the reconstruction methods described in this dissertation can apply to the case that $r_i = 0$ for emission tomography by using the modified objective function, for simplicity we henceforth assume that $r_i > 0$ for all i unless stated otherwise.

⁹For strict concavity, Lemma 3.1 still applies to $\tilde{\Phi}$. If $\boldsymbol{\lambda} \in \mathbb{R}_+^p \setminus \mathcal{E}$, then its corresponding diagonal element of $\mathbf{W}(\boldsymbol{\lambda})$ in (3.16) would change to $-\ddot{h}_i(\epsilon) = y_i/\epsilon^2$, and this leads to the same conclusion.

¹⁰A function f is said to be *Lipschitz continuous* on D if there exists some $L > 0$ such that $\|f(\mathbf{x}) - f(\mathbf{y})\| \leq L \|\mathbf{x} - \mathbf{y}\|$ for all $\mathbf{x}, \mathbf{y} \in D$. A differentiable function is Lipschitz continuous if its derivatives are bounded. Conversely, the derivatives of a Lipschitz continuous function are bounded when they exist. Therefore, Lipschitz continuity conditions on the gradients of a function imply that the curvatures of the function are bounded if they exist.

3.2 Transmission Tomography

3.2.1 Statistical Model

We assume the following Poisson statistical model for (monoenergetic) transmission measurements:

$$Y_i \sim \text{Poisson}\left\{b_i e^{-[\mathbf{A}\boldsymbol{\mu}^{\text{true}}]_i} + r_i\right\}, \quad i = 1, \dots, N \quad (3.23)$$

where Y_i denotes the transmission measurement random variable, b_i denotes the blank scan counts of the i th detector unit (detector pair for PET), r_i denotes the mean number of background counts, and $[\mathbf{A}\boldsymbol{\mu}]_i = \sum_{j=1}^p a_{ij}\mu_j$ represents the i th line integral of the attenuation map. The goal is to estimate an unknown parameter vector $\boldsymbol{\mu} = [\mu_1, \dots, \mu_p]'$ from a realization $\mathbf{y} = [y_1, \dots, y_N]'$ of the random vector $\mathbf{Y} = [Y_1, \dots, Y_N]'$ where μ_j is the mean attenuation coefficient in the j th pixel. We assume that a_{ij} and r_i are known nonnegative constants, and that b_i are known positive constants.

3.2.2 Penalized-Likelihood Reconstruction

We focus on penalized-likelihood (PL) reconstruction where one must find a maximizer of the PL objective function,¹¹ over \mathbb{R}_+^p ,

$$\Phi(\boldsymbol{\mu}) = L(\boldsymbol{\mu}) - R(\boldsymbol{\mu}) \quad (3.24)$$

which includes the log-likelihood function

$$L(\boldsymbol{\mu}) = \sum_{i=1}^N h_i([\mathbf{A}\boldsymbol{\mu}]_i) \quad (3.25)$$

$$h_i(l) = y_i \log(b_i e^{-l} + r_i) - (b_i e^{-l} + r_i) \quad (3.26)$$

and the roughness penalty R is defined in (3.9).

¹¹Throughout this dissertation, we use the same notations (Φ, L, h_i, U, \dots) in different contexts. One can distinguish them by the argument variable of functions: $\boldsymbol{\lambda}$ is used for emission tomography, $\boldsymbol{\mu}$ for transmission tomography, and \boldsymbol{x} for both cases or a general optimization problem unless stated otherwise.

Suppose $r_i = 0$ for all i . If $\mathbf{A}'\mathbf{y} > \mathbf{0}$, then one can show that the objective function Φ is strictly concave for R with strictly convex potential functions ψ , and that there exists a maximizer of Φ . However, in a physically realistic case that $r_i > 0$, the PL objective can be nonconcave [1], complicating maximization.

We apply a box constraint

$$\mathcal{B} = \{\boldsymbol{\mu} \in \mathbb{R}^p : 0 \leq \mu_j \leq U, \forall j\} \quad (3.27)$$

as in the emission case. The nonnegativity restriction is desirable on physical grounds, and the upper bound $U > 0$ is set by the user to be a value that is larger than the maximum attenuation coefficient conceivable for the object being scanned whereas the upper bound for the emission case was analytically computable in Section 3.1.5. The reason for using the box constraint \mathcal{B} rather than the nonnegativity constraint \mathbb{R}_+^p is that the convergence proofs for most reconstruction algorithms need the iterates to be bounded. However, imposing upper bounds is not overly restrictive in a sense that one can choose a physically meaningful upper bound for attenuation coefficients, and the PL image estimate is unlikely to be affected by U if one chooses an arbitrarily large U . In practice, if the upper bound happens to be hit by some iterate, then the user could rerun the algorithm with a larger bound.

3.3 Optimization Transfer Algorithms

Most iterative image reconstruction algorithms to find a maximizer of an objective function like (3.7) or (3.24) are based on or have their roots in *optimization transfer principles* [70]. Many algorithms can be more easily understood in the framework of the optimization transfer approach although their original motivation is different. In this section we describe the optimization transfer approach and provide two representative families of optimization transfer algorithms particularly useful for tomographic image reconstruction.

3.3.1 Optimization Transfer Principles

The problem is to find the maximizer of an objective function:

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x} \in \mathcal{D}} \Phi(\mathbf{x}).$$

In the optimization transfer approach, when the objective function $\Phi(\cdot)$ is difficult to maximize, one finds a surrogate function $\phi(\cdot; \mathbf{x}^n)$, which is easier to maximize than Φ , at a current iterate \mathbf{x}^n , and then maximizes the surrogate to obtain a new iterate \mathbf{x}^{n+1} . Repeating this process, a sequence $\{\mathbf{x}^n\}_{n=0}^{\infty}$ is generated for some initial estimate $\mathbf{x}^0 \in \mathcal{D}$. This approach can be summarized as follows:

$$\mathbf{x}^{n+1} = \arg \max_{\mathbf{x} \in \mathcal{D}} \phi(\mathbf{x}; \mathbf{x}^n), \quad (3.28)$$

for $n \in \mathbb{N}$. If we choose $\phi(\cdot; \cdot)$ appropriately, the sequence $\{\mathbf{x}^n\}$ should converge to a solution $\hat{\mathbf{x}}$.

Suppose that we choose surrogate functions satisfying the following *minorization conditions*:

$$\phi(\mathbf{x}^n; \mathbf{x}^n) = \Phi(\mathbf{x}^n) \quad (3.29)$$

$$\phi(\mathbf{x}; \mathbf{x}^n) \leq \Phi(\mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{D} \quad (3.30)$$

$$\nabla \Phi(\mathbf{x})|_{\mathbf{x}=\mathbf{x}^n} = \nabla^{10} \phi(\mathbf{x}; \mathbf{x}^n)|_{\mathbf{x}=\mathbf{x}^n} \quad (3.31)$$

for all n , where ∇^{10} is the column gradient operator with respect to the first argument.

The following key property of optimization transfer algorithms is derived from (3.29) and (3.30):

$$\Phi(\mathbf{x}) - \Phi(\mathbf{x}^n) \geq \phi(\mathbf{x}; \mathbf{x}^n) - \phi(\mathbf{x}^n; \mathbf{x}^n), \quad \forall \mathbf{x} \in \mathcal{D}, \quad (3.32)$$

which ensures that the update (3.28) *monotonically* increases Φ . Fig. 3.1 illustrates the optimization transfer principles in 1D.

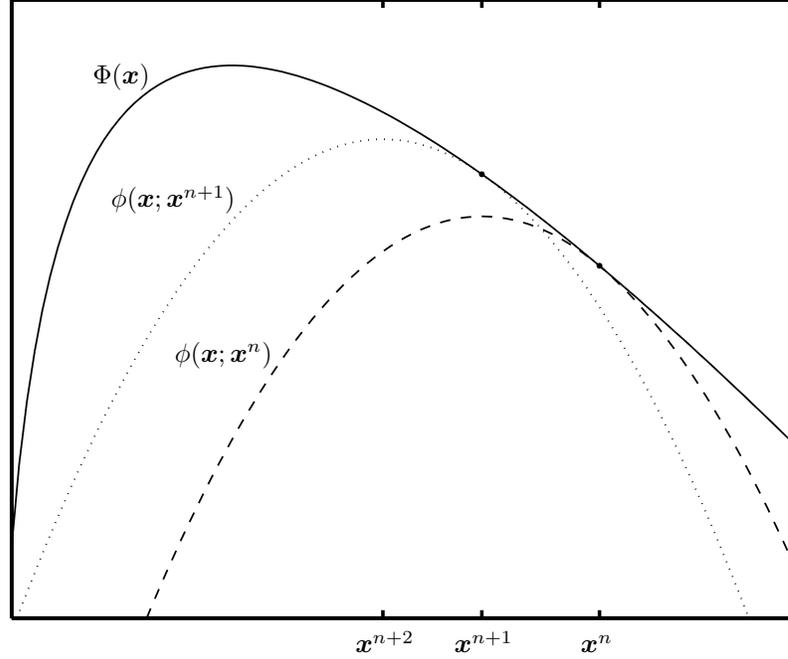


Figure 3.1: Illustration of optimization transfer in 1D. In this figure, Φ and ϕ denote an objective function and a surrogate function, respectively.

Optimization transfer algorithms can be shown globally convergent¹² under mild conditions as follows (cf [70]).

Theorem 3.2. *Suppose Φ and ϕ are differentiable functions that satisfy (3.29)–(3.31). Let \mathcal{M} be a point-to-set mapping defined on $\mathcal{D} \in \mathbb{R}^p$ that assigns to every point $\mathbf{x} \in \mathcal{D}$ a subset of \mathcal{D} such that $\mathcal{M}(\bar{\mathbf{x}}) = \{\mathbf{x}^* \in \mathcal{D} : \phi(\mathbf{x}^*; \bar{\mathbf{x}}) \geq \phi(\mathbf{x}; \bar{\mathbf{x}}), \forall \mathbf{x} \in \mathcal{D}\}$. Assume $\mathcal{M}(\bar{\mathbf{x}}) \neq \emptyset$ for all $\bar{\mathbf{x}} \in \mathcal{D}$. Suppose that, given $\mathbf{x}^0 \in \mathcal{D}$, the sequence $\{\mathbf{x}^n\}_{n=0}^\infty$ is generated satisfying $\mathbf{x}^{n+1} \in \mathcal{M}(\mathbf{x}^n)$; in other words, the sequence is obtained by (3.28). Define a set of stationary points (in the context of maximization) of Φ by*

$$\Gamma \triangleq \{\mathbf{x}^* \in \mathcal{D} : \nabla \Phi(\mathbf{x}^*)'(\mathbf{x} - \mathbf{x}^*) \leq 0, \quad \forall \mathbf{x} \in \mathcal{D}\}.$$

If $\{\mathbf{x}^n\}$ is bounded (e.g., \mathcal{D} is bounded or a level set defined by $\{\mathbf{x} \in \mathcal{D} : \Phi(\mathbf{x}) \geq$

¹²Some authors define *global convergence* as the property that an algorithm has stationary points that are limit points [78, p. 228] or that are limits [13, p. 312], irrespective of starting points. We adopt the former convention in this dissertation unless stated otherwise.

$\Phi(\mathbf{x}^0)\}$ is bounded), and the mapping \mathcal{M} is closed [78, p. 185], then any limit point of the sequence $\{\mathbf{x}^n\}_{n=0}^\infty$ is an element of Γ .

Proof. By (3.32), if $\mathbf{x} \in \mathcal{D}$, then $\Phi(\mathbf{x}^*) \geq \Phi(\mathbf{x})$ for all $\mathbf{x}^* \in \mathcal{M}(\mathbf{x})$. If $\bar{\mathbf{x}} \in \mathcal{D}$ but $\bar{\mathbf{x}} \notin \Gamma$, then $\bar{\mathbf{x}}$ cannot be a (local) maximizer of $\phi(\cdot; \bar{\mathbf{x}})$ [15, p. 194] and so $\phi(\mathbf{x}^*; \bar{\mathbf{x}}) > \phi(\bar{\mathbf{x}}; \bar{\mathbf{x}})$ for $\mathbf{x}^* \in \mathcal{M}(\bar{\mathbf{x}})$. By (3.32), it follows that $\Phi(\mathbf{x}^*) > \Phi(\bar{\mathbf{x}})$ for all $\mathbf{x}^* \in \mathcal{M}(\bar{\mathbf{x}})$. Using Zangwill's Convergence Theorem [136, p. 91] leads to the conclusion. \square

The ‘‘matched gradient’’ condition (3.31) is implied by (3.29) and (3.30) when \mathbf{x}^n is an interior point of \mathcal{D} as follows.

Lemma 3.3. *Suppose $\Phi(\cdot)$ and $\phi(\cdot; \mathbf{x}^n)$ are differentiable. If \mathbf{x}^n is an interior point of \mathcal{D} , then (3.29) and (3.30) imply (3.31).*

Proof. Assume there exists $\mathbf{z} \in \mathbb{R}^p$ such that $\|\mathbf{z}\| = 1$ and $\langle \nabla \Phi(\mathbf{x}^n) - \nabla^{10} \phi(\mathbf{x}^n; \mathbf{x}^n), \mathbf{z} \rangle < 0$ where $\langle \cdot, \cdot \rangle$ is the usual inner product. If the inequality is reversed, then take $-\mathbf{z}$. Define

$$f(\alpha) \triangleq \frac{\phi(\mathbf{x}^n + \alpha \mathbf{z}; \mathbf{x}^n) - \phi(\mathbf{x}^n; \mathbf{x}^n) - \langle \nabla^{10} \phi(\mathbf{x}^n; \mathbf{x}^n), \alpha \mathbf{z} \rangle}{\alpha} - \frac{\Phi(\mathbf{x}^n + \alpha \mathbf{z}) - \Phi(\mathbf{x}^n) - \langle \nabla \Phi(\mathbf{x}^n), \alpha \mathbf{z} \rangle}{\alpha}$$

for $\alpha > 0$. By the differentiability assumption, $\lim_{\alpha \searrow 0} f(\alpha) = 0$ [77, p. 172]. However,

$$\begin{aligned} f(\alpha) &= \frac{\phi(\mathbf{x}^n + \alpha \mathbf{z}; \mathbf{x}^n) - \Phi(\mathbf{x}^n + \alpha \mathbf{z})}{\alpha} - (\langle \nabla^{10} \phi(\mathbf{x}^n; \mathbf{x}^n), \mathbf{z} \rangle - \langle \nabla \Phi(\mathbf{x}^n), \mathbf{z} \rangle) \\ &\leq -(\langle \nabla^{10} \phi(\mathbf{x}^n; \mathbf{x}^n), \mathbf{z} \rangle - \langle \nabla \Phi(\mathbf{x}^n), \mathbf{z} \rangle) < 0 \end{aligned}$$

for $\alpha > 0$ by (3.29) and (3.30). This is a contradiction, and so (3.31) must hold. \square

However, for \mathbf{x}^n that lies at the boundary of \mathcal{D} (when \mathcal{D} is closed), (3.29) and (3.30) do not necessarily ensure (3.31). Fig. 3.2 shows a 1D example in which (3.31) does not hold at $\mathbf{x}^n = 0$ while (3.29) and (3.30) hold. In this example, $\mathbf{x} = 0$ is a fixed point of the optimization transfer algorithm whereas it is not a stationary point of the objective

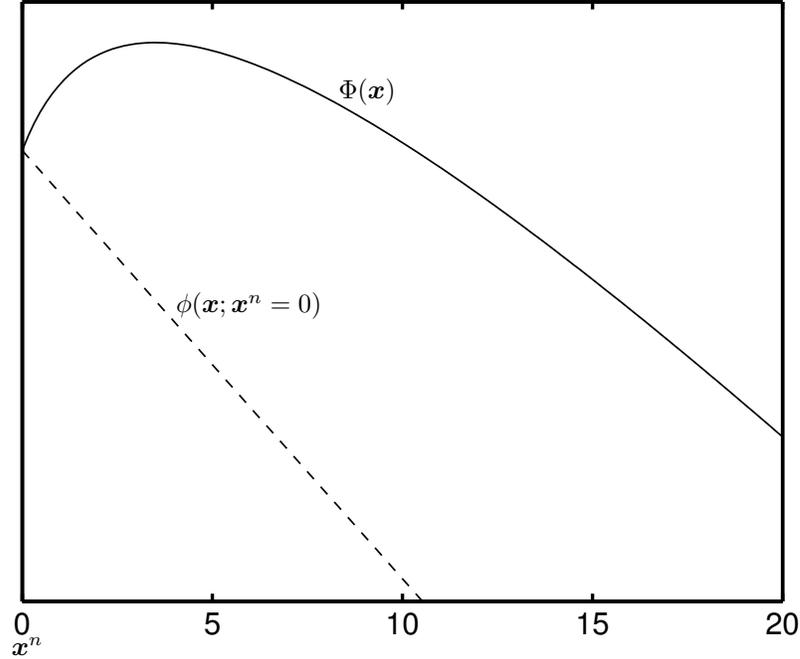


Figure 3.2: This is a 1D example, whose constraint is $x \geq 0$, showing that the “matched gradient” condition (3.31) does not hold at $x^n = 0$ while the “matched Φ value” condition (3.29) and “lies below” condition (3.30) hold.

function. Therefore, we need to state explicitly the condition (3.31) particularly when the constraint set is closed and is not \mathbb{R}^p . The condition can be slightly weaker, for example, as in [58] by using directional derivatives.

One can show that the local convergence rate of the optimization transfer algorithm in the neighborhood of an optimal point \hat{x} is determined by the following spectral radius:

$$\rho(\mathbf{I}_p - [\nabla^{20} \phi(\hat{x}; \hat{x})]^{-1} \nabla^2 \Phi(\hat{x})) \quad (3.33)$$

where \mathbf{I}_p is a $p \times p$ identity matrix, $\rho(\cdot)$ is the spectral radius, and ∇^{20} is the Hessian operator with respect to the first argument [70]. This suggests that as a surrogate $\phi(\cdot; x^n)$ approximates the objective $\Phi(\cdot)$ more accurately, or the curvature of the surrogate becomes lower, the convergence rate becomes faster. A trivial extreme case is that one takes the objective $\Phi(\cdot)$ as a surrogate $\phi(\cdot; x^n)$. However, a surrogate function should be chosen so that it is computationally efficient to maximize.

The optimization transfer approach is a general framework to yield a very broad family of monotonic algorithms by choosing application-dependent surrogate functions. Next we review two families of optimization transfer algorithms useful for tomographic reconstruction.

3.3.2 EM Surrogate Algorithms

The expectation maximization (EM) algorithms [32] provide methods for computing the ML estimate

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x} \in \mathcal{D}} \Phi(\mathbf{x}), \quad \Phi(\mathbf{x}) = \log f(\mathbf{y}; \mathbf{x})$$

where $\mathcal{D} \subset \mathbb{R}^p$ is a set of feasible parameters, $\mathbf{y} \in \mathbb{R}^N$ denotes a realization of an observable random vector \mathbf{Y} with probability distribution $f(\mathbf{y}; \mathbf{x}^{\text{true}})$, and $\mathbf{x}^{\text{true}} \in \mathbb{R}^p$ is the true value of the unknown parameter. Assume that we identify an admissible complete-data¹³ random vector \mathbf{Z} for $f(\mathbf{y}; \mathbf{x})$. Define

$$\phi(\mathbf{x}; \mathbf{x}^n) \triangleq E[\log f(\mathbf{Z}; \mathbf{x}) | \mathbf{Y} = \mathbf{y}; \mathbf{x}^n], \quad (3.34)$$

which is called an *EM surrogate function*. One obtains \mathbf{x}^{n+1} by maximizing (3.34) as in (3.28). By computing the EM surrogate in (3.34) (“E-step”) and maximizing the EM surrogate (“M-step”) repeatedly, one generates a sequence $\{\mathbf{x}^n\}_{n=0}^{\infty}$ for an initial estimate $\mathbf{x}^0 \in \mathcal{D}$. The EM surrogate functions in (3.34) ensure the monotonicity condition (3.32) [32]. The EM algorithms can be viewed as a special case of optimization transfer algorithms using the particular EM surrogates.

The EM algorithms successfully apply to the emission reconstruction problem in Section 3.1 to yield a simple form of update equations whereas the EM surrogate does not have a closed-form maximizer for the transmission problem in Section 3.2 [68, 112]. For

¹³A random vector \mathbf{Z} with probability distribution $f(\mathbf{z}; \mathbf{x})$ is called an admissible complete-data vector for $f(\mathbf{y}; \mathbf{x})$ if $f(\mathbf{y}, \mathbf{z}; \mathbf{x}) = f(\mathbf{y}|\mathbf{z})f(\mathbf{z}; \mathbf{x})$ [39, 40]. A special case is that \mathbf{Y} is a deterministic function of \mathbf{Z} .

the emission problem we review the EM algorithm on which some algorithms we develop are based. We derive the EM surrogates in (3.34) by using an alternate algebraic trick which is simpler rather than computing the conditional expectation of the complete-data log-likelihood.

EM for ML emission reconstruction

First, we derive the EM surrogate function ϕ^L for the emission log-likelihood L in (3.2) with (3.3) by using concavity of the logarithm [29]. Assume $r_i > 0$ for all i for simplicity although the derivation below can be generalized to a case that $r_i = 0$. For $\boldsymbol{\lambda}^n > \mathbf{0}$, we have

$$\begin{aligned}
L(\boldsymbol{\lambda}) &= \sum_{i=1}^N y_i \left[\log \left(\sum_{j=1}^p a_{ij} \lambda_j + r_i \right) - \left(\sum_{j=1}^p a_{ij} \lambda_j + r_i \right) \right] \\
&= \sum_{i=1}^N \left[y_i \log \left(\sum_{j=1}^p \frac{a_{ij} \lambda_j^n}{\bar{y}_i(\boldsymbol{\lambda}^n)} \frac{\lambda_j}{\lambda_j^n} \bar{y}_i(\boldsymbol{\lambda}^n) + \frac{r_i}{\bar{y}_i(\boldsymbol{\lambda}^n)} \bar{y}_i(\boldsymbol{\lambda}^n) \right) - \left(\sum_{j=1}^p a_{ij} \lambda_j + r_i \right) \right] \\
&\geq \sum_{i=1}^N \left[y_i \left\{ \sum_{j=1}^p \frac{a_{ij} \lambda_j^n}{\bar{y}_i(\boldsymbol{\lambda}^n)} \log \left(\frac{\lambda_j}{\lambda_j^n} \bar{y}_i(\boldsymbol{\lambda}^n) \right) + \frac{r_i}{\bar{y}_i(\boldsymbol{\lambda}^n)} \log \bar{y}_i(\boldsymbol{\lambda}^n) \right\} \right. \\
&\quad \left. - \left(\sum_{j=1}^p a_{ij} \lambda_j + r_i \right) \right] \tag{3.35}
\end{aligned}$$

where \bar{y}_i is defined in (3.1). By ignoring constants independent of $\boldsymbol{\lambda}$ in (3.35), one can define the following minorizing surrogate function

$$\phi^L(\boldsymbol{\lambda}; \boldsymbol{\lambda}^n) = \sum_{j=1}^p \left\{ \left(\sum_{i=1}^N \frac{a_{ij} \lambda_j^n}{\bar{y}_i(\boldsymbol{\lambda}^n)} y_i \right) \log \lambda_j - a_j \lambda_j \right\}, \tag{3.36}$$

where $a_j \triangleq \sum_{i=1}^N a_{ij}$, which is equivalent to the EM surrogate (3.34) for usual complete data in the emission problem (to within constants with respect to $\boldsymbol{\lambda}$), and satisfies the minorization conditions (3.29)–(3.31). By setting its derivative to zero, one can obtain the maximizer of the separable EM surrogate $\phi^L(\boldsymbol{\lambda}; \boldsymbol{\lambda}^n)$ over $\boldsymbol{\lambda} > \mathbf{0}$ as follows:

$$\lambda_j^{n+1} = \frac{\lambda_j^n}{a_j} \sum_{i=1}^N \frac{a_{ij} y_i}{\bar{y}_i(\boldsymbol{\lambda}^n)} > 0 \text{ for } j = 1, \dots, p, \tag{3.37}$$

which is the EM algorithm for ML reconstruction. The above EM update can be written equivalently using the gradient of the log-likelihood as follows:

$$\lambda_j^{n+1} = \lambda_j^n + \frac{\lambda_j^n}{a_j} \frac{\partial L}{\partial \lambda_j}(\boldsymbol{\lambda}^n) \quad (3.38)$$

for all j .

De Pierro's modified EM for PL emission reconstruction

Next, we derive a separable surrogate ϕ^R for the penalty function R in (3.9) by using the convexity trick proposed in [30]. For a potential function ψ satisfying (3.10)–(3.15), we have

$$\begin{aligned} \psi(\lambda_j - \lambda_k) &= \psi\left(\frac{1}{2}[2\lambda_j - \lambda_j^n - \lambda_k^n] + \frac{1}{2}[-2\lambda_k + \lambda_j^n + \lambda_k^n]\right) \\ &\leq \frac{1}{2}\psi(2\lambda_j - \lambda_j^n - \lambda_k^n) + \frac{1}{2}\psi(-2\lambda_k + \lambda_j^n + \lambda_k^n) \end{aligned} \quad (3.39)$$

by the convexity of ψ [30, 69]. But there exists the following majorizing quadratic surrogate function for ψ [54, p. 184]:

$$q(t; s) \triangleq \psi(s) + \dot{\psi}(s)(t - s) + \frac{1}{2}\omega_\psi(s)(t - s)^2 \geq \psi(t). \quad (3.40)$$

Combining (3.39) and (3.40), one can obtain a separable quadratic surrogate ϕ^R for the penalty function as follows:

$$\begin{aligned} -R(\boldsymbol{\lambda}) &\geq -\frac{\beta}{2} \sum_{j=1}^p \sum_{k \in \mathcal{N}_j} \frac{w_{jk}}{2} \{ \psi(2\lambda_j - \lambda_j^n - \lambda_k^n) + \psi(-2\lambda_k + \lambda_j^n + \lambda_k^n) \} \\ &= -\frac{\beta}{2} \sum_{j=1}^p \sum_{k \in \mathcal{N}_j} w_{jk} \psi(2\lambda_k - \lambda_j^n - \lambda_k^n) \end{aligned} \quad (3.41)$$

$$\geq -\frac{\beta}{2} \sum_{j=1}^p \sum_{k \in \mathcal{N}_j} w_{jk} q(2\lambda_j - \lambda_j^n - \lambda_k^n; \lambda_j^n - \lambda_k^n) \quad (3.42)$$

$$\begin{aligned} &= -\sum_{j=1}^p \left[\frac{r_j(\boldsymbol{\lambda}^n)}{2} (\lambda_j - \lambda_j^n)^2 + \beta \sum_{k \in \mathcal{N}_j} w_{jk} \dot{\psi}(\lambda_j - \lambda_j^n) + \frac{\beta}{2} w_{jk} \psi(\lambda_j^n - \lambda_k^n) \right] \\ &= -\sum_{j=1}^p \left[\frac{r_j(\boldsymbol{\lambda}^n)}{2} (\lambda_j - \lambda_j^n)^2 + \frac{\partial R}{\partial \lambda_j}(\boldsymbol{\lambda}^n) (\lambda_j - \lambda_j^n) \right] - R(\boldsymbol{\lambda}^n) \\ &\triangleq \phi^R(\boldsymbol{\lambda}; \boldsymbol{\lambda}^n) \end{aligned} \quad (3.43)$$

where (3.41) is due to symmetry of ψ and w_{jk} , (3.42) is due to (3.40), and

$$r_j(\boldsymbol{\lambda}) = 2\beta \sum_{k \in \mathcal{N}_j} w_{jk} \omega_\psi(\lambda_j - \lambda_k). \quad (3.44)$$

Now we have a global surrogate function by combining (3.36) and (3.43):

$$\phi(\boldsymbol{\lambda}; \boldsymbol{\lambda}^n) = \phi^L(\boldsymbol{\lambda}; \boldsymbol{\lambda}^n) + \phi^R(\boldsymbol{\lambda}; \boldsymbol{\lambda}^n). \quad (3.45)$$

By setting the derivative of (3.45) to zero, one obtains the following De Pierro's modified EM (DPEM) algorithm for PL reconstruction in the emission problem [30]:

$$\lambda_j^{n+1} = \text{root} \left(r_j(\boldsymbol{\lambda}^n), \frac{1}{2} \left(a_j + \frac{\partial R}{\partial \lambda_j}(\boldsymbol{\lambda}^n) - r_j(\boldsymbol{\lambda}^n) \lambda_j^n \right), \sum_{i=1}^N \frac{a_{ij} \lambda_j^n}{\bar{y}_i(\boldsymbol{\lambda}^n)} y_i \right) \quad (3.46)$$

for $j = 1, \dots, p$, where $\text{root}(\alpha, \beta, \gamma)$ denotes the nonnegative root of $\alpha x^2 + 2\beta x - \gamma = 0$.

3.3.3 Quadratic Surrogate Algorithms

It is natural to use quadratic surrogate functions which are usually easy to maximize. For both emission and transmission problems, one can find the following 1D quadratic

function which is a *proper* surrogate for a 1D marginal log-likelihood h_i in (3.3) or (3.26) for some choice of the curvature $c_i > 0$ in the sense that $q_i(l; \bar{l}) \leq h_i(l)$, $\forall l \geq 0$, and $q_i(\bar{l}; \bar{l}) = h_i(\bar{l})$ [1, 33]:

$$q_i(l; \bar{l}) = h_i(\bar{l}) + \dot{h}_i(\bar{l})(l - \bar{l}) - \frac{c_i(\bar{l})}{2}(l - \bar{l})^2. \quad (3.47)$$

There are at least two types of such curvatures: “maximum curvature (MC)” and “optimum curvature (OC).” The maximum curvature is given by

$$c_i^{\text{MC}}(\bar{l}) \triangleq \max_{l \geq 0} \left\{ -\ddot{h}_i(l) \right\}, \quad (3.48)$$

which is a constant, and can be precomputed and stored. However, this curvature is usually large and the convergence rate is slow in light of (3.33). On the other hand, the optimum curvature is the smallest curvature the quadratic surrogate in (3.47) can have while satisfying the minorization conditions. One can compute it analytically for both emission and transmission marginal log-likelihoods h_i in (3.3) and (3.26) as follows [1, 33]:

$$\begin{aligned} c_i^{\text{OC}}(\bar{l}) &\triangleq \min \left\{ c \geq 0 : h_i(\bar{l}) + \dot{h}_i(\bar{l})(l - \bar{l}) - \frac{c_i(\bar{l})}{2}(l - \bar{l})^2 \leq h_i(l), \forall l \geq 0 \right\} \\ &= \begin{cases} \left[-2 \frac{h_i(0) - h_i(\bar{l}) + \dot{h}_i(\bar{l}) \cdot \bar{l}}{\bar{l}^2} \right]_+, & \bar{l} > 0 \\ \left[-\ddot{h}_i(0) \right]_+, & \bar{l} = 0. \end{cases} \end{aligned} \quad (3.49)$$

Summing the 1D surrogate functions in (3.47), a quadratic surrogate Q for the log-likelihood function L in (3.2) or (3.25) is obtained as follows:

$$\begin{aligned} Q(\mathbf{x}; \mathbf{x}^n) &\triangleq \sum_{i=1}^N q_i([\mathbf{A}\mathbf{x}]_i; [\mathbf{A}\mathbf{x}^n]_i) \\ &= -\frac{1}{2}(\mathbf{x} - \mathbf{x}^n)' \mathbf{A}' \text{diag}\{c_i(\mathbf{x}^n)\} \mathbf{A}(\mathbf{x} - \mathbf{x}^n) + \\ &\quad \nabla L(\mathbf{x}^n)'(\mathbf{x} - \mathbf{x}^n) + L(\mathbf{x}^n). \end{aligned} \quad (3.50)$$

where $\text{diag}\{\cdot\}$ denotes a diagonal matrix appropriately formed. The above surrogate function is nonseparable with respect to \mathbf{x} , and so it is difficult to maximize. One can perform

a sequential update by using a coordinate ascent approach [1]. However, this method requires column access of the system matrix \mathbf{A} , and it is difficult to implement unless the system matrix is precomputed and stored.¹⁴ Fortunately, one can apply the concavity trick in [30] again to Q and obtain the following separable quadratic surrogate for the log-likelihood (see [2] for detailed derivation):

$$\phi^L(\mathbf{x}; \mathbf{x}^n) = -\frac{1}{2}(\mathbf{x} - \mathbf{x}^n)' \text{diag}\{\tilde{c}_j(\mathbf{x}^n)\} (\mathbf{x} - \mathbf{x}^n) + \nabla L(\mathbf{x}^n)'(\mathbf{x} - \mathbf{x}^n) + L(\mathbf{x}^n) \quad (3.51)$$

where

$$\tilde{c}_j(\mathbf{x}) = \sum_{i=1}^N a_{ij} a_i c_i(\mathbf{x}). \quad (3.52)$$

with $a_i \triangleq \sum_{j=1}^p a_{ij}$. Now, by combining (3.51) and (3.43), we have a separable quadratic global surrogate function for the PL objective function in (3.7) or (3.24). By maximizing the surrogate function, the following separable paraboloidal surrogates (SPS) update is obtained [2]:

$$x_j^{n+1} = \left[x_j^n + \frac{\nabla \Phi(\mathbf{x}^n)}{\tilde{c}_j(\mathbf{x}^n) + r_j(\mathbf{x}^n)} \right]_+ \quad (3.53)$$

where $[x]_+ = \max\{x, 0\}$, and \tilde{c}_j and r_j are defined in (3.52) and (3.44) respectively.

We derived the separable surrogate in (3.51) for the log-likelihood function from the nonseparable surrogate in (3.50) by using De Pierro's concavity tricks in [30]. The separable surrogate is easier to maximize than the nonseparable one, that is, less computation per iteration is required. However, the curvatures of the separable surrogate are higher, that is, the convergence rate is slower. There is a trade-off between a convergence rate and a computational efficiency in choosing surrogates.

¹⁴Consider the following example where the column access of the system matrix is difficult. Suppose that the system matrix (for 2D) is factored as follows: $\mathbf{A} = \mathbf{B}\mathbf{G}$ where \mathbf{G} denotes geometric projection and \mathbf{B} represents the blurring caused by the detector responses including the effects of crystal penetration and inter-crystal scattering. Let $\mathbf{G} = [\mathbf{G}_1, \dots, \mathbf{G}_{n_a}]'$ where \mathbf{G}_i denotes the projection at the i th angle and n_a denotes the number of angles in the sinogram. Assuming the blurring is applied only in the radial direction, the blurring matrix can be written as $\mathbf{B} = \mathbf{I} \otimes \mathbf{B}_{1D}$ where \mathbf{I} is the $n_a \times n_a$ identity matrix and \mathbf{B}_{1D} is a 1D blurring matrix. Now the system matrix can be written as: $\mathbf{A} = [\mathbf{B}_{1D}\mathbf{G}_1, \dots, \mathbf{B}_{1D}\mathbf{G}_{n_a}]'$. One can see that it is computationally inefficient to access a column of \mathbf{A} whereas the row access is convenient.

CHAPTER 4

Convergent Relaxed Ordered Subsets Algorithms

4.1 Introduction

Statistical image reconstruction methods lead to improved image quality over conventional filtered backprojection (FBP) methods by using accurate physical models, taking the stochastic nature of noise into account, and enforcing object constraints like nonnegativity. However, iterative algorithms for achieving maximum likelihood (ML) or penalized-likelihood (PL) reconstruction require considerable computation, and there have been ongoing efforts to develop fast algorithms.

Ordered subsets (OS) algorithms, also known as block iterative or incremental gradient methods, have been very popular in the medical imaging community for tomographic image reconstruction due to their fast convergence rates [2, 4, 18, 19, 31, 55, 60, 65, 66, 76]. The incremental gradient type algorithms are also found in convex programming [61, 63, 88, 89].

The classical “algebraic reconstruction technique” (ART) [42, 47] can be considered to be a type of “ordered subsets” method in which each subset consists of a single measurement. However, most ART methods formulate the reconstruction problem as one of finding the solution to a system of equations that involves the imaging physics but not the measurement statistics. Some ART algorithms can be made to converge by introducing

relaxation, but the limiting solution has a geometric interpretation in terms of distances to hyperplanes, rather than arising from statistical considerations [23–25]. In this work we focus on OS algorithms that are designed to maximize an objective function that captures the statistical properties of the measurements.

The OS algorithms apply successfully to problems where an objective function of interest is a sum of a large number of component functions. Because of the assumed statistical independence of tomographic data, such sums arise in statistical tomographic reconstruction problems like PL or ML reconstruction. Typically, the OS methods decompose the sum of component functions into several subsums or subobjective functions, each corresponding to a subset of the projection views, and update the image estimate by using, in a specified cyclic order (or sometimes randomly [88, 89]), the gradient of a subobjective function as an approximate gradient of the original objective function. When the noise is independent, each subobjective function is associated with a subset of the measurement data. If the subset gradients are suitably “balanced,” then the gradient approximation can be quite reasonable when the initial estimates are far from an optimal point (a PL or ML solution). Thus OS methods initially accelerate “convergence” in the sense that less computation is required to achieve nearly the same level of the objective increase as with non-OS gradient based methods. The OS methods can be viewed as incremental gradient methods in which incremental gradients, that is, gradients of subobjective functions, are used for each update [89].

For emission image reconstruction, the OS idea was applied to the classical EM algorithm [32, 68, 112] to generate several OS algorithms. ML reconstruction algorithms include the ordered subsets expectation maximization (OS-EM) [55], the rescaled block-iterative expectation maximization maximum likelihood (RBI-EMML) [19], and the row-action maximum likelihood algorithm (RAMLA) [18]. PL reconstruction algorithms in-

clude the block sequential regularized expectation maximization (BSREM) [31] (RAMLA is a special unregularized case of BSREM). The paraboloidal surrogates (PS) methods [1, 33] also adopted the OS idea to construct the ordered subsets separable paraboloidal surrogates (OS-SPS) [2], originally named the ordered subsets transmission (OSTR) algorithm in the context of transmission tomography.

Although OS algorithms, including OS-EM, RBI-EMML, and OS-SPS, successfully accelerate “convergence,” they generally do not converge to an optimal solution but rather become eventually stuck at a suboptimal limit cycle that consists of as many points as there are subsets. In fact, OS-EM and RBI-EM in their original forms [19, 55] usually do not converge to the optimal point even if relaxed due to their subset-dependent scaling (or preconditioning) matrices [4].

Convergence to an optimal solution is important for any algorithm for optimization problems, particularly in medical applications where reliability and stability are essential. For PL (or MAP) reconstruction, the convergence issue is more critical than ML for which we do not usually run algorithms to convergence. It is desirable to achieve both fast initial convergence rates (typical of OS algorithms) and global convergence.

One method for making OS algorithms globally convergent is to use relaxation parameters, that is, diminishing stepsizes. This comes from the intuition that the size of a limit cycle should be proportional to the stepsize. BSREM uses relaxation parameters, and De Pierro and Yamagishi [31] provided convergence proofs for BSREM after imposing a few “a posteriori” assumptions: the convergence of the objective sequence, and the positivity and boundedness of each iterate. We relax these assumptions by making some modifications to BSREM.

Kudo, Nakazawa, and Saito [65, 66] also used a relaxation scheme in their block-gradient method applied to penalized weighted least-squares image reconstruction for

emission tomography; however, they ignored the nonnegativity constraint. Their method appears to be a special case of incremental gradient methods [61, 88, 89]. Nedić and Bertsekas analyzed the incremental gradient methods and obtained many useful results about convergence properties [88, 89]. We observe that OS-SPS is a special case of diagonally-scaled version of incremental gradient methods with a constant stepsize, and then prove the global convergence of diagonally-scaled incremental gradient methods with diminishing stepsizes, thereby establishing global convergence of relaxed OS-SPS.

In this chapter we focus on emission image reconstruction although the methods presented here are quite general. We present two types of relaxed OS algorithms [3]: modified BSREM and relaxed OS-SPS. We also prove the global convergence of the algorithms. Both of them use diagonally-scaled gradient ascent for each update to maximize a PL objective function. Although the main difference between these two methods is the form of scaling functions, the approaches of convergence proofs are quite different. These algorithms are parallelizable, that is, they are able to update all pixels simultaneously and independently, so they are computationally convenient like EM.

In Section 4.2, we review the OS algorithms. Section 4.3 provides our modified BSREM and relaxed OS-SPS algorithms. Section 4.4 gives simulation results including discussion of relaxation parameters as related to convergence rate.

4.2 Ordered Subsets Algorithms

In many applications, an objective function of interest can be expressed as a sum of several subobjective functions:

$$\Phi(\mathbf{x}) = \sum_{m=1}^M \Phi_m(\mathbf{x}).$$

Assume that we find a surrogate function ϕ_m for each subobjective function Φ_m such that ϕ_m satisfy the minorization conditions (3.29) and (3.30) for Φ_m . Given a current iterate

\mathbf{x}^n , maximizing one of $\phi_m(\cdot; \mathbf{x}^n)$ in a cyclic order leads to an “ordered subsets” algorithm:

$$\mathbf{x}^{n+1} = \arg \max_{\mathbf{x} \in \mathcal{D}} \phi_{1+(n \bmod M)}(\mathbf{x}; \mathbf{x}^n) \quad (4.1)$$

where \mathcal{D} denotes a constraint set. When the data are independent, each subobjective function involves a subset of data: hence the name “ordered subsets.” When the number M of subobjective functions is one, the OS algorithm reduces to an ordinary optimization transfer algorithm in (3.28). In Section 4.2.2 an alternate geometric interpretation on OS algorithms is given using gradients.

4.2.1 PL Emission Reconstruction Problem

In this chapter, we focus on PL reconstruction in emission tomography described in Section 3.1. We rewrite the problem for completeness. The goal is to find the following PL estimate:

$$\hat{\boldsymbol{\lambda}} = \arg \max_{\boldsymbol{\lambda} \in \mathcal{B}} \Phi(\boldsymbol{\lambda}), \quad \Phi(\boldsymbol{\lambda}) = L(\boldsymbol{\lambda}) - R(\boldsymbol{\lambda}) \quad (4.2)$$

where the box constraint \mathcal{B} is defined in (3.18), the penalty function R is given in (3.3), and the log-likelihood L is given as follows:

$$L(\boldsymbol{\lambda}) = \sum_{i=1}^N h_i([\mathbf{A}\boldsymbol{\lambda}]_i), \quad h_i(l) = y_i \log(l + r_i) - (l + r_i). \quad (4.3)$$

In this chapter for simplicity we focus on a quadratic penalty function using

$$\psi(t) = \frac{t^2}{2}.$$

We assume for simplicity that background contribution $r_i > 0$ for all i although one can apply the methods presented here to the case that $r_i = 0$ by modifying the objective function as discussed in Section 3.1.6.

4.2.2 Ordered Subsets (OS) Algorithms

Most iterative algorithms¹ involve the gradient of the objective function, e.g., EM in (3.38) or SPS in (3.53), and many “parallelizable” algorithms—able to update all the pixels simultaneously like EM or SPS—can be written in the following form²:

$$\lambda_j^{n+1} = \lambda_j^n + \alpha_n d_j(\boldsymbol{\lambda}^n) \frac{\partial \Phi}{\partial \lambda_j}(\boldsymbol{\lambda}^n), \quad j = 1, 2, \dots, p, \quad (4.4)$$

where $\alpha_n > 0$ is a relaxation parameter³ (or stepsize), and $d_j(\boldsymbol{\lambda})$ is a nonnegative scaling function. We call the nonnegative function $d_j(\boldsymbol{\lambda})$ a *scaling function* to emphasize that it *scales* the derivative. Likewise, in vector form,

$$\boldsymbol{\lambda}^{n+1} = \boldsymbol{\lambda}^n + \alpha_n \mathbf{D}(\boldsymbol{\lambda}^n) \nabla \Phi(\boldsymbol{\lambda}^n), \quad (4.5)$$

we call the $p \times p$ matrix $\mathbf{D}(\boldsymbol{\lambda})$ a *scaling matrix*.⁴ The partial derivative of Φ is given by:

$$\frac{\partial \Phi}{\partial \lambda_j}(\boldsymbol{\lambda}) = \sum_{i=1}^N a_{ij} \dot{h}_i([\mathbf{A}\boldsymbol{\lambda}]_i) - \frac{\partial R}{\partial \lambda_j}(\boldsymbol{\lambda}). \quad (4.6)$$

The derivative involves a sum over sinogram index i , that is, backprojection. Let $\{S_m\}_{m=1}^M$ be disjoint subsets of $\{1, 2, \dots, N\}$ such that $\bigcup_{m=1}^M S_m = \{1, 2, \dots, N\}$, and define sub-objective functions as

$$\Phi_m(\boldsymbol{\lambda}) \triangleq \sum_{i \in S_m} h_i([\mathbf{A}\boldsymbol{\lambda}]_i) - \gamma_m R(\boldsymbol{\lambda}), \quad (4.7)$$

resulting in

$$\Phi = \sum_{m=1}^M \Phi_m, \quad (4.8)$$

¹DPEM in (3.46) is an exception.

²Although for some algorithms we need to enforce the constraint each iteration, we ignore this detail in this subsection to simplify explanation of OS principles. We do consider this important detail in the convergence proofs, however.

³When an update equation is written as (new iterate) = (current iterate) + $\alpha_n \times$ (correction term), one speaks of *underrelaxation* if $\alpha_n < 1$, and *overrelaxation* if $\alpha_n > 1$ [125, p. 546].

⁴The matrix \mathbf{D} is also called a *preconditioner* or a *preconditioning matrix*.

where the regularization term R is included in one or more of the Φ_m 's by choosing $\gamma_m \geq 0$ and $\sum_{m=1}^M \gamma_m = 1$. (Typically we choose $\gamma_m = 1/M$.) Suppose that the following “subset gradient balance” conditions hold:

$$\nabla\Phi_1(\boldsymbol{\lambda}) \cong \nabla\Phi_2(\boldsymbol{\lambda}) \cong \dots \cong \nabla\Phi_M(\boldsymbol{\lambda}), \quad (4.9)$$

or equivalently,

$$\nabla\Phi(\boldsymbol{\lambda}) \cong M\nabla\Phi_m(\boldsymbol{\lambda}), \quad \forall m. \quad (4.10)$$

Then one can replace $\partial\Phi/\partial\lambda_j$ with $M\partial\Phi_m/\partial\lambda_j$ in (4.4), and using $\partial\Phi_m/\partial\lambda_j$ in a cyclic order leads to an ordered subsets version of (4.4):⁵

$$\begin{aligned} \lambda_j^{n,0} &= \lambda_j^n \\ \lambda_j^{n,m} &= \lambda_j^{n,m-1} + \alpha_n d_j(\boldsymbol{\lambda}^{n,m-1}) \frac{\partial\Phi_m}{\partial\lambda_j}(\boldsymbol{\lambda}^{n,m-1}), \quad \text{for } m = 1, \dots, M \\ \lambda_j^{n+1} &= \lambda_j^{n,M} \end{aligned} \quad (4.11)$$

where the factor M in $M\partial\Phi_m/\partial\lambda_j$ is absorbed into d_j (or α_n). We refer to each update in (4.11) as the m th subiteration of the n th iteration. In the tomography context, the partition $\{S_m\}_{m=1}^M$ is naturally chosen so that projections within one subset correspond to projections with downsampled projection angles. It is desirable to order the subsets such that projections corresponding to one subset are as “perpendicular” as possible to previously used angles at each subiteration [47]. This strategy has a long history; Hamaker and Solomon [46] analyzed quantitatively the relationship between the convergence rate of ART and ordering in terms of the angles between the null spaces of each projection.

⁵One could use relaxation parameters $\alpha_{n,m}$ which depends on m . In this case, for global convergence, the variations of $\alpha_{n,m}$ over each cycle must be sufficiently small asymptotically (as n goes to ∞). For example, see [61]. However, to avoid undue complexity in convergence analysis, we focus on relaxation parameters that are held constant during each iteration, as is widely used [18, 31, 65, 89].

Fig. 4.1 illustrates the behavior of an OS algorithm for a toy example with the following objective function:

$$\Phi(\mathbf{x}) = \sum_{i=1}^3 \left(-\frac{1}{2} \mathbf{x}' \mathbf{Q}_i \mathbf{x} + \mathbf{b}'_i \mathbf{x} \right)$$

where $\mathbf{Q}_1 = \begin{bmatrix} 1 & 1 \\ 1 & 2 \end{bmatrix}$, $\mathbf{Q}_2 = \begin{bmatrix} 2 & -1 \\ -1 & 1 \end{bmatrix}$, $\mathbf{Q}_3 = \begin{bmatrix} 3 & 0 \\ 0 & 1 \end{bmatrix}$, $\mathbf{b}_1 = [1.25 \ 2.5]'$, $\mathbf{b}_2 = [-1.25 \ 0.25]'$, $\mathbf{b}_3 = [3 \ -0.75]'$, and the maximizer is $\hat{\mathbf{x}} = [0.50 \ 0.5]'$. We compare an ordinary gradient ascent method:

$$\mathbf{x}^{n+1} = \mathbf{x}^n + \alpha \nabla \Phi(\mathbf{x}^n)$$

with $\alpha = .05$, and its ordered subsets version with 3 subsets:

$$\mathbf{x}^{n,m} = \mathbf{x}^{n,m-1} + 3\alpha \nabla \Phi_m(\mathbf{x}^{n,m-1}) \text{ for } m = 1, 2, 3$$

where $\Phi_m = -\frac{1}{2} \mathbf{x}' \mathbf{Q}_m \mathbf{x} + \mathbf{b}'_m \mathbf{x}$. As can be seen in the figure, the OS algorithm is about three times faster initially far from the optimal point, but it converges to a limit cycle.

Although it is hard to prove the existence of such a limit cycle, one can expect that a set of limit points $\{\boldsymbol{\lambda}^{*,m}\}_{m=1}^M$ of a sequence $\{\boldsymbol{\lambda}^{n,m}\}$ generated by (4.11) would satisfy:

$$\lambda_j^{*,m} = \lambda_j^{*,m-1} + \alpha d_j(\boldsymbol{\lambda}^{*,m-1}) \frac{\partial \Phi_m}{\partial \lambda_j}(\boldsymbol{\lambda}^{*,m-1}), \quad \forall m$$

$$\lambda_j^{*,M} = \lambda_j^{*,0}.$$

These conditions generally differ from the true optimality conditions, that is,

$$\frac{\partial \Phi}{\partial \lambda_j}(\boldsymbol{\lambda}^*) = \sum_{m=1}^M \frac{\partial \Phi_m}{\partial \lambda_j}(\boldsymbol{\lambda}^*) = 0$$

in the context of unconstrained optimization.

One may need to use relaxation parameters α_n such that $\lim_{n \rightarrow \infty} \alpha_n = 0$ to suppress the limit cycle of an OS algorithm. Intuitively, as the stepsizes are diminishing, the size of

the limit cycle would decrease to zero, that is, a point. However, even if such relaxation make the OS algorithm converge to some λ^* , we must still ensure that the limit λ^* is a PL solution. The next section describes appropriate choices of $d_j(\cdot)$ and α_n that ensure global convergence.

4.3 Relaxed OS Algorithms

The algorithms described in this section (and the accompanying convergence proofs in the appendices) are applicable to a broad family of objective functions that have the same general properties as the emission tomography case. Specifically, the properties that we exploit are the following:

- Φ is concave (or strictly concave) and differentiable.
- The maximizers of Φ lie in a bounded set defined by $0 \leq x_j \leq U$ where U is a computable upper bound.
- Φ has the summation form (4.8) where each Φ_m is concave and differentiable.

In addition, in the convergence proofs we assume that the gradients of the Φ_m functions are Lipschitz continuous; however, this is automatically ensured by the continuity of Φ_m over the bounded and closed constraint set \mathcal{B} . Collectively these are fairly unrestrictive assumptions so the algorithms should have broad applicability.

To achieve the goal of maximizing Φ over \mathcal{B} , we present two types of relaxed OS algorithms that are globally convergent: modified BSREM methods and diagonally-scaled incremental gradient methods of which relaxed OS-SPS is a special case for emission tomography. For both of these OS algorithms, we use the subobjective functions given in (4.7). The main difference is in the form of $d_j(\cdot)$ in (4.4).

4.3.1 Modified BSREM

De Pierro and Yamagishi [31] presented the BSREM algorithm and proved its global convergence under the following assumptions: the sequence $\{\boldsymbol{\lambda}^n\}$ generated by the algorithm is positive and bounded; and the objective sequence $\{\Phi(\boldsymbol{\lambda}^n)\}$ converges. These conditions are not automatically ensured by the form of the original BSREM. We eliminate those assumptions in our convergence analysis by modifying the $d_j(\cdot)$ functions.

The basic idea of the modification is to ensure that all iterates lie in the interior of the constraint set \mathcal{B} by choosing suitable scaling functions $d_j(\cdot)$ and relaxation parameters α_n . For EM-like algorithms including BSREM, we observe that using the form $d_j(\boldsymbol{\lambda}) = (\text{some term}) \times \lambda_j$ can help each iterate keep positivity and avoid crossing the lower boundary $\lambda_j = 0$. We enforce the upper bound U similarly. Consider the following algorithm called “modified BSREM-I” in vector notation:

$$\begin{aligned}\boldsymbol{\lambda}^{n,0} &= \boldsymbol{\lambda}^n \\ \boldsymbol{\lambda}^{n,m} &= \boldsymbol{\lambda}^{n,m-1} + \alpha_n \mathbf{D}(\boldsymbol{\lambda}^{n,m-1}) \nabla \Phi_m(\boldsymbol{\lambda}^{n,m-1}), \quad \text{for } m = 1, \dots, M \\ \boldsymbol{\lambda}^{n+1} &= \boldsymbol{\lambda}^{n,M}\end{aligned}\tag{4.12}$$

where $\alpha_n > 0$ and $\mathbf{D}(\boldsymbol{\lambda}) = \text{diag}\{d_j(\boldsymbol{\lambda})\}$ with

$$d_j(\boldsymbol{\lambda}) \triangleq \begin{cases} \frac{\lambda_j}{p_j} & \text{for } 0 \leq \lambda_j < \frac{U}{2} \\ \frac{U - \lambda_j}{p_j} & \text{for } \frac{U}{2} \leq \lambda_j \leq U \end{cases}\tag{4.13}$$

for some $p_j > 0$. (The original BSREM used $d_j(\boldsymbol{\lambda}) = \lambda_j$.)

The convergence analysis of this type of algorithm for a strictly concave objective function is given in Appendix C. The first part (Lemma C.1) of the analysis states that if

$$\alpha_n < \alpha, \quad \forall n \quad \text{for some small } \alpha > 0,\tag{4.14}$$

$$\text{and } \boldsymbol{\lambda}^0 \in \text{Int } \mathcal{B},\tag{4.15}$$

then the iterates generated by (4.12) automatically stay in the interior of \mathcal{B} , where $\text{Int } \mathcal{B}$ denotes the interior of \mathcal{B} . The second part (Lemmas C.2–C.5) is about convergence: the iterates generated by (4.12) converge to the solution $\hat{\boldsymbol{\lambda}} = \arg \max_{\boldsymbol{\lambda} \in \mathcal{B}} \Phi(\boldsymbol{\lambda})$ if

$$\sum_{n=0}^{\infty} \alpha_n = \infty, \quad (4.16)$$

$$\sum_{n=0}^{\infty} \alpha_n^2 < \infty, \quad (4.17)$$

$$\text{and } \boldsymbol{\lambda}^{n,m} \in \text{Int } \mathcal{B}, \forall n, m. \quad (4.18)$$

But the first part says that (4.18) is guaranteed if (4.14) and (4.15) hold. So, combining two parts, one can conclude (Theorem C.6 and Corollary C.7) that the modified BSREM-I is globally convergent if (4.14)–(4.17) hold.

A practical and critical issue is how small the relaxation parameter should be in (4.14) for ensuring (4.18). If some iterate hits the boundary, then all subsequent iterates remain stuck at the boundary because the scaling function is zero on the boundary. As shown in Lemma C.1, one may compute a bound ensuring (4.18) and use relaxation parameters smaller than the bound. However, a conservatively small bound will adversely affect convergence rate. So the convergence theorem for BSREM-I still leaves users with practical dilemmas. To overcome these limitations of BSREM-I, we propose to add the following step after (4.12) for each update:

$$\boldsymbol{\lambda}^{n,m} = \begin{cases} \mathcal{P}_{\mathcal{T}}(\boldsymbol{\lambda}^{n,m}) & \text{for } \boldsymbol{\lambda}^{n,m} \notin \text{Int } \mathcal{B} \\ \boldsymbol{\lambda}^{n,m} & \text{otherwise} \end{cases} \quad (4.19)$$

where $\mathcal{P}_{\mathcal{T}}(\boldsymbol{\lambda})$ is the projection⁶ of $\boldsymbol{\lambda} \in \mathbb{R}^p$ onto $\mathcal{T} \triangleq \{\boldsymbol{\lambda} \in \mathbb{R}^p : t \leq \lambda_j \leq U - t, \forall j\}$ for some small $t > 0$. Consider this modified algorithm (4.12) with (4.19), called “mod-

⁶For a Hilbert space \mathcal{H} , a projection $\mathcal{P}_K(\boldsymbol{x})$ of $\boldsymbol{x} \in \mathcal{H}$ onto a nonempty closed convex subset $K \subset \mathcal{H}$ is defined by $\mathcal{P}_K(\boldsymbol{x}) = \arg \min_{\boldsymbol{y} \in K} \|\boldsymbol{x} - \boldsymbol{y}\|$. Here the projection $\mathcal{P}_{\mathcal{T}}(\boldsymbol{\lambda})$ is easily calculated componentwise as $[\mathcal{P}_{\mathcal{T}}(\boldsymbol{\lambda})]_j = \text{median}\{t, \lambda_j, U - t\}$.

ified BSREM-II,” and suppose that conditions (4.16) and (4.17) hold. Then (4.18) is always satisfied by (4.19) regardless of whether (4.14) and/or (4.15) hold. Since (4.16) implies [110, p. 70] that $\lim_{n \rightarrow \infty} \alpha_n = 0$, there exists $K \in \mathbb{N}$ such that α_n satisfies (4.14) for $n \geq K$. Treating $\boldsymbol{\lambda}^K \in \text{Int } \mathcal{B}$ as a “new” starting point, one can see that the iterates after K iterations never hit the boundary by the first part of the analysis mentioned in the previous paragraph. This implies that the step (4.19) becomes vacuous and in subsequent iterations the modified BSREM-II becomes equivalent to the modified BSREM-I. So by the second part of the analysis the modified BSREM-II is globally convergent. The addition of step (4.19) removes the conditions (4.14) and (4.15) while retaining global convergence.

In (4.13), any $p_j > 0$ can be used for global convergence. But we want to choose p_j such that stepsize selection becomes convenient, akin to the appropriateness of a unity stepsize in Newton’s methods due to the scaling by the Hessian’s inverse. Motivated by the EM algorithm for emission tomography, a reasonable choice for p_j is:

$$p_j = \frac{1}{M} \sum_{i=1}^N a_{ij}. \quad (4.20)$$

If $M = 1$ (one subset), $\alpha_n = 1$ (unrelaxed), and $R = 0$ (unregularized), then (4.12) with (4.20) reduces to ML-EM except the term $(U - \lambda_j)/p_j$ in (4.13). Although (4.20) ignores the regularization term, it seems to work well for the regularized case unless the regularization term is too large compared to the log-likelihood part. This is verified experimentally in Section 4.4.

If we take larger and larger U , then $\mathcal{B} \rightarrow \mathbb{R}_+^p$ and $d_j(\boldsymbol{\lambda}) \rightarrow \lambda_j/p_j$. So the modified BSREM should behave quite similarly to the original BSREM for large U in practice except for scaling by p_j . The upper bound U seems to be more important for convergence analysis than for practical implementation.

4.3.2 Diagonally-Scaled Incremental Gradient Method

As an alternative to the BSREM methods, we consider next a family of OS algorithms with *constant* scaling functions $d_j(\cdot) = d_j$ as follows:

$$\begin{aligned}\boldsymbol{\lambda}^{n,0} &= \boldsymbol{\lambda}^n \\ \boldsymbol{\lambda}^{n,m} &= \mathcal{P}_{\mathcal{B}}(\boldsymbol{\lambda}^{n,m-1} + \alpha_n \mathbf{D} \nabla \Phi_m(\boldsymbol{\lambda}^{n,m-1})) \quad \text{for } m = 1, \dots, M \\ \boldsymbol{\lambda}^{n+1} &= \boldsymbol{\lambda}^{n,M}\end{aligned}\tag{4.21}$$

where $\alpha_n > 0$ and $\mathbf{D} = \text{diag}\{d_j\}$ with $d_j > 0, \forall j$, and $\mathcal{P}_{\mathcal{B}}(\boldsymbol{\lambda})$ is the projection⁷ of $\boldsymbol{\lambda} \in \mathbb{R}^p$ onto \mathcal{B} . We call these algorithms “diagonally-scaled incremental gradient methods” since if we choose $\mathbf{D} = \mathbf{I}$, the algorithm (4.21) becomes an incremental gradient method [89]. Appendix D presents the convergence analysis of this type of algorithm for a concave objective function (possibly having multiple solutions). The iterates generated by (4.21) converge to a maximizer if $\sum_{n=0}^{\infty} \alpha_n = \infty$ and $\sum_{n=0}^{\infty} \alpha_n^2 < \infty$ as shown in Theorem D.3 and Corollary D.4. The global convergence holds regardless of \mathbf{D} as long as it is diagonal with positive elements.

A practical issue is how to choose \mathbf{D} for fast convergence rate and easy stepsize selection. Fortunately, some hints are given by observing that the ordered subsets separable paraboloidal surrogates (OS-SPS) method, which showed fairly fast convergence [2], is a special case of (4.21). In particular, (4.21) becomes quadratically-penalized OS-SPS for a likelihood of the form (4.3) if $\alpha_n = 1$ and the scaling constants are chosen as follows:

$$d_j = M \left(\sum_{i=1}^N a_{ij} a_i c_i^{\text{PC}} + 2\beta \sum_{k \in \mathcal{N}_j} \omega_{jk} \right)^{-1}, \quad \forall j\tag{4.22}$$

⁷The projection is readily computed componentwise as $[\mathcal{P}_{\mathcal{B}}(\boldsymbol{\lambda})]_j = \text{median}\{0, \lambda_j, U\}$.

where $a_i \triangleq \sum_{j=1}^p a_{ij}$, M is the number of subsets, and

$$c_i^{\text{PC}} \triangleq \begin{cases} -\ddot{h}_i([y_i - r_i]_+) & \text{for } y_i > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (4.23)$$

The OS-SPS is an OS version of SPS described in Section 3.3.3. Here we use “precomputed curvature (PC)” [1, 2] instead of MC in (3.48) or OC in (3.49) because in this OS algorithm we forgo monotonicity. The precomputed curvatures in (4.23) are approximated Newton’s curvatures:

$$-\ddot{h}_i(l) \approx -\ddot{h} \left(\arg \max_{\tilde{l} \geq 0} h_i(\tilde{l}) \right) = c_i^{\text{PC}}.$$

The above approximation is reasonable since the projections $[\mathbf{A}\boldsymbol{\lambda}^n]_i$ do not change much during iterations in tomography. The precomputed curvatures in (4.23) can be precomputed and stored like MC but are smaller than MC, that is, yield faster convergence rates than MC.

For nonquadratic penalties, the second term in the parenthesis of (4.22) could be substituted with the curvatures of the penalty function at an initial point or at a uniform image. Although OS-SPS is not globally convergent in general, by using relaxation parameters, we obtain a relaxed OS-SPS that is readily shown to be globally convergent as a special member of the family (4.21). Interestingly, whereas the original quadratic surrogate based methods [33] for emission tomography required $r_i > 0$ for monotonicity and convergence, we eliminate this requirement here by the modification of the objective function in Section 3.1.6.

One of required conditions for the global convergence proofs of diagonally-scaled incremental gradient methods is the boundedness of $\nabla\Phi_m$ on \mathcal{B} . If the gradient ∇R of the penalty function is bounded on \mathbb{R}_+^p , then we can take $\mathcal{B} = \mathbb{R}_+^p$ while retaining global convergence since the gradient of the log-likelihood is bounded on \mathbb{R}_+^p . Such penalties include

the Huber penalty [91]:

$$\psi(x) = \begin{cases} x^2/2, & \text{for } |x| \leq \delta \\ \delta|x| - \delta^2/2, & \text{otherwise,} \end{cases}$$

for some $\delta > 0$.

4.3.3 Regularization into OS Algorithms

There are two typical ways of distributing the regularization term into subobjective functions, *i.e.*, how to choose γ_m in (4.7). One way is to include regularization in every Φ_m as in [2]:

$$\gamma_m = \frac{|S_m|}{N}, \quad \forall m \quad (4.24)$$

where $|S_m|$ is the number of elements in S_m . ($\gamma_m = 1/M$ for equally sized subsets.) Another way is to take the regularization term as a separate subobjective function as in [31]:

$$\gamma_m = 0 \text{ for } m = 1, \dots, M, \text{ and } \gamma_{M+1} = 1 \quad (4.25)$$

where we have $(M + 1)$ subobjective functions and take $S_{M+1} = \emptyset$. Both cases satisfy the condition $\Phi = \sum_m \Phi_m$. However, the convergence *rates* of the two choices can differ if the regularization parameter β is not small. Recalling the motivations of OS algorithms, (4.9) and (4.10), one can expect that (4.24) will yield faster initial convergence rates since (4.25) may cause poor “subset gradient balance.” In other words, the amplitude of a limit cycle that is supposed to be suppressed by relaxation is larger for (4.25) due to significant dissimilarities between the subobjective functions. On the other hand, (4.24) requires more computation since the gradient of the regularization part should be computed every subiteration. This additional computational cost is proportional to the number of subsets; however, usually it is relatively small compared to the computation of the log-likelihood part. In experiments not shown, we have observed that the choice (4.24) usually makes

algorithms faster and more stable, so we focus on (4.24) in Section 4.4. Nevertheless, our convergence results apply to any choices for the γ_m 's.

4.3.4 Subiteration-Independent Scaling Matrices Are Essential

Both algorithms in (4.12) and (4.21) belong to the class (4.11) where the functions $d_j(\cdot)$ are independent of subiteration index m . Classical OS-EM [55] does not belong to this class. As pointed out by Browne and De Pierro [18], OS-EM in general does not converge to a solution even if relaxed. We generalize their argument. One could write a more general form of OS algorithms by allowing different scaling matrices over subiterations:

$$\boldsymbol{\lambda}^{n,m} = \boldsymbol{\lambda}^{n,m-1} + \alpha_n \mathbf{D}_m(\boldsymbol{\lambda}^{n,m-1}) \nabla \Phi_m(\boldsymbol{\lambda}^{n,m-1}), \quad \forall m \quad (4.26)$$

where $\alpha_n > 0$, $\forall n$ and $\mathbf{D}_m(\boldsymbol{\lambda})$ is some nonnegative definite diagonal matrix (function). When we choose $\alpha_n = 1$ and $\mathbf{D}_m(\boldsymbol{\lambda}) = \text{diag}\{\lambda_j / \sum_{i \in S_m} a_{ij}\}$, the algorithm (4.26) becomes OS-EM for $R = 0$. Now consider a relaxed version by assuming $\lim_{n \rightarrow \infty} \alpha_n = 0$ and⁸ $\sum_{n=0}^{\infty} \alpha_n = \infty$. Following [18], one can write the following expression for $\boldsymbol{\lambda}^{n+1}$:

$$\begin{aligned} \boldsymbol{\lambda}^{n+1} &= \boldsymbol{\lambda}^n + \alpha_n \sum_{m=1}^M \mathbf{D}_m(\boldsymbol{\lambda}^{n,m-1}) \nabla \Phi_m(\boldsymbol{\lambda}^{n,m-1}) \\ &= \boldsymbol{\lambda}^0 + \sum_{k=0}^n \alpha_k \sum_{m=1}^M \mathbf{D}_m(\boldsymbol{\lambda}^{k,m-1}) \nabla \Phi_m(\boldsymbol{\lambda}^{k,m-1}). \end{aligned}$$

Now suppose that the sequence $\{\boldsymbol{\lambda}^{n,m}\}$ generated by (4.26) converges to some $\boldsymbol{\lambda}^*$. Assuming that $\mathbf{D}_m \nabla \Phi_m$ is continuous, we have:

$$\lim_{k \rightarrow \infty} \mathbf{D}_m(\boldsymbol{\lambda}^{k,m-1}) \nabla \Phi_m(\boldsymbol{\lambda}^{k,m-1}) = \mathbf{D}_m(\boldsymbol{\lambda}^*) \nabla \Phi_m(\boldsymbol{\lambda}^*).$$

⁸If we take a diminishing stepsize ($\lim_{n \rightarrow \infty} \alpha_n = 0$), we need the assumption: $\sum_{n=0}^{\infty} \alpha_n = \infty$. Suppose that $\sum_{n=0}^{\infty} \alpha_n < \infty$. Since $\|\boldsymbol{\lambda}^{n+1} - \boldsymbol{\lambda}^n\| = O(\alpha_n)$ (by assuming that $\mathbf{D}_m \nabla \Phi_m$ is bounded), we will never get to the optimum point if an initial point is sufficiently far from it.

If $\sum_{m=1}^M \mathbf{D}_m(\boldsymbol{\lambda}^*) \nabla \Phi_m(\boldsymbol{\lambda}^*) \neq \mathbf{0}$, then $\{\boldsymbol{\lambda}^n\}$ diverges since $\sum_{n=0}^{\infty} \alpha_n = \infty$. So it must be the case that:

$$\sum_{m=1}^M \mathbf{D}_m(\boldsymbol{\lambda}^*) \nabla \Phi_m(\boldsymbol{\lambda}^*) = \mathbf{0}. \quad (4.27)$$

However, if the \mathbf{D}_m 's are different, then (4.27) is generally different from the true optimality conditions, *e.g.*, $\nabla \Phi(\boldsymbol{\lambda}^*) = \sum_{m=1}^M \nabla \Phi_m(\boldsymbol{\lambda}^*) = \mathbf{0}$ for unconstrained optimization. So, in general, OS algorithms with subiteration-dependent scaling matrices, including OS-EM and RBI-EM [19], do not converge *to the desired optimum point* even if they become convergent due to relaxation.

Table 4.1: Outline for the algorithms presented in this chapter

Compute a bound U on a solution by (A.1) in Appendix A. Compute ϵ by (B.1) in Appendix B if $\mathcal{I} \neq \emptyset$, that is, $r_i = 0$ but $y_i > 0$ for some i . Precompute $p_j = \sum_{i=1}^N a_{ij}/M$ for modified BSREM, or precompute d_j^P for relaxed OS-SPS. Use (4.22) for quadratic penalty. for each iteration $n = 1, \dots, \text{niter}$ for each subset $m = 1, \dots, M$ $\hat{l}_i = \sum_{j=1}^p a_{ij} \hat{\lambda}_j + r_i$ for $i \in S_m$ $\dot{h}_i = \begin{cases} \dot{h}_i(\epsilon) + \ddot{h}_i(\epsilon)(\hat{l}_i - \epsilon) & \text{for } i \in \mathcal{I} \text{ and } \hat{l}_i \leq \epsilon, \text{ where } h_i(l) = y_i \log l - l \\ (y_i/\hat{l}_i) - 1 & \text{otherwise} \end{cases}$ $\boldsymbol{\lambda}^{\text{old}} = \hat{\boldsymbol{\lambda}}$ for $j = 1, \dots, p$ $\dot{\Phi}_j = \sum_{i \in S_m} a_{ij} \dot{h}_i - \beta \sum_{k \in \mathcal{N}_j} \omega_{jk} \dot{\psi}(\lambda_j^{\text{old}} - \lambda_k^{\text{old}})/M$ Update $\hat{\lambda}_j$. (See Table 4.2.) end end end

4.4 Results

The outline of modified BSREM and relaxed OS-SPS algorithms for a Poisson penalized-likelihood in emission tomography are summarized in Tables 4.1 and 4.2. In addition to those conditions in Table 4.2, for a general objective function, modified BSREM requires that Φ is strictly concave, and $\nabla \Phi_m(\boldsymbol{\lambda})$ and $\mathbf{D}(\boldsymbol{\lambda}) \nabla \Phi_m(\boldsymbol{\lambda})$ are Lipschitz continuous on

Table 4.2: Comparison of algorithms

Algorithm	Update in Table 4.1	Sufficient conditions for convergence
Modified BSREM-I	$d_j^B = \begin{cases} \hat{\lambda}_j/p_j & \text{for } \hat{\lambda}_j < U/2 \\ (U - \hat{\lambda}_j)/p_j & \text{for } \hat{\lambda}_j \geq U/2 \end{cases}$ $\hat{\lambda}_j := \hat{\lambda}_j + \alpha_n d_j^B \Phi_j$	(i) $\sum_n \alpha_n = \infty$ (ii) $\sum_n \alpha_n^2 < \infty$ (iii) α_n is sufficiently small (vi) $\hat{\lambda}^{\text{initial}} \in \text{Int } \mathcal{B}$ or, instead of (iii) and (iv), (v) All iterates lie in the interior of \mathcal{B}
Modified BSREM-II	$\hat{\lambda}_j := \hat{\lambda}_j + \alpha_n d_j^B \Phi_j$ same as above $\hat{\lambda}_j := \begin{cases} t & \text{if } \hat{\lambda}_j \leq 0 \\ U - t & \text{if } \hat{\lambda}_j \geq U \end{cases}$	$\sum_n \alpha_n = \infty$ and $\sum_n \alpha_n^2 < \infty$
Relaxed OS-SPS	$\hat{\lambda}_j := \hat{\lambda}_j + \alpha_n d_j^B \Phi_j$ $\hat{\lambda}_j := \begin{cases} 0 & \text{if } \hat{\lambda}_j \leq 0 \\ U & \text{if } \hat{\lambda}_j \geq U \end{cases}$	$\sum_n \alpha_n = \infty$ and $\sum_n \alpha_n^2 < \infty$

t is a small value, say, $0.001 \max_j \hat{\lambda}_j^{\text{FBP}}$.

\mathcal{B} . Diagonally-scaled incremental gradient methods including relaxed OS-SPS require that $\nabla \Phi_m$ are bounded on \mathcal{B} and Φ_m are concave. We focus on modified BSREM-II rather than modified BSREM-I in this section. A critical issue in practice will be how to determine relaxation parameters to get close to a solution within a few iterations. The sufficient conditions on a relaxation sequence for global convergence are the following: $\sum_{n=0}^{\infty} \alpha_n = \infty$ and $\sum_{n=0}^{\infty} \alpha_n^2 < \infty$. One may try to optimize a finite number of relaxation parameters by training [18,31,47] if a reasonable training set is given for a particular task. The relaxation parameters obtained thus might not seem to satisfy those conditions. However, it may not be relevant since those conditions are only *sufficient* and, moreover, *asymptotic*.

One simple choice of relaxation parameters satisfying those conditions is:

$$\alpha_n = \frac{\alpha_0}{\gamma n + 1}, \quad \forall n \quad (4.28)$$

for $\gamma > 0$ and $\alpha_0 > 0$. We run simulations using these simple relaxation parameters. Our goal here is not to try to find the best relaxation but to get some insight into the effects of relaxation parameters on convergence rate through some experiments. By design, our modified BSREM and relaxed OS-SPS are properly scaled, meaning that even a constant

$\alpha_n = 1$ works fairly well. So we could obtain reasonably good results by setting $\alpha_0 = 1$ and tuning experimentally only γ .

We performed image reconstruction using two-dimensional SPECT simulation data generated with the Shepp-Logan digital phantom. The projection space was 128 radial bins with 3.6 mm ray spacing and 120 angles over 360 degrees, and the reconstructed images were 128×128 with 3.6 mm pixel size. The distance from the center of rotation to the detector plane was 288 mm. The system matrix \mathbf{A} was generated by ASPIRE 3.0 [35] and it assumed a Gaussian shaped point spread function with the following model for the depth-dependent full-width at half-maximum (FWHM):

$$\text{FWHM} = \sqrt{(0.0868056 \cdot z)^2 + (3\text{mm})^2}$$

where z is the distance from a pixel's center to the detector. We did not consider attenuation in this simulation. The total counts were 5×10^5 , and r_i corresponded to a uniform field of 10% of background events, a very crude approximation of the effects of scatter. We regularized the log-likelihood using the first-order quadratic penalty $\psi(x) = x^2/2$ with $\beta = 1.5$, and we took a FBP reconstruction as a starting image for PL reconstruction. Because the relaxed OS algorithms are additive updates, the scaling of the initial image can affect the initial convergence rate, so we implemented the FBP algorithm carefully with respect to the global scale factor. In contrast, the classical ML-EM and OS-EM methods for emission tomography are multiplicative, so the initial scaling is unimportant.

Fig. 4.2 compares two nonOS algorithms: SPS with optimum curvature (see Section 3.3.3) and DPEM (see Section 3.3.2); and two unrelaxed OS algorithms: unrelaxed OS-SPS and unrelaxed modified BSREM with $\alpha_n = 1$ and with 8 subsets and 40 subsets. The OS algorithms initially increase the objective function much faster than the non-OS ones, but they get stuck at suboptimal points. The figure shows the normalized Φ difference $(\Phi(\hat{\boldsymbol{\lambda}}) - \Phi(\boldsymbol{\lambda}^n))/(\Phi(\hat{\boldsymbol{\lambda}}) - \Phi(\boldsymbol{\lambda}^0))$ versus iteration number where $\hat{\boldsymbol{\lambda}}$ is the solution

estimated by 5,000 iterations of DPEM, a globally convergent method [30]. One can see that the scaling factors (4.13) with (4.20), and (4.22) for the OS algorithms are reasonable since the stepsize of unity worked fairly well. For both unrelaxed OS-SPS and unrelaxed modified BSREM, using more subsets accelerated “convergence” but made the algorithms reach a limit cycle earlier. Roughly speaking, in early iterations more subsets are desirable but in later iterations fewer subsets would be preferable in the unrelaxed case.

Now we see how relaxation improves convergence. Fig. 4.3 compares unrelaxed modified BSREM and relaxed modified BSREM. As can be seen in the figure, the unrelaxed modified BSREM algorithms converged to a limit cycle whereas the relaxed ones showed better performance in increasing the objective function by suppressing the amplitude of the cycle (note the logarithmic scale). We chose $\alpha_n = 1/(\frac{1}{15}n + 1)$ for relaxed modified BSREM-8, and $\alpha_n = 1/(n + 1)$ for relaxed modified BSREM-40. In this experiment the second part of the scaling function in (4.13) was never invoked due to the very large bound U used; the scaling matrix we used was effectively the same as that of original BSREM except for p_j . Fig. 4.4 shows results for relaxed OS-SPS that are similar to those for modified BSREM. We chose $\alpha_n = 1/(\frac{1}{5}n + 1)$ for relaxed OS-SPS-8, and $\alpha_n = 1/(n + 1)$ for relaxed OS-SPS-40. Fig. 4.5 summarizes Fig. 4.3 and Fig. 4.4. The reconstructed images are shown in Fig. 4.6.

We observed, from experiments with relaxation parameters, that applying relaxation (less than unity) before an algorithm reaches a limit cycle far from the optimum point does not improve convergence rate because it slows down the algorithm’s progress toward the optimum point. Apparently relaxation is most helpful when an algorithm is nearing a limit cycle. Generally speaking, rapidly diminishing stepsizes are preferable for an algorithm using many subsets since such algorithms tend to reach a limit cycle quickly. But relaxation should be applied gradually in cases where it takes many iterations for an algorithm

to reach a limit cycle, *e.g.*, unregularized ML reconstruction or when few subsets are used.

4.5 Conclusion

We presented two types of convergent relaxed OS algorithms: modified BSREM and relaxed OS-SPS which differ in their scaling functions $d_j(\cdot)$. We proved global convergence of both algorithms without “a posteriori” assumptions. A natural subsequent question is about convergence rate. This is related to how to determine the relaxation parameters. For relaxation parameters, we showed through experiments that relaxation improves the OS algorithms convergence rates when the algorithms are approaching a limit cycle. Hopefully, future work on quantitative convergence rate analysis will provide more useful rules for determining relaxation parameters, perhaps adaptively.

The practical question of whether it is preferable to achieve convergence by using relaxation or by reducing the number of subsets with iteration remains open, and may simply be a matter of preference. The approach of progressively decreasing the number of subsets will guarantee global convergence as long as it eventually takes a convergent nonOS (one subset version) algorithm. However, for fast initial convergence rates, one would need to determine a (nearly) optimal schedule of reducing the number of subsets, which is as inconvenient as tuning the relaxation parameters for relaxed OS algorithms. As a rough rule of thumb, one should reduce the current number of subsets when the algorithms are nearing a limit cycle. However, the absence of automatic rules to determine relaxation parameters and a schedule of reducing the number of subsets makes it difficult to compare the two approaches. When iterative algorithms become implemented in special purpose hardware, the consistent data flow provided by the relaxation approach may be beneficial.

We have not tried to evaluate the relative merits of modified BSREM and relaxed OS-SPS. Both algorithms are globally convergent, and simulation results showed that appro-

priate relaxation accelerates convergence similarly for both of them. Finding better scaling functions in terms of convergence speed and computational efficiency could also be interesting future work.

The algorithms presented in this chapter are easily adapted to transmission tomography for zero backgrounds ($r_i = 0$). However, for a nonzero background case, the penalized-likelihood objective function can become nonconcave [1]. It will also be interesting future work to investigate whether the relaxed OS algorithms can be proved to converge to local maxima in nonconcave cases.

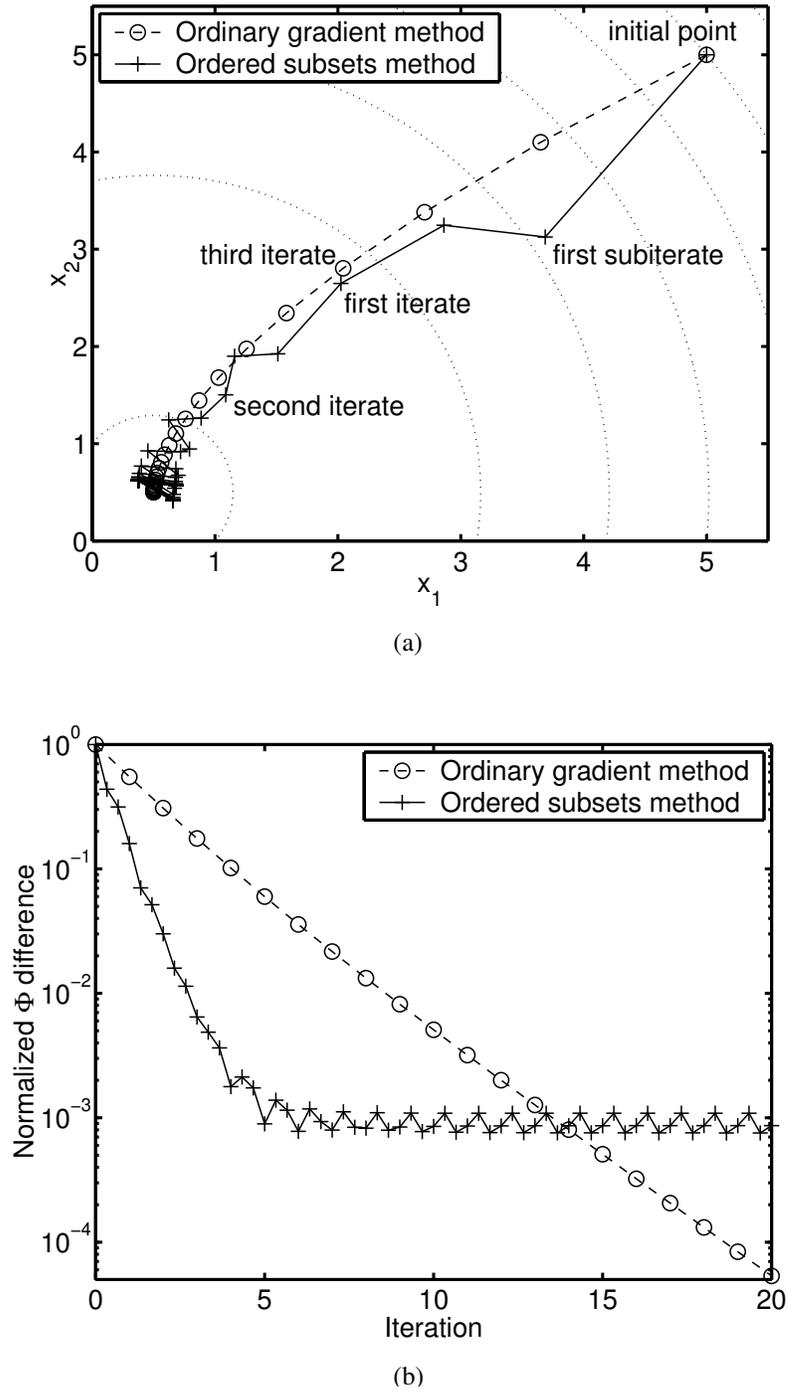


Figure 4.1: Toy example of OS algorithms. (a) Trajectory of iterates of a nonOS gradient method with a constant stepsize and its OS version with 3 subsets. The optimal point is $\hat{x} = (0.5, 0.5)$ and the initial point is $x^0 = (5, 5)$. (b) $(\Phi(\hat{x}) - \Phi(x^{n,m})) / (\Phi(\hat{x}) - \Phi(x^0))$ versus iteration number. For the OS method, each subiterate is denoted.

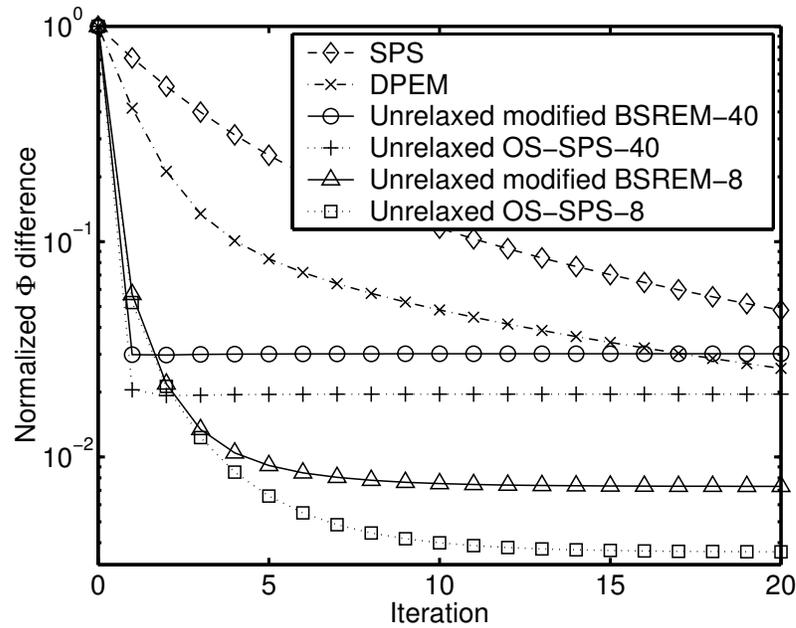


Figure 4.2: Comparison of normalized Φ difference $(\Phi(\hat{\lambda}) - \Phi(\lambda^n)) / (\Phi(\hat{\lambda}) - \Phi(\lambda^0))$ versus iteration number for non-OS algorithms including SPS and DPEM; and unrelaxed OS algorithms including unrelaxed OS-SPS and unrelaxed modified BSREM (with 8 and 40 subsets).

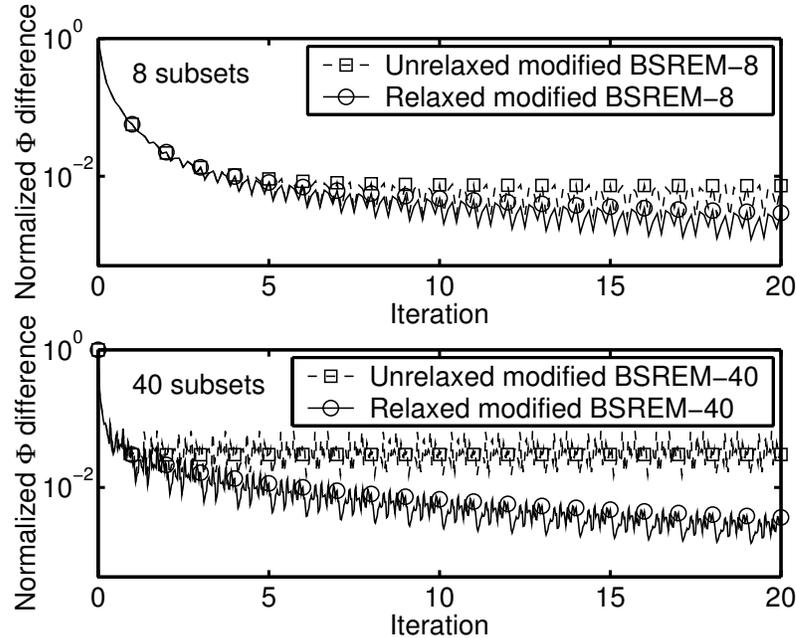


Figure 4.3: Comparison of normalized Φ difference $(\Phi(\hat{\lambda}) - \Phi(\lambda^{n,m})) / (\Phi(\hat{\lambda}) - \Phi(\lambda^0))$ versus iteration number for unrelaxed modified BSREM and relaxed modified BSREM with 8 and 40 subsets. For relaxed modified BSREM-8 (top) and relaxed modified BSREM-40 (bottom), $\alpha_n = 1/(\frac{1}{15}n + 1)$ and $\alpha_n = 1/(n + 1)$ are used, respectively. This figure shows every subiterate.

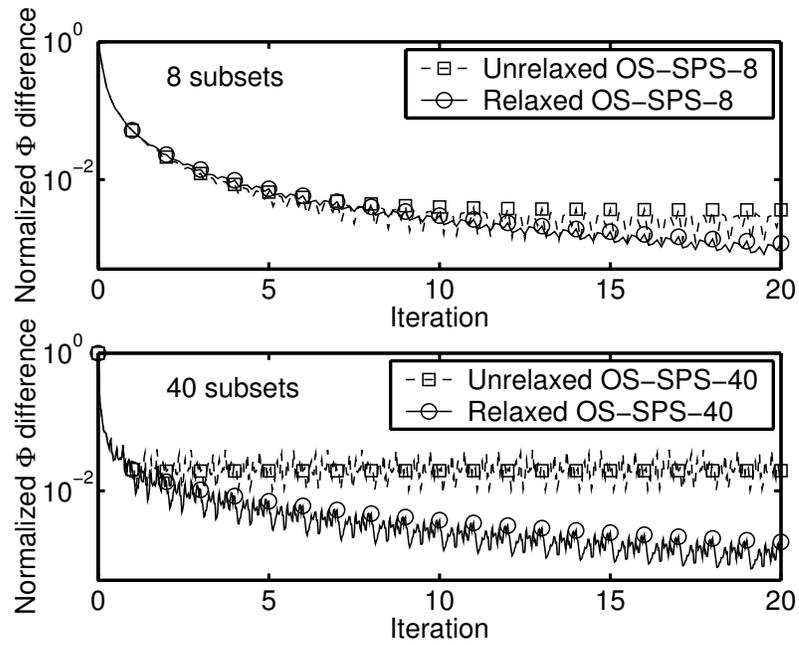
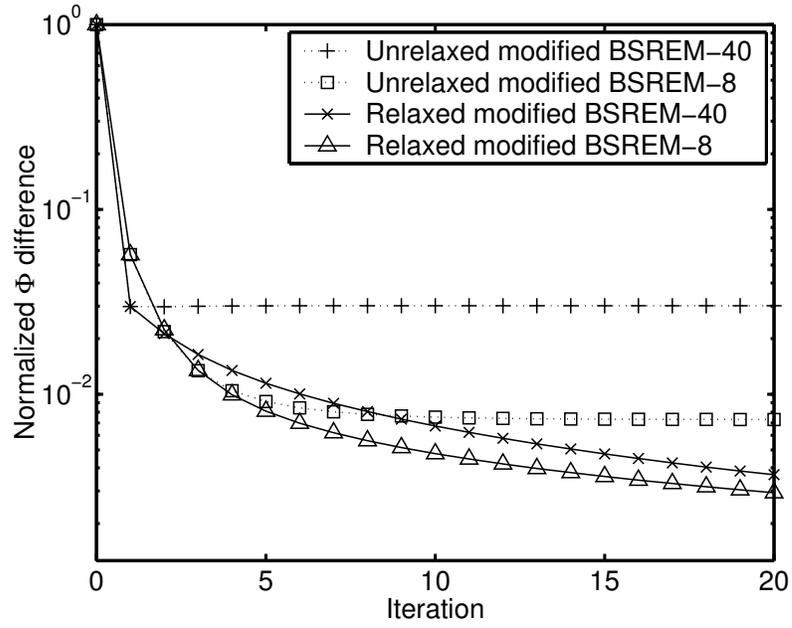
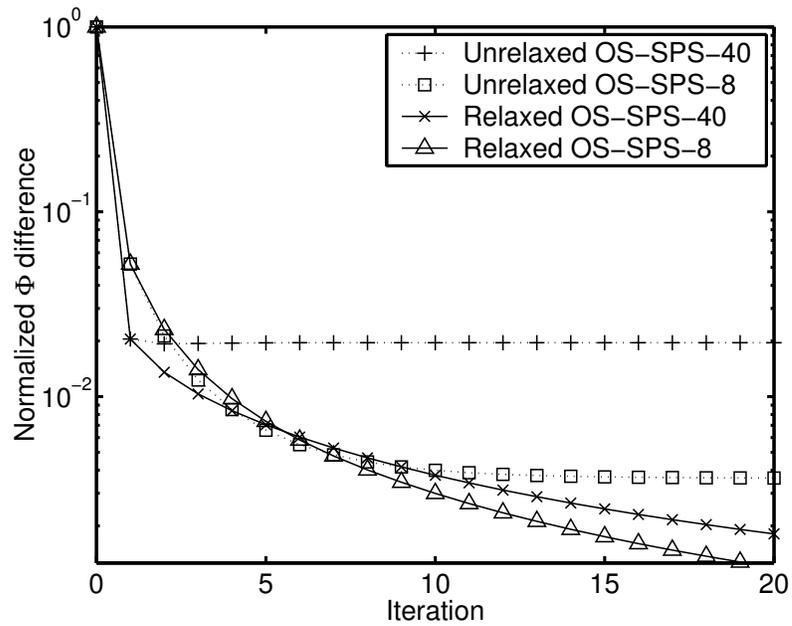


Figure 4.4: Comparison of normalized Φ difference $(\Phi(\hat{\lambda}) - \Phi(\lambda^{n,m})) / (\Phi(\hat{\lambda}) - \Phi(\lambda^0))$ versus iteration number for unrelaxed OS-SPS and relaxed OS-SPS with 8 and 40 subsets. For relaxed OS-SPS-8 (top) and relaxed OS-SPS-40 (bottom), $\alpha_n = 1/(\frac{1}{5}n + 1)$ and $\alpha_n = 1/(n + 1)$ are used, respectively. This figure shows every subiterate.



(a)



(b)

Figure 4.5: Comparison of normalized Φ difference $(\Phi(\hat{\lambda}) - \Phi(\lambda^n)) / (\Phi(\hat{\lambda}) - \Phi(\lambda^0))$ versus iteration number for unrelaxed OS algorithms and relaxed ones. (a) Unrelaxed modified BSREM and relaxed modified BSREM. This figure is the same as Fig. 4.3 except that it shows only each iterate. (b) Unrelaxed OS-SPS and relaxed OS-SPS. This figure is the same as Fig. 4.4 except that it shows only each iterate.

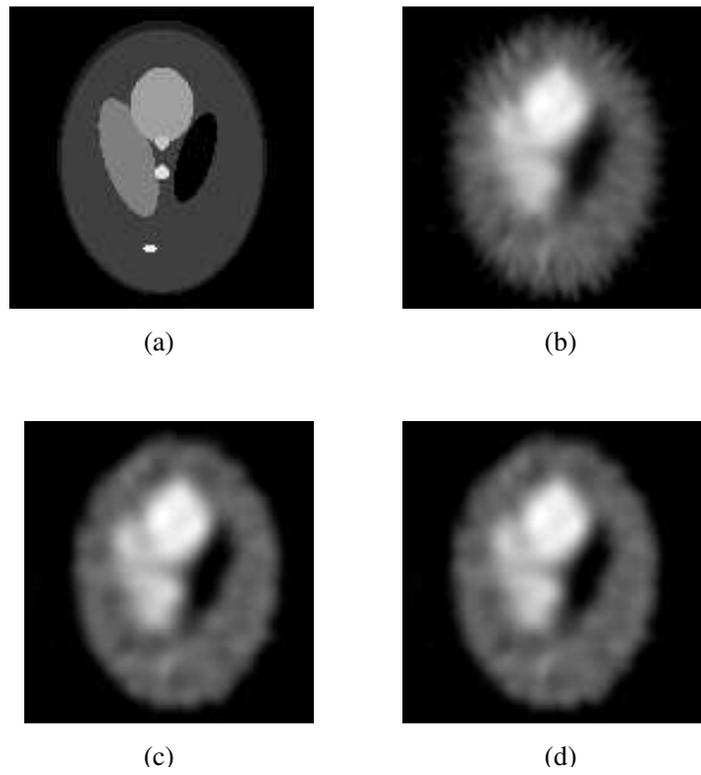


Figure 4.6: (a) Shepp-Logan digital phantom (true image). (b) FBP reconstruction (starting image). (c) PL reconstruction using 20 iterations of relaxed modified BSREM with 8 subsets. (d) PL reconstruction using 20 iterations of relaxed OS-SPS with 8 subsets.

CHAPTER 5

Convergent Incremental Optimization Transfer Algorithms

5.1 Introduction

In the previous chapter, we studied convergent relaxed OS algorithms. In this chapter, we develop and investigate another family of convergent incremental (or OS type) algorithms: incremental optimization transfer algorithms.

Relaxation parameters are used widely to render OS algorithms convergent, and suitably relaxed algorithms can be shown to converge to an optimal solution under mild regularity conditions [4, 18, 23, 24, 31, 61, 65, 66, 76, 88, 89, 119]. However, since relaxation parameters should be scheduled to converge to zero for global convergence, the asymptotic convergence rates of relaxed OS algorithms are very slow. Also, inappropriately chosen (*e.g.*, too rapidly decreasing) relaxation parameters could make initial convergence rates even worse than those of non-OS algorithms. On the other hand, overly large relaxation parameters can lead to unstable or divergent behavior. Given a system and an object (or data), to determine optimal (or at least suboptimal) relaxation parameters in terms of convergence rates, one may need some experimentation and trial-and-error; as a rule of thumb, for properly scaled OS algorithms such as modified BSREM and relaxed OS-SPS, one should initialize the relaxation parameter near unity and decrease it gradually as convergence to a limit cycle nears [4]. One may optimize a few initial relaxation parameters

by training when a training set is available for a particular task [18, 47]. Or one could use the dynamic stepsize rule in [88, 89], but that method needs to compute the objective value at every update, which is computationally expensive in tomographic reconstruction problems. So the use of relaxation parameters can be inconvenient and we would like to find alternative methods.

In contrast, incremental EM algorithms [87] do not require user-specified relaxation parameters. They are convergent [44] yet faster than ordinary EM algorithms although slower than nonconvergent OS-EM type algorithms [50, 52]. Such incremental EM algorithms have been applied to emission tomography [45, 50, 52, 60].

Recently, Blatt *et al.* [16] proposed a convergent incremental gradient method that does not require a relaxation parameter, called incremental aggregated gradient (IAG). The IAG method computes a single subset gradient per iteration but aggregates it with the stale subsets gradients that were computed in previous iterations. The use of the aggregated gradient to approximate the full gradient of the objective function leads to the convergence property. Similarly, as discussed below, the use of the sum of surrogate functions (rather than a single surrogate function) to approximate a minorizing function yields convergent algorithms.

In this chapter we generalize the incremental EM algorithms by introducing an approach we call “incremental optimization transfer”; this is akin to the generalization of the EM algorithms [32] by the optimization transfer principles [70]. In fact, the broad family of “incremental optimization transfer algorithms” includes the ordinary optimization transfer algorithms (*e.g.*, EM), also referred to as MM (minorize-maximize or majorize-minimize) algorithms in [57], as a special case where the objective function consists of only one subobjective function.

In the incremental optimization transfer approach, for *each* subobjective function, we

define an augmented vector which has the same size as the parameter vector to be estimated. The augmented vector plays a role as an expansion point with which a minorizing surrogate function is defined for the subobjective function (see Section 5.2 for details). Then the sum of the surrogate functions defines an augmented objective that is a function of the parameter vector and the augmented vectors. With surrogate functions satisfying usual minorization conditions [1, 70], a solution to the problem of maximizing the original objective can be found by maximizing the augmented objective instead. Applying a block coordinate ascent approach to the augmented problem leads to a new class of “incremental optimization transfer algorithms.” By using the block coordinate ascent approach, incremental optimization transfer algorithms are monotonic in the augmented objective value though not in the original objective; nevertheless, global convergence is ensured under mild regularity conditions. Incremental optimization transfer algorithms show faster convergence rates than their nonincremental counterparts like EM [50, 52, 87].

The difference between incremental gradient (or ordinary OS) methods and incremental optimization transfer methods is as follows. Incremental gradient methods use only a partial gradient at every (sub)iteration whereas incremental optimization transfer methods use the whole gradient which is incrementally updated.

Incremental optimization transfer is a general framework in which one can develop many different optimization algorithms by using a broad family of application-dependent surrogate functions. These methods are particularly useful for large-scale problems where the objective function is expressed as a sum of several subobjective functions. However, in this chapter, we focus on PL image reconstruction for transmission tomography, which is a challenging nonconcave maximization problem. We propose a particular incremental optimization transfer algorithm that uses separable paraboloidal surrogates (SPS) [2]. Such quadratic surrogates simplify the maximization. In contrast, the EM surrogates for

transmission tomography do not have a closed-form maximizer in the ‘‘M-step’’ [68].

The proposed ‘‘transmission incremental optimization transfer (TRIOT)’’ algorithm is convergent yet converges faster than ordinary SPS [2]; it can be further accelerated by the method in [53] or by initializing through a few iterations of OS-SPS (see Section 5.3 for details). It is parallelizable, and the nonnegativity constraint is naturally enforced. In addition, it is easily implemented for many system models whereas pixel-grouped coordinate ascent based methods require column access of the system matrix [39, 40, 58, 111].

Section 5.2 describes the incremental optimization transfer algorithms in a general framework and discusses their convergence properties. Section 5.3 develops incremental optimization transfer algorithms for transmission tomography, and addresses some acceleration methods. Section 5.4 provides simulation and real PET data results.

5.2 Incremental Optimization Transfer

5.2.1 Incremental Optimization Transfer Algorithms

Most objective functions of interest in image reconstruction can be expressed as a sum of subobjective functions:¹

$$\Phi(\mathbf{x}) = \sum_{m=1}^M \Phi_m(\mathbf{x}), \quad (5.1)$$

where $\Phi_m : \mathcal{D} \subset \mathbb{R}^p \rightarrow \mathbb{R}$ is a continuously differentiable function whose domain \mathcal{D} is a nonempty, convex and closed set. We consider the following optimization problem:

$$\text{maximize } \Phi(\mathbf{x}) \text{ subject to } \mathbf{x} \in \mathcal{D}. \quad (5.2)$$

Since usually there exists no closed-form solution to the above problem, one must apply iterative algorithms. Assume that for each subobjective function Φ_m , we find a surrogate

¹Such functions are said to be *additive-separable* in [61]; and to be *partially separable* when each $\Phi_m(\mathbf{x})$ is a function of fewer components of $\mathbf{x} \in \mathbb{R}^p$ than p in [92].

function $\phi_m : \mathcal{D}^2 \subset \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}$ that is easier to maximize than Φ_m and that satisfies the following usual minorization conditions as in Section 3.3.1:

$$\begin{aligned}\phi_m(\mathbf{x}; \mathbf{x}) &= \Phi_m(\mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{D} \\ \phi_m(\mathbf{x}; \bar{\mathbf{x}}) &\leq \Phi_m(\mathbf{x}), \quad \forall \mathbf{x}, \bar{\mathbf{x}} \in \mathcal{D},\end{aligned}\tag{5.3}$$

where \mathcal{D}^n denotes the n -ary Cartesian product over the set \mathcal{D} for $n \in \mathbb{N}$ throughout this chapter. It follows from the above conditions that

$$\Phi_m(\mathbf{x}) - \Phi_m(\bar{\mathbf{x}}) \geq \phi_m(\mathbf{x}; \bar{\mathbf{x}}) - \phi_m(\bar{\mathbf{x}}; \bar{\mathbf{x}}), \quad \forall \mathbf{x}, \bar{\mathbf{x}} \in \mathcal{D}.$$

In other words, choosing \mathbf{x} such that $\phi_m(\mathbf{x}; \bar{\mathbf{x}}) \geq \phi_m(\bar{\mathbf{x}}; \bar{\mathbf{x}})$ ensures that $\Phi_m(\mathbf{x}) \geq \Phi_m(\bar{\mathbf{x}})$.

Define the following ‘‘divergence’’ functions:

$$D_m(\mathbf{x} \parallel \bar{\mathbf{x}}) \triangleq \Phi_m(\mathbf{x}) - \phi_m(\mathbf{x}; \bar{\mathbf{x}}).$$

Then by (5.3), we have the following properties:²

$$D_m(\mathbf{x} \parallel \bar{\mathbf{x}}) \geq 0 \text{ and } D_m(\mathbf{x} \parallel \mathbf{x}) = 0.\tag{5.4}$$

Now we define the following ‘‘augmented’’ objective function:

$$F(\mathbf{x}; \bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M) = \Phi(\mathbf{x}) - \sum_{m=1}^M D_m(\mathbf{x} \parallel \bar{\mathbf{x}}_m)\tag{5.5}$$

$$= \sum_{m=1}^M \phi_m(\mathbf{x}; \bar{\mathbf{x}}_m).\tag{5.6}$$

Since

$$\min_{(\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M) \in \mathcal{D}^M} \sum_{m=1}^M D_m(\mathbf{x} \parallel \bar{\mathbf{x}}_m) = 0, \quad \forall \mathbf{x} \in \mathcal{D},$$

²When there exists $\check{\mathbf{x}} \neq \bar{\mathbf{x}}$ such that $D_m(\check{\mathbf{x}} \parallel \bar{\mathbf{x}}) = 0$, using a modified surrogate $\phi_m^{\text{new}}(\mathbf{x}; \bar{\mathbf{x}}) = \phi_m(\mathbf{x}; \bar{\mathbf{x}}) - \epsilon \|\mathbf{x} - \bar{\mathbf{x}}\|^2$ for any fixed $\epsilon > 0$ would lead to the following property: $D_m^{\text{new}}(\mathbf{x} \parallel \bar{\mathbf{x}}) \geq 0$ where equality holds *if and only if* $\mathbf{x} = \bar{\mathbf{x}}$. Although this modification might provide a more natural definition of divergence, it is not needed for our convergence proofs so we allow the less restrictive conditions in (5.3) and (5.4).

that is,

$$\max_{(\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M) \in \mathcal{D}^M} F(\mathbf{x}; \bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M) = \Phi(\mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{D},$$

one can rewrite the optimization problem (5.2) equivalently as follows:

$$\begin{aligned} & \text{maximize } F(\mathbf{x}; \bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M) \\ & \text{subject to } (\mathbf{x}; \bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M) \in \mathcal{D}^{M+1}, \end{aligned} \tag{5.7}$$

in a sense that $\mathbf{x}^* \in \mathcal{D}$ is an optimal solution of (5.2) if and only if $(\mathbf{x}^*; \bar{\mathbf{x}}_1^*, \dots, \bar{\mathbf{x}}_M^*) \in \mathcal{D}^{M+1}$ is an optimal solution of (5.7) for some $(\bar{\mathbf{x}}_1^*, \dots, \bar{\mathbf{x}}_M^*) \in \mathcal{D}^M$. Therefore we can find a solution to problem (5.2) by maximizing F with respect to $(\mathbf{x}; \bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M)$.

By alternating between updating \mathbf{x} and one of the $\bar{\mathbf{x}}_m$'s, we obtain an ‘‘incremental optimization transfer algorithm’’ outlined in Table 5.1, where we assume that there exists one or possibly more maximizers in (5.8), and ‘‘arg max’’ denotes one of those maximizers.

Table 5.1: Outline for incremental optimization transfer algorithms. The right side of (5.9) is due to (5.4) and (5.5).

Initialize $\mathbf{x}^0, \bar{\mathbf{x}}_1^0, \dots, \bar{\mathbf{x}}_M^0 \in \mathcal{D}$	
for $n = 0, \dots, n_{\text{iter}}$	
for $m = 1, \dots, M$	
$\mathbf{x}^{\text{new}} = \arg \max_{\mathbf{x} \in \mathcal{D}} F(\mathbf{x}; \bar{\mathbf{x}}_1^{n+1}, \dots, \bar{\mathbf{x}}_{m-1}^{n+1}, \bar{\mathbf{x}}_m^n, \bar{\mathbf{x}}_{m+1}^n, \dots, \bar{\mathbf{x}}_M^n)$	(5.8)
$\bar{\mathbf{x}}_m^{n+1} = \mathbf{x}^{\text{new}} = \arg \max_{\bar{\mathbf{x}}_m \in \mathcal{D}} F(\mathbf{x}^{\text{new}}; \bar{\mathbf{x}}_1^{n+1}, \dots, \bar{\mathbf{x}}_{m-1}^{n+1}, \bar{\mathbf{x}}_m, \bar{\mathbf{x}}_{m+1}^n, \dots, \bar{\mathbf{x}}_M^n)$	(5.9)
end	
$\mathbf{x}^{n+1} = \bar{\mathbf{x}}_M^{n+1}$	(5.10)
end	

The incremental optimization transfer algorithm shown in Table 5.1 can be viewed as a block coordinate ascent algorithm for maximizing F with respect to $(\mathbf{x}; \bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M)$ [15, p. 270]. It monotonically increases the augmented objective function F , but not necessarily the original objective function Φ [20]. If one has only one subobjective function

in (5.1), that is, $M = 1$, then the incremental optimization transfer algorithm reduces to an ordinary optimization transfer algorithm as in Section 3.3.1. The incremental approach ($M > 1$) usually leads to faster convergence rates than nonincremental methods ($M = 1$) [87]. The incremental EM algorithms [45, 87] including COSEM [50, 52] are a special case where the surrogates ϕ_m are constructed by EM principles (see Section 3.3.2).

If one were to maximize one of ϕ_m 's instead of the sum of ϕ_m 's in (5.6), or in other words, if one were to maximize $\phi_m(\cdot; \bar{\mathbf{x}}_m^n)$ instead of $F(\cdot; \bar{\mathbf{x}}_1^{n+1}, \dots, \bar{\mathbf{x}}_{m-1}^{n+1}, \bar{\mathbf{x}}_m^n, \bar{\mathbf{x}}_{m+1}^n, \dots, \bar{\mathbf{x}}_M^n)$ in (5.8), then one would have ordinary OS type algorithms as described in Section 4.2. Although this greedy approach usually yields faster initial convergence rates than incremental optimization transfer algorithms, the OS type algorithms are not monotonic in F nor in Φ . In the previous chapter, we achieved global convergence by introducing relaxation into the nonconvergent ordinary OS type algorithms. In contrast, incremental optimization transfer methods presented in this chapter update the surrogate function incrementally and keep monotonicity (in an augmented function), consequently achieving convergence.

For incremental optimization transfer algorithms one must store M vectors $\{\bar{\mathbf{x}}_m\}_{m=1}^M$, so one needs more memory compared to ordinary OS algorithms; however, this is not a severe limitation unless M is overly large.

5.2.2 Special Case: Incremental EM Algorithms

We briefly review the incremental EM algorithms within a framework of the incremental optimization transfer principles given in the preceding subsection.

For maximum likelihood (ML) estimation, one must maximize a log-likelihood function

$$\Phi(\mathbf{x}) = \log f(\mathbf{y}; \mathbf{x})$$

with respect to parameter $\mathbf{x} \in \mathbb{R}^p$ over a feasible set $\mathcal{D} \subset \mathbb{R}^p$ where $\mathbf{y} \in \mathbb{R}^N$ denotes a realization of an observable random vector \mathbf{Y} with probability distribution $f(\mathbf{y}; \mathbf{x}^{\text{true}})$, and $\mathbf{x}^{\text{true}} \in \mathbb{R}^p$ is the true value of the unknown parameter. But in many applications including imaging problems with independent data, the log-likelihood objective is additive-separable, that is,

$$\Phi(\mathbf{x}) = \sum_{m=1}^M \Phi_m(\mathbf{x}), \quad \Phi_m(\mathbf{x}) = \log f(\mathbf{y}_m; \mathbf{x}),$$

where $\mathbf{Y} = (\mathbf{Y}_1, \dots, \mathbf{Y}_M)$ is some decomposition of the data. Assume we identify admissible complete-data random vectors \mathbf{Z}_m for $f(\mathbf{y}_m; \mathbf{x})$ as in Section 3.3.2. Then, for each $\Phi_m(\mathbf{x})$, one can obtain an incremental EM surrogate

$$\phi_m(\mathbf{x}; \bar{\mathbf{x}}) = E[\log f(\mathbf{Z}_m; \mathbf{x}) | \mathbf{Y}_m = \mathbf{y}_m; \bar{\mathbf{x}}] \quad (5.11)$$

that also satisfies the minorization conditions in (5.3). Defining the augmented objective function as in (5.6) and then alternating between updating \mathbf{x} and one of the $\bar{\mathbf{x}}_m$'s as in Table 5.1 leads to the incremental EM algorithms [45,87]. The COSEM algorithm [50,52], a special case of the incremental EM for emission tomography, can be readily derived.

In some applications, using surrogates other than (5.11) can lead to more convenient implementation or faster convergence (*e.g.*, see Section 5.3.2).

5.2.3 Convergence Properties

Global Convergence

Since the incremental optimization transfer algorithms monotonically increase the augmented objective F , the sequence of augmented objective values converges to some value in the usual case where F has an upper bound. However, the question of whether the algorithms really converge to a maximizer of (5.2) is addressed next.

Define a solution set as the collection of stationary points of (5.2):

$$\Gamma \triangleq \{\mathbf{x}^* \in \mathcal{D} : \nabla \Phi(\mathbf{x}^*)'(\mathbf{x} - \mathbf{x}^*) \leq 0, \quad \forall \mathbf{x} \in \mathcal{D}\}, \quad (5.12)$$

where $'$ denotes matrix or vector transpose, and we assume $\Gamma \neq \emptyset$. Each element of the solution set Γ satisfies the first-order necessary condition for a local maximizer of Φ over \mathcal{D} [15, p. 194]. We want algorithms to converge to some point in Γ . If the objective function Φ is concave, then Γ is the set of (possibly multiple) global maximizers of Φ over \mathcal{D} [15, p. 194]. If Φ is strictly concave, then Γ is the singleton of a unique global maximizer [15, p. 685]. On the other hand, for a nonconcave objective function Φ (as in Section 5.3), the solution set Γ could contain even a local minimizer as well as a local maximizer. In fact, it is difficult to find a global maximizer of a nonconcave objective function that may have multiple local maxima. However, the hope is that, with an initial point reasonably close to a global maximizer, the iterates generated by our monotonic algorithms will be attracted to the global maximizer (see [58] for discussion about convergence to a globally optimal point).

In Appendix E, we show that every limit point³ of the sequence generated by the incremental optimization transfer algorithm is an element of the solution set Γ of stationary points regardless of initial points when the following general sufficient conditions hold: (i) each Φ_m and $\phi_m(\cdot; \cdot)$ is continuously differentiable, (ii) the iterates are bounded (*e.g.*, \mathcal{D} is a bounded set), (iii) the surrogates ϕ_m satisfy the minorization conditions in (5.3), (iv) the gradients of Φ_m and $\phi_m(\cdot; \bar{\mathbf{x}})$ match at $\bar{\mathbf{x}}$ (see Condition E.4 in Appendix E), and (v) the maximizer in (5.8) is defined uniquely (*e.g.*, $\phi_m(\cdot; \bar{\mathbf{x}}_m)$ is strictly concave). Consequently, if the objective function Φ is strictly concave, then the algorithm converges to the global maximizer. For a nonconcave objective function Φ , if the points in Γ are isolated, the algorithm will still converge to some stationary point in Γ that we hope is a global maximizer or at least a local maximizer (see Appendix E). It is an open question whether

³Recall the distinction between a limit and a limit point. A point $\tilde{\mathbf{x}}$ is called a *limit* of a sequence $\{\mathbf{x}^n\}$ if $\forall \epsilon > 0$, $\exists N$ such that $\forall n > N$, $\|\tilde{\mathbf{x}} - \mathbf{x}^n\| < \epsilon$. On the other hand, a point $\bar{\mathbf{x}}$ is called a *limit point* of a sequence $\{\mathbf{x}^n\}$ if $\forall \epsilon > 0$, $\forall N$, $\exists n > N$ such that $\|\bar{\mathbf{x}} - \mathbf{x}^n\| < \epsilon$, in other words, if there exists a subsequence $\{\mathbf{x}^{n_k}\}$ whose limit is $\bar{\mathbf{x}}$.

optimization transfer algorithms converge to nonisolated stationary points [58].

Also we analyze the asymptotic local convergence rate of the incremental optimization transfer algorithms, and provide an illustrative one-parameter example for a comparison of the convergence rates of incremental and nonincremental algorithms (see Appendix F).

5.3 Application to Transmission Tomography

In this section we develop a particular incremental optimization transfer algorithm for transmission tomographic reconstruction. We use quadratic surrogates as in Section 3.3.3 rather than EM surrogates in (5.11) because the standard complete-data proposed in [68] for transmission tomography does not yield a closed-form M-step [38]. Of course, using quadratic surrogates is not limited to the transmission case [6, 33, 118]; the incremental optimization transfer algorithms using quadratic surrogates developed in this section are easily extended to other applications including emission tomography.

5.3.1 PL Attenuation Map Reconstruction Problem

For completeness, we briefly rewrite the problem described in Section 3.2. The goal is to find the following PL estimate of attenuation coefficients:

$$\hat{\boldsymbol{\mu}}^{\text{PL}} = \arg \max_{\boldsymbol{\mu} \in \mathcal{B}} \Phi(\boldsymbol{\mu}), \quad \Phi(\boldsymbol{\mu}) = L(\boldsymbol{\mu}) - R(\boldsymbol{\mu}) \quad (5.13)$$

where the box constraint \mathcal{B} is defined in (3.27), the penalty function R is given in (3.3), and the log-likelihood L is given as follows:

$$L(\boldsymbol{\mu}) = \sum_{i=1}^N h_i([\mathbf{A}\boldsymbol{\mu}]_i), \quad h_i(l) = y_i \log(b_i e^{-l} + r_i) - (b_i e^{-l} + r_i).$$

We used the following edge-preserving nonquadratic potential function in our PL reconstruction results [67]:

$$\psi(t) = \delta^2[|t/\delta| - \log(1 + |t/\delta|)] \quad (5.14)$$

for some $\delta > 0$.

5.3.2 Transmission Incremental Optimization Transfer (TRIOT)

We decompose the objective function Φ into the following subobjective functions as in the previous chapter:

$$\Phi_m(\boldsymbol{\lambda}) = \sum_{i \in S_m} h_i([\mathbf{A}\boldsymbol{\mu}]_i) - \frac{\beta}{M} R(\boldsymbol{\mu})$$

where $\{S_m\}_{m=1}^M$ is a partition of $\{1, \dots, N\}$. We use the usual subsets corresponding to downsampled projection angles [55]. Consider the following *separable* quadratic surrogate ϕ_m for the subobjective function Φ_m :

$$\phi_m(\boldsymbol{\mu}; \bar{\boldsymbol{\mu}}) = \Phi_m(\bar{\boldsymbol{\mu}}) + \nabla \Phi_m(\bar{\boldsymbol{\mu}})'(\boldsymbol{\mu} - \bar{\boldsymbol{\mu}}) - \frac{1}{2}(\boldsymbol{\mu} - \bar{\boldsymbol{\mu}})' \check{\mathbf{C}}_m(\bar{\boldsymbol{\mu}})(\boldsymbol{\mu} - \bar{\boldsymbol{\mu}}) \quad (5.15)$$

with

$$\check{\mathbf{C}}_m(\boldsymbol{\mu}) = \text{diag}\{\check{c}_{mj}(\boldsymbol{\mu})\} \quad (5.16)$$

where $\check{c}_{mj}(\cdot) > 0$ and $\text{diag}\{\cdot\}$ denotes a diagonal matrix appropriately formed. The surrogates ϕ_m in (5.15) satisfy Conditions E.4 and E.5 in Appendix E.

As described in Section 3.3.3, to make ϕ_m additionally satisfy the minorization conditions in (5.3), one has at least two choices for \check{c}_{mj} : “optimum curvature (OC)” and “maximum curvature (MC).” Those curvatures \check{c}_{mj} have the following form:

$$\check{c}_{mj}(\boldsymbol{\mu}) = \max \left\{ \sum_{i \in S_m} a_{ij} a_i c_i([\mathbf{A}\boldsymbol{\mu}]_i) + \frac{2\beta}{M} \sum_{k \in \mathcal{N}_j} w_{jk} \omega_\psi(\mu_j - \mu_k), \epsilon \right\} \quad (5.17)$$

for some small value $\epsilon > 0$ where $a_i = \sum_{j=1}^p a_{ij}$ and $\omega_\psi(t) = \dot{\psi}(t)/t$. The functionals $c_i(\cdot)$ are defined as follows. For OC, we define

$$c_i^{\text{OC}}(l) \triangleq \begin{cases} \left[\frac{2h_i(0) - h_i(l) + \dot{h}_i(l) \cdot l}{l^2} \right]_+, & l > 0 \\ \left[\ddot{h}_i(0) \right]_+, & l = 0, \end{cases} \quad (5.18)$$

and for MC,

$$c_i^{\text{MC}}(l) \triangleq \left[\ddot{h}_i(0) \right]_+, \quad (5.19)$$

where $[x]_+ = \max\{x, 0\}$. Detailed derivations of (5.17)–(5.19) can be found in [1]. We leave the second term in (5.17) as a function of $\boldsymbol{\mu}$ even for MC since its computation is cheap compared to projection and backprojection operations unless M is too large; in contrast, we used constant curvatures for the penalty function for relaxed OS-SPS in the preceding chapter.

The augmented objective function F defined in (5.6) with (5.15) is readily maximized with respect to $\boldsymbol{\mu}$ over the box constraint \mathcal{B} as follows:

$$\hat{\boldsymbol{\mu}} = \mathcal{P}_{\mathcal{B}} \left(\left[\sum_{m=1}^M \check{\mathbf{C}}_m(\bar{\boldsymbol{\mu}}_m) \right]^{-1} \sum_{m=1}^M \left[\check{\mathbf{C}}_m(\bar{\boldsymbol{\mu}}_m) \bar{\boldsymbol{\mu}}_m + \nabla \Phi_m(\bar{\boldsymbol{\mu}}_m) \right] \right) \quad (5.20)$$

where $\mathcal{P}_{\mathcal{B}}(\boldsymbol{\mu})$ is the orthogonal projection of $\boldsymbol{\mu} \in \mathbb{R}^p$ onto \mathcal{B} and is easily computed componentwise as follows: $[\mathcal{P}_{\mathcal{B}}(\boldsymbol{\mu})]_j = \text{median}\{0, \mu_j, U\}$ for all j . Using (5.20) in the step (5.8) leads to a new “transmission incremental optimization transfer (TRIOT)” algorithm, which is outlined in Table 5.2. When $M = 1$, this TRIOT reduces to ordinary SPS in (3.53). The TRIOT update begins after $n_{\text{iter}}^{\text{OS}} (\geq 1)$ iteration(s) of OS-SPS (see the next subsection for OS-SPS in detail). Running at least one iteration of fast, though not convergent, OS-SPS is more effective than initializing all $\bar{\boldsymbol{\mu}}_m$ ’s to be the same image (*e.g.*, FBP or uniform image) since in any case one needs to compute partial gradients $\nabla \Phi_m(\bar{\boldsymbol{\mu}}_m)$ (and curvatures) for all m to perform the TRIOT update.

In Table 5.2, a TRIOT using MC in (5.19), we call TRIOT-MC⁴, is shown; however, OC in (5.18) can be easily included. The two steps (5.8) and (5.9) in Table 5.1 are combined in Table 5.2. In (5.22), one can avoid the sum $\sum_{l=1}^M$ at every subiteration by maintaining that sum as a state vector that is updated incrementally as in [50, 52, 53]. One could slightly modify the algorithm to perform (5.22) more than one time at every subiteration so that one additionally updates the surrogate for the penalty part with fixing the surrogate for the

⁴The second part denotes a specific curvature used (*e.g.*, SPS-OC).

likelihood part as in [1]. One iteration, indexed by n , of TRIOT-MC requires one projection and one backprojection operation while TRIOT-OC needs an extra backprojection [see (5.17) and (5.18)].

Table 5.2: Outline for transmission incremental optimization transfer (TRIOT) algorithm using maximum curvature (MC).

Initialize: $\hat{\boldsymbol{\mu}} = \hat{\boldsymbol{\mu}}^0 = \left[\text{FBP} \left\{ \log \left(\frac{b_i}{y_i - r_i} \right) \right\}_{i=1}^N \right]^+$
Precompute: $d_{mj}^{\text{MC}} = \sum_{i \in S_m} a_{ij} a_i \left[\left(1 - \frac{y_i r_i}{(b_i + r_i)^2} \right) b_i \right]_+$ and $d_j^{\text{PC}} = \frac{1}{M} \sum_{i=1}^N a_{ij} a_i c_i^{\text{PC}}$, $\forall m, j$
for each iteration $n = 1, \dots, n_{\text{iter}}$
for each subset (subiteration) $m = 1, \dots, M$
$\hat{l}_i = \sum_{j=1}^p a_{ij} \hat{\mu}_j, \quad \hat{h}_i = \left(1 - \frac{y_i}{b_i e^{-\hat{l}_i} + r_i} \right) b_i e^{-\hat{l}_i}, \quad \forall i \in S_m$
$\dot{L}_{mj} = \sum_{i \in S_m} a_{ij} \dot{h}_i, \quad r_{mj} = \frac{2\beta}{M} \sum_{k \in \mathcal{N}_j} w_{jk} \omega_\psi(\hat{\mu}_j - \hat{\mu}_k), \quad \forall j$
$\bar{\mu}_{mj} = \hat{\mu}_j, \quad \forall j$
if $n \leq n_{\text{iter}}^{\text{OS}}$, perform the following OS-SPS update:
$\hat{\mu}_j = \left[\bar{\mu}_{mj} + \frac{\dot{L}_{mj} - \frac{\beta}{M} \sum_{k \in \mathcal{N}_j} w_{jk} \dot{\psi}(\bar{\mu}_{mj} - \bar{\mu}_{mk})}{\max \{ d_j^{\text{PC}} + r_{mj}, \epsilon \}} \right]^+, \quad \forall j \quad (5.21)$
else , perform the following TRIOT-MC update:
$\hat{\mu}_j = \left[\frac{\sum_{l=1}^M \left[\bar{\mu}_{lj} \max \{ d_{lj}^{\text{MC}} + r_{lj}, \epsilon \} + \left(\dot{L}_{lj} - \frac{\beta}{M} \sum_{k \in \mathcal{N}_j} w_{jk} \dot{\psi}(\bar{\mu}_{lj} - \bar{\mu}_{lk}) \right) \right]}{\sum_{l=1}^M \max \{ d_{lj}^{\text{MC}} + r_{lj}, \epsilon \}} \right]^+, \quad \forall j \quad (5.22)$
end
end
$\hat{\boldsymbol{\mu}}^n = \hat{\boldsymbol{\mu}}$
end

Here ϵ is some small positive value; c_i^{PC} is defined in (5.24); and $[\mu]^+ \triangleq \text{median}\{0, \mu, U\}$, which should not be confused with $[\mu]_+ \triangleq \max\{\mu, 0\}$.

The discussion and proofs for global convergence given in Section 5.2.3 and Appendix E apply to TRIOT. When $r_i = 0$ for all i , under mild conditions, since the PL objective for transmission tomography is strictly concave, the algorithm converges to the

optimal solution [69]. In the case $r_i \neq 0$, the objective function is not necessarily concave [1], and we have a weaker conclusion that every limit point of a sequence generated by TRIOT is a stationary point. However, in our practical experience, we obtained the same limit in all experiments with different initializations, suggesting that suboptimal local maxima are rare, or are far from reasonable starting images.

5.3.3 OS-SPS

Since we use *unrelaxed* OS-SPS in initializing and accelerating TRIOT, we briefly review OS-SPS for completeness (cf Section 4.3.2). For each subiteration, indexed by m , maximizing the m th subobjective $\phi_m(\cdot; \bar{\boldsymbol{\mu}}_m)$ in (5.15) instead of the augmented objective $F(\cdot; \bar{\boldsymbol{\mu}}_1, \dots, \bar{\boldsymbol{\mu}}_M)$ in (5.6) leads to the following OS-SPS update:

$$\bar{\boldsymbol{\mu}}_{(m \bmod M)+1} = \mathcal{P}_{\mathcal{B}} \left(\bar{\boldsymbol{\mu}}_m + \left[\check{\mathbf{C}}_m(\bar{\boldsymbol{\mu}}_m) \right]^{-1} \nabla \Phi_m(\bar{\boldsymbol{\mu}}_m) \right) \quad (5.23)$$

for $m = 1, \dots, M$ where $\check{\mathbf{C}}_m(\cdot)$ is based on (5.17). This greedy approach does not ensure monotonicity, in neither the augmented objective nor the PL objective, so we need not insist that the curvatures satisfy the minorization conditions. A natural choice for $c_i(\cdot)$ is the Newton's curvatures $-\ddot{h}_i(\cdot)$; these can be approximated as follows:

$$\begin{aligned} -\ddot{h}_i(l) \approx c_i^{\text{PC}} &\triangleq -\ddot{h}_i \left(\arg \max_{\tilde{l} \geq 0} h_i(\tilde{l}) \right) \\ &= \begin{cases} \frac{(y_i - r_i)^2}{y_i}, & y_i > r_i \\ 0, & \text{otherwise.} \end{cases} \end{aligned} \quad (5.24)$$

This choice is called ‘‘precomputed curvature (PC)’’ [1, 2]. Usually, for OS-SPS, the following subset-independent preconditioning matrix using PC is used in place of $\check{\mathbf{C}}_m(\bar{\boldsymbol{\mu}}_m)$ in (5.23):

$$\begin{aligned} \check{\mathbf{C}}^{\text{PC}}(\boldsymbol{\mu}) &= \text{diag}_j \{ \check{c}_j^{\text{PC}}(\boldsymbol{\mu}) \} \\ \check{c}_j^{\text{PC}}(\boldsymbol{\mu}) &= \max \left\{ \frac{1}{M} \sum_{i=1}^N a_{ij} a_i c_i^{\text{PC}} + \frac{2\beta}{M} \sum_{k \in \mathcal{N}_j} w_{jk} \omega_\psi(\mu_j - \mu_k), \epsilon \right\} \end{aligned} \quad (5.25)$$

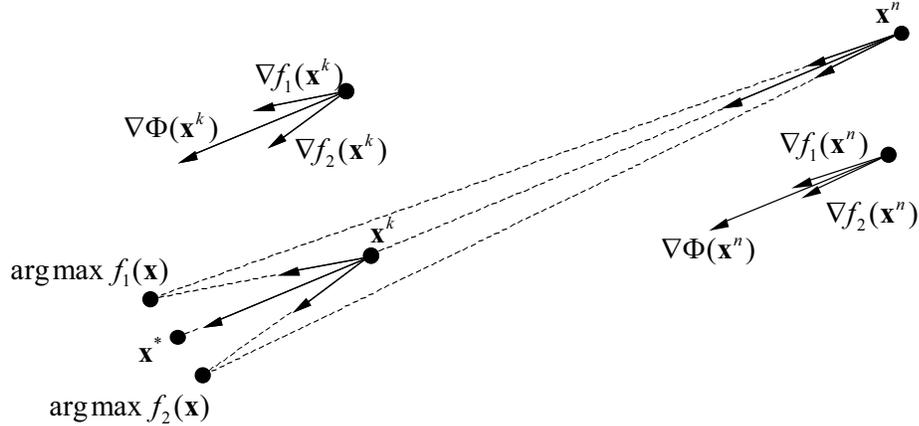


Figure 5.1: Illustration of a geometric view of OS algorithms. In this example, there are two subobjective functions f_1 and f_2 such that $\Phi = f_1 + f_2$. When \mathbf{x}^n is far from the optimal point \mathbf{x}^* , all $\nabla\Phi(\mathbf{x}^n)$, $\nabla f_1(\mathbf{x}^n)$, and $\nabla f_2(\mathbf{x}^n)$ point roughly towards \mathbf{x}^* . However, when \mathbf{x}^k is near \mathbf{x}^* , they do not.

where c_i^{PC} is given in (5.24). The first term on the right side in (5.25) can be precomputed and stored like the maximum curvatures (MC). The benefit of using PC is that it leads to faster convergence rates than MC. The update for OS-SPS is shown in (5.21) in Table 5.2.

The OS-SPS shows very fast initial convergence rates but becomes eventually stuck at a limit cycle. Using more subsets leads to a faster initial convergence rate but causes the points in the limit cycle to be farther from the optimal solution.

Fig. 5.1 illustrates a geometric view of OS algorithms. Suppose a current iterate is far from an optimal point. Then even partial gradients should point roughly to the optimal point. This explains initial accelerated convergence rates of typical OS algorithms. However, when a current iterate is near the optimal point, the subset gradient balance conditions (4.9) or (4.10) never hold, and OS algorithms generally exhibit limit cycle behavior particularly with a constant stepsize $\alpha_n = \alpha$.

5.3.4 Acceleration

TRIOT-OC/MC is convergent yet faster than nonincremental ordinary SPS [2], but it is still slower initially than OS-SPS which is not convergent unless relaxed. We discuss

methods to accelerate TRIOT.

Enhanced Incremental Optimization Transfer Algorithms

Hsiao *et al.* proposed E-COSEM, an accelerated version of COSEM [53]. The idea is to choose for each subiteration a convex combination of a fast yet nonconvergent OS algorithm and a slow yet convergent incremental optimization transfer algorithm such that the combination both ensures monotonicity in the augmented objective and is as close to the OS algorithm as possible. This approach usually accelerates incremental optimization transfer algorithms without destroying the monotonicity in the augmented objective.

Switch from OS-SPS to TRIOT

It is a popular idea to switch from a nonconvergent yet fast OS type algorithm to a convergent yet slow non-OS algorithm at some point to take advantage of both fast initial convergence rates of OS methods and fast asymptotic convergence rates (or global convergence) of non-OS methods. Alternatively, one could decrease the number of subsets as iterations proceed or could use a hybrid class of methods that combine OS and non-OS algorithms [14]. However, the switching schedule or the parameters for the hybrid class are as inconvenient to determine as relaxation parameters in relaxed OS algorithms.

We observed that it is very effective to switch to TRIOT from OS-SPS at the point where the OS-SPS algorithm nearly gets to a limit cycle; even one single subiteration of TRIOT moves the iterate from the limit cycle very close to the optimal solution. The reason is as follows: a group of the points in the limit cycle would be roughly centered around the optimal point and the update for TRIOT includes a weighted average of the points [see the first term on the right side in (5.20) or (5.22)].

To obtain further insight into this property, consider a simple unconstrained quadratic

problem where the objective function and the subobjective functions are

$$\Phi(\mathbf{x}) = -\frac{1}{2}\mathbf{x}'\mathbf{Q}\mathbf{x} + \mathbf{b}'\mathbf{x}, \quad \Phi_m(\mathbf{x}) = -\frac{1}{2}\mathbf{x}'\mathbf{Q}_m\mathbf{x} + \mathbf{b}'_m\mathbf{x}$$

respectively where $\sum_{m=1}^M \mathbf{Q}_m = \mathbf{Q}$ and $\sum_{m=1}^M \mathbf{b}_m = \mathbf{b}$. Assume that each surrogate function $\phi_m(\mathbf{x}; \bar{\mathbf{x}})$ is equal to its corresponding subobjective $\Phi_m(\mathbf{x})$ and so it has a closed-form maximizer $\hat{\mathbf{x}}_m = \mathbf{Q}_m^{-1}\mathbf{b}_m$ where we assume \mathbf{Q}_m are invertible. Then the OS approach will generate a limit cycle that consists of those $\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_M$. Now applying just one iteration of the incremental optimization transfer method as in (5.20) leads to

$$\begin{aligned} \hat{\mathbf{x}} &= \left(\sum_{m=1}^M \mathbf{Q}_m \right)^{-1} \sum_{m=1}^M \mathbf{Q}_m \hat{\mathbf{x}}_m = \mathbf{Q}^{-1} \sum_{m=1}^M \mathbf{Q}_m (\mathbf{Q}_m^{-1} \mathbf{b}_m) \\ &= \mathbf{Q}^{-1} \mathbf{b}, \end{aligned}$$

which is the maximizer of the original objective [the second term on the right side in (5.20) equals zero]. This example suggests that the built-in averaging operation in TRIOT helps iterates escape from a limit cycle, generated by nonconvergent OS algorithms, towards the optimal solution.

However, when OS-SPS is still far from the limit cycle and is making progress towards the optimal point, that is, at early iterations, TRIOT is usually slower than OS-SPS due to the averaging of the past subiterates; or one can say the reason is that the incremental optimization transfer approach updates the surrogates incrementally, that is, conservatively to ensure monotonicity. So it is desirable to get to a limit cycle quickly using OS-SPS with *many* subsets and then switch to TRIOT. In a 2D reconstruction case in Section 5.4, the use of 64 subsets is sufficient to reach nearly a limit cycle within a couple of iterations.

Precomputed Curvatures

Forgoing monotonicity (in the augmented objective) and (probably) global convergence, one can use for TRIOT the “precomputed curvatures (PC)” in (5.24) usually used

for nonconvergent OS-SPS. TRIOT-PC is faster than provably convergent TRIOT-OC/MC. It is an open question whether TRIOT-PC converges to an optimal solution. However, in our experiments, TRIOT-PC yielded the same limit as convergent algorithms like SPS-OC within numerical precision!

5.4 Results

To assess the performance and to investigate the behavior of the proposed algorithms, we performed 2D attenuation map reconstructions.

5.4.1 Simulation

We used a slice of the (modified) Zubal phantom [137] in Fig. 5.5(a). The true attenuation coefficient levels for the image were 0.002 (lungs), 0.0096 (soft tissue and water), and 0.012 mm^{-1} (bone). The sinogram had 160 radial bins and 192 angles, and the reconstructed images were 128×128 with 4.2 mm pixels. The system geometry was approximated with 3.375 mm wide strip integrals and 3.375 mm ray spacing; the system matrix was generated using ASPIRE [35]. The known r_i factor corresponded to a uniform field of 10% backgrounds, and the total counts amounted to 10^6 . We used the edge-preserving nonquadratic penalty (5.14) with $\delta = 4 \times 10^{-4} \text{ mm}^{-1}$ and $\beta = 2^{17.5}$, chosen by visual inspection. Pseudo-random independent Poisson variates were drawn according to (3.23). We used the FBP reconstruction shown in Fig. 5.5(b) as a starting image.

Images were reconstructed using convergent yet slow SPS-MC/OC, fast yet nonconvergent OS-SPS-PC, TRIOT algorithms, and enhanced TRIOT algorithms. For enhanced TRIOT algorithms, we used the same method as given in [53]. For OS-SPS and TRIOT methods, we used 16 subsets (a moderate number) and 64 subsets (a little larger number than usual).

Fig. 5.2 shows normalized Φ difference versus iteration for different algorithms using

16 subsets. The normalized Φ difference is defined by $(\Phi(\hat{\boldsymbol{\mu}}^{\text{PL}}) - \Phi(\hat{\boldsymbol{\mu}}^n))/(\Phi(\hat{\boldsymbol{\mu}}^{\text{PL}}) - \Phi(\hat{\boldsymbol{\mu}}^0))$ where $\hat{\boldsymbol{\mu}}^{\text{PL}}$ is a maximizer of the penalized likelihood; and a small value means a better image (closer to the optimal image $\hat{\boldsymbol{\mu}}^{\text{PL}}$). The optimal image $\hat{\boldsymbol{\mu}}^{\text{PL}}$ [shown in Fig. 5.5(c)] was estimated by 30 iterations of OS-SPS followed by 800 iterations of the monotonic SPS-OC algorithm. As described in Section 5.3.2, TRIOT methods were initialized by running one iteration of OS-SPS. At iteration 2, all TRIOT methods decreased Φ slightly, apparently because the subiterates generated by OS-SPS had not reached a limit cycle yet. This shows nonmonotonicity in the objective despite monotonicity in the augmented objective. Although OS-SPS shows a fast initial convergence rate, it become stuck at a suboptimal point whereas other methods continue to improve in terms of objective values. TRIOT-OC lead to faster convergence rates than TRIOT-MC; however, TRIOT-OC requires an extra backprojection per iteration. The enhancement proposed in [53] is very effective here. Although global convergence is not provably ensured for TRIOT-PC, the limit of TRIOT-PC (say, obtained by 1000 iterations) was the same as that of SPS-OC (obtained similarly) within numerical precision (not shown here), which suggests TRIOT-PC also has desirable convergence properties.

Fig. 5.3 shows the behavior of TRIOT algorithms when they are initialized by running 3 iterations of OS-SPS. Since OS-SPS nearly reached a limit cycle after 3 iterations, the built-in weighted-averaging in the TRIOT algorithms leads to considerable improvement at iteration 4. This shows that it is effective to switch from OS-SPS to TRIOT around when OS-SPS almost reaches a limit cycle. However, it is inconvenient to predict how many iterations are required for OS-SPS to arrive at a limit cycle; and if the required iterations are many, then there is no room to improve the initial convergence rates during OS-SPS being run. In Fig. 5.3, at early iterations, the enhancement method did slow TRIOT algorithms. Further investigation is needed.

As shown in Fig. 5.4, when we increase the number of subsets to 64, the initial convergence rate of OS-SPS becomes much faster (even one iteration leads to a limit cycle) but OS-SPS stagnates at a worse image. Meanwhile, the TRIOT algorithms and their enhanced versions are quite effective even they used only one iteration of OS-SPS as initialization. TRIOT-PC (with enhancement) shows very nice properties: fast initial convergence rates and, in this study, convergence to the same limit as the convergent SPS-OC [the reconstructed image is shown in Fig. 5.5(d)].

5.4.2 Real Data

We acquired real PET data using a Siemens/CTI ECAT EXACT 921 PET scanner with rotating rod transmission sources [128]. We used an anthropomorphic thorax phantom (Data Spectrum, Chapel Hill, NC). The sinogram size (160×190) and the image size (128×128) were the same as the previous simulation study, and the system geometry was similar. The total counts amounted to 9.2×10^5 . We used the edge-preserving nonquadratic penalty (5.14) with $\delta = 4 \times 10^{-4} \text{ mm}^{-1}$ and $\beta = 2^{18.5}$. The results were very similar to the previous simulation case (see Figs. 5.6–5.8). Fig. 5.9(c) shows an image of the limit cycle generated by OS-SPS with 64 subsets (see Fig. 5.8), The figure looks different from the true PL solution in Fig. 5.9(b) and its image quality is not very good. On the other hand, Fig. 5.9(d) reconstructed by TRIOT-PC looks similar to the PL solution image and shows good image quality. In light of these similar results, our strategy—switch to TRIOT after few iterations of OS-SPS with many subsets—seems robust whereas relaxed OS algorithms are sensitive to relaxation parameters.

5.5 Conclusion

We presented a broad family of incremental optimization transfer algorithms by generalizing the incremental EM algorithms. The incremental optimization transfer algorithms

show faster convergence rates than ordinary optimization transfer methods like EM, but are globally convergent.

We also developed a particular incremental optimization transfer algorithm for transmission tomography by using separable quadratic surrogates: TRIOT algorithms. We found that it is very effective to switch from OS-SPS to TRIOT when OS-SPS nearly reaches a limit cycle. The enhancement method in [53] worked effectively. When reasonably many subsets are used, as few as one iteration of OS-SPS can be enough to approach a limit cycle (although it would depend on the degree of regularization applied and the problem size). This switching strategy is much more convenient than relaxed OS algorithms that require determining relaxation parameters. Also, TRIOT is preferable to reducing the number of subsets with iteration since the consistent data flow in OS-SPS and TRIOT could be beneficial and it would be inconvenient to determine an optimal schedule for reducing the number of subsets. The switching idea is also found in [104].

One iteration of TRIOT-MC/PC or OS-SPS requires computing one projection and one backprojection plus the penalty related gradients and curvatures (the use of OC needs an extra backprojection); so the computational cost is almost the same as classic ML-EM except for the contribution of the penalty part. As the number of subsets increases, required computation per iteration also increases due to the penalty part being updated for each subiteration. Although the computational contribution of the penalty function is usually small compared to projection/backprojection particularly for a large-scale problem like 3D, further investigation could help reduce this computation further, *e.g.*, by subsetizing the penalty part.

Our recommended method for transmission tomography is to switch from OS-SPS to (enhanced) TRIOT-PC when the OS-SPS nearly reaches a limit cycle (in our 2D case, one iteration was enough for 64 subsets). Although the TRIOT-PC was numerically found to

be convergent, if one really wants provable convergence, one could switch to TRIOT-MC or OC at some point.

We expect that incremental optimization transfer methods with quadratic surrogates like TRIOT will work similarly for emission image reconstruction since the update equation is of the same form as (5.20). The use of EM surrogates leads to COSEM algorithms [50, 52, 53], and the comparison of quadratic surrogates and EM surrogates in incremental optimization transfer methods for emission tomography will need investigation.

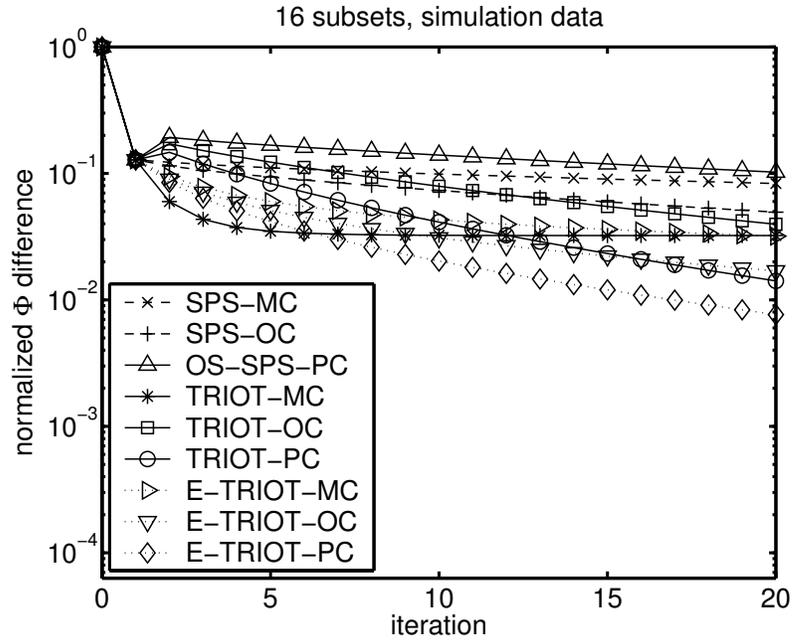


Figure 5.2: Comparison of non-OS algorithms (SPS-MC/OC), an OS algorithm (OS-SPS), incremental optimization transfer algorithms (TRIoT-MC/OC/PC), and enhanced TRIoT (E-TRIoT-MC/OC/PC) for 2D attenuation map reconstruction using simulation data. The normalized Φ difference is defined by $(\Phi(\hat{\mu}^{\text{PL}}) - \Phi(\hat{\mu}^n)) / (\Phi(\hat{\mu}^{\text{PL}}) - \Phi(\hat{\mu}^0))$ where $\hat{\mu}^{\text{PL}}$ is a PL estimate image. OS-SPS and TRIoT algorithms used 16 subsets, and TRIoT algorithms include one iteration of OS-SPS initially. The starting image was a FBP image for all cases.

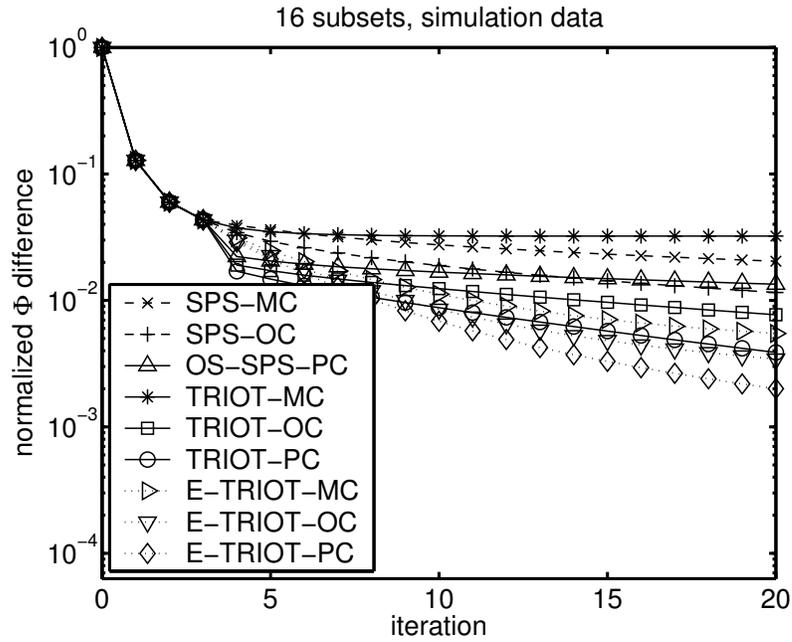


Figure 5.3: Same as Fig. 5.2, but TRIOT algorithms include *three* iterations of OS-SPS initially.

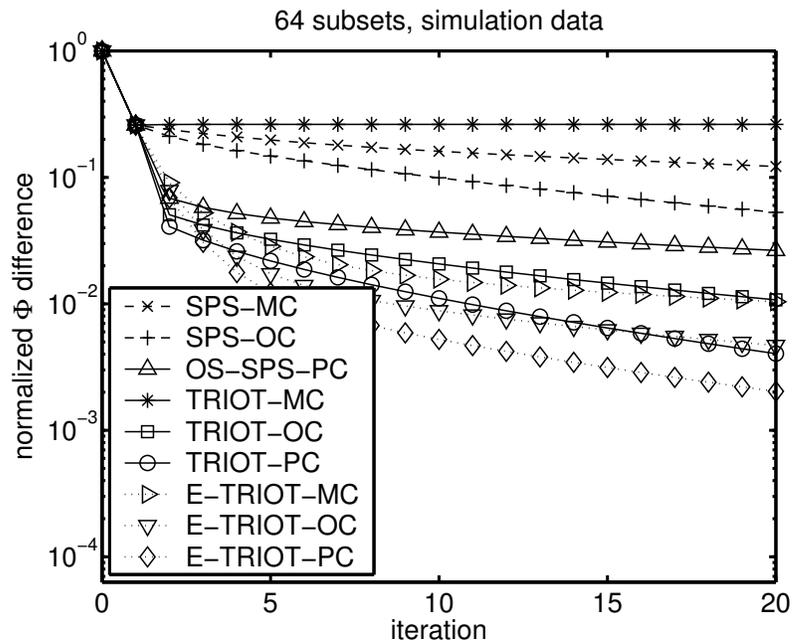


Figure 5.4: Same as Fig. 5.2, but 64 subsets are used for OS-SPS and TRIOT algorithms.

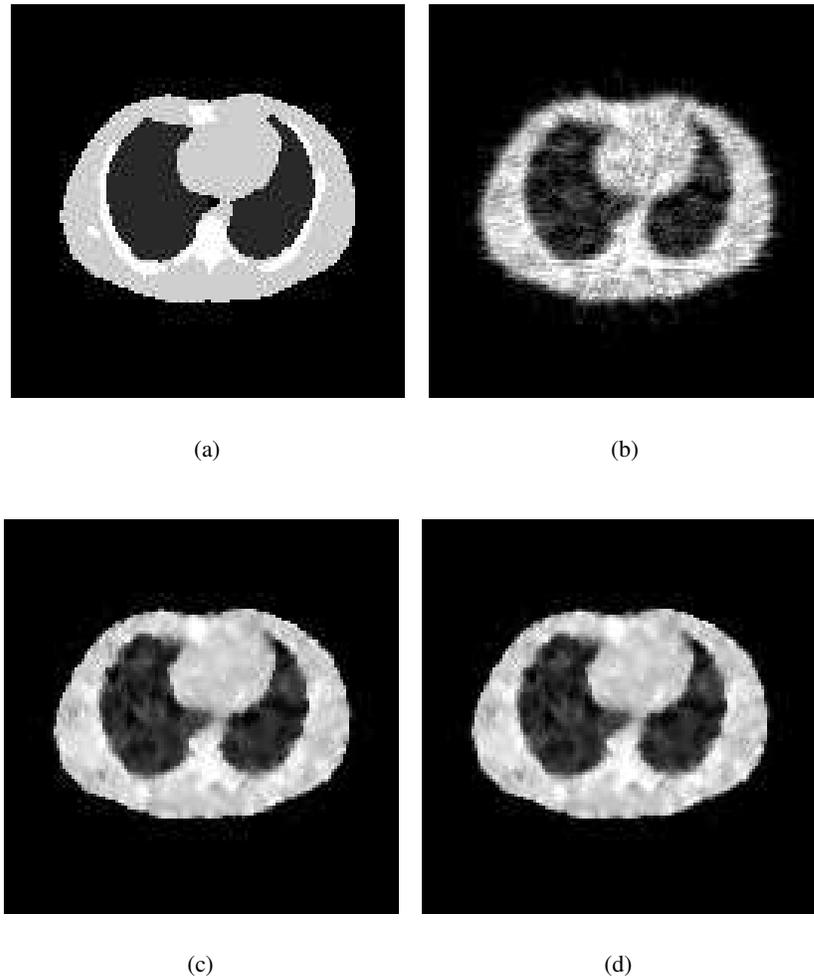


Figure 5.5: (a) A slice of the (modified) Zubal phantom (true image). (b) FBP reconstruction (starting image). (c) PL estimate image $\hat{\mu}^{\text{PL}}$ obtained using 30 iterations of OS-SPS-PC with 16 subsets followed by 800 iterations of SPS-OC. (d) PL reconstruction using 1 iteration of OS-SPS-PC and 19 iterations of TRIOT-PC with 64 subsets.

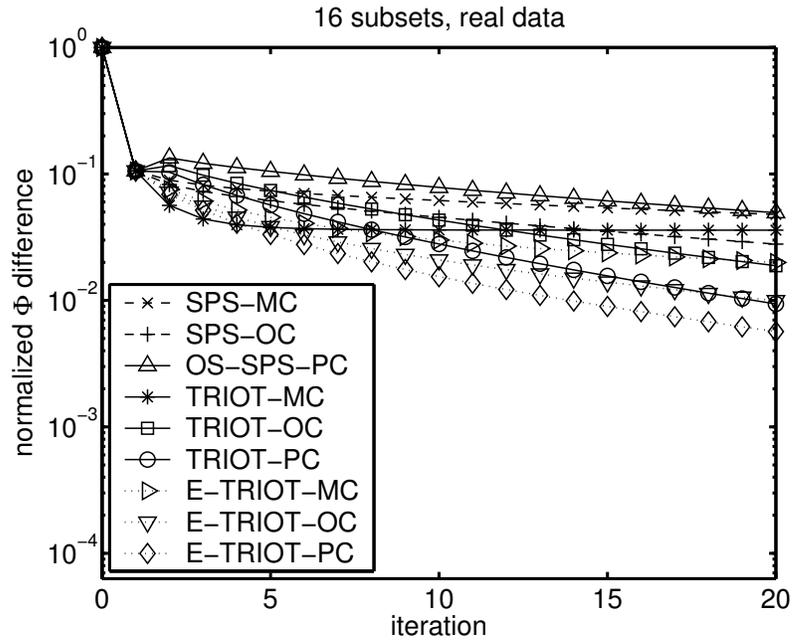


Figure 5.6: Same as Fig. 5.2, but real PET data are used.

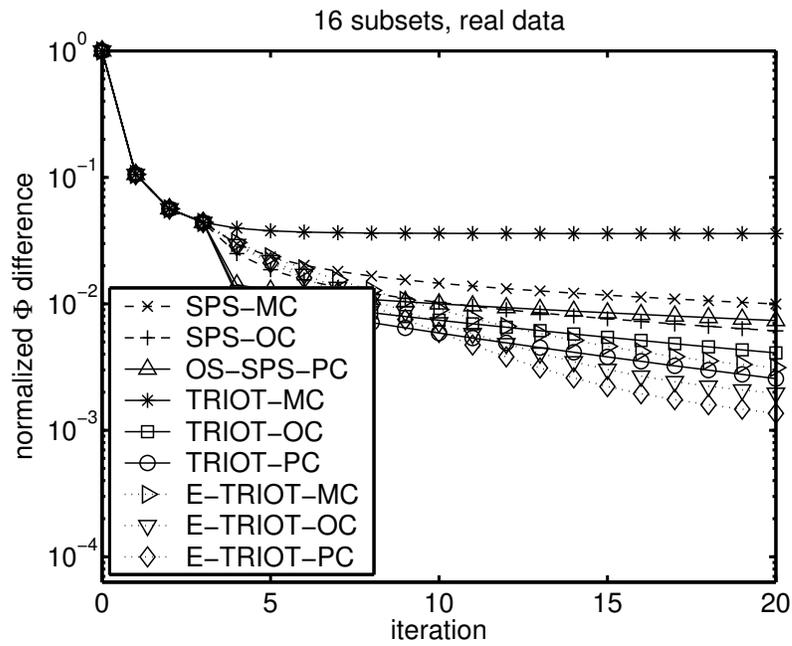


Figure 5.7: Same as Fig. 5.3, but real PET data are used.

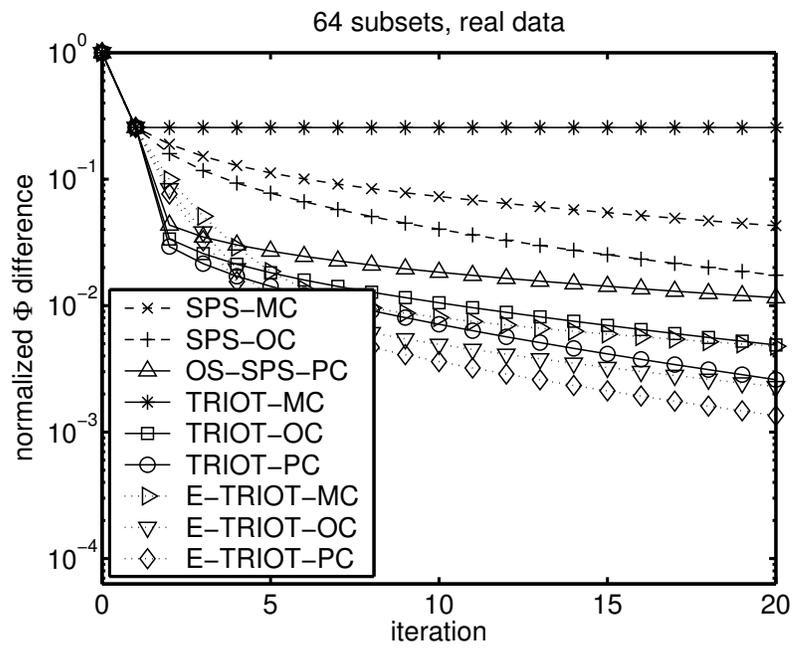


Figure 5.8: Same as Fig. 5.4, but real PET data are used.

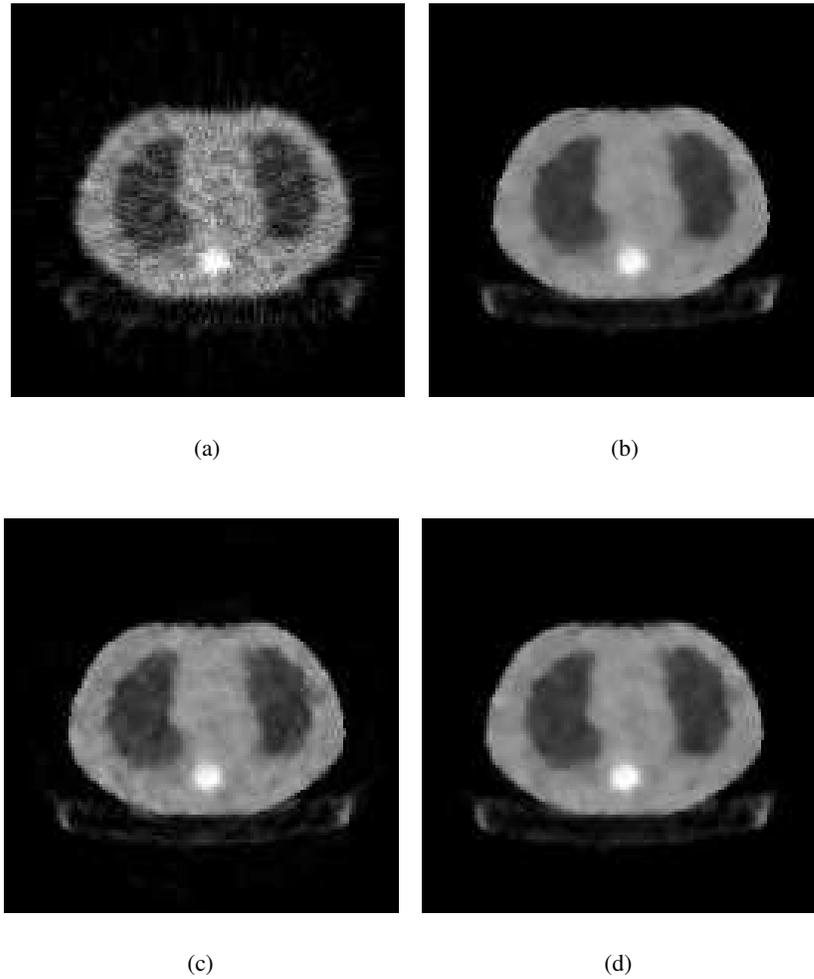


Figure 5.9: Reconstructed attenuation maps from real PET data. (a) FBP reconstruction (starting image). (b) PL estimate image $\hat{\mu}^{\text{PL}}$ obtained using 30 iterations of OS-SPS-PC with 16 subsets followed by 800 iterations of SPS-OC. (c) PL reconstruction using 20 iterations of OS-SPS-PC with 64 subsets (an image that is one point of a limit cycle). (d) PL reconstruction using 1 iteration of OS-SPS-PC and 19 iterations of TRIOT-PC with 64 subsets.

CHAPTER 6

Statistical Image Reconstruction for Randoms-Precorrected PET Scans

6.1 Introduction

Accidental coincidence (AC) events, also known as randoms, are a primary source of background noise in positron emission tomography (PET) [48]. AC events occur when two photons that arise from separate positron emissions are detected within a coincidence timing window and recorded as having originated from the same emission [95,99]. Quantitative PET studies require correction for AC events.

Usually, PET systems detect coincidence events during “prompt” windows and “delayed” windows as described in Section 2.3.1 [120, 131]. The delayed coincidences represent AC events (or randoms), and the prompt coincidences represent true coincidences contaminated by AC events (plus Compton scatter events). In most PET scans, the prompt data are precorrected for the effects of AC events by real-time subtraction of the delayed coincidences [48]. The subtraction compensates for the AC events in terms of the mean but increases the variance of the data [86]. Ideally, scanners would maintain both prompt and randoms sinograms. One could then estimate the mean of AC events from the randoms sinogram and incorporate these estimates into an appropriate model for the prompt measurement to estimate unknown parameters (radioactivity for emission

scans and attenuation coefficients for transmission scans) [22, 49, 85, 86, 99]. However, because of data storage limitations and historical momentum, most PET centers store the randoms-precorrected data only [131]. This chapter focuses on the problem of reconstructing emission images by considering the measurement statistics based on *only* randoms-precorrected data *without* access to separate prompt and randoms sinograms. We do assume that a rough estimate of the randoms contribution is available, such as can be computed from the block singles rates that are often available [86].

Whereas both (prompt and randoms) sinograms are well approximated as being Poisson distributed [135], the randoms-precorrected data do not follow Poisson statistics. The exact log-likelihood of precorrected data is inconvenient to maximize. Several practical approximations to the exact log-likelihood have been investigated [129–133]. A shifted Poisson (SP) model and a saddle-point (SD) model are such approximations [130, 131]. For *transmission scans*, both SP and SD models have been shown to outperform conventional ordinary Poisson (OP) and weighted least squares (WLS) models in terms of systematic bias and variance [130–132]. In *transmission* image reconstruction, the SP model seems more attractive than the SD model since its implementation is simpler but their performance is comparable.

For *emission scans*, the SP and SD models again lead to lower variance than the OP model [133]. However, SP suffers from a positive systematic bias for low counts per ray, albeit generally less than OP, while SD seems to be free of such a bias [133]. The systematic bias is caused by zeroing negative sinogram values (note that randoms-precorrected data can be negative) [133]; it can affect contrast and quantitative studies adversely. The zero-thresholding of negative values for SP and OP was a natural choice since those models are based on Poisson approximations. Moreover, negative sinogram values can cause reconstruction algorithms like classic ML-EM to diverge. Negative sinogram values also

cause the Poisson log-likelihood to become nonconcave, and it is difficult to develop algorithms that globally maximize a nonconcave objective function. By contrast, in a *transmission* case, negative values do not cause reconstruction algorithms to diverge; in fact, they help ensure concavity of the log-likelihood that otherwise could be nonconcave [129, Sec. 4.6].

To eliminate the positive systematic bias in *emission scans*, we propose new SP and OP models that allow negative sinogram values, departing from the conventional tendency to zero-threshold them [83, 84, 102, 133]. We will henceforth call our new methods “SP⁻” and “OP⁻” to differentiate from the conventional ones with zero-thresholding that will be called “SP⁺” and “OP⁺” in this dissertation. In contrast to some previous methods that allow the pixel values to be negative [93], here we enforce the usual nonnegativity constraint in the *image domain* but allow the *sinogram* values to be negative for SP⁻ and OP⁻. We show that negative sinogram values in emission scans need *not* cause divergence of appropriate algorithms for SP⁻ and OP⁻. Although negative values can cause the likelihood for SP⁻ or OP⁻ to be nonconcave, one can achieve at least a locally optimal reconstruction by employing algorithms that increase the objective function monotonically. We use the “optimization transfer principles” [70] to derive two monotonic algorithms that allow negative values: separable paraboloidal surrogates (SPS) and a variant of maximum likelihood expectation maximization (ML-EM). Our practical experience shows that the locally optimal reconstruction obtained by monotonic algorithms is very good regardless of initializations.

We show analytically that our new SP⁻ model is nearly free of systematic bias (as is the new OP⁻ model) and leads to less variance than other methods including OP⁻ and filtered backprojection (FBP); this is corroborated by simulation results in Section 6.7. In other words, the new SP⁻ model, our recommended method for randoms-precorrected emission image reconstruction, is comparable, in spite of its simpler implementation, to SD which

has shown the best performance in terms of systematic bias and variance.

Section 6.2 reviews the statistical model for precorrected measurements and its exact log-likelihood. Section 6.3 describes conventional approximation models and our new ones for emission scans. We analyze the systematic bias due to zero-thresholding and the asymptotic variances in Sections 6.4 and 6.5. Section 6.6 provides monotonic algorithms for the new models and Section 6.7 gives simulation results.

6.2 Measurement Model and Exact Log-likelihood

Let $\mathbf{Y} = [Y_1, \dots, Y_N]'$ denote the *precorrected* measurements for PET emission scans, where $'$ denotes vector and matrix transpose. The precorrected measurement for the i th bin is

$$Y_i = Y_i^{\text{prompt}} - Y_i^{\text{delay}} \quad (6.1)$$

where Y_i^{prompt} and Y_i^{delay} are the number of coincidences detected within the prompt and delayed windows, respectively. The prompts and delays can be modeled reasonably as independent Poisson random variables [135] as follows:

$$Y_i^{\text{prompt}} \sim \text{Poisson}\{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + r_i + s_i\} \quad (6.2)$$

$$Y_i^{\text{delay}} \sim \text{Poisson}\{r_i\} \quad (6.3)$$

where $a_{ij} \geq 0$ is the entry in the system matrix \mathbf{A} incorporating scan geometry, attenuation, detector efficiencies, *etc.*; $\lambda_j^{\text{true}} \geq 0$ is the activity at the j th voxel; and $r_i \geq 0$ and $s_i \geq 0$ are the means of AC events and scatters, respectively. To focus on the problem of estimating the unknown activity $\boldsymbol{\lambda}^{\text{true}} = [\lambda_1^{\text{true}}, \dots, \lambda_p^{\text{true}}]'$, we assume that $\mathbf{r} = [r_1, \dots, r_N]'$ and $\mathbf{s} = [s_1, \dots, s_N]'$ are *known*.¹ In other words, we investigate the

¹Even in a case where one does not have access to the delayed events separately, the total number of AC events or the block singles rates are often available at the end of the scan and can be used to estimate AC rates [86, 103]. Indeed, approximate models like SP and SD are known to be robust to errors in estimating AC rates [130]. Regarding scatter estimation and correction, see [86, 94, 127] for example.

“upper bound” of performance of each reconstruction method that needs estimates of \mathbf{r} and \mathbf{s} . In [49], the effects of random estimates on bias for various reconstruction methods were investigated. We also assume $r_i > 0$ for all i for simplicity; the analysis and algorithms are easily adopted to include rays where $r_i = 0$.

Let $\mathbf{y} = [y_1, \dots, y_N]'$ be an observed realization of \mathbf{Y} . Recall that $\mathbf{Y}^{\text{prompt}}$ and $\mathbf{Y}^{\text{delay}}$ are not separately accessible. Since the measurements are independent, one can obtain the exact log-likelihood, ignoring constants independent of $\boldsymbol{\lambda}$, as in [130, 131]:

$$L(\boldsymbol{\lambda}; \mathbf{Y}) = \sum_{i=1}^N h_i^{\text{EX}}([\mathbf{A}\boldsymbol{\lambda}]_i) \quad (6.4)$$

with

$$h_i^{\text{EX}}(l) = \log \left(\sum_{m=[-y_i]_+}^{\infty} \frac{(l + r_i + s_i)^{y_i+m} r_i^m}{(y_i + m)! m!} \right) - (l + 2r_i + s_i) \quad (6.5)$$

where $[x]_+ = \max\{x, 0\}$. For notational simplicity, we omit an argument indicating the dependence of h_i^{EX} on y_i in (6.4) and (6.5).

For penalized-likelihood (PL) reconstruction, one must find a maximizer of the objective function

$$\Phi(\boldsymbol{\lambda}; \mathbf{Y}) = L(\boldsymbol{\lambda}; \mathbf{Y}) - R(\boldsymbol{\lambda}) \quad (6.6)$$

over a nonnegativity constraint on the *image* $\boldsymbol{\lambda}$, where R is a regularization term that controls a trade-off of resolution and noise in the reconstructed image as in (3.9). The exact log-likelihood function (6.4) is inconvenient to maximize although it can be expressed without the infinite summations in (6.5) using Bessel functions [129, Sec. 3.2]. The next section describes practical approximations to the exact log-likelihood.

6.3 Approximations to Exact Log-likelihood

6.3.1 Ordinary Poisson (OP) Approximation

A simple approach that does not need an estimate of AC events r is to approximate the measurements as Poisson random variables as follows:

$$Y_i \stackrel{\text{OP}}{\underset{\text{approx.}}{\sim}} \text{Poisson}\{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i\}. \quad (6.7)$$

This model matches the first moment of Y_i only. The log-likelihood L^{OP^-} corresponding to this “OP⁻” approximation² is of the form (6.4) with

$$h_i^{\text{OP}^-}(l) = y_i \log(l + s_i) - (l + s_i). \quad (6.8)$$

We assume $s_i > 0$ in (6.8); otherwise, negative values y_i would cause reconstruction algorithms to diverge since $h_i^{\text{OP}^-}(0) = +\infty$ for $y_i < 0$ and $s_i = 0$. To avoid such divergence, past studies of the OP approach for emission scans have used zero-thresholded values as follows [83, 133]:

$$h_i^{\text{OP}^+}(l) = [y_i]_+ \log(l + s_i) - (l + s_i), \quad (6.9)$$

called the “OP⁺” approximation in this dissertation. (Note the slightly different use of terms from [133].) The zero-thresholding is natural in view of the nonnegative nature of Poisson random variables in (6.7). Moreover, it guarantees the concavity of $h_i^{\text{OP}^+}$, and hence the existence and uniqueness of the penalized-likelihood reconstruction under mild conditions [4]. However, zero-thresholding destroys the first moment matching in (6.7), and the increase of the precorrected data by zero-thresholding causes the estimators to have a positive systematic bias since emission data is linearly related to activity in the mean. Section 6.4 shows that the seemingly unnatural use of negative sinogram values in the Poisson framework can alleviate the systematic bias problem of OP⁺.

²The minus sign signifies that this approximation allows negative precorrected data $y_i < 0$.

6.3.2 Shifted Poisson (SP) Approximation

An improved approximation is to match both the first and second moments as follows:

$$Y_i + 2r_i \stackrel{\text{SP}}{\underset{\text{approx.}}{\sim}} \text{Poisson}\{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i + 2r_i\}, \quad (6.10)$$

where in practice one must use an estimate \hat{r}_i . This ‘‘SP⁻’’ approximation³ leads to a log-likelihood function L^{SP^-} of the form (6.4) with

$$h_i^{\text{SP}^-}(l) = (y_i + 2r_i) \log(l + s_i + 2r_i) - (l + s_i + 2r_i). \quad (6.11)$$

Similarly, its conventional zero-thresholded version L^{SP^+} uses [133]

$$h_i^{\text{SP}^+}(l) = [y_i + 2r_i]_+ \log(l + s_i + 2r_i) - (l + s_i + 2r_i). \quad (6.12)$$

The zero-thresholding again ensures the concavity of L^{SP^+} but also causes positive systematic bias, albeit generally less than that of OP^+ since it is more likely that $y_i < 0$ than $y_i + 2r_i < 0$. Section 6.4 describes the details.

6.3.3 Saddle-Point (SD) Approximation

Another approach is to make a second order Taylor series approximation in the z -transform domain to the probability generating function and then carry out the inverse transform [108,116]. The log-likelihood L^{SD} corresponding to this SD approximation [133] is of the form (6.4) with

$$h_i^{\text{SD}}(l) = y_i \log\left(\frac{l + s_i + r_i}{z_i + u_i(l)}\right) - l + u_i(l) - \frac{1}{2} \log u_i(l) \quad (6.13)$$

where

$$z_i = \begin{cases} y_i + 1, & \text{for } y_i \geq 0 \\ y_i - 1, & \text{for } y_i < 0 \end{cases}$$

³The minus sign signifies that this approximation allows $y_i + 2r_i < 0$.

and

$$u_i(l) = \sqrt{z_i^2 + 4(l + r_i + s_i)r_i}.$$

The SD model for emission image reconstruction is free of systematic bias and leads to lower variance than OP⁺ [133]. Indeed, in all cases studied to date, the SD model has shown the best performance for randoms-precorrected PET emission reconstruction. We observe those properties empirically in Section 6.7. However, the new SP⁻, despite its simpler implementation, performs comparably to SD.

6.3.4 Log-likelihood for Prompt Data

If one has access to the prompt data $\mathbf{Y}^{\text{prompt}}$, then one can use the log-likelihood for the prompt data in the form (6.4) with

$$h_i^{\text{PR}}(l) = y_i^{\text{prompt}} \log(l + s_i + r_i) - (l + s_i + r_i).$$

We include this PR model for comparing the bias and variance of the methods for randoms-precorrected data in Section 6.7. Since $\mathbf{Y}^{\text{prompt}}$ has lower variance than \mathbf{Y} , it serves as a baseline for comparing algorithms.

6.4 Effects of Zero-thresholding on Bias

The sinogram zero-thresholding in (6.9) and (6.12) increases the mean values of the data. This section analyzes the effects of this shift.

First, we focus on a single ray to investigate the properties of OP⁺ and SP⁺. Let Y be a precorrected measurement modeled as the difference of two independent Poisson random variables as follows:

$$Y \triangleq \text{Poisson}\{\theta + r\} - \text{Poisson}\{r\} \tag{6.14}$$

where θ and r denote the mean number of trues (possibly including scatters) and AC

events, respectively. The normalized effective means of trues are given by

$$m^{\text{OP}^+} = \frac{E\{[Y]_+\}}{\theta} \quad \text{for OP}^+,$$

and

$$m^{\text{SP}^+} = \frac{E\{[Y + 2r]_+\} - 2r}{\theta} \quad \text{for SP}^+,$$

as a function of θ . Note that m^{OP^+} and m^{SP^+} would be unity without zero-thresholding.

We calculated these expectations using the Bessel function expression for the probability mass function for Y [129, Sec. 3.2] as follows:

$$\begin{aligned} P(Y = y; \theta) &= \sum_{m=[-y]_+}^{\infty} \frac{(\theta + r)^{y+m} e^{-(\theta+r)} r^m e^{-r}}{(y+m)! m!} \\ &= \frac{e^{-(\theta+2r)}}{i^{|y|}} \left(\sqrt{\frac{\theta+r}{r}} \right)^y J_{|y|} \left(2i\sqrt{(\theta+r)r} \right) \end{aligned}$$

where $i = \sqrt{-1}$ and $J_n(\cdot)$ is the Bessel function of the first kind of order n [10, p. 575].

Fig. 6.1 shows the results, from which we infer that 1) for counts per ray higher than 10 (or 1), there is little effect of zero-thresholding for OP^+ (or SP^+) as long as randoms amount to less than 100% of trues, and 2) OP^+ generally leads to a higher bias than SP^+ . For extremely low counts or low AC rates, OP^+ and SP^+ yield similar results.

Next, we consider a one-parameter example for the OP approach to illustrate how allowing negative values can be helpful for reducing systematic positive bias. Let the measurements be the difference of two independent Poisson random variables as follows:

$$Z_i \triangleq \text{Poisson}\{a_i\theta^{\text{true}} + s_i + r_i\} - \text{Poisson}\{r_i\} \quad (6.15)$$

for $i = 1, \dots, N$. Setting $s_i = 0$, ML estimates based on OP^- and OP^+ models [see (6.4) with (6.8) and (6.9)] with an image nonnegativity constraint have the following analytical solutions:

$$\hat{\theta}^{\text{OP}^-} = \frac{\left[\sum_{i=1}^N z_i \right]_+}{\sum_{i=1}^N a_i} \quad (6.16)$$

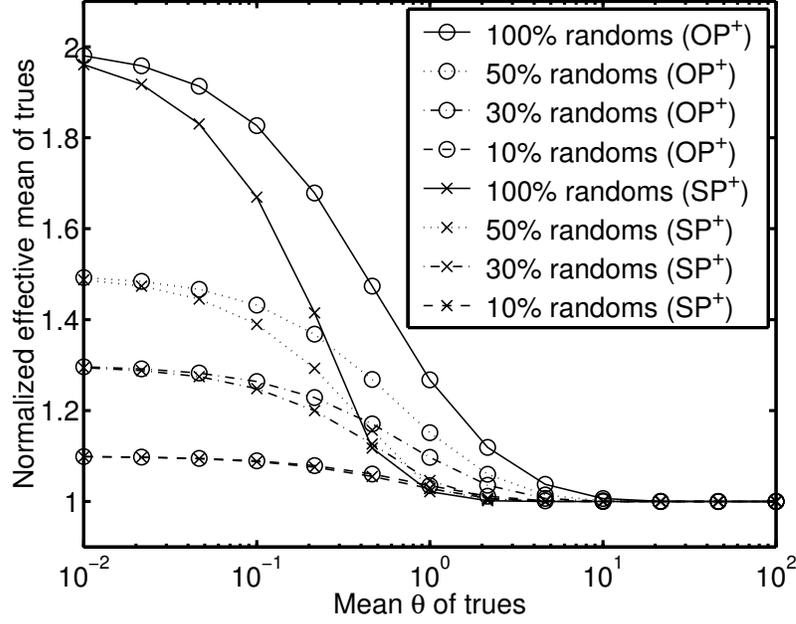


Figure 6.1: Effective means of trues increased by zero-thresholding for OP^+ and SP^+ . In this figure, randoms fractions in % denote $(\text{mean of randoms})/(\text{mean of trues})$.

and

$$\hat{\theta}^{OP^+} = \frac{\sum_{i=1}^N [z_i]_+}{\sum_{i=1}^N a_i}. \quad (6.17)$$

Note the zero-thresholding in (6.16) is due to the *image-domain* nonnegativity constraint and not a primary source of the positive bias, whereas the zero-thresholding $[z_i]_+$ in (6.17) is in the *sinogram domain* from (6.9). In view of Fig. 6.1, $\hat{\theta}^{OP^+}$ is more biased than $\hat{\theta}^{OP^-}$. For instance, if $\theta^{\text{true}} = 1$, $a_i = 1$, $r_i = 0.5$ for all i , and $N = 10$, then the estimator biases can be computed using Fig. 6.1 with (6.16) and (6.17) as follows:

$$b^{OP^-} = E\{\hat{\theta}^{OP^-}\} - \theta^{\text{true}} = 1.014 - 1 = 0.014$$

$$b^{OP^+} = E\{\hat{\theta}^{OP^+}\} - \theta^{\text{true}} = 1.152 - 1 = 0.152.$$

So the OP^- model reduces significantly the positive bias in OP^+ . This example suggests that when the rays passing through a particular voxel have low counts but high AC rates, OP^+ will yield a higher positive systematic bias than OP^- . The comparison of SP^- and SP^+ would be similar although there are no closed-form estimators for SP^- and SP^+ like

(6.16) and (6.17).

For *high counts* per ray cases, sinogram zero-thresholding is not problematic since the probability of negative values is greatly reduced. The next section investigates the asymptotic behavior of the estimators for high counts.

6.5 Asymptotic Analysis

This section analyzes the *asymptotic* bias and covariance of OP^- , SP^- , and SD estimators for *high counts* (per ray) cases. The purpose of the analysis is to compare the estimator properties rather than to accurately predict estimator behavior. (The prediction of the mean and covariance of PL or ML estimators could be conducted following [36]; see [132] for such analysis for randoms-precorrected PET *transmission* scans.) We focus on ML estimators for simplicity. We do not consider OP^+ and SP^+ since they should behave quite similarly to OP^- and SP^- , respectively, for high counts per ray.

6.5.1 Asymptotic Bias and Asymptotic Covariance

Let the *precorrected* measurement $\mathbf{Y}^n = [Y_1^n, \dots, Y_N^n]'$ be such that

$$Y_i^n \triangleq \text{Poisson}\{n([\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + r_i + s_i)\} - \text{Poisson}\{nr_i\} \quad (6.18)$$

for $i = 1, \dots, N$ where $n \in \mathbb{N}$ represents a factor proportional to the number of total counts or the scan time. Define $\bar{\mathbf{Y}}^n \triangleq \mathbf{Y}^n/n$, then⁴

$$\bar{\mathbf{Y}}^n \xrightarrow{p} \bar{\mathbf{Y}} \quad (6.19)$$

as $n \rightarrow \infty$ by the weak law of large numbers [105, p. 112] where

$$\bar{\mathbf{Y}} = E\{\mathbf{Y}^1\} = \mathbf{A}\boldsymbol{\lambda}^{\text{true}} + \mathbf{s} \quad (6.20)$$

⁴Note that \mathbf{Y}^n is identical to the sum of n iid random vectors each of which is identical to \mathbf{Y}^1 .

and “ \xrightarrow{p} ” denotes convergence in probability. Also,

$$\sqrt{n}(\bar{\mathbf{Y}}^n - \bar{\mathbf{Y}}) \xrightarrow{\mathcal{L}} \mathcal{N}(\mathbf{0}, \text{Cov}\{\mathbf{Y}^1\}) \quad (6.21)$$

as $n \rightarrow \infty$ by the central limit theorem [74, p. 61] where

$$\text{Cov}\{\mathbf{Y}^1\} = \text{diag}\{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i + 2r_i\} \quad (6.22)$$

and “ $\xrightarrow{\mathcal{L}}$ ” denotes convergence in law (or distribution). Because $h_i^{\text{OP}^-}$ in (6.8) is affine in \mathbf{Y} , we can write the OP^- estimate based on \mathbf{Y}^n as follows:

$$\begin{aligned} \arg \max_{\boldsymbol{\lambda} \geq \mathbf{0}} L^{\text{OP}^-}(\boldsymbol{\lambda}; \mathbf{Y}^n) &= \arg \max_{\boldsymbol{\lambda} \geq \mathbf{0}} \left\{ \frac{L^{\text{OP}^-}(\boldsymbol{\lambda}; \mathbf{Y}^n)}{n} \right\} \\ &= \arg \max_{\boldsymbol{\lambda} \geq \mathbf{0}} L^{\text{OP}^-}(\boldsymbol{\lambda}; \bar{\mathbf{Y}}^n) \end{aligned} \quad (6.23)$$

$$\triangleq \hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\bar{\mathbf{Y}}^n) \quad (6.24)$$

where L^{OP^-} is of the form (6.4) with (6.8). One can show $L^{\text{OP}^-}(\boldsymbol{\lambda}^{\text{true}}; \bar{\mathbf{Y}}) \geq L^{\text{OP}^-}(\boldsymbol{\lambda}; \bar{\mathbf{Y}})$ for all $\boldsymbol{\lambda} \geq \mathbf{0}$ where $\bar{\mathbf{Y}}$ is defined in (6.20) since $h_i^{\text{OP}^-}(l)$ attains a maximum over $l \geq 0$ at $l = [y_i - s_i]_+$. We assume that the $N \times p$ system matrix \mathbf{A} has full column rank, ensuring uniqueness of the noiseless reconstruction $\hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\bar{\mathbf{Y}}) = \boldsymbol{\lambda}^{\text{true}}$.

One can easily show that

$$\nabla^{10} L^{\text{OP}^-}(\boldsymbol{\lambda}^{\text{true}}; \bar{\mathbf{Y}}) = \mathbf{0}$$

and that

$$\nabla^{20} L^{\text{OP}^-}(\boldsymbol{\lambda}^{\text{true}}; \bar{\mathbf{Y}}) = \mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i} \right\} \mathbf{A}$$

is positive definite since \mathbf{A} has full column rank where $\nabla^{10} = [\frac{\partial}{\partial \lambda_1}, \dots, \frac{\partial}{\partial \lambda_p}]'$ and ∇^{20} denote the column gradient operator and the Hessian operator, respectively. Then $\hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\cdot)$ is continuously differentiable at $\bar{\mathbf{Y}}$ by the implicit function theorem [15, p. 668].

Since $\hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\cdot)$ is continuous at $\bar{\mathbf{Y}}$, one can show [105, p. 124]

$$\hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\bar{\mathbf{Y}}^n) \xrightarrow{p} \hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\bar{\mathbf{Y}}) = \boldsymbol{\lambda}^{\text{true}} \quad (6.25)$$

as $n \rightarrow \infty$, in view of (6.19). In other words, the OP^- estimator (and OP^+ as well) is *asymptotically unbiased*.

Next, we investigate the asymptotic variance for OP^- . Since $\hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\cdot)$ is continuously differentiable in a neighborhood of $\bar{\mathbf{Y}}$, it can be shown by the Delta method [74, p. 61], in view of (6.21), that

$$\sqrt{n} \left(\hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\bar{\mathbf{Y}}^n) - \boldsymbol{\lambda}^{\text{true}} \right) \xrightarrow{\mathcal{L}} \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}^{\text{OP}^-})$$

with

$$\boldsymbol{\Sigma}^{\text{OP}^-} = \nabla \hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\bar{\mathbf{Y}}) \text{Cov}\{\mathbf{Y}^1\} [\nabla \hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\bar{\mathbf{Y}})]' \quad (6.26)$$

where $\nabla = [\frac{\partial}{\partial Y_1}, \dots, \frac{\partial}{\partial Y_N}]$ denotes the row gradient operator. The gradient $\nabla \hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\bar{\mathbf{Y}})$ of the implicitly defined function (6.24) can be computed as in [36]. Assuming that⁵ $\boldsymbol{\lambda}^{\text{true}} > \mathbf{0}$, one obtains

$$\begin{aligned} \nabla \hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\bar{\mathbf{Y}}) &= \left[-\nabla^{20} L^{\text{OP}^-}(\hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\bar{\mathbf{Y}}); \bar{\mathbf{Y}}) \right]^{-1} \nabla^{11} L^{\text{OP}^-}(\hat{\boldsymbol{\lambda}}^{\text{OP}^-}(\bar{\mathbf{Y}}); \bar{\mathbf{Y}}) \quad (6.27) \\ &= \left[-\nabla^{20} L^{\text{OP}^-}(\boldsymbol{\lambda}^{\text{true}}; \bar{\mathbf{Y}}) \right]^{-1} \nabla^{11} L^{\text{OP}^-}(\boldsymbol{\lambda}^{\text{true}}; \bar{\mathbf{Y}}) \\ &= - \left[\mathbf{A}' \text{diag} \left\{ \frac{\bar{Y}_i}{([\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i)^2} \right\} \mathbf{A} \right]^{-1} \mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i} \right\} \\ &= - \left[\mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i} \right\} \mathbf{A} \right]^{-1} \mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i} \right\} \end{aligned}$$

due to (6.20) and (6.25) where the (j, i) th element of the $p \times N$ operator ∇^{11} is $\frac{\partial^2}{\partial \lambda_j \partial Y_i}$.

Now, (6.26) can be written as

$$\boldsymbol{\Sigma}^{\text{OP}^-} = \mathbf{F}_{\text{OP}^-}^{-1} \mathbf{A}' \text{diag} \left\{ \frac{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i + 2r_i}{([\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i)^2} \right\} \mathbf{A} \mathbf{F}_{\text{OP}^-}^{-1} \quad (6.28)$$

where

$$\mathbf{F}_{\text{OP}^-} = \mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i} \right\} \mathbf{A}.$$

⁵Although the analysis method here does not apply to a case where some $\lambda_j^{\text{true}} = 0$, it provides reasonably accurate prediction of the covariance of an implicitly defined estimator [36].

Similarly, one can show that the SP^- method (and SP^+ as well) is also *asymptotically unbiased* and that its asymptotic covariance is

$$\Sigma^{\text{SP}^-} = \left[\mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i + 2r_i} \right\} \mathbf{A} \right]^{-1}. \quad (6.29)$$

To analyze the SD case, one needs the following approximation that, from (6.13), is valid for large n :

$$\frac{L^{\text{SD}}(\boldsymbol{\lambda}; \mathbf{Y}^n)}{n} \approx L^{\text{SD}'}(\boldsymbol{\lambda}; \bar{\mathbf{Y}}^n) \quad (6.30)$$

with

$$L^{\text{SD}'}(\boldsymbol{\lambda}; \mathbf{Y}) = \sum_{i=1}^N h_i^{\text{SD}'}(l_i(\boldsymbol{\lambda}))$$

where

$$h_i^{\text{SD}'}(l) = y_i \log \left(\frac{l + s_i + r_i}{y_i + \tilde{u}_i(l)} \right) - l + \tilde{u}_i(l)$$

and

$$\tilde{u}_i(l) = \sqrt{y_i^2 + 4(l + r_i + s_i)r_i}.$$

The SD estimate can be written as follows:

$$\begin{aligned} \arg \max_{\boldsymbol{\lambda} \geq \mathbf{0}} L^{\text{SD}}(\boldsymbol{\lambda}; \mathbf{Y}^n) &= \arg \max_{\boldsymbol{\lambda} \geq \mathbf{0}} \left\{ \frac{L^{\text{SD}}(\boldsymbol{\lambda}; \mathbf{Y}^n)}{n} \right\} \\ &\approx \arg \max_{\boldsymbol{\lambda} \geq \mathbf{0}} L^{\text{SD}'}(\boldsymbol{\lambda}; \bar{\mathbf{Y}}^n) \\ &\triangleq \hat{\boldsymbol{\lambda}}^{\text{SD}'}(\bar{\mathbf{Y}}^n). \end{aligned}$$

Since 1) $L^{\text{SD}'}(\boldsymbol{\lambda}; \bar{\mathbf{Y}})$ has the unique maximizer (over $\boldsymbol{\lambda} \geq \mathbf{0}$), 2) $\hat{\boldsymbol{\lambda}}^{\text{SD}'}(\bar{\mathbf{Y}}) = \boldsymbol{\lambda}^{\text{true}}$ (note $L^{\text{SD}'}$ can be shown to be strictly concave) and 3) the approximation (6.30) becomes more accurate as n increases, it can be shown that the SD method is also *asymptotically unbiased*. By similar manipulations, one can obtain the asymptotic covariance for SD as follows:

$$\Sigma^{\text{SD}} \approx \Sigma^{\text{SD}'} = \left[\mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i + 2r_i} \right\} \mathbf{A} \right]^{-1}, \quad (6.31)$$

which is equal to (6.29).

Both SP^- and SD are *asymptotically efficient* in the following sense. Noting that $\bar{\mathbf{Y}}^n$ is asymptotically normal with mean $\bar{\mathbf{Y}}$ and covariance $\text{Cov}\{\mathbf{Y}^1\}/n$ from (6.21), one can obtain the Cramér-Rao bound (CRB) from the asymptotic normal likelihood as follows:

$$\begin{aligned} \mathbf{B}(\boldsymbol{\lambda}^{\text{true}}) &\approx \frac{1}{n} [\mathbf{A}' \text{Cov}\{\mathbf{Y}^1\}^{-1} \mathbf{A}]^{-1} \\ &= \frac{1}{n} \left[\mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + s_i + 2r_i} \right\} \mathbf{A} \right]^{-1}. \end{aligned}$$

Now one can see SP^- and SD asymptotically achieve this bound from (6.29) and (6.31).

Note that the exact CRB appears intractable due to form of (6.5).

The reasons that SP^- and SD are asymptotically efficient are as follows. First, for SP^- , the precorrected data are modeled as the Poisson approximation in (6.10) that matches the first and second moments, so the SP model approaches the asymptotic normal distribution in (6.21) of the precorrected data in (6.18) for large n . Intuitively, this suggests that SP estimators approach ML estimators for large n , and consequently, they are asymptotically efficient. Next, noting that the saddle-point approximation of a normal variate is exact, one could also expect SD estimators to approach ML estimators asymptotically; so they should also be asymptotically efficient.

To summarize, we have shown that all OP^- , SP^- and SD are asymptotically unbiased, and we have derived their asymptotic covariances (6.28), (6.29) and (6.31). We have also shown that SP^- and SD are asymptotically efficient.

6.5.2 Comparison with Weighted Least Squares Methods

We compare SP^- and SD with some weighted least squares (WLS) methods⁶ in a way similar to the approach in Section 6.5.1. As an alternative to Poisson models, in view of

⁶We are not concerned about implementation issues, nonnegativity constraints, or rigorous convergence proofs for WLS methods. Our aim here is to obtain some insights into the asymptotic behavior of SP^- and SD compared to those of WLS methods.

$\bar{\mathbf{Y}}^n$ being (asymptotically normal) with mean $\bar{\mathbf{Y}}$ in (6.20) and covariance $\text{Cov}\{\mathbf{Y}^1\}/n$ in (6.22), a natural “model-weighted” least squares (MWLS) objective function is

$$\Phi^{\text{MWLS}}(\boldsymbol{\lambda}; \bar{\mathbf{Y}}^n) = -n (\bar{\mathbf{Y}}^n - (\mathbf{A}\boldsymbol{\lambda} + \mathbf{s}))' \mathbf{W}(\boldsymbol{\lambda}) (\bar{\mathbf{Y}}^n - (\mathbf{A}\boldsymbol{\lambda} + \mathbf{s})) \quad (6.32)$$

where

$$\mathbf{W}(\boldsymbol{\lambda}) = \text{diag} \left\{ \frac{1}{[\mathbf{A}\boldsymbol{\lambda}]_i + 2r_i + s_i} \right\}.$$

Define

$$\hat{\boldsymbol{\lambda}}^{\text{MWLS}}(\mathbf{Y}) \triangleq \arg \max_{\boldsymbol{\lambda}} \Phi^{\text{MWLS}}(\boldsymbol{\lambda}; \mathbf{Y}).$$

Assuming that \mathbf{A} is of full column rank, one can compute the unique noiseless estimate $\hat{\boldsymbol{\lambda}}^{\text{MWLS}}(\bar{\mathbf{Y}}) = \boldsymbol{\lambda}^{\text{true}}$ where $\bar{\mathbf{Y}}$ is defined in (6.20). Then it follows that

$$\hat{\boldsymbol{\lambda}}^{\text{MWLS}}(\bar{\mathbf{Y}}^n) \xrightarrow{p} \boldsymbol{\lambda}^{\text{true}}$$

when $n \rightarrow \infty$, as in (6.25). That is, MWLS is also asymptotically unbiased. Similar to (6.26), the asymptotic covariance for MWLS is given by

$$\boldsymbol{\Sigma}^{\text{MWLS}} = \nabla \hat{\boldsymbol{\lambda}}^{\text{MWLS}}(\bar{\mathbf{Y}}) \text{Cov}\{\mathbf{Y}^1\} [\nabla \hat{\boldsymbol{\lambda}}^{\text{MWLS}}(\bar{\mathbf{Y}})]'. \quad (6.33)$$

But one can compute

$$\begin{aligned} \nabla \hat{\boldsymbol{\lambda}}^{\text{MWLS}}(\bar{\mathbf{Y}}) &= \left[-\nabla^2 \Phi^{\text{MWLS}}(\hat{\boldsymbol{\lambda}}^{\text{MWLS}}(\bar{\mathbf{Y}}); \bar{\mathbf{Y}}) \right]^{-1} \nabla^1 \Phi^{\text{MWLS}}(\hat{\boldsymbol{\lambda}}^{\text{MWLS}}(\bar{\mathbf{Y}}); \bar{\mathbf{Y}}) \\ &= - \left[\mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + 2r_i + s_i} \right\} \mathbf{A} \right]^{-1} \\ &\quad \mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + 2r_i + s_i} \right\} \end{aligned} \quad (6.34)$$

as in (6.27). Combining (6.22), (6.33), and (6.34) yields

$$\boldsymbol{\Sigma}^{\text{MWLS}} = \left[\mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + 2r_i + s_i} \right\} \mathbf{A} \right]^{-1}, \quad (6.35)$$

which is equal to (6.29) and (6.31).

Now we consider an iterative reweighted least squares (IRLS) method [43]. Since the objective function in (6.32) is difficult to maximize due to its nonlinearity, one could take the IRLS approach by using the previous iterate for determining the current weight:

$$\boldsymbol{\lambda}^{k+1} = \arg \min_{\boldsymbol{\lambda}} (\bar{\mathbf{Y}}^n - (\mathbf{A}\boldsymbol{\lambda} + \mathbf{s}))' \mathbf{W}(\boldsymbol{\lambda}^k) (\bar{\mathbf{Y}}^n - (\mathbf{A}\boldsymbol{\lambda} + \mathbf{s})).$$

Ignoring the nonnegativity constraint, one has the following analytical solution to the above optimization problem:

$$\boldsymbol{\lambda}^{k+1} = (\mathbf{A}'\mathbf{W}(\boldsymbol{\lambda}^k)\mathbf{A})^{-1} \mathbf{A}'\mathbf{W}(\boldsymbol{\lambda}^k)(\bar{\mathbf{Y}}^n - \mathbf{s}). \quad (6.36)$$

Assuming a sequence $\{\boldsymbol{\lambda}^k\}$ generated by (6.36) converges to $\hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\bar{\mathbf{Y}}^n)$, one obtains the following relationship:

$$\hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\bar{\mathbf{Y}}^n) = (\mathbf{A}'\mathbf{W}(\hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\bar{\mathbf{Y}}^n))\mathbf{A})^{-1} \mathbf{A}'\mathbf{W}(\hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\bar{\mathbf{Y}}^n))(\bar{\mathbf{Y}}^n - \mathbf{s}). \quad (6.37)$$

Since $\hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\bar{\mathbf{Y}}) = \boldsymbol{\lambda}^{\text{true}}$, IRLS is asymptotically unbiased, too. Define

$$\Psi(\hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\mathbf{Y}); \mathbf{Y}) \triangleq \mathbf{A}'\mathbf{W}(\hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\mathbf{Y})) (\mathbf{Y} - (\mathbf{A}\hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\mathbf{Y}) + \mathbf{s})) = \mathbf{0} \quad (6.38)$$

where the equality is due to (6.37). Differentiating (6.38) with respect to \mathbf{Y} yields

$$\begin{aligned} \nabla \hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\mathbf{Y}) &= \left[-\nabla^{10} \Psi(\hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\mathbf{Y}); \mathbf{Y}) \right]^{-1} \nabla^{01} \Psi(\hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\mathbf{Y}); \mathbf{Y}) \\ &= - \left[\mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\mathbf{Y})]_i + 2r_i + s_i} \right\} \mathbf{A} \right]^{-1} \\ &\quad \mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\mathbf{Y})]_i + 2r_i + s_i} \right\} \end{aligned} \quad (6.39)$$

where ∇^{10} and ∇^{01} denote the row gradient operator with respect to the first and the second argument of $\Psi(\cdot; \cdot)$, respectively. Now one can compute the asymptotic covariance for IRLS by using (6.22) and (6.39) as follows:

$$\begin{aligned} \boldsymbol{\Sigma}^{\text{IRLS}} &= \nabla \hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\bar{\mathbf{Y}}) \text{Cov}\{\mathbf{Y}^1\} [\nabla \hat{\boldsymbol{\lambda}}^{\text{IRLS}}(\bar{\mathbf{Y}})]' \\ &= \left[\mathbf{A}' \text{diag} \left\{ \frac{1}{[\mathbf{A}\boldsymbol{\lambda}^{\text{true}}]_i + 2r_i + s_i} \right\} \mathbf{A} \right]^{-1}, \end{aligned}$$

which is equal to (6.29) and (6.31) as well as (6.35).

To summarize, SP^- and SD show the same asymptotic behavior as WLS methods such as MWLS and IRLS.

6.5.3 Comparison of Covariances

We compare the asymptotic variances of OP^- and SP^- (equivalently, SD) estimators for a simple one-parameter example introduced in (6.15) (see [130, 131] for a similar comparison in a *transmission* case). Using (6.28) and (6.29), one obtains asymptotic variances (or approximate variances for high counts) for OP^- and SP^- :

$$\text{Var}\{\hat{\theta}^{\text{OP}^-}\} = \frac{\sum_{i=1}^N a_i^2 \frac{a_i \theta^{\text{true}} + s_i + 2r_i}{(a_i \theta^{\text{true}} + s_i)^2}}{\left(\sum_{i=1}^N \frac{a_i^2}{a_i \theta^{\text{true}} + s_i}\right)^2}$$

and

$$\text{Var}\{\hat{\theta}^{\text{SP}^-}\} = \left(\sum_{i=1}^N \frac{a_i^2}{a_i \theta^{\text{true}} + s_i + 2r_i}\right)^{-1}.$$

Using the Schwartz inequality [126, p. 107], one can show

$$\text{Var}\{\hat{\theta}^{\text{OP}^-}\} \geq \text{Var}\{\hat{\theta}^{\text{SP}^-}\} \quad (6.40)$$

where equality holds if and only if the $(a_i \theta^{\text{true}} + s_i)/(a_i \theta^{\text{true}} + s_i + 2r_i)$ ratios are equal, which is impossible as long as $r_i > 0$ and $s_i > 0$. Therefore, the inequality (6.40) is strict; the variance of OP^- is higher than that of SP^- (or SD). This is corroborated by empirical results for a multi-parameter case in Section 6.7.

6.6 Reconstruction Algorithms

After choosing a suitable likelihood approximation, one needs an algorithm to maximize the corresponding objective function for ML or PL estimation. It is straightfor-

ward to use globally convergent (and monotonic) algorithms⁷ such as SPS in Section 3.3.3 and ML-EM in Section 3.3.2 for PR, OP⁺, SP⁺ and SD, all of which have concave log-likelihoods. However, the new OP⁻ and SP⁻ models can have nonconcave log-likelihood functions when negative sinogram values are present. The algorithms need some modifications to ensure monotonicity for the nonconcave case as well. Monotonicity is one of the most desirable properties to enable at least locally optimal reconstruction.

A large class of monotonic iterative algorithms (including SPS and ML-EM) are based on the “optimization transfer principles” in Section 3.3: at each iteration we choose a surrogate function that is easier to maximize than the original objective function, and then maximize the surrogate. To ensure monotonicity, the surrogate function is chosen so that increasing the surrogate guarantees the increase of the original objective function.

The idea for extending the algorithms to allow negative y_i 's is to choose a linear surrogate when a marginal log-likelihood h_i is convex. That is, for OP⁻, if $y_i < 0$, a tangent line to $h_i^{\text{OP}^-}$ at a current iterate l_i^n (in projection domain)

$$q_i^{\text{OP}^-}(l; l_i^n) = \dot{h}_i^{\text{OP}^-}(l_i^n)(l - l_i^n) + h_i^{\text{OP}^-}(l_i^n) \quad (6.41)$$

is a proper surrogate for $h_i^{\text{OP}^-}$ in light of [1, Eq. 7] since $q_i^{\text{OP}^-}$ lies below for all $l \geq 0$ due to convexity of $h_i^{\text{OP}^-}$, as illustrated in Fig. 6.2.

The same principle applies to SP⁻ when $y_i + 2r_i < 0$. We derive modified SPS and ML-EM applicable to OP⁻ and SP⁻, using a linear surrogate (6.41) when needed.

⁷Ordered subsets algorithms [2,55] can also be used with the aim of accelerating convergence speeds at the expense of monotonicity or global convergence.

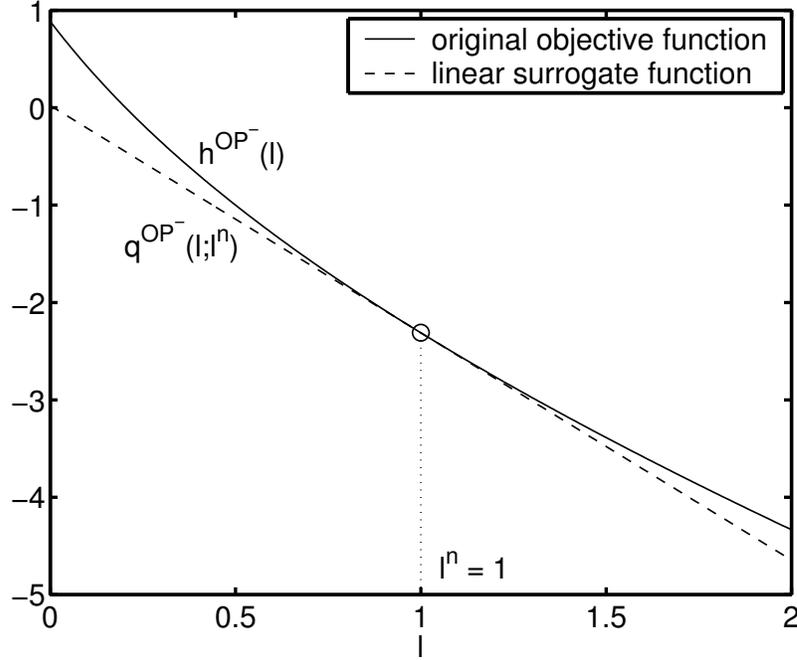


Figure 6.2: Illustration of a linear surrogate q^{OP^-} [see (6.41)] at $l^n = 1$ for an OP^- log-likelihood h^{OP^-} for a negative value $y_i < 0$. The concave surrogate q^{OP^-} lies below the objective h^{OP^-} that is convex. One can see that $q^{\text{OP}^-}(l; l^n) \geq q^{\text{OP}^-}(l^n; l^n)$ implies that $h^{\text{OP}^-}(l) \geq h^{\text{OP}^-}(l^n)$ for $l \geq 0$.

6.6.1 SPS for OP^- , SP^- , and SD

We consider the PL objective function Φ in (6.6) with a quadratic penalty for simplicity:

$$R(\boldsymbol{\lambda}) = \frac{\beta}{2} \sum_{j=1}^p \sum_{k \in \mathcal{N}_j} \omega_{jk} \frac{(\lambda_j - \lambda_k)^2}{2}. \quad (6.42)$$

A monotonic SPS method for OP^- and SP^- is readily derived following Section 3.3.3 with (6.41). The resulting algorithm differs only slightly from the ordinary SPS algorithm in (3.53), and uses the following iteration:

$$\lambda_j^{n+1} = \left[\lambda_j^n + \frac{1}{d_j^{\text{SPS}}(\boldsymbol{\lambda}^n)} \frac{\partial \Phi(\boldsymbol{\lambda}^n)}{\partial \lambda_j} \right]_+ \quad (6.43)$$

with

$$d_j^{\text{SPS}}(\boldsymbol{\lambda}) = \sum_{i=1}^N a_{ij} a_i c_i ([\mathbf{A}\boldsymbol{\lambda}]_i) + 2\beta \sum_{k \in \mathcal{N}_j} \omega_{jk} \quad (6.44)$$

where $a_i = \sum_{j=1}^p a_{ij}$ and the optimum curvatures are

$$c_i(l) = \begin{cases} \frac{2[h_i(l) - h_i(0) - l\dot{h}_i(l)]}{l^2}, & l > 0, x_i > 0 \\ -\ddot{h}_i(0), & l = 0, x_i > 0 \\ 0, & x_i \leq 0, \end{cases} \quad (6.45)$$

in which h_i represents $h_i^{\text{OP}^-}$ for OP^- and $h_i^{\text{SP}^-}$ for SP^- , and we define

$$x_i \triangleq \begin{cases} y_i, & \text{for } \text{OP}^- \\ y_i + 2r_i, & \text{for } \text{SP}^-. \end{cases} \quad (6.46)$$

The only difference from the ordinary SPS method with optimum curvatures in Section 3.3.3 is that here we set c_i to zero for $x_i < 0$ in (6.45) (note that $x_i < 0$ never occurs for previous zero-thresholding or prompt models). So one can easily modify existing codes to apply OP^- or SP^- . Nonquadratic penalties are included as in [2].

Being constructed by the optimization transfer principle, the iteration (6.43) increases monotonically the objective function Φ every iteration. Since the step (6.44) requires an “extra” backprojection, we often forego strict monotonicity by replacing the curvatures c_i with the following “precomputed curvatures,”

$$c_i(l) = \begin{cases} -\ddot{h}_i(\hat{l}_i), & x_i > 0 \\ 0, & x_i \leq 0 \end{cases}$$

where $\hat{l}_i = \arg \max_{l \geq 0} h_i(l) = [y_i - s_i]_+$. This allows d_j^{SPS} to be computed prior to iterating.

Paraboloidal surrogates algorithms for SD were developed in [129, Sec. 5.6]. A monotonic SPS version has the form of (6.43) and (6.44) with the following curvatures,

$$c_i(l) = \begin{cases} -\ddot{h}_i^{\text{SD}}(l_i^*), & y_i = 0, l_i^* > 0 \\ -\ddot{h}_i^{\text{SD}}(l_i^{**}), & y_i = -1, l_i^{**} > 0 \\ t_i(l), & \text{otherwise} \end{cases}$$

and

$$t_i(l) = \begin{cases} \frac{2[h_i^{\text{SD}}(l) - h_i^{\text{SD}}(0) - l\dot{h}_i^{\text{SD}}(l)]}{l^2}, & l > 0 \\ -\ddot{h}_i^{\text{SD}}(0), & l = 0 \end{cases}$$

with $l_i^* = 7/9 - 4r_i(r_i + s_i)$ and $l_i^{**} = x_0^2 - 1 - r_i(r_i + s_i)$ where $x_0 \approx -1.1193219$ is a root of a polynomial [129, Appendix E]. In this case for SD, the following precomputed curvatures can be used for saving computation at the expense of monotonicity,

$$c_i(l) = -\ddot{h}_i^{\text{SD}}(\hat{l}_i),$$

where $\hat{l}_i = \arg \max_{l \geq 0} h_i^{\text{SD}}(l)$, or one could use a simple estimate $\hat{l}_i \approx [y_i - s_i]_+$.

6.6.2 Variation of ML-EM for OP⁻ and SP⁻

Following the derivation of ML-EM in Section 3.3.2 leads to a variation of ML-EM for OP⁻ and SP⁻. Although we used SPS rather than ML-EM variants for the results in Section 6.7, we provide the ML-EM variants for completeness as follows:

$$\lambda_j^{n+1} = \frac{\lambda_j^n}{d_j^{\text{EM}}(\boldsymbol{\lambda}^n)} \sum_{i=1}^N \frac{a_{ij}[x_i]_+}{\bar{x}_i(\boldsymbol{\lambda}^n)} \quad (6.47)$$

or

$$\lambda_j^{n+1} = \lambda_j^n + \frac{\lambda_j^n}{d_j^{\text{EM}}(\boldsymbol{\lambda}^n)} \frac{\partial L(\boldsymbol{\lambda}^n)}{\partial \lambda_j}$$

with

$$d_j^{\text{EM}}(\boldsymbol{\lambda}) = \sum_{i=1}^N a_{ij} \left(1 + \frac{[-x_i]_+}{\bar{x}_i(\boldsymbol{\lambda})} \right), \quad (6.48)$$

where x_i is defined in (6.46), and

$$\bar{x}_i(\boldsymbol{\lambda}) \triangleq [\mathbf{A}\boldsymbol{\lambda}]_i + b_i \quad (6.49)$$

in which

$$b_i \triangleq \begin{cases} s_i, & \text{for OP} \\ s_i + 2r_i, & \text{for SP.} \end{cases} \quad (6.50)$$

This variation of ML-EM reverts to classic ML-EM in (3.37) as a special case for non-negative sinogram values. Regularization can also be incorporated as for DPEM in Section 3.3.2. The steps (6.48) require an extra backprojection each iteration compared to classic ML-EM in (3.37).

We present the derivation of the variation of ML-EM in (6.47). Define $L^+(\boldsymbol{\lambda}; \mathbf{Y}) \triangleq \sum_{i:x_i \geq 0} h_i([\mathbf{A}\boldsymbol{\lambda}]_i)$ and $L^-(\boldsymbol{\lambda}; \mathbf{Y}) \triangleq \sum_{i:x_i < 0} h_i([\mathbf{A}\boldsymbol{\lambda}]_i)$ with $h_i(l) = x_i \log(l+b_i) - (l+b_i)$ where x_i and b_i are defined in (6.46), and (6.50), respectively. Then, by concavity of \log [30],

$$\begin{aligned} L^+(\boldsymbol{\lambda}; \mathbf{Y}) &= \sum_{i:x_i \geq 0} x_i \log \left(\sum_{j=1}^p \frac{a_{ij} \lambda_j^n}{\bar{x}_i(\boldsymbol{\lambda}^n)} \frac{\lambda_j}{\lambda_j^n} \bar{x}_i(\boldsymbol{\lambda}^n) + \frac{b_i}{\bar{x}_i(\boldsymbol{\lambda}^n)} \bar{x}_i(\boldsymbol{\lambda}^n) \right) - \bar{x}_i(\boldsymbol{\lambda}) \\ &\geq \sum_{i:x_i \geq 0} \sum_{j=1}^p \frac{x_i a_{ij} \lambda_j^n}{\bar{x}_i(\boldsymbol{\lambda}^n)} \log \lambda_j - a_{ij} \lambda_j + C^+ \\ &\triangleq Q^+(\boldsymbol{\lambda}; \boldsymbol{\lambda}^n) \end{aligned}$$

where \bar{x}_i is defined in (6.49) and C^+ is a constant with respect to $\boldsymbol{\lambda}$. On the other hand, since $h_i(\cdot)$ is convex for $x_i < 0$,

$$\begin{aligned} L^-(\boldsymbol{\lambda}; \mathbf{Y}) &\geq \sum_{i:x_i < 0} \dot{h}_i([\mathbf{A}\boldsymbol{\lambda}^n]_i)([\mathbf{A}\boldsymbol{\lambda}]_i - [\mathbf{A}\boldsymbol{\lambda}^n]_i) + h_i([\mathbf{A}\boldsymbol{\lambda}^n]_i) \\ &= \sum_{i:x_i < 0} \sum_{j=1}^p \left(\frac{x_i}{\bar{x}_i(\boldsymbol{\lambda}^n)} - 1 \right) a_{ij} \lambda_j + C^- \\ &\triangleq Q^-(\boldsymbol{\lambda}; \boldsymbol{\lambda}^n) \end{aligned}$$

where C^- is a constant with respect to $\boldsymbol{\lambda}$. Since

$$\begin{aligned} Q(\boldsymbol{\lambda}; \boldsymbol{\lambda}^n) &\triangleq Q^+(\boldsymbol{\lambda}; \boldsymbol{\lambda}^n) + Q^-(\boldsymbol{\lambda}; \boldsymbol{\lambda}^n) \\ &\leq L^+(\boldsymbol{\lambda}; \mathbf{Y}) + L^-(\boldsymbol{\lambda}; \mathbf{Y}) = L(\boldsymbol{\lambda}; \mathbf{Y}) \end{aligned}$$

and $Q(\boldsymbol{\lambda}^n; \boldsymbol{\lambda}^n) = L(\boldsymbol{\lambda}^n; \mathbf{Y})$, one can show that $Q(\cdot; \boldsymbol{\lambda}^n)$ is a minorizing surrogate for $L(\cdot; \mathbf{Y})$. The surrogate $Q(\cdot; \boldsymbol{\lambda}^n)$ can be maximized by setting its derivative to zero, and as

a result, its maximizer λ^{n+1} is calculated as (6.47). Because of the optimization transfer principle, this derivation ensures monotonicity.

The following is an alternative “intuitive” but not rigorous derivation. The partial derivatives of the log-likelihood function L^{SP^-} or L^{OP^-} at a nonnegative maximizer $\hat{\lambda}$ are, by the Karush-Kuhn-Tucker conditions [15, p. 310],

$$\frac{\partial}{\partial \lambda_j} L(\hat{\lambda}) = \sum_{i=1}^N a_{ij} \left(\frac{x_i}{\bar{x}_i(\hat{\lambda})} - 1 \right) \begin{cases} = 0, & \hat{\lambda}_j > 0 \\ \leq 0, & \hat{\lambda}_j = 0 \end{cases}$$

where x_i and \bar{x}_i are defined in (6.46) and (6.49), respectively. So, for $\hat{\lambda}_j > 0$,

$$\sum_{i=1}^N a_{ij} = \sum_{i=1}^N \frac{a_{ij}}{\bar{x}_i(\hat{\lambda})} x_i = \sum_{i=1}^N \frac{a_{ij}}{\bar{x}_i(\hat{\lambda})} ([x_i]_+ - [-x_i]_+).$$

Moving the subtracted term to the other side (*cf.* [71, 72]) leads to the following:

$$\sum_{i=1}^N a_{ij} \left(1 + \frac{[-x_i]_+}{\bar{x}_i(\hat{\lambda})} \right) = \sum_{i=1}^N \frac{a_{ij} [x_i]_+}{\bar{x}_i(\hat{\lambda})}.$$

The ratio of these terms yields the multiplicative update (6.47). In other words, $\hat{\lambda}$ is a fixed point of the iteration (6.47).

6.7 Simulations

6.7.1 Methods

To compare the bias and variance properties of the estimators (OP⁻, OP⁺, SP⁻, SP⁺ and SD), we simulated 2D PET emission scans. The PR model was also included for comparison purposes since in this simulation we had access to Y_i^{prompt} and Y_i^{delay} separately.

The synthetic emission phantom shown in Fig. 6.3 was used; its warm background, left cold disc, and right hot disc had relative emission activities of 2, 0.5, and 4, respectively. The sinograms had 192 radial bins and 120 angles uniformly sampled over 180 degrees. The system matrix was generated using ASPIRE [35]; the system geometry was approximated with 3 mm wide strip integrals and 3 mm ray spacing. We simulated nonuniform

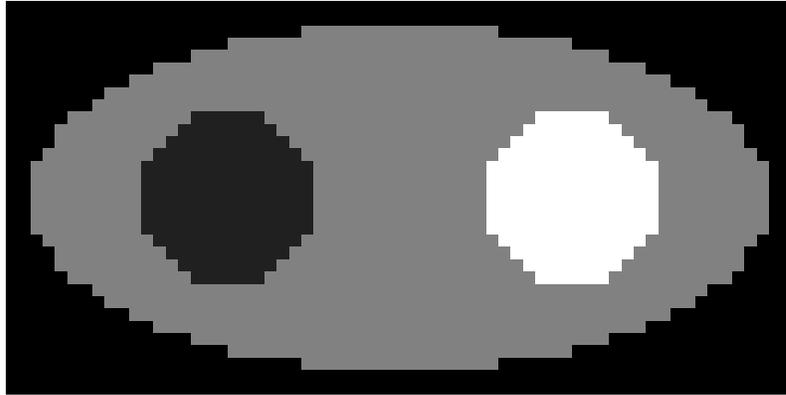


Figure 6.3: Digital phantom used in simulations. The background, left cold disc, and right hot disc have relative emission activities of 2, 0.5, and 4, respectively.

detector efficiencies using pseudo-random log-normal variates with standard deviation of 0.3. Attenuation was not considered in this simulation. The reconstructed images were 64 by 32 with 9 mm pixels. The *known* r_i and s_i factors corresponded to a uniform field of 60% randoms and 10% scatters, respectively.⁸

The specific aim of the simulation was to compare biases for low counts and to compare variances for high counts. We performed two studies with 2K and 2M total counts. We generated 500 realizations of pseudo-random emission measurements according to (6.1) with (6.2) and (6.3). For each realization, images were reconstructed using 100 iterations of the SPS method (with optimum curvature) for 2K counts, and using 40 iterations of SPS (with optimum curvature) after 10 iterations of ordered subsets SPS (with 8 subsets and precomputed curvature) [2] for 2M counts. The FBP reconstruction for each realization served as an initial image for the iterations. The number of iterations was determined by looking at objective function values over iteration for a few realizations to ensure that convergence was reasonably achieved. For initial FBP reconstructions, a Hanning filter was used with such a cut-off frequency that their impulse responses were of 3 pixels full-

⁸The fractions in % denote (mean of randoms)/(mean of trues) and (mean of scatters)/(mean of trues), respectively.

width half-maximum (FWHM).

For regularization, we used a second-order quadratic penalty in (6.42). It is important to match the spatial resolution in reconstructed images for a fair comparison of different estimators. Penalty functions can be designed to achieve spatially uniform resolution [41, 122, 123]. However, in this simulation, we used a simpler hybrid technique [122] consisting of two steps: 1) for each method, we adjusted a global regularization parameter β so that the local impulse response⁹ at the center pixel was of 1.5 pixels FWHM, and then performed PL reconstructions; 2) we applied a 2D Gaussian post-smoothing filter to the PL reconstructions so that the overall local impulse response (at the center pixel), which is the convolution of the post-smoothing filter and the original local impulse response (of 1.5 pixels FWHM), achieved a target resolution of 3 pixels FWHM.

This technique enables us to obtain reconstructions with various target resolutions by simply changing the post-smoothing filter. As the post-smoothing filter becomes wider (higher FWHM), the overall resolution becomes more uniform spatially since post-smoothing dominates the overall response. To check the spatial uniformity, the overall resolutions at every third pixel were investigated and it was found that, except the 2 pixel wide strip along the phantom boundary, each pixel achieved the target resolution (3 pixels FWHM) within 5% errors for all estimators—reasonably uniform resolution.

⁹The approximate expression for the local impulse response, which could be interpreted as the *point spread function*, of an implicitly defined estimator was given in [41, Eq. 14]. It can be computed efficiently using 2D fast Fourier Transforms by assuming local shift-invariance as in [123, Eq. 9]. All resolutions in this chapter (except those of simple linear FBP reconstruction) were computed as FWHM of the local impulse response (at a specific pixel) obtained using the methods in [41, 123]. One might doubt the feasibility of the approximate expressions in a low-counts-per-ray case where nonnegativity constraints are often active; however, in Fig. 6.4(a), the sample means of reconstructed images seem to have reasonably matched resolutions for different methods even for extremely low counts per ray (*cf.* Fig. 6.5(a) for high counts).

6.7.2 Results

Fig. 6.4(a) shows the profiles through the sample mean images of different estimators for 2K counts—very low counts. Both OP^+ and SP^+ showed large positive systematic biases particularly in the cold spot (pixels 12–25) and near the ends (phantom boundary). Zero-thresholding in sinogram domain contributes to the positive bias since the rays passing through those regions (cold spot and boundary) have low counts, as discussed in Section 6.4. Overall the systematic bias of OP^+ was slightly larger than that of SP^+ , as predicted in Section 6.4.

On the other hand, other methods (OP^- , SP^- , SD, and PR) seem reasonably free of such a bias. However, some positive biases are present in the cold spot (pixels 12–25) for OP^- , SP^- , SD and PR. The positive bias in the cold spot is mainly due to the interaction of the image-domain nonnegativity constraints and the large variances, which causes the nonnegativity constraints to be active frequently and, consequently, increases image mean values. Note that the bias is not due to zero-thresholding in sinogram domain since PR does not require any such thresholding. The coefficients of variation for those methods are larger than 100% in the cold spot in Fig. 6.4(b). In fact, the positive biases in the cold spot for OP^+ and SP^+ are caused by both sinogram-domain zero-thresholding and image-domain nonnegativity combined with large variances. Also, note that small negative biases in background and hot regions for OP^- , SP^- , SD, and PR in Fig. 6.4(a). Our hypothesis is that the positive bias in the cold spot tends to decrease image values in other regions since the reconstruction methods try to make projections of image values close to given sinogram data. However, it is hard to analyze the effects of image-domain nonnegativity constraints or to study them even experimentally since we would need, for comparison purposes, new models and algorithms (like NEG-ML in [93]) allowing negative image values. Further investigation is needed, and it would be interesting future work. To summarize, the

results showed that both SP^- and SD were comparable to PR, the baseline reconstruction, and that they were free of systematic bias *caused by sinogram-domain zero-thresholding* which appeared in SP^+ and OP^+ .

As shown in Fig. 6.4(a), FBP was nearly unbiased since image-domain nonnegativity constraints are not imposed on the FBP reconstruction. However, it showed significantly large variances in Fig. 6.4(b). Although not shown here, we found that even if FBP is constrained by image nonnegativity, it still shows larger bias and variance than other methods [5].

Fig. 6.5(a) shows the profiles through the sample mean images of different estimators for 2M counts—high counts. All of the methods are seen to be unbiased, as predicted from the analysis in Section 6.5. Fig. 6.5(b) shows profiles through the sample standard deviation images. FBP again showed the highest standard deviation and PR showed the lowest as expected. SP^- and SD showed similar performance, and OP^- led to higher standard deviation than both of them. These empirical results corroborate the analysis of asymptotic variance in Section 6.5. For each pixel, we computed the ratios of the sample standard deviation of different methods to the sample standard deviation of PR (baseline method) as shown in Fig. 6.6; and the means (over the entire image) of the ratios were 1.20 for FBP, 1.16 for OP^- , 1.11 for SP^- , and 1.12 for SD. This also supports the claim that both SP^- and SD lead to less variance than OP^- (and FBP).

The reason we did not see as clear differences in sample standard deviations of the estimators in Fig. 6.4(b) as in Fig. 6.5(b) is as follows. For low counts, the nonnegativity constraint frequently becomes active, and consequently, our asymptotic analysis based on the high count assumption will become inaccurate. In other words, there is no guarantee that PR (SD or SP) should yield less variance than OP. Additionally, the standard deviations are coupled with the biases that are distinctive in the low count case. (Generally,

there is a trade-off between bias and standard deviation.) The effects of the bias and the nonnegativity constraint seem dominant in Fig. 6.4.

Whereas SP^- performed comparably to SD, the computation time for SP^- reconstruction was shorter than SD by 3–20% (depending on curvature type and counts) for the image and sinogram size here in our C and MATLAB implementation. However, as the image and sinogram size increases, the difference in computation would become smaller since projection and backprojection operations will contribute more significantly to the computational cost.

6.8 Conclusion

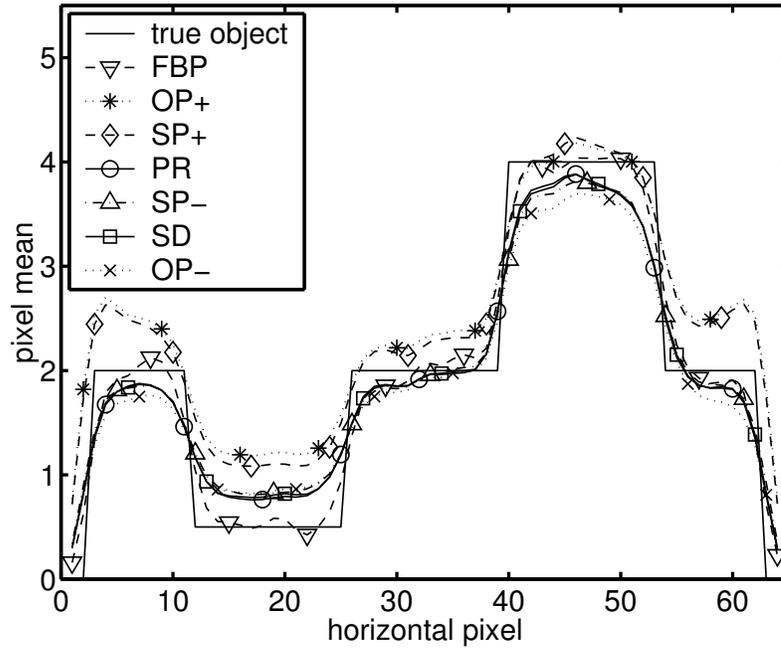
We proposed new log-likelihood approximations (SP^- and OP^-) for randoms-pre-corrected PET emission image reconstruction by allowing negative sinogram values, and we also developed algorithms (SPS and ML-EM variants) for the new models. The new methods are free of the positive systematic bias that degrades SP^+ and OP^+ images. The positive biases appearing in SP^+ and OP^+ are more distinguishable in low counts per ray regions such as cold spots, the boundary of an object, or high attenuation regions rather than depending solely on total counts. Our new models seem particularly promising for fully 3D PET emission scans where AC rates are high and photon counts per ray can be low, essentially for newer scanners with small crystals.

The new SP^- model yields less variance (than OP^- and FBP). Its performance is comparable to SD in terms of systematic bias and variance; yet its implementation is simpler. Indeed, when implemented with the usual ordered-subsets approach [4], the modified OS-SPS algorithm presented for the new SP^- model has essentially the same compute complexity as the popular OS-EM method for PET.

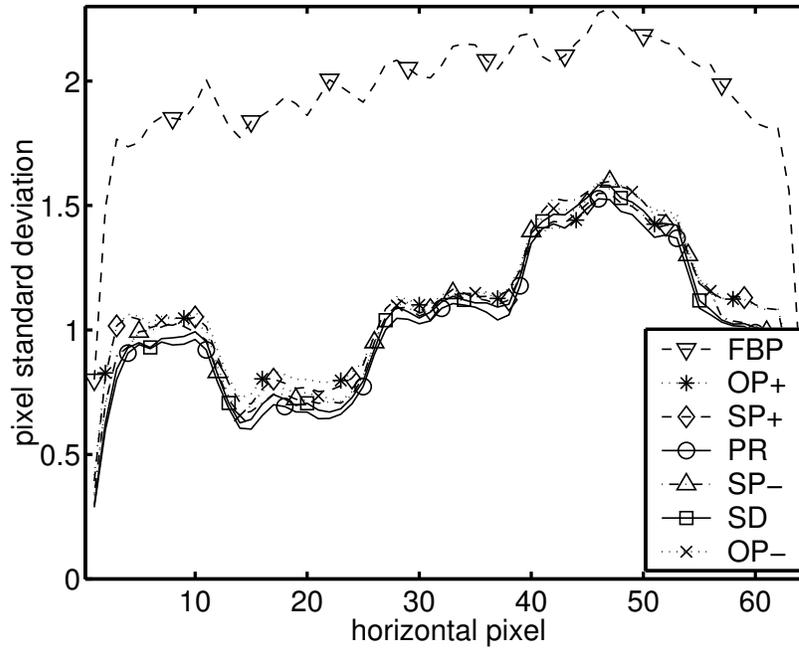
We recommend the PR method if the prompt and the randoms data are accessible sep-

arately; however, if only randoms-precorrected data are available, the new SP^- is our recommended method.

horizontal profile 17 through sample mean for 2K counts

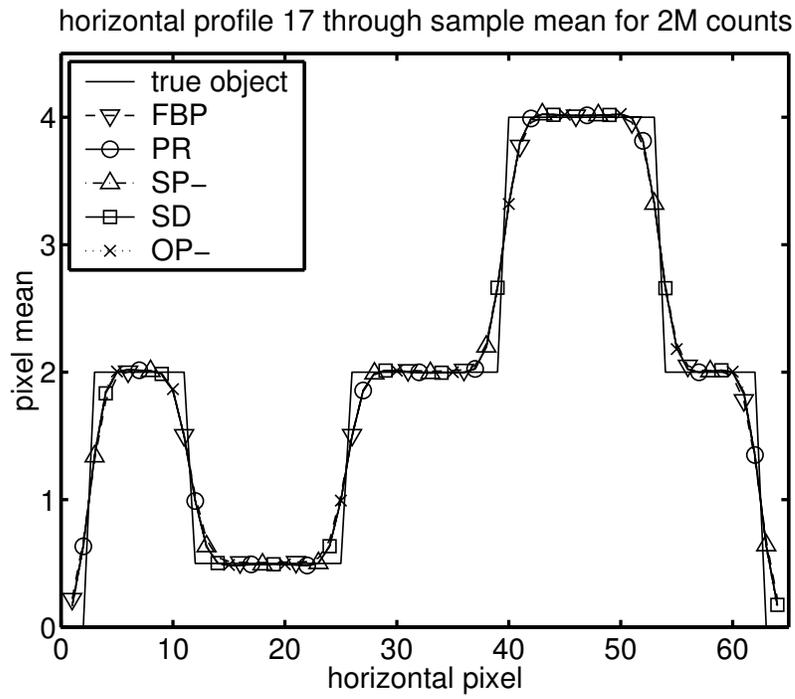


(a)

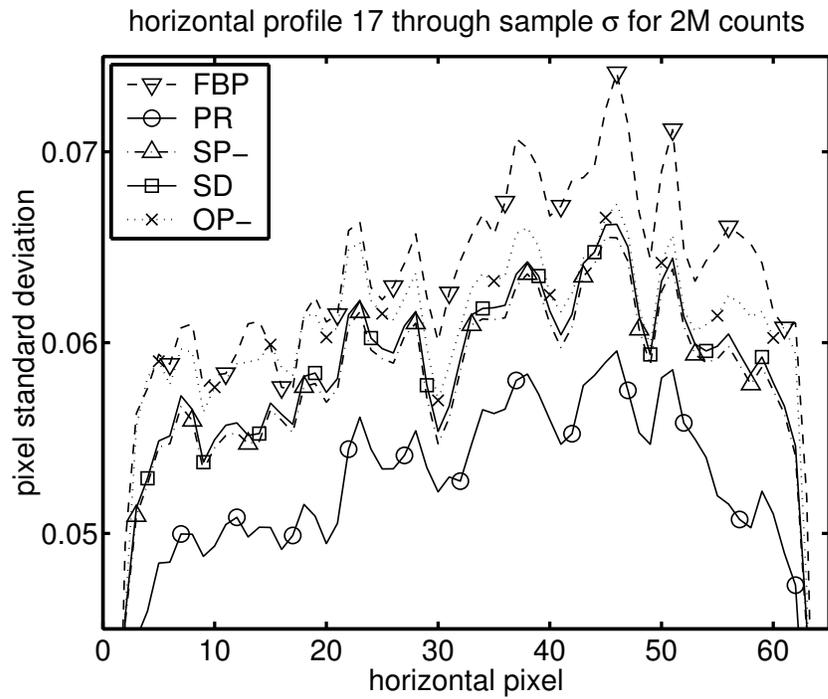
horizontal profile 17 through sample σ for 2K counts

(b)

Figure 6.4: (a) Horizontal profile through sample mean of estimators for 2K counts. (b) Horizontal profile through sample standard deviation of estimators for 2K counts.



(a)



(b)

Figure 6.5: (a) Horizontal profile through sample mean of estimators for 2M counts. (b) Horizontal profile through sample standard deviation of estimators for 2M counts.

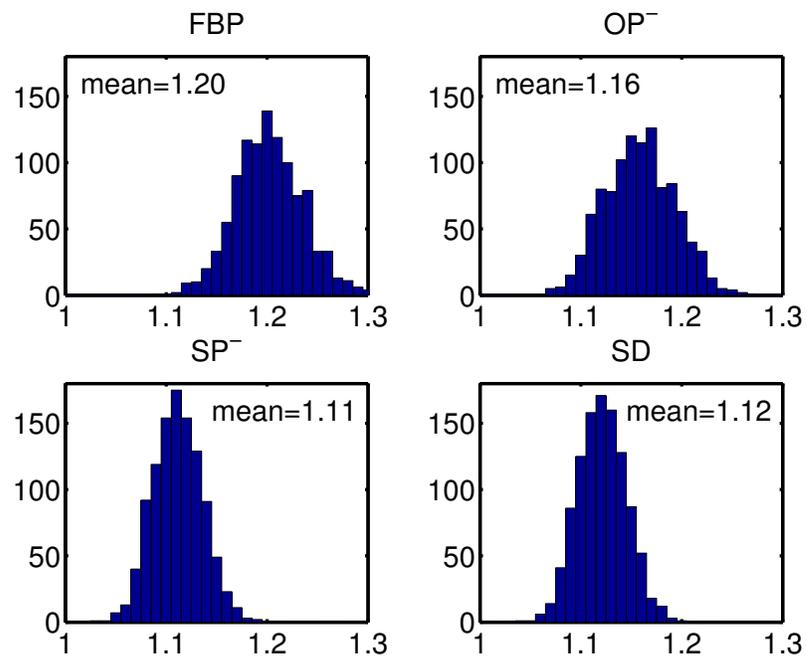


Figure 6.6: Histogram of the ratio of the standard deviation of different reconstruction methods to the standard deviation of PR method for 2M counts.

CHAPTER 7

Covariance of Kinetic Parameter Estimators Based on Time Activity Curve Reconstructions: 1D Study

7.1 Introduction

A primary application of dynamic PET or SPECT imaging is to quantify parameters of tracer kinetic models or compartmental models representing specific physiological processes (see Sec. 2.5). The goal is to estimate the kinetic parameters of the model for each region of interest (ROI) or voxel. Kinetic parameters are conventionally estimated as follows [56]: a series of images are reconstructed frame-by-frame, ROIs are identified and then kinetic parameters are obtained by fitting a compartmental model (with a measured or estimated blood input function) to spatially-averaged reconstructed image values for each ROI. Although there have been attempts to estimate parameters directly from sinogram data, they are computationally challenging [21, 26].

Recently, spatio-temporal reconstruction methods have been proposed to reconstruct time activity curves (TACs) by modeling each TAC as a *linear* combination of cubic B-splines [90]. Also, TAC reconstructions for each ROI obtained using B-spline temporal basis functions have been used to estimate kinetic parameters [106]. The performance, such as bias and variance, of the kinetic parameter estimators is affected by the choice of temporal basis functions for TACs (*e.g.*, the order of B-splines [106] and their knot loca-

tions). Although the effects of basis functions on TAC reconstructions have been studied in [11, 107], the effects on kinetic parameter estimators have little been analyzed [81, 106].

In this chapter we provide approximate analytical expressions for the covariance of kinetic parameter estimators in a simple 1D temporal “imaging” case. We do not analyze bias since we estimate the kinetic parameters from TAC reconstructions by (asymptotically) *unbiased* maximum likelihood (ML) estimators as opposed to widely-used (data-weighted) least squares estimators. The approximation formulas are very useful tools since they enable one to assess and optimize temporal basis functions in terms of complexity and variance without exhaustive simulations. They also show the effects of temporal regularization in TAC reconstruction.

Our approximations apply to list-mode data as well as (temporal) bin-mode data. List-mode acquisitions are more attractive than conventional frame-by-frame scans since all temporal information is contained in the event list. Our expressions can also be used to compute how much information is lost through temporal binning compared to list-mode data.

7.2 Problem

To focus on temporal aspects rather than interactions with spatial distributions, we consider a single-voxel or single-ROI object (containing a radiotracer) and a single detector unit, recording list-mode data (the arrival times of detected photons), or temporal bin-mode data. The model is not an unrealistically simple one; for example, in planar dynamic imaging, one could take a ROI and investigate the (average or dominant) dynamic tracer behavior using corresponding data. The goal is to estimate tracer kinetic parameters governing dynamic activity changes.

The photon emissions in the object can be modeled as an inhomogeneous Poisson pro-

cess whose rate function $\eta(t; \boldsymbol{\theta})$ corresponds to a TAC parameterized by kinetic parameters $\boldsymbol{\theta} = [\theta_1, \dots, \theta_p]'$ [90]. Suppose $\{\tau_k\}_{k=1}^K$ denotes list-mode data, that is, event arrival times.

Then the log-likelihood of $\boldsymbol{\theta}$ given the list-mode data is [115, p. 57]

$$L(\boldsymbol{\theta}, \{\tau_k\}_{k=1}^K) = \sum_{k=1}^K \log\{\alpha(\eta(\tau_k; \boldsymbol{\theta}) + r(\tau_k))\} - \int_0^T \alpha(\eta(t; \boldsymbol{\theta}) + r(t)) dt$$

where $r(t)$ is the rate function of the background process such as scatters and randoms, T denotes the total scan time, and α denotes a constant factor proportional to a radioisotope dosage. Although the background process $r(t)$ can be a function of α , we neglect the dependence for simplicity. One can obtain the Fisher information matrix $\mathbf{I}_\tau(\boldsymbol{\theta})$ for estimating $\boldsymbol{\theta}$ from $\{\tau_k\}_{k=1}^K$ as [115, p. 81]

$$[\mathbf{I}_\tau(\boldsymbol{\theta})]_{ij} = \alpha \int_0^T \frac{\partial \eta(t; \boldsymbol{\theta})}{\partial \theta_i} \frac{\partial \eta(t; \boldsymbol{\theta})}{\partial \theta_j} \frac{1}{\eta(t; \boldsymbol{\theta}) + r(t)} dt.$$

The inverse of $\mathbf{I}_\tau(\boldsymbol{\theta})$ can serve as an approximation to the baseline covariance of the direct estimator of $\boldsymbol{\theta}$ (without TAC reconstruction) based on list-mode data. However, in cases where the kinetic model is under development, it can be preferable to first estimate a TAC, and then fit various kinetic models to the TAC reconstruction.

Next, we describe the procedure of TAC reconstruction using temporal basis functions followed by kinetic parameter estimation.

7.2.1 TAC Reconstruction

We model the rate function as a linear combination of temporal basis functions $\{B_l(t)\}_{l=1}^L$, which for example can be B-splines, as

$$\eta(t) \cong \sum_{l=1}^L w_l B_l(t),$$

and we reconstruct the coefficients by PL estimation.

For simplicity we consider temporal bin-mode data $\mathbf{y} = [y_1, \dots, y_N]'$ where y_n is the number of events detected in the n th temporal bin (note $N \geq L$ or possibly $N \gg L$); the

list-mode data is a limiting case where $N \rightarrow \infty$ and the bin widths approach zero [11]. The bin-mode data \mathbf{y} are independent Poisson random variables, and the mean of each element is given by

$$\bar{y}_n(\boldsymbol{\theta}) \triangleq E[y_n] = \alpha p_n(\boldsymbol{\theta}) \quad (7.1)$$

$$p_n(\boldsymbol{\theta}) = \int_{t_{n-1}}^{t_n} \eta(t; \boldsymbol{\theta}) dt + r_n \quad (7.2)$$

where t_{n-1} and t_n are the end points of the n th temporal bin, and r_n represents background contributions. The log-likelihood of \mathbf{w} given \mathbf{y} can be obtained, ignoring constants independent of \mathbf{w} , as

$$L(\mathbf{w}, \mathbf{y}) = \sum_{n=1}^N \{y_n \log(\alpha \tilde{p}_n(\mathbf{w})) - \alpha \tilde{p}_n(\mathbf{w})\}$$

where

$$\tilde{p}_n(\mathbf{w}) = [\mathbf{B}\mathbf{w}]_n + r_n.$$

The $N \times L$ matrix \mathbf{B} has the (n, l) th entry as

$$b_{nl} = \int_{t_{n-1}}^{t_n} B_l(t) dt.$$

We assume that the $\{r_n\}$ are known (see [90] for methods of estimating randoms and scatters).

A PL estimate of \mathbf{w} is obtained finding the following maximizer:

$$\hat{\mathbf{w}}(\mathbf{y}) = \arg \max_{\mathbf{w} \in \mathcal{W}} \Phi(\mathbf{w}, \mathbf{y}) \quad (7.3)$$

where

$$\Phi(\mathbf{w}, \mathbf{y}) = L(\mathbf{w}, \mathbf{y}) - \frac{\beta}{2} \mathbf{w}' \mathbf{R} \mathbf{w} \quad (7.4)$$

and

$$\mathcal{W} = \left\{ \mathbf{w} : \sum_{l=1}^L w_l B_l(t) \geq 0, \forall t \in [0, T] \right\}. \quad (7.5)$$

The last term in (7.4) represents a roughness penalty encouraging temporal smoothness [11, 90], \mathbf{R} is a symmetric nonnegative definite matrix, and β is a regularization parameter. The set in (7.5) represents the nonnegativity constraint on reconstructed TACs, $\hat{\eta}(t) = \sum_{l=1}^L \hat{w}_l B_l(t)$.

7.2.2 Kinetic Parameter Estimation

To estimate kinetic parameters θ from \hat{w} in (7.3), we assume \hat{w} is Gaussian-distributed:

$$\hat{w} \sim \mathcal{N}(\boldsymbol{\mu}_{\hat{w}}(\theta), \mathbf{K}_{\hat{w}}(\theta)) \quad (7.6)$$

where $\boldsymbol{\mu}_{\hat{w}}$ and $\mathbf{K}_{\hat{w}}$ are the mean and the covariance matrix of the estimator \hat{w} , respectively. The higher counts per time (or temporal bin), the more the Gaussian assumption becomes accurate. Then one can compute a ML estimate of θ as follows:

$$\hat{\theta}(\hat{w}) = \arg \max_{\theta \in \Theta} \Psi(\theta, \hat{w}) \quad (7.7)$$

where Θ is a set of feasible θ , and the log-likelihood of θ given \hat{w} can be obtained, neglecting constants independent of θ , as

$$\Psi(\theta, \hat{w}) = -\frac{1}{2}(\hat{w} - \boldsymbol{\mu}_{\hat{w}}(\theta))'[\mathbf{K}_{\hat{w}}(\theta)]^{-1}(\hat{w} - \boldsymbol{\mu}_{\hat{w}}(\theta)) - \frac{1}{2} \log |\mathbf{K}_{\hat{w}}(\theta)| \quad (7.8)$$

where $|\cdot|$ denotes determinant. Generally, the TAC estimator $\hat{\eta}(t) = \sum_{l=1}^L \hat{w}_l B_l(t)$ from (7.3) is not consistent since it can be a case that $\hat{\eta}(t; \theta^{\text{true}}) \neq \sum_{l=1}^L w_l B_l(t)$ for all w_l 's; even in such a case, however, the ML kinetic parameter estimator $\hat{\theta}$ in (7.7) can be (nearly) unbiased as shown in Sec. 7.4. Therefore, the Cramér-Rao bound, that is, the inverse of the Fisher information matrix, which is shown in the next section, can serve as an approximation to the covariance of $\hat{\theta}$.

7.3 Covariance of kinetic parameter estimators

7.3.1 Derivation

First, we need approximate expressions for $\boldsymbol{\mu}_{\hat{\boldsymbol{w}}}$ and $\mathbf{K}_{\hat{\boldsymbol{w}}}$ in (7.6). Using a first-order Taylor approximation of $\hat{\boldsymbol{w}}(\boldsymbol{y})$ at $\bar{\boldsymbol{y}} = [\bar{y}_1, \dots, \bar{y}_N]'$ in (7.1), the chain rule and the implicit function theorem with some reasonable assumptions [36], one can obtain the following approximations:

$$\boldsymbol{\mu}_{\hat{\boldsymbol{w}}}(\boldsymbol{\theta}) \cong \hat{\boldsymbol{w}}(\bar{\boldsymbol{y}}(\boldsymbol{\theta})) \triangleq \check{\boldsymbol{w}}(\boldsymbol{\theta}) \quad (7.9)$$

and

$$\begin{aligned} \mathbf{K}_{\hat{\boldsymbol{w}}}(\boldsymbol{\theta}) &\cong \frac{1}{\alpha} \left[\mathbf{F}_{\text{bin}}(\boldsymbol{\theta}) + \frac{\beta}{\alpha} \mathbf{R} \right]^{-1} \mathbf{F}_{\text{bin}}(\boldsymbol{\theta}) \left[\mathbf{F}_{\text{bin}}(\boldsymbol{\theta}) + \frac{\beta}{\alpha} \mathbf{R} \right]^{-1} \\ &\triangleq [\tilde{\mathbf{F}}(\boldsymbol{\theta})]^{-1} \end{aligned} \quad (7.10)$$

where

$$\mathbf{F}_{\text{bin}}(\boldsymbol{\theta}) = \mathbf{B}' \text{diag} \left\{ \frac{p_n(\boldsymbol{\theta})}{\tilde{p}_n^2(\check{\boldsymbol{w}}(\boldsymbol{\theta}))} \right\} \mathbf{B}.$$

Now one can compute the Fisher information matrix from (7.8) by replacing $\boldsymbol{\mu}_{\hat{\boldsymbol{w}}}$ and $\mathbf{K}_{\hat{\boldsymbol{w}}}$ with their approximations in (7.9) and (7.10). Some manipulation leads to our final expression for the Fisher information matrix for estimating $\boldsymbol{\theta}$ from $\hat{\boldsymbol{w}}$,

$$\begin{aligned} \mathbf{I}_{\hat{\boldsymbol{w}}, \text{bin}}(\boldsymbol{\theta}) &= E_{\boldsymbol{\theta}}[-\nabla_{\boldsymbol{\theta}}^2 \Psi(\boldsymbol{\theta}, \hat{\boldsymbol{w}})] \\ &\cong [\nabla_{\boldsymbol{\theta}} \check{\boldsymbol{w}}(\boldsymbol{\theta})]' \tilde{\mathbf{F}}(\boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}} \check{\boldsymbol{w}}(\boldsymbol{\theta}) \\ &= \alpha [\nabla_{\boldsymbol{\theta}} \mathbf{p}(\boldsymbol{\theta})]' \text{diag} \left\{ \frac{1}{\tilde{p}_n(\check{\boldsymbol{w}}(\boldsymbol{\theta}))} \right\} \mathbf{B} [\mathbf{F}_{\text{bin}}(\boldsymbol{\theta})]^{-1} \mathbf{B}' \cdot \\ &\quad \text{diag} \left\{ \frac{1}{\tilde{p}_n(\check{\boldsymbol{w}}(\boldsymbol{\theta}))} \right\} \nabla_{\boldsymbol{\theta}} \mathbf{p}(\boldsymbol{\theta}) \end{aligned} \quad (7.11)$$

where $\nabla_{\boldsymbol{\theta}} = [\frac{\partial}{\partial \theta_1}, \dots, \frac{\partial}{\partial \theta_p}]$ denotes the row gradient operators, $\nabla_{\boldsymbol{\theta}}^2$ denotes the Hessian operator, and $\mathbf{p} = [p_1, \dots, p_N]'$ is defined in (7.2). The information matrix $\mathbf{I}_{\hat{\boldsymbol{w}}, \text{bin}}(\boldsymbol{\theta})$ depends implicitly on temporal regularization only through $\tilde{p}_n(\check{\boldsymbol{w}}(\boldsymbol{\theta}))$ [see (7.3), (7.4) and (7.9)].

7.3.2 Information Matrix for List-Mode Data

By increasing the number of bins N to ∞ and decreasing the bin widths to 0 in (7.11), one can obtain the following information matrix for list-mode data:

$$\begin{aligned}
 [\mathbf{I}_{\dot{\mathbf{w}},\text{list}}(\boldsymbol{\theta})]_{ij} &\cong \alpha \sum_{p=1}^L \sum_{q=1}^L [\mathbf{F}_{\text{list}}(\boldsymbol{\theta})^{-1}]_{pq} \\
 &\int_0^T \frac{\partial \eta(t; \boldsymbol{\theta})}{\partial \theta_i} \frac{B_p(t)}{\sum_{l=1}^L \check{w}_l(\boldsymbol{\theta}) B_l(t) + r(t)} dt \cdot \\
 &\int_0^T \frac{\partial \eta(t; \boldsymbol{\theta})}{\partial \theta_j} \frac{B_q(t)}{\sum_{l=1}^L \check{w}_l(\boldsymbol{\theta}) B_l(t) + r(t)} dt
 \end{aligned} \tag{7.12}$$

where

$$[\mathbf{F}_{\text{list}}(\boldsymbol{\theta})]_{ij} = \int_0^T B_i(t) B_j(t) \frac{\eta(t; \boldsymbol{\theta}) + r(t)}{\left(\sum_{l=1}^L \check{w}_l(\boldsymbol{\theta}) B_l(t) + r(t) \right)^2} dt.$$

If temporal basis functions are constant B-splines as

$$B_l(t) = I_{[t_{l-1}, t_l]}(t)$$

where $I_{[t_{l-1}, t_l]}$ is an indicator function, then the information matrix in (7.12) becomes

$$\mathbf{I}_{\dot{\mathbf{w}},\text{list}}(\boldsymbol{\theta}) \cong \alpha [\nabla_{\boldsymbol{\theta}} \mathbf{p}(\boldsymbol{\theta})]' \text{diag} \left\{ \frac{1}{p_n(\boldsymbol{\theta})} \right\} \nabla_{\boldsymbol{\theta}} \mathbf{p}(\boldsymbol{\theta}). \tag{7.13}$$

This information matrix *is independent of temporal regularization!* One can also obtain the same result as (7.13) from (7.11) by making temporal bins agree with the constant B-spline basis functions $\{B_l(t)\}$. In this case \mathbf{B} and \mathbf{F} are diagonal, and the regularization-related terms $\tilde{p}_n(\check{\mathbf{w}}(\boldsymbol{\theta}))$ are canceled out in (7.11). In fact, the equality happens to hold in (7.13) [115, p. 81].

7.4 Results

To assess the accuracy of the approximation for the covariance of kinetic parameter estimators given by the inverse of (7.11), we simulated dynamic imaging data. The simulated

TAC was given by

$$\eta(t; \boldsymbol{\theta}) = b(t) \star h(t; \boldsymbol{\theta})$$

where \star denotes the convolution, the *known* blood input function is

$$b(t) = te^{-1.4t}u(t)$$

and the impulse response of a one tissue compartment model is

$$h(t; \boldsymbol{\theta}) = \theta_2 \exp(-\theta_1 t)u(t) + f_v \delta(t). \quad (7.14)$$

The unit step function is denoted by $u(t)$, and the delta function is denoted by $\delta(t)$. In (7.14), f_v denotes an unknown fractional volume of blood to be also estimated. The total scan time T was set to 15 min, and the true kinetic parameters were set as $\theta_1^{\text{true}} = 0.15$ and $\theta_2^{\text{true}} = 0.7$. The total counts were 100K, and r_n corresponded to a temporally uniform field of 10% of background events. The data were acquired using 30 uniform temporal bins ($N = 30$). We used the uniform quadratic penalty with $\beta = 0.1$, and 10 B-spline basis functions of different orders (constant, linear, quadratic, and cubic) with uniformly spaced knots for TAC reconstruction ($L = 10$). Given temporal basis functions and given (simulated) noisy data, $\hat{\mathbf{w}}$ was estimated by (7.3) with linear constraints $\mathbf{B}\mathbf{w} \geq \mathbf{0}$ as a reasonable approximation to (7.5), and then $\boldsymbol{\theta}$ and f_v was estimated by (7.7) with (7.8) using (7.9) and (7.10) with nonnegativity constraints $\boldsymbol{\theta} \geq \mathbf{0}$ and $f_v \geq 0$. We computed the sample mean and the sample covariance of $\hat{\boldsymbol{\theta}}$ from 100 realizations for each set of basis functions. For maximization in (7.3) and (7.7), we used separable surrogates [2], and the `fmincon` and `quadprog` functions of Matlab.

Table 7.1 shows that the sample means of the parameter estimators $\hat{\boldsymbol{\theta}}$ agreed with true values $\boldsymbol{\theta}^{\text{true}}$ for all basis function sets used in this paper; this shows that the estimators are unbiased. Table 7.2 shows that predicted standard deviations of $\hat{\boldsymbol{\theta}}$ obtained from the inverse of (7.11) were reasonably close to empirical sample standard deviations.

Table 7.1: Comparison of predicted and empirical sample means of kinetic parameter estimators

basis	parameter	true	sample mean
cubic	θ_1	0.15	0.15 ± 0.0001
B-splines	θ_2	0.70	0.70 ± 0.0005
quadratic	θ_1	0.15	0.15 ± 0.0001
B-splines	θ_2	0.70	0.70 ± 0.0005
linear	θ_1	0.15	0.15 ± 0.0001
B-splines	θ_2	0.70	0.70 ± 0.0005
constant	θ_1	0.15	0.15 ± 0.0001
B-splines	θ_2	0.70	0.70 ± 0.0006

Table 7.2: Comparison of predicted and empirical standard deviations of kinetic parameter estimators

basis	parameter	predicted std dev ($\times 10^{-3}$)	sample std dev ($\times 10^{-3}$)
cubic	θ_1	1.15	1.22 ± 0.09
B-splines	θ_2	4.89	5.20 ± 0.37
quadratic	θ_1	1.15	1.22 ± 0.09
B-splines	θ_2	4.89	5.22 ± 0.37
linear	θ_1	1.16	1.23 ± 0.09
B-splines	θ_2	4.91	5.27 ± 0.37
constant	θ_1	1.20	1.24 ± 0.09
B-splines	θ_2	5.16	5.48 ± 0.39

7.5 Conclusion

We derived the covariance matrix (the inverse of the Fisher information matrix) of kinetic parameter estimators based on TAC reconstructions using temporal basis functions for list-mode data as well as bin-mode data in a 1D temporal problem. We demonstrated the covariance approximation predicts the empirical covariance reasonably well for high counts.

CHAPTER 8

Summary and Future Work

8.1 Summary

Global convergence is an important property for any optimization algorithm. One may argue that it is not critical because in practice we do not run algorithms to convergence. However, in medical imaging, reliability is essential, that is, there must be a guarantee that medical imaging systems provide a solution to the problem we formulated.¹ In addition, it is not impractical to run algorithms to convergence for PL reconstruction because of better conditioning by regularization, growing computing powers, and fast algorithms, whereas we usually stop algorithms long before convergence for ill-posed ML reconstruction. Finally, when one develops new statistical models or objective functions, and investigates their properties, one needs to obtain maximizers of given objective functions by running algorithms to convergence.

We have developed two families of fast and convergent algorithms: relaxed OS algorithms and incremental optimization transfer algorithms. Relaxed OS algorithms include

¹As in Sections 3.1 and 3.2, a non-obvious part in problem formulation is how to choose a penalty function (and a regularization parameter). Noise-resolution trade-offs and edge-preserving properties in reconstructed images are related to penalty functions used. For example, there are methods of designing quadratic penalty functions to achieve target resolution in reconstructed images [123, 124]. Suppose we have designed such a penalty function. Then, to achieve the target resolution, one should obtain a maximizer to the corresponding PL objective function, and accordingly one needs a convergent algorithm.

modified BSREM and diagonally-scaled incremental gradient algorithm (or relaxed OS-SPS as a special case). Incremental optimization transfer algorithms include TRIOT as a special case for transmission tomography. They are provably convergent, and the results show they also have fast convergence rates. They are quite general, so they have broad applicability. As discussed in Chapter 5, it seems that incremental optimization transfer algorithms, particularly TRIOT algorithms for transmission tomography, are more convenient than relaxed OS algorithms because determining relaxation parameters is inconvenient.

We also developed new statistical models for randoms-precorrected PET by allowing negative sinogram values. Analysis and simulation results show that the new shifted Poisson (SP^-) model is nearly free of systematic bias yet keeps low variance. Despite its simpler implementation, the new SP^- performs comparably to the saddle-point (SD) model which has shown the best performance (as to systematic bias and variance) in randoms-precorrected PET emission reconstruction. We recommend the prompt (PR) method if the prompt and the randoms data are accessible separately; however, if only randoms-precorrected data are available, the new SP^- is our recommended method.

Finally, we developed approximate expressions for the covariance matrix of kinetic parameter estimators based on time activity curve (TAC) reconstructions when TACs are modeled as a linear combination of temporal basis functions such as B-splines. Simulation results show that the expression is reasonably accurate, and the approximate formulas are useful tools for assessing and optimizing the reconstruction methods for dynamic imaging.

8.2 Future Work

- We have proven global convergence of relaxed OS algorithms, including modified BSREM and diagonally-scaled incremental gradient methods, and incremental optimization transfer algorithms; and we have investigated their convergence rates qual-

itatively through experiments. However, a quantitative analysis of convergence rates would give insights into how to accelerate algorithms more systematically. For example, although we analyzed convergence rates for incremental optimization transfer algorithms and showed that incremental optimization transfer algorithms are faster than nonincremental ones for a 1D case, whether it is generally true is an open problem.

- For relaxed OS algorithms, it will be an interesting question whether the sufficient conditions for convergence can be relaxed, for example, whether they would still converge when they apply to nonconcave objective functions, or whether modified BSREM would converge when an objective function is (non strictly) concave and has multiple solutions.
- For relaxed OS algorithms, it would be desirable to develop some automatic step-size rules that do not require expensive computation, removing the inconvenience to determine the relaxation parameters.
- Modified BSREM and diagonally-scaled incremental gradient methods (or relaxed OS-SPS as a special case) have similar forms, and the major difference is in the form of $d_j(\cdot)$. It is an open problem to prove convergence of relaxed OS algorithms when general scaling functions $d_j(\cdot)$ are used.
- For OS or incremental algorithms, we distributed the penalty function into all sub-objective function, that is, we computed gradients and/or curvatures of the penalty function at every subiteration. As the number of subsets increases, the computation involving the penalty function also increases. The effectiveness of the use of partial penalty functions at every subiteration needs to be examined.
- For the (incremental) optimization transfer framework, we used separable surrogate functions derived using De Pierro's concavity tricks in [30]. The separable surro-

gates satisfy the minorization conditions, consequently ensuring monotonicity in the (augmented) objective function. Whether better surrogate functions can be found is an open problem. Separable surrogates are usually easier to maximize than nonseparable ones, but their curvatures are higher and accordingly their convergence rates are slower. In other words, there is a trade-off between computational efficiency and convergence rates. There might exist some nonseparable surrogates yielding better performance in terms of objective increase per CPU time. Even if we are interested in only separable surrogates, it will be an interesting problem what the best separable surrogate is in terms of convergence rates. Another challenging question is whether we can relax the monotonicity or the minorization conditions to obtain better surrogate functions and under what conditions such surrogates will ensure convergence.

- In Chapter 5, we found that switching from OS-SPS to TRIOT can be effective. As Bertsekas unified incremental gradient methods and ordinary gradient methods in [14], OS-SPS and TRIOT could be unified in a more general framework.
- Throughout this dissertation, we assumed the system matrix is known without any error. It is an open problem how a model mismatch will affect the reconstruction methods presented here, for example, the convergence properties of reconstruction algorithms, the ϵ in Appendix B, and so on.
- We developed new reconstruction algorithms in this dissertation; however, we are still seeking for better algorithms that are convergent, simpler, and faster.
- We analyzed the covariance of kinetic parameter estimators for 1D in Chapter 7. It is natural to extend the analysis to 3D or 4D by incorporating spatial aspects.

APPENDICES

APPENDIX A

Computable Upper Bounds for PL Solutions in Emission Tomography

In this appendix, we provide two types of computable upper bounds for PL solutions in emission tomography (see Section 3.1.5).

A.1 Computable Upper Bound: Example 1

Define

$$U \triangleq \max_i \left\{ \frac{y_i}{\min_{j: a_{ij} \neq 0} a_{ij}} \right\}. \quad (\text{A.1})$$

Suppose $\boldsymbol{\lambda}$ is a vector in \mathcal{D} for which the index set of “too large” elements $\mathcal{J} = \{j = 1, \dots, p : \lambda_j > U\}$ is *nonempty* where \mathcal{D} is defined in (3.8). Define $\tilde{\boldsymbol{\lambda}}$ by clipping as

$$\tilde{\lambda}_j = \begin{cases} U, & j \in \mathcal{J} \\ \lambda_j, & j \notin \mathcal{J}. \end{cases}$$

To prove the U in (A.1) is an upper bound for PL solutions, it suffices to show that $\Phi(\boldsymbol{\lambda}) < \Phi(\tilde{\boldsymbol{\lambda}}) \leq \max_{\boldsymbol{\xi} \in \mathcal{D}} \Phi(\boldsymbol{\xi})$. First, note that, for any i , if there exists $j_i \in \mathcal{J}$ such that $a_{ij_i} > 0$, then:

$$\begin{aligned} \bar{y}_i(\boldsymbol{\lambda}) &= a_{ij_i} \lambda_{j_i} + \sum_{j \neq j_i} a_{ij} \lambda_j + r_i \\ &> a_{ij_i} U + \sum_{j \neq j_i} a_{ij} \tilde{\lambda}_j + r_i = \bar{y}_i(\tilde{\boldsymbol{\lambda}}) \end{aligned} \quad (\text{A.2})$$

$$\geq a_{ij_i} U \geq y_i \quad (\text{A.3})$$

where in (A.2) we used the fact that $\lambda_j \geq \tilde{\lambda}_j, \forall j$ and that $\lambda_{j_i} > U = \tilde{\lambda}_{j_i}$; and (A.3) is due to our construction of U in (A.1). Therefore, $h_i([\mathbf{A}\boldsymbol{\lambda}]_i) < h_i([\mathbf{A}\tilde{\boldsymbol{\lambda}}]_i)$ by (3.5). Second, if such j_i does not exist, then $h_i([\mathbf{A}\boldsymbol{\lambda}]_i) = h_i([\mathbf{A}\tilde{\boldsymbol{\lambda}}]_i)$ since $[\mathbf{A}\boldsymbol{\lambda}]_i = [\mathbf{A}\tilde{\boldsymbol{\lambda}}]_i$. Third, one can verify that there exists some i for which such j_i exists by the assumption of nonzero sensitivity factors. Combining these, we have $L(\boldsymbol{\lambda}) = \sum_i^N h_i([\mathbf{A}\boldsymbol{\lambda}]_i) < \sum_i^N h_i([\mathbf{A}\tilde{\boldsymbol{\lambda}}]_i) = L(\tilde{\boldsymbol{\lambda}})$. One can also show that “clipping” all elements of $\boldsymbol{\lambda}$ greater than U will always decrease the roughness penalty R in (3.9) due to our assumption that the potential function $\psi(t)$ is nondecreasing in $|t|$, implied by (3.11) and (3.13). We have established that $\phi(\boldsymbol{\lambda}) = L(\boldsymbol{\lambda}) - R(\boldsymbol{\lambda}) < L(\tilde{\boldsymbol{\lambda}}) - R(\tilde{\boldsymbol{\lambda}}) = \Phi(\tilde{\boldsymbol{\lambda}})$.

A.2 Computable Upper Bound: Example 2

Alternatively, we construct an upper bound for a broader family of penalty functions more general than those based on differences of neighboring pixels. Pick any $\boldsymbol{\nu} \in \mathcal{D}$, e.g., $\boldsymbol{\nu} = \mathbf{1}$. Define $i_j = \arg \max_i a_{ij}, \forall j$. By the assumption of nonzero sensitivity factors, we have $a_{i_j j} > 0, \forall j$. Choose any $c > 0$ for all j , and define

$$U_j \triangleq \frac{1}{a_{i_j j}} \max \left\{ y_{i_j} - r_{i_j}, \frac{\sum_{i \neq i_j} h_i(y_i - r_i) + y_{i_j} \log(y_{i_j} + c) - \Phi(\boldsymbol{\nu})}{c/(y_{i_j} + c)}, 1 \right\}, \forall j. \quad (\text{A.4})$$

The last term, 1, in the brace ensures that $U_j > 0$. Let $U \triangleq \max_j U_j$. We prove (3.18) for the U . Suppose that $\boldsymbol{\lambda} \in \mathcal{D} \setminus \mathcal{B}$, that is, $\lambda_q > U$ for some q , where \mathcal{B} is defined in (3.18). To prove the U is an upper bound for PL solution, it suffices to show $\Phi(\boldsymbol{\lambda}) < \Phi(\boldsymbol{\nu}) \leq$

$\max_{\boldsymbol{\xi} \in \mathcal{D}} \Phi(\boldsymbol{\xi})$. One can show this as follows:

$$\begin{aligned} \Phi(\boldsymbol{\lambda}) &= \sum_{i=1}^N h_i([\mathbf{A}\boldsymbol{\lambda}]_i) - R(\boldsymbol{\lambda}) \\ &\leq h_{i_q}([\mathbf{A}\boldsymbol{\lambda}]_{i_q}) + \sum_{i \neq i_q} h_i([\mathbf{A}\boldsymbol{\lambda}]_i) \end{aligned} \quad (\text{A.5})$$

$$< h_{i_q}(a_{i_q q}U - r_{i_q}) + \sum_{i \neq i_q} h_i(y_i - r_i) \quad (\text{A.6})$$

$$\begin{aligned} &\leq h_{i_q}(y_{i_q} - r_{i_q} + c) + \dot{h}_{i_q}(y_{i_q} - r_{i_q} + c)(a_{i_q q}U - r_{i_q} - (y_{i_q} - r_{i_q} + c)) \\ &\quad + \sum_{i \neq i_q} h_i(y_i - r_i) \end{aligned} \quad (\text{A.7})$$

$$\begin{aligned} &= y_{i_q} \log(y_{i_q} + c) - (y_{i_q} + c) + \left(\frac{y_{i_q}}{y_{i_q} + c} - 1 \right) (a_{i_q q}U - (y_{i_q} + c)) \\ &\quad + \sum_{i \neq i_q} h_i(y_i - r_i) \end{aligned} \quad (\text{A.8})$$

$$\begin{aligned} &\leq y_{i_q} \log(y_{i_q} + c) - \frac{c}{y_{i_q} + c} a_{i_q q}U + \sum_{i \neq i_q} h_i(y_i - r_i) \\ &< \Phi(\boldsymbol{\nu}) \leq \Phi(\hat{\boldsymbol{\lambda}}). \end{aligned} \quad (\text{A.9})$$

where (A.5) is due to (3.11); (A.6) is due to (3.4), (3.5) and (A.4); (A.7) is due to the concavity of h_{i_q} from (3.6); (A.8) is computed from the definition of h_{i_q} in (3.3); and (A.9) is due to (A.4).

APPENDIX B

Determination of ϵ for Modified Objective Function in Emission Tomography

In this appendix, we determine ϵ such that (3.22) holds (see Section 3.1.6). If $\mathcal{I} = \emptyset$, then we do not need to modify the PL objective function Φ , where \mathcal{I} is defined in (3.19).

Suppose $\mathcal{I} \neq \emptyset$. Pick any $\boldsymbol{\nu} \in \mathcal{D}$, e.g., $\boldsymbol{\nu} = \mathbf{1}$. Define

$$\epsilon \triangleq \min_{q \in \mathcal{I}} \left\{ y_q, \exp\left(\frac{\Phi(\boldsymbol{\nu}) - \sum_{i \neq q} h_i(y_i - r_i)}{y_q}\right) \right\} > 0. \quad (\text{B.1})$$

We show that $\Lambda^* \subset \mathcal{E}$ where Λ^* and \mathcal{E} are defined in (3.17) and (3.21), respectively (recall \mathcal{E} is the region on which Φ is not modified). Let $\boldsymbol{\lambda} \in \mathcal{D} \setminus \mathcal{E}$, that is, $0 \leq [\mathbf{A}\boldsymbol{\lambda}]_q \leq \epsilon$ for some $q \in \mathcal{I}$. It suffices to show that $\Phi(\boldsymbol{\lambda}) < \Phi(\boldsymbol{\nu}) \leq \arg \max_{\boldsymbol{\xi} \in \mathcal{D}} \Phi(\boldsymbol{\xi})$. One can show this as follows:

$$\begin{aligned} \Phi(\boldsymbol{\lambda}) &= \sum_{i=1}^N h_i([\mathbf{A}\boldsymbol{\lambda}]_i) - R(\boldsymbol{\lambda}) \\ &\leq h_q([\mathbf{A}\boldsymbol{\lambda}]_q) + \sum_{i \neq q} h_i(y_i - r_i) \end{aligned} \quad (\text{B.2})$$

$$\leq h_q(\epsilon) + \sum_{i \neq q} h_i(y_i - r_i) \quad (\text{B.3})$$

$$< y_q \log \epsilon + \sum_{i \neq q} h_i(y_i - r_i) \quad (\text{B.4})$$

$$\leq \Phi(\boldsymbol{\nu}) \quad (\text{B.5})$$

where (B.2) is a consequence of (3.4) and (3.10); (B.3) is due to (3.5) and (B.1); (B.4) is due to the definition (3.3) (recall $r_q = 0$ since $q \in \mathcal{I}$); and (B.5) is a consequence of

(B.1). Therefore, $\Lambda^* \subset \mathcal{E}$. Similarly, one can verify that $\tilde{\Lambda}^{**} \triangleq \{\boldsymbol{\lambda}^* \in \mathbb{R}_+^p : \tilde{\Phi}(\boldsymbol{\lambda}^*) \geq \tilde{\Phi}(\boldsymbol{\lambda}), \forall \boldsymbol{\lambda} \in \mathbb{R}_+^p\} \subset \mathcal{E}$. But since $\Phi(\boldsymbol{\lambda}) = \tilde{\Phi}(\boldsymbol{\lambda})$ for $\boldsymbol{\lambda} \in \mathcal{E}$, we have $\tilde{\Lambda}^{**} = \Lambda^*$. Now since $\Lambda^* \subset \mathcal{B}$ from the preceding section, we have $\tilde{\Lambda}^{**} = \tilde{\Lambda}^{**} \cap \mathcal{B} = \tilde{\Lambda}^*$. We have shown that the maximizers of the modified objective function $\tilde{\Phi}$ are the same as those of the original objective function Φ over \mathcal{D} or \mathcal{B} .

APPENDIX C

Convergence Proofs for modified BSREM

We prove that the modified BSREM-I in (4.12) with (4.13) is globally convergent. The required assumptions on the objective function Φ as follows: $\Phi(\boldsymbol{\lambda})$ is strictly concave on \mathcal{B} ; and $\nabla\Phi_m(\boldsymbol{\lambda})$ and $\mathbf{D}(\boldsymbol{\lambda})\nabla\Phi_m(\boldsymbol{\lambda})$ are Lipschitz continuous (and thus bounded) on \mathcal{B} . They are satisfied by the emission PL objective function.

Lemma C.1. *Suppose that $\{\boldsymbol{\lambda}^{n,m}\}$ is a sequence generated by (4.12) with $\boldsymbol{\lambda}^0 \in \text{Int } \mathcal{B}$, where $\text{Int } \mathcal{B}$ denotes the interior of \mathcal{B} . Then there exists $\alpha > 0$ such that if $0 < \alpha_n \leq \alpha$, $\forall n$, then $\boldsymbol{\lambda}^{n,m} \in \text{Int } \mathcal{B}$, $\forall n, m$.*

Proof. Since $(\partial\Phi_m/\partial\lambda_j)(\boldsymbol{\lambda})$ is bounded over \mathcal{B} for all j and m , one can choose $\alpha > 0$ such that

$$\alpha \left| \frac{1}{p_j} \frac{\partial\Phi_m}{\partial\lambda_j}(\boldsymbol{\lambda}) \right| < 1, \quad \forall \boldsymbol{\lambda} \in \mathcal{B} \text{ and } \forall j, m.$$

Suppose that $0 < \alpha_n \leq \alpha$, $\forall n$ and that $\boldsymbol{\lambda}^{n,m-1} \in \text{Int } \mathcal{B}$. If $0 < \lambda_j^{n,m-1} < U/2$, one can show that $0 < \lambda_j^{n,m} < U$, using the following expression for $\lambda_j^{n,m}$:

$$\begin{aligned} \lambda_j^{n,m} &= \lambda_j^{n,m-1} + \alpha_n \frac{\lambda_j^{n,m-1}}{p_j} \frac{\partial\Phi_m}{\partial\lambda_j}(\boldsymbol{\lambda}^{n,m-1}) \\ &= \lambda_j^{n,m-1} \left(1 + \alpha_n \frac{1}{p_j} \frac{\partial\Phi_m}{\partial\lambda_j}(\boldsymbol{\lambda}^{n,m-1}) \right). \end{aligned}$$

If $U/2 \leq \lambda_j^{n,m-1} < U$, one can also show that $0 < \lambda_j^{n,m} < U$, using the following

expression for $\lambda_j^{n,m}$, in view of (4.13):

$$U - \lambda_j^{n,m} = (U - \lambda_j^{n,m-1}) \left(1 - \alpha_n \frac{1}{p_j} \frac{\partial \Phi_m}{\partial \lambda_j}(\boldsymbol{\lambda}^{n,m-1}) \right).$$

This implies that $\boldsymbol{\lambda}^{n,m} \in \text{Int } \mathcal{B}$. □

Lemma C.2. *Suppose that $\{\boldsymbol{\lambda}^n\}$ is a sequence generated by (4.12) with $\alpha_n > 0$ such that $\sum_{n=0}^{\infty} \alpha_n = \infty$ and $\sum_{n=0}^{\infty} \alpha_n^2 < \infty$. If $\boldsymbol{\lambda}^{n,m} \in \mathcal{B}$, $\forall n, m$, then $\{\Phi(\boldsymbol{\lambda}^n)\}$ converges in \mathbb{R} and there exists a limit point $\boldsymbol{\lambda}^* \in \mathcal{B}$ of $\{\boldsymbol{\lambda}^n\}$ such that $\mathbf{D}(\boldsymbol{\lambda}^*) \nabla \Phi(\boldsymbol{\lambda}^*) = \mathbf{0}$.*

Proof. Using the definition of the sequence $\{\boldsymbol{\lambda}^n\}$, we have:

$$\begin{aligned} \boldsymbol{\lambda}^{n+1} &= \boldsymbol{\lambda}^n + \alpha_n \sum_{m=1}^M \mathbf{D}(\boldsymbol{\lambda}^{n,m-1}) \nabla \Phi_m(\boldsymbol{\lambda}^{n,m-1}) \\ &= \boldsymbol{\lambda}^n + \alpha_n \mathbf{D}(\boldsymbol{\lambda}^n) \nabla \Phi(\boldsymbol{\lambda}^n) + \\ &\quad \alpha_n \sum_{m=1}^M (\mathbf{D}(\boldsymbol{\lambda}^{n,m-1}) \nabla \Phi_m(\boldsymbol{\lambda}^{n,m-1}) - \mathbf{D}(\boldsymbol{\lambda}^n) \nabla \Phi_m(\boldsymbol{\lambda}^n)) \\ &= \boldsymbol{\lambda}^n + \alpha_n \mathbf{D}(\boldsymbol{\lambda}^n) \nabla \Phi(\boldsymbol{\lambda}^n) + O(\alpha_n^2) \end{aligned} \tag{C.1}$$

where the last equality is obtained using the Lipschitz continuity of $\mathbf{D}(\boldsymbol{\lambda}) \nabla \Phi_m(\boldsymbol{\lambda})$ on \mathcal{B} .

In particular, for some positive $L \in \mathbb{R}$, we have:

$$\begin{aligned} &\left\| \sum_{m=1}^M (\mathbf{D}(\boldsymbol{\lambda}^{n,m-1}) \nabla \Phi_m(\boldsymbol{\lambda}^{n,m-1}) - \mathbf{D}(\boldsymbol{\lambda}^n) \nabla \Phi_m(\boldsymbol{\lambda}^n)) \right\| \\ &\leq \sum_{m=1}^M \left\| \mathbf{D}(\boldsymbol{\lambda}^{n,m-1}) \nabla \Phi_m(\boldsymbol{\lambda}^{n,m-1}) - \mathbf{D}(\boldsymbol{\lambda}^n) \nabla \Phi_m(\boldsymbol{\lambda}^n) \right\| \\ &\leq L \sum_{m=1}^M \left\| \boldsymbol{\lambda}^{n,m-1} - \boldsymbol{\lambda}^n \right\| \\ &\leq \alpha_n L \sum_{m=1}^M \sum_{k=1}^{m-1} \left\| \mathbf{D}(\boldsymbol{\lambda}^{n,k-1}) \nabla \Phi_k(\boldsymbol{\lambda}^{n,k-1}) \right\| \\ &\leq \alpha_n L M^2 \max_{m, \boldsymbol{\lambda} \in \mathcal{B}} \left\| \mathbf{D}(\boldsymbol{\lambda}) \nabla \Phi_m(\boldsymbol{\lambda}) \right\|. \end{aligned}$$

Consider the objective sequence $\{\Phi(\boldsymbol{\lambda}^n)\}$. Since $\nabla \Phi(\boldsymbol{\lambda})$ is Lipschitz continuous on \mathcal{B} , we have [100, p. 6]:

$$\Phi(\boldsymbol{\lambda}^{n+1}) = \Phi(\boldsymbol{\lambda}^n) + \nabla \Phi(\boldsymbol{\lambda}^n)' (\boldsymbol{\lambda}^{n+1} - \boldsymbol{\lambda}^n) + O(\|\boldsymbol{\lambda}^{n+1} - \boldsymbol{\lambda}^n\|^2). \tag{C.2}$$

Using (C.1) and (C.2), for large n , we establish:

$$\Phi(\boldsymbol{\lambda}^{n+1}) = \Phi(\boldsymbol{\lambda}^n) + \alpha_n \nabla \Phi(\boldsymbol{\lambda}^n)' \mathbf{D}(\boldsymbol{\lambda}^n) \nabla \Phi(\boldsymbol{\lambda}^n) + O(\alpha_n^2). \quad (\text{C.3})$$

In view of (i) $\sum_{n=0}^{\infty} \alpha_n^2 < \infty$, (ii) the boundedness of $\Phi(\boldsymbol{\lambda})$ on \mathcal{B} , (iii) the nonnegative definiteness of $\mathbf{D}(\boldsymbol{\lambda}^n)$, and (C.3), one can show that:

$$\sum_{n=l}^k \alpha_n \nabla \Phi(\boldsymbol{\lambda}^n)' \mathbf{D}(\boldsymbol{\lambda}^n) \nabla \Phi(\boldsymbol{\lambda}^n) < q < \infty, \quad \forall k > l$$

for some $q \in \mathbb{R}$ and some large l . This implies that $\sum_{n=0}^{\infty} \alpha_n \nabla \Phi(\boldsymbol{\lambda}^n)' \mathbf{D}(\boldsymbol{\lambda}^n) \nabla \Phi(\boldsymbol{\lambda}^n) < \infty$. Given any $\epsilon > 0$, suppose that there exists k such that $\nabla \Phi(\boldsymbol{\lambda}^n)' \mathbf{D}(\boldsymbol{\lambda}^n) \nabla \Phi(\boldsymbol{\lambda}^n) > \epsilon$, $\forall n > k$. Then since $\sum_{n=0}^{\infty} \alpha_n = \infty$, we have $\sum_{n=0}^{\infty} \alpha_n \nabla \Phi(\boldsymbol{\lambda}^n)' \mathbf{D}(\boldsymbol{\lambda}^n) \nabla \Phi(\boldsymbol{\lambda}^n) = \infty$, which is a contradiction. So it must be the case that there exists a subsequence $\{\boldsymbol{\lambda}^{n_k}\}$ of $\{\boldsymbol{\lambda}^n\}$ such that $\lim_{k \rightarrow \infty} \boldsymbol{\lambda}^{n_k} = \boldsymbol{\lambda}^* \in \mathcal{B}$ and $\nabla \Phi(\boldsymbol{\lambda}^*)' \mathbf{D}(\boldsymbol{\lambda}^*) \nabla \Phi(\boldsymbol{\lambda}^*) = 0$. Therefore, $\mathbf{D}(\boldsymbol{\lambda}^*) \nabla \Phi(\boldsymbol{\lambda}^*) = \mathbf{0}$ because $\mathbf{D}(\boldsymbol{\lambda}^*)$ is a nonnegative definite diagonal matrix.

On the other hand, from (C.3), one can show that $\{\Phi(\boldsymbol{\lambda}^n)\}$ is a Cauchy sequence in \mathbb{R} in view of $\sum_{n=0}^{\infty} \alpha_n^2 < \infty$ and $\sum_{n=0}^{\infty} \alpha_n \nabla \Phi(\boldsymbol{\lambda}^n)' \mathbf{D}(\boldsymbol{\lambda}^n) \nabla \Phi(\boldsymbol{\lambda}^n) < \infty$. This implies that $\{\Phi(\boldsymbol{\lambda}^n)\}$ converges [110, p. 46]. \square

Lemma C.3. *Suppose that $\{\boldsymbol{\lambda}^n\}$ is a sequence generated by (4.12) with $\alpha_n > 0$ such that $\lim_{n \rightarrow \infty} \alpha_n = 0$. If $\boldsymbol{\lambda}^{n,m} \in \mathcal{B}$, $\forall n, m$, then $\lim_{n \rightarrow \infty} (\boldsymbol{\lambda}^{n,m} - \boldsymbol{\lambda}^n) = \mathbf{0}$, $\forall m$.*

Proof. Since $\mathbf{D}(\boldsymbol{\lambda}) \nabla \Phi_m(\boldsymbol{\lambda})$ is bounded on \mathcal{B} , $\forall m$, and $\lim_{n \rightarrow \infty} \alpha_n = 0$, we have:

$$\boldsymbol{\lambda}^{n,m} - \boldsymbol{\lambda}^n = \alpha_n \sum_{k=1}^m \mathbf{D}(\boldsymbol{\lambda}^{n,k-1}) \nabla \Phi_k(\boldsymbol{\lambda}^{n,k-1}) \rightarrow \mathbf{0}$$

as $n \rightarrow \infty$. \square

Corollary C.4. $\lim_{n \rightarrow \infty} (\boldsymbol{\lambda}^{n+1} - \boldsymbol{\lambda}^n) = \mathbf{0}$.

Lemma C.5. *The limit point $\boldsymbol{\lambda}^* \in \mathcal{B}$ in Lemma C.2 such that $\mathbf{D}(\boldsymbol{\lambda}^*) \nabla \Phi(\boldsymbol{\lambda}^*) = \mathbf{0}$ is a maximizer of $\Phi(\boldsymbol{\lambda})$ over \mathcal{B} if $\boldsymbol{\lambda}^{n,m} \in \text{Int } \mathcal{B}$, $\forall n, m$.*

Proof. We extend the proof of Proposition 3 of [18]. It is clear that $(\partial\Phi/\partial\lambda_j)(\boldsymbol{\lambda}^*) = 0$ if $0 < \lambda_j^* < U$. Considering the optimality conditions [100, p. 203], we need to prove that $\lambda_j^* = 0$ implies $(\partial\Phi/\partial\lambda_j)(\boldsymbol{\lambda}^*) \leq 0$, and $\lambda_j^* = U$ implies $(\partial\Phi/\partial\lambda_j)(\boldsymbol{\lambda}^*) \geq 0$. Define $\mathcal{J} = \mathcal{J}_1 \cup \mathcal{J}_2$ where $\mathcal{J}_1 = \{j = 1, \dots, p : \lambda_j^* = 0 \text{ and } (\partial\Phi/\partial\lambda_j)(\boldsymbol{\lambda}^*) > 0\}$ and $\mathcal{J}_2 = \{j = 1, \dots, p : \lambda_j^* = U \text{ and } (\partial\Phi/\partial\lambda_j)(\boldsymbol{\lambda}^*) < 0\}$. We show that $\mathcal{J} = \emptyset$. Since $\nabla\Phi$ is continuous on \mathcal{B} , there exists $0 < \delta < U/2$ such that if $\boldsymbol{\lambda} \in \mathcal{B}_\delta$ then $(\partial\Phi/\partial\lambda_j)(\boldsymbol{\lambda}) > 0, \forall j \in \mathcal{J}_1$ and $(\partial\Phi/\partial\lambda_j)(\boldsymbol{\lambda}) < 0, \forall j \in \mathcal{J}_2$ where $\mathcal{B}_\delta \triangleq \{\boldsymbol{\lambda} \in \mathcal{B} : \|\boldsymbol{\lambda} - \boldsymbol{\lambda}^*\| < \delta\}$.

Suppose that $\boldsymbol{\lambda}^n \in \mathcal{B}_\delta$ where n is sufficiently large. Then, by using Lemma C.3, we have $\boldsymbol{\lambda}^{n,m} \in \mathcal{B}_\delta, \forall m$ since n is large. For $j \in \mathcal{J}_2$, since $d_j(\boldsymbol{\lambda}^{n,m}) = (U - \lambda_j^{n,m})/p_j, \forall m$, one can show:

$$U - \lambda_j^{n,m} = (U - \lambda_j^{n,m-1}) \left(1 - \frac{\alpha_n}{p_j} \frac{\partial\Phi_m}{\partial\lambda_j}(\boldsymbol{\lambda}^{n,m-1}) \right).$$

Then using the Lipschitz continuity of $\nabla\Phi_m$, we have:

$$\begin{aligned} U - \lambda_j^{n+1} &= (U - \lambda_j^n) \prod_{m=1}^M \left(1 - \frac{\alpha_n}{p_j} \frac{\partial\Phi_m}{\partial\lambda_j}(\boldsymbol{\lambda}^{n,m-1}) \right) \\ &= (U - \lambda_j^n) \left(1 - \frac{\alpha_n}{p_j} \sum_{m=1}^M \frac{\partial\Phi_m}{\partial\lambda_j}(\boldsymbol{\lambda}^{n,m-1}) + O(\alpha_n^2) \right) \\ &= (U - \lambda_j^n) \left(1 - \frac{\alpha_n}{p_j} \frac{\partial\Phi}{\partial\lambda_j}(\boldsymbol{\lambda}^n) + O(\alpha_n^2) \right). \end{aligned}$$

Now we have $U - \lambda_j^{n+1} > U - \lambda_j^n$, that is, $\lambda_j^{n+1} < \lambda_j^n$ since $(\partial\Phi/\partial\lambda_j)(\boldsymbol{\lambda}^n) < 0$. Similarly, one can show that $\lambda_j^{n+1} > \lambda_j^n$ for $j \in \mathcal{J}_1$.

Let $\{\boldsymbol{\lambda}^{n_k}\}$ be a subsequence of $\{\boldsymbol{\lambda}^n\}$ such that $\lim_{k \rightarrow \infty} \boldsymbol{\lambda}^{n_k} = \boldsymbol{\lambda}^*$. Let $t_k = \max\{q < n_k : \boldsymbol{\lambda}^q \notin \mathcal{B}_\delta\}$. If $\boldsymbol{\lambda}^q \in \mathcal{B}_\delta, \forall q < n_k$ for some k , set $t_k = 0$. Then $\{t_k\}$ is a monotone increasing sequence of nonnegative integers such that $\boldsymbol{\lambda}^q \in \mathcal{B}_\delta$ for $t_k + 1 \leq q \leq n_k$ for large k . Suppose that $\lim_{k \rightarrow \infty} t_k = t < \infty$, that is, $\boldsymbol{\lambda}^n$ stays in \mathcal{B}_δ for large n . Then $\lambda_j^k > \lambda_j^l > 0, \forall k > l, \forall j \in \mathcal{J}_1$ and $\lambda_j^k < \lambda_j^l < U, \forall k > l, \forall j \in \mathcal{J}_2$ for some large l . This is a contradiction since we have assumed that $\{\boldsymbol{\lambda}^n\}$ has a limit point $\boldsymbol{\lambda}^*$ such that

$\lambda_j^* = 0, \forall j \in \mathcal{J}_1$ and $\lambda_j^* = U, \forall j \in \mathcal{J}_2$. So it must be the case that $\lim_{k \rightarrow \infty} t_k = \infty$. Now we have $\lambda_j^{n_k} > \lambda_j^{t_k+1} \geq 0, \forall j \in \mathcal{J}_1$ and $\lambda_j^{n_k} < \lambda_j^{t_k+1} \leq U, \forall j \in \mathcal{J}_2$ for large k . Since $\lim_{k \rightarrow \infty} \lambda_j^{n_k} = 0, \forall j \in \mathcal{J}_1$ and $\lim_{k \rightarrow \infty} \lambda_j^{n_k} = U, \forall j \in \mathcal{J}_2$, we have $\lim_{k \rightarrow \infty} \lambda_j^{t_k+1} = 0, \forall j \in \mathcal{J}_1$ and $\lim_{k \rightarrow \infty} \lambda_j^{t_k+1} = U, \forall j \in \mathcal{J}_2$. By Corollary C.4, $\lim_{k \rightarrow \infty} \lambda_j^{t_k} = 0, \forall j \in \mathcal{J}_1$ and $\lim_{k \rightarrow \infty} \lambda_j^{t_k} = U, \forall j \in \mathcal{J}_2$. Now one can construct a subsequence $\{\lambda^{t_{k_l}}\}$ of $\{\lambda^{t_k}\}$, which is also a subsequence of $\{\lambda^n\}$, such that $\lim_{l \rightarrow \infty} \lambda^{t_{k_l}} = \lambda^{**}$ with $\lambda_j^{**} = 0, \forall j \in \mathcal{J}_1$ and $\lambda_j^{**} = U, \forall j \in \mathcal{J}_2$ but $\lambda^{**} \neq \lambda^*$ (since $\lambda^{t_k} \notin \mathcal{B}_\delta$ and thus $\lambda^{**} \notin \mathcal{B}_\delta$). Then $\Phi(\lambda^*) = \Phi(\lambda^{**})$ by Lemma C.2. We have two different maximizers λ^* and λ^{**} of Φ over $\{\lambda \in \mathcal{B} : \lambda_j = 0, \forall j \in \mathcal{J}_1 \text{ and } \lambda_j = U, \forall j \in \mathcal{J}_2\}$. This is a contradiction since Φ is strictly concave. So it must be the case that $\mathcal{J} = \emptyset$. \square

Theorem C.6. *A sequence $\{\lambda^n\}$ generated by (4.12) with sufficiently small $\alpha_n > 0$ such that $\sum_{n=0}^{\infty} \alpha_n = \infty$ and $\sum_{n=0}^{\infty} \alpha_n^2 < \infty$, converges to $\hat{\lambda} = \arg \max_{\lambda \in \mathcal{B}} \Phi(\lambda)$.*

Proof. By Lemmas C.1, C.2, and C.5, the maximizer $\hat{\lambda}$ is a limit point of $\{\lambda^n\}$. Suppose that λ^{**} is a limit point of $\{\lambda^n\}$. Then $\Phi(\lambda^{**}) = \Phi(\hat{\lambda})$ by Lemma C.2. This implies that λ^{**} is also a maximizer. By the uniqueness of the maximizer, $\lambda^{**} = \hat{\lambda}$. So $\{\lambda^n\}$ has a unique limit point $\hat{\lambda}$. This implies that the bounded sequence $\{\lambda^n\}$ converges to $\hat{\lambda}$ by Proposition A.5 of [15, p. 652]. \square

Corollary C.7. $\lim_{n \rightarrow \infty} \lambda^{n,m} = \hat{\lambda}, \forall m$.

Proof. Use Lemma C.3 and Theorem C.6. \square

APPENDIX D

Convergence Proofs for Diagonally-Scaled Incremental Gradient Method

We prove the global convergence of the diagonally-scaled incremental gradient method in (4.21). The required assumptions on the objective function are the following: $\nabla\Phi_m$ are bounded on \mathcal{B} , and Φ_m are concave. They are satisfied by the emission PL objective function. Define a norm $\|\cdot\|_{D^{-1}}$ on \mathbb{R}^p by $\|\lambda\|_{D^{-1}} = (\lambda'D^{-1}\lambda)^{1/2}$ for $\lambda \in \mathbb{R}^p$. Suppose that $\Phi^* = \sup_{\lambda \in \mathcal{B}} \Phi(\lambda)$.

Lemma D.1. *Let $\{\lambda^n\}$ be a sequence generated by (4.21). Then for any $\lambda \in \mathcal{B}$, one can show:*

$$\|\lambda^{n+1} - \lambda\|_{D^{-1}}^2 \leq \|\lambda^n - \lambda\|_{D^{-1}}^2 - 2\alpha_n(\Phi(\lambda) - \Phi(\lambda^n)) + \alpha_n^2 C$$

for all n and some $C > 0$.

Proof. One can verify that the algorithm in (4.21) is equivalent to the following:

$$\eta^{n,m} = \mathcal{P}_{\mathcal{B}'}(\eta^{n,m-1} + \alpha_n \nabla g_m(\eta^{n,m-1}))$$

for $m = 1, \dots, M$, where $\eta^{n,m} \triangleq D^{-1/2}\lambda^{n,m}$, $g_m(\eta) = \Phi_m(D^{1/2}\eta)$, and $\mathcal{B}' = \{\eta \in \mathbb{R}^p : 0 \leq \eta_j \leq Ud_j^{-1/2}\}$. Use Lemma 2.1 of [89]. □

Lemma D.2. *Suppose that $\{\lambda^n\}$ is a sequence generated by (4.21) with $\alpha_n > 0$ such that $\lim_{n \rightarrow \infty} \alpha_n = 0$ and $\sum_{n=0}^{\infty} \alpha_n = \infty$. Then $\limsup_{n \rightarrow \infty} \Phi(\lambda^n) = \Phi^*$.*

Proof. The proof is due to Proposition 1.2 of [28]. Assume for contradiction that there are $\delta > 0$, $N \in \mathbb{N}$, and $\nu \in \mathcal{B}$ such that $\Phi(\nu) > \Phi(\lambda^n) + \delta$ for all $n \geq N$. Since $\lim_{n \rightarrow \infty} \alpha_n = 0$, one can assume that N is so large that $\alpha_n C < \delta$ where $C > 0$ is a constant from Lemma D.1. Using Lemma D.1, one obtains:

$$\begin{aligned} \|\lambda^{n+1} - \nu\|_{D^{-1}}^2 &\leq \|\lambda^n - \nu\|_{D^{-1}}^2 + \alpha_n(\alpha_n C - 2\delta) \\ &\leq \|\lambda^n - \nu\|_{D^{-1}}^2 - \alpha_n \delta \end{aligned}$$

for all $n \geq N$. Summing up, this gives:

$$0 \leq \|\lambda^n - \nu\|_{D^{-1}}^2 \leq \|\lambda^N - \nu\|_{D^{-1}}^2 - \delta \sum_{k=N}^{n-1} \alpha_k$$

for all $n > N$. This is a contradiction since $\sum_{n=0}^{\infty} \alpha_n = \infty$. \square

Theorem D.3. *Let $\{\lambda^n\}$ be the sequence generated by (4.21) with $\alpha_n > 0$ such that $\sum_{n=0}^{\infty} \alpha_n = \infty$ and $\sum_{n=0}^{\infty} \alpha_n^2 < \infty$. Then $\{\lambda^n\}$ converges to some $\lambda^* \in \Lambda^* = \{\nu \in \mathcal{B} : \Phi(\nu) \geq \Phi(\lambda), \forall \lambda \in \mathcal{B}\}$.*

Proof. Using Lemma D.1 with some $\nu \in \Lambda^*$, we have:

$$\|\lambda^{n+1} - \nu\|_{D^{-1}}^2 \leq \|\lambda^0 - \nu\|_{D^{-1}}^2 - 2 \sum_{k=0}^n \alpha_k (\Phi^* - \Phi(\lambda^k)) + \sum_{k=0}^n \alpha_k^2 C \quad (\text{D.1})$$

for all n . Since $\sum_{n=0}^{\infty} \alpha_n^2 < \infty$, we have

$$2 \sum_{k=0}^n \alpha_k (\Phi^* - \Phi(\lambda^k)) \leq \|\lambda^0 - \nu\|_{D^{-1}}^2 + \sum_{k=0}^n \alpha_k^2 C < q < \infty$$

for all n and some q where C is a constant from Lemma D.1. This implies that $\sum_{k=0}^{\infty} \alpha_k (\Phi^* - \Phi(\lambda^k)) < \infty$ since $\Phi^* - \Phi(\lambda^k) \geq 0, \forall k$. Therefore, (D.1) implies that $\{\lambda^n\}$ is bounded. By Lemma D.2, there exists a subsequence $\{\lambda^{n_k}\}$ of $\{\lambda^n\}$ such that $\lim_{k \rightarrow \infty} \Phi(\lambda^{n_k}) = \Phi^*$. Since $\{\lambda^{n_k}\}$ is bounded, there exists a subsequence $\{\lambda^{n_{k_l}}\}$ of $\{\lambda^{n_k}\}$ such that $\{\lambda^{n_{k_l}}\}$ converges to some $\lambda^* \in \mathcal{B}$ [15, p. 652]. By the continuity of Φ , we have $\Phi(\lambda^*) = \Phi^*$,

that is, $\boldsymbol{\lambda}^* \in \Lambda^*$. We have obtained a limit point $\boldsymbol{\lambda}^* \in \Lambda^*$ of $\{\boldsymbol{\lambda}^n\}$. Now we follow the line of the proof of Proposition 1.3 of [28]. For any $\delta > 0$, take $N \in \mathbb{N}$ such that $\|\boldsymbol{\lambda}^N - \boldsymbol{\lambda}^*\|_{\mathcal{D}^{-1}}^2 \leq \delta/2$ and $\sum_{k=N}^{\infty} (-2\alpha_k(\Phi^* - \Phi(\boldsymbol{\lambda}^k)) + \alpha_k^2 C) \leq \delta/2$. Using Lemma D.1, one obtains:

$$\|\boldsymbol{\lambda}^{n+1} - \boldsymbol{\lambda}^*\|_{\mathcal{D}^{-1}}^2 \leq \|\boldsymbol{\lambda}^N - \boldsymbol{\lambda}^*\|_{\mathcal{D}^{-1}}^2 + \sum_{k=N}^n (-2\alpha_k(\Phi^* - \Phi(\boldsymbol{\lambda}^k)) + \alpha_k^2 C) \leq \delta$$

for all $n \geq N$. □

Corollary D.4. $\lim_{n \rightarrow \infty} \boldsymbol{\lambda}^{n,m} = \boldsymbol{\lambda}^* \in \Lambda^*$, $\forall m$.

Proof. Use $\lim_{n \rightarrow \infty} \alpha_n = 0$ with the assumption that $\nabla \Phi_m$ is bounded on \mathcal{B} . □

APPENDIX E

Convergence Proofs for Incremental Optimization Transfer Algorithms

In this appendix we prove the convergence of the incremental optimization transfer algorithm given in Table 5.1. Define $\mathbf{z} \triangleq (\mathbf{x}; \bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M) \in \mathcal{D}^{M+1}$, and define a mapping $\mathcal{M} : \mathcal{D}^{M+1} \rightarrow \mathcal{D}^{M+1}$ such that $\mathcal{M}(\mathbf{z}^n) = \mathbf{z}^{n+1}$ where $\mathbf{z}^{n+1} = (\mathbf{x}^{n+1}; \bar{\mathbf{x}}_1^{n+1}, \dots, \bar{\mathbf{x}}_M^{n+1})$ is computed by (5.8)–(5.10) for $\mathbf{z}^n = (\mathbf{x}^n; \bar{\mathbf{x}}_1^n, \dots, \bar{\mathbf{x}}_M^n)$. Suppose that the algorithm generates a sequence $\{\mathbf{z}^n\}$ (or a sequence $\{\mathbf{x}^n\}$ by taking the first component of \mathbf{z}^n), given some initial point $\mathbf{z}^0 \in \mathcal{D}^{M+1}$. Define an augmented solution set as follows:

$$\Lambda \triangleq \{\mathbf{z} = (\mathbf{x}; \mathbf{x}, \dots, \mathbf{x}) \in \mathcal{D}^{M+1} : \mathbf{x} \in \Gamma\} \quad (\text{E.1})$$

where Γ is defined in (5.12). We impose the following assumptions.

Assumption E.1. Each Φ_m and $\phi_m(\cdot; \cdot)$ is continuously differentiable on a nonempty, closed, and convex set $\mathcal{D} \subset \mathbb{R}^p$ and $\mathcal{D}^2 \subset \mathbb{R}^p \times \mathbb{R}^p$, respectively.

Assumption E.2. The iterates $\{\mathbf{z}^n\}$ are bounded where $\mathbf{z}^n = (\mathbf{x}^n; \bar{\mathbf{x}}_1^n, \dots, \bar{\mathbf{x}}_M^n)$.

Assumption E.2 is ensured by either of the following sufficient conditions.

Assumption E.2'. The feasible set \mathcal{D} is bounded.

Assumption E.2''. A level set defined by $\{\mathbf{z} \in \mathcal{D}^{M+1} : F(\mathbf{z}) \geq F(\mathbf{z}^0)\}$ is bounded.

We assume that the surrogates ϕ_m satisfy the following conditions.

Condition E.3. The functionals ϕ_m satisfy the minorization conditions in (5.3).

Condition E.4. The following derivatives match for all m and $\mathbf{x} \in \mathcal{D}$:

$$\nabla\Phi_m(\mathbf{x}) = \nabla^{10}\phi_m(\mathbf{x}; \mathbf{x}) \quad (\text{E.2})$$

where ∇^{10} is the column gradient operator with respect to the first argument (see [58] for less restrictive conditions).

Condition E.5. There exists a *unique* maximizer in (5.8).

The following is sufficient for Condition E.5.

Condition E.5'. Each $\phi_m(\cdot; \bar{\mathbf{x}}_m)$ is strictly concave for all $\bar{\mathbf{x}}_m \in \mathcal{D}$, and there exists a maximizer of $F(\cdot; \bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M)$ over \mathcal{D} for all $\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M \in \mathcal{D}$.

Using the above assumptions and conditions, we prove a series of lemmas necessary for proving convergence.

Lemma E.6. *The iterates $\{\mathbf{z}^n\}$ generated by (5.8)–(5.10) yield monotonic increases in F , that is, $F(\mathbf{z}^{n+1}) \geq F(\mathbf{z}^n)$ for all n .*

Proof. It follows from the cyclic block coordinate ascent updates in (5.8) and (5.9). \square

Lemma E.7. *Suppose that $\mathbf{z}^* \in \mathcal{D}^{M+1}$ is a fixed point of \mathcal{M} , that is, $\mathcal{M}(\mathbf{z}^*) = \mathbf{z}^*$. Then $\mathbf{z}^* \in \Lambda$ where Λ is defined in (E.1).*

Proof. For the fixed point $\mathbf{z}^* = (\mathbf{x}^*; \bar{\mathbf{x}}_1^*, \dots, \bar{\mathbf{x}}_M^*)$, in view of Condition E.5, one can show that $\mathbf{x}^* = \bar{\mathbf{x}}_1^* = \dots = \bar{\mathbf{x}}_M^*$. Since \mathbf{x}^* is a maximizer of $\sum_{m=1}^M \phi_m(\cdot; \mathbf{x}^*)$ over \mathcal{D} , it follows that $\sum_{m=1}^M \nabla^{10}\phi_m(\mathbf{x}^*; \mathbf{x}^*)'(\mathbf{x} - \mathbf{x}^*) \leq 0$ for all $\mathbf{x} \in \mathcal{D}$ [15, p. 194]. Therefore, by Condition E.4, $\nabla\Phi(\mathbf{x}^*)'(\mathbf{x} - \mathbf{x}^*) \leq 0$ for all $\mathbf{x} \in \mathcal{D}$, and it follows that $\mathbf{x}^* \in \Gamma$. \square

Lemma E.8. *If $\mathbf{z} \notin \Lambda$, then $F(\mathcal{M}(\mathbf{z})) > F(\mathbf{z})$.*

Proof. The $z \notin \Lambda$ is not a fixed point of \mathcal{M} by Lemma E.7. Combining Condition E.5 and Lemma E.6 leads to the conclusion. \square

Now we prove the following theorem on the convergence of the incremental optimization transfer algorithm.

Theorem E.9. *Suppose that $\{z^n\}$ is a sequence generated by (5.8)–(5.10) with $z^0 \in \mathcal{D}^{M+1}$ and that Assumptions E.1 and E.2 and Conditions E.3–E.5 hold. Then any limit point of $\{z^n\}$ is an element of Λ .*

Proof. Following [78, p. 209 and p. 228], one can show that the mapping \mathcal{M} is closed, in other words, \mathcal{M} is continuous. The conclusion then follows from Zangwill’s Convergence Theorem [136, p. 91] with Assumption E.2, Lemmas E.6 and E.8, and the closedness of \mathcal{M} . \square

The following corollaries and lemmas also hold when “ $\{z^n\}$ ” is replaced with “ $\{\bar{x}_m^n\}$ ” for all m .

Corollary E.10. *Suppose $\{x^n\}$ is a sequence obtained by taking the first component from z^n in Theorem E.9. Then any limit point of $\{x^n\}$ is an element of Γ .*

Proof. Use Theorem E.9, Assumption E.2, and the definition of Λ in (E.1). \square

Corollary E.11. *If Φ is concave, then any limit point of $\{x^n\}$ is a global maximizer of Φ over \mathcal{D} . Moreover, if Φ is strictly concave, then $\{x^n\}$ converges to the global maximizer of Φ over \mathcal{D} .*

Proof. Use Corollary E.10 and [15, Proposition 2.1.2]. \square

When Φ is not strictly concave, there is no guarantee that the algorithm converges to a limit. However, convergence can be established by additionally assuming that the solution set Γ is discrete.

Lemma E.12. *Suppose $\{\mathbf{x}^n\}$ is a sequence from Corollary E.10. Then $\|\mathbf{x}^{n+1} - \mathbf{x}^n\| \rightarrow 0$.*

Proof. It follows from [82, Theorem 3.1] that $\|\mathbf{z}^{n+1} - \mathbf{z}^n\| \rightarrow 0$. Since $\|\mathbf{z}^{n+1} - \mathbf{z}^n\|^2 = \|\mathbf{x}^{n+1} - \mathbf{x}^n\|^2 + \sum_{m=1}^M \|\bar{\mathbf{x}}_m^{n+1} - \bar{\mathbf{x}}_m^n\|^2$, it must be a case that $\|\mathbf{x}^{n+1} - \mathbf{x}^n\| \rightarrow 0$. \square

Lemma E.13. *Suppose $\{\mathbf{x}^n\}$ is a sequence from Corollary E.10. Additionally, suppose that the set Γ is discrete. Then $\{\mathbf{x}^n\}$ converges to an element in Γ .*

Proof. Let S be a set of limit points of $\{\mathbf{x}^n\}$. Then $S \subset \Gamma$ by Theorem E.9. But, by Lemma E.12, S is connected [97, p. 173]. Since S is both discrete and connected, it is a singleton. \square

The above lemma implies that if stationary points of (5.2) are isolated, then the algorithm converges to one of them.

APPENDIX F

Local Convergence Rate Analysis for Incremental Optimiazaition Transfer Algorithms

F.1 Asymptotic Convergence Rate

We analyze the asymptotic convergence rate of the incremental optimization transfer algorithm given in Table 5.1. As in usual local convergence analysis, we assume that a sequence $\{\bar{\mathbf{x}}_m^n\}_{n=1}^\infty$ generated by the algorithm converges to an optimal point $\hat{\mathbf{x}}$ of (5.2) for all m , and that every iterate $\bar{\mathbf{x}}_m^n$ and the limit $\hat{\mathbf{x}}$ lie in the interior of \mathcal{D} .

Consider the following first-order Taylor's expansion of $\nabla^{10}\phi_m(\cdot; \bar{\mathbf{x}}_m^n)$ with respect to the first argument about $\bar{\mathbf{x}}_m^n$:

$$\nabla^{10}\phi_m(\mathbf{x}; \bar{\mathbf{x}}_m^n) \approx \nabla^{10}\phi_m(\bar{\mathbf{x}}_m^n; \bar{\mathbf{x}}_m^n) + \nabla^{20}\phi_m(\bar{\mathbf{x}}_m^n; \bar{\mathbf{x}}_m^n)(\mathbf{x} - \bar{\mathbf{x}}_m^n) \quad (\text{F.1})$$

where ∇^{20} is the Hessian operator with respect to the first argument. The first term on the right hand side can be further approximated as

$$\begin{aligned} \nabla^{10}\phi_m(\bar{\mathbf{x}}_m^n; \bar{\mathbf{x}}_m^n) &= \nabla\Phi_m(\bar{\mathbf{x}}_m^n) \\ &\approx \nabla\Phi_m(\hat{\mathbf{x}}) + \nabla^2\Phi_m(\hat{\mathbf{x}})(\bar{\mathbf{x}}_m^n - \hat{\mathbf{x}}) \end{aligned} \quad (\text{F.2})$$

where the equality is due to (E.2). Because of the assumption of $\hat{\mathbf{x}}$ and $\bar{\mathbf{x}}_1^{n+1}$ being in the interior of \mathcal{D} , and the construction of $\bar{\mathbf{x}}_1^{n+1} = \arg \max_{\mathbf{x} \in \mathcal{D}} F(\mathbf{x}; \bar{\mathbf{x}}_1^n, \dots, \bar{\mathbf{x}}_M^n)$ [see (5.8)

and (5.9)], it follows that

$$\begin{aligned} \sum_{m=1}^M \nabla \Phi_m(\hat{\mathbf{x}}) &= \nabla \Phi(\hat{\mathbf{x}}) = \mathbf{0} \\ \sum_{m=1}^M \nabla^{10} \phi_m(\bar{\mathbf{x}}_1^{n+1}; \bar{\mathbf{x}}_m^n) &= \mathbf{0}. \end{aligned}$$

Now combining (F.1) and (F.2) yields the following approximation:

$$\mathbf{e}_1^{n+1} \approx \left[\sum_{m=1}^M \nabla^{20} \phi_m(\bar{\mathbf{x}}_m^n; \bar{\mathbf{x}}_m^n) \right]^{-1} \sum_{m=1}^M [\nabla^{20} \phi_m(\bar{\mathbf{x}}_m^n; \bar{\mathbf{x}}_m^n) - \nabla^2 \Phi_m(\hat{\mathbf{x}})] \mathbf{e}_m^n$$

where $\mathbf{e}_m^n \triangleq \bar{\mathbf{x}}_m^n - \hat{\mathbf{x}}$ for all m and n . Similarly, one can obtain the following approximation for all m :

$$\begin{aligned} \mathbf{e}_m^{n+1} \approx & \left[\sum_{k=1}^{m-1} \nabla^{20} \phi_k(\bar{\mathbf{x}}_k^{n+1}; \bar{\mathbf{x}}_k^{n+1}) + \sum_{k=m}^M \nabla^{20} \phi_k(\bar{\mathbf{x}}_k^n; \bar{\mathbf{x}}_k^n) \right]^{-1} \cdot \\ & \left(\sum_{k=1}^{m-1} [\nabla^{20} \phi_k(\bar{\mathbf{x}}_k^{n+1}; \bar{\mathbf{x}}_k^{n+1}) - \nabla^2 \Phi_k(\hat{\mathbf{x}})] \mathbf{e}_k^{n+1} + \right. \\ & \left. \sum_{k=m}^M [\nabla^{20} \phi_k(\bar{\mathbf{x}}_k^n; \bar{\mathbf{x}}_k^n) - \nabla^2 \Phi_k(\hat{\mathbf{x}})] \mathbf{e}_k^n \right). \end{aligned} \quad (\text{F.3})$$

Assuming that $\nabla^{20} \phi_m(\cdot; \cdot)$ is continuous, it will converge to $\nabla^{20} \phi_m(\hat{\mathbf{x}}; \hat{\mathbf{x}})$ as $\lim_{n \rightarrow \infty} \bar{\mathbf{x}}_m^n = \hat{\mathbf{x}}$. For notational convenience, define $\mathbf{D}_m \triangleq \nabla^{20} \phi_m(\hat{\mathbf{x}}; \hat{\mathbf{x}})$, $\mathbf{H}_m \triangleq \nabla^2 \Phi_m(\hat{\mathbf{x}})$, and $\mathbf{T}_m \triangleq (\sum_{k=1}^M \mathbf{D}_k)^{-1} (\mathbf{D}_m - \mathbf{H}_m)$ for all m . Then one can write the *asymptotic* approximation of (F.3) in matrix form as follows:

$$\boldsymbol{\mathcal{E}}^{n+1} \approx (\mathbf{I}_{pM} - \boldsymbol{\Gamma}_l)^{-1} \boldsymbol{\Gamma}_u \boldsymbol{\mathcal{E}}^n$$

where $\boldsymbol{\mathcal{E}}^n \triangleq [(\mathbf{e}_1^n)', \dots, (\mathbf{e}_M^n)']'$ is a $pM \times 1$ column vector, \mathbf{I}_k is a $k \times k$ identity matrix,

and

$$\mathbf{\Gamma}_l = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{T}_1 & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{T}_1 & \mathbf{T}_2 & \cdots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & & \vdots & \vdots \\ \mathbf{T}_1 & \mathbf{T}_2 & \cdots & \mathbf{T}_{M-1} & \mathbf{0} \end{bmatrix} \quad (\text{F.4})$$

$$\mathbf{\Gamma}_u = \begin{bmatrix} \mathbf{T}_1 & \mathbf{T}_2 & \cdots & \mathbf{T}_{M-1} & \mathbf{T}_M \\ \mathbf{0} & \mathbf{T}_2 & \cdots & \mathbf{T}_{M-1} & \mathbf{T}_M \\ \vdots & \vdots & & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{T}_{M-1} & \mathbf{T}_M \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{T}_M \end{bmatrix} \quad (\text{F.5})$$

with $\mathbf{0}$ being a $p \times p$ zero matrix. Thus, the root-convergence factor [96, p. 288] of the sequence $\{[(\bar{\mathbf{x}}_1^n)', \dots, (\bar{\mathbf{x}}_M^n)']'\}_{n=1}^\infty$ for the incremental optimization transfer algorithm is given by the spectral radius

$$\rho_M = \rho((\mathbf{I}_{pM} - \mathbf{\Gamma}_l)^{-1} \mathbf{\Gamma}_u) \quad (\text{F.6})$$

where $\rho(\cdot)$ denotes spectral radius. One can show that the root-convergence factor of the sequence $\{\bar{\mathbf{x}}_m^n\}_{n=1}^\infty$ is also governed by the above spectral radius for all m . For ordinary optimization transfer algorithms, that is, when $M = 1$, the spectral radius (F.6) reduces to

$$\rho_1 = \rho(\mathbf{I}_p - [\nabla^2 \phi(\hat{\mathbf{x}}; \hat{\mathbf{x}})]^{-1} \nabla^2 \Phi(\hat{\mathbf{x}})), \quad (\text{F.7})$$

as is well known [70]. To compare ρ_1 and ρ_M for $M > 1$, we provide an illustrative example in the following subsection.

F.2 One-Parameter Example

We consider a simple one-parameter transmission problem. Suppose the measurement model is:

$$y_i \sim \text{Poisson}\{b_i e^{-ax}\}, \quad i = 1, \dots, N$$

where $a > 0$ and $b_i > 0, \forall i$. Assuming $\sum_{i=1}^N y_i > 0$, the ML estimate is given by

$$\hat{x} = \left[\frac{1}{a} \log \frac{\sum_{i=1}^N b_i}{\sum_{i=1}^N y_i} \right]_+$$

where $[x]_+ = \max\{x, 0\}$. Assuming $\hat{\mu} > 0$, which is very likely for high SNR data, the root-convergence factor ρ_M for an incremental optimization transfer algorithm, TRIOT-MC (see Section 5.3.2 for details), is given by (F.6) with substituting

$$\mathbf{T}_m = \frac{\sum_{i \in S_m} b_i}{\sum_{i=1}^N b_i} \left(1 - \frac{\sum_{i=1}^N y_i}{\sum_{i=1}^N b_i} \right), \quad m = 1, \dots, M \quad (\text{F.8})$$

in (F.4) and (F.5), where $\{S_m\}_{m=1}^M$ is a partition of $\{1, \dots, N\}$. Fig. F.1 shows the mean root-convergence factor $E[\rho_M]$ as a function of the number M of subsets for an example where $x^{\text{true}} = 0.7$, $N = 128$, $a = 1$, and b_i was simulated using pseudorandom uniform variates with mean of 0.5. The mean was approximately computed by replacing y_i in (F.8) with its mean $b_i e^{-ax^{\text{true}}}$; this approximation is reasonably accurate for high SNR. For example, for $M = 1$, that is, for a nonincremental version, ordinary SPS-MC, the mean of the root-convergence factor is given by

$$E[\rho_1] \approx 1 - e^{-ax^{\text{true}}}.$$

As shown in Fig. F.1, for this one-parameter example, the asymptotic convergence rates of incremental optimization transfer algorithms ($M > 1$) are faster than that of the nonincremental one ($M = 1$), and the convergence rate of the incremental one becomes faster as the number M of subsets increases.

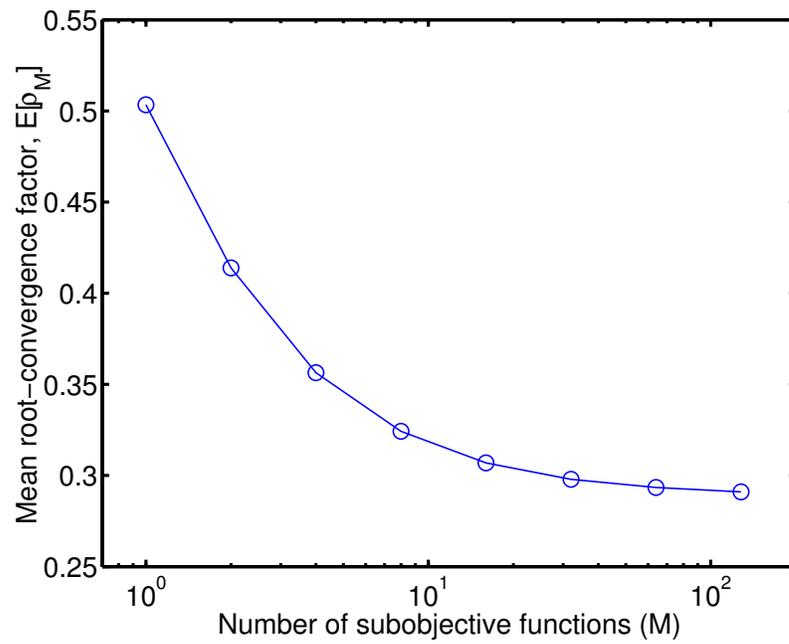


Figure F.1: Comparison of mean root-convergence factors of incremental SPS with different numbers M of subobjective functions for a one-parameter transmission problem. Nonincremental ordinary SPS corresponds to the case $M = 1$.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] H. Erdoğan and J. A. Fessler. Monotonic algorithms for transmission tomography. *IEEE Trans. Med. Imag.*, 18(9):801–814, September 1999.
- [2] H. Erdoğan and J. A. Fessler. Ordered subsets algorithms for transmission tomography. *Phys. Med. Biol.*, 44(11):2835–2851, November 1999.
- [3] S. Ahn and J. A. Fessler. Globally convergent ordered subsets algorithms: Application to tomography. In *Proc. IEEE Nuclear Science Symp. Medical Imaging Conf.*, volume 2, pages 1064–1068, 2001.
- [4] S. Ahn and J. A. Fessler. Globally convergent image reconstruction for emission tomography using relaxed ordered subsets algorithms. *IEEE Trans. Med. Imag.*, 22(3):613–626, May 2003.
- [5] S. Ahn and J. A. Fessler. Statistical emission image reconstruction for randoms-precorrected PET scans using negative sinogram values. In *Proc. IEEE Nuclear Science Symp. Medical Imaging Conf.*, 2003. To appear.
- [6] S. Ahn and J. A. Fessler. Emission image reconstruction for randoms-precorrected PET allowing negative sinogram values. *IEEE Trans. Med. Imag.*, 23(5):501–601, May 2004.
- [7] S. Ahn, J. A. Fessler, D. Blatt, and A. O. Hero. Incremental optimization transfer algorithms: Application to tomography. *IEEE Trans. Med. Imag.*, 2004. Submitted.
- [8] S. Ahn, J. A. Fessler, D. Blatt, and A. O. Hero. Incremental optimization transfer algorithms: application to transmission tomography. In *Proc. IEEE Nuclear Science Symp. Medical Imaging Conf.*, 2004. To appear.
- [9] S. Ahn, J. A. Fessler, T. E. Nichols, and R. A. Koeppe. Covariance of kinetic parameter estimators based on time activity curve reconstructions: Preliminary study on 1D dynamic imaging. In *Proc. IEEE Intl. Symp. Biomedical Imaging*, pages 368–371, 2004.
- [10] G. Arfken. *Mathematical methods for physicists*. Academic Press, Orlando, FL, 3 edition, 1985.
- [11] E. Asma, T. E. Nichols, and R. M. Leahy. Temporal resolution properties of dynamic PET reconstructions. In *Proc. IEEE Nuclear Science Symp. Medical Imaging Conf.*, volume 2, pages 778–781, 2002.
- [12] B. Bai, A. Ruangma, R. Laforest, Y. C. Tai, and R. M. Leahy. Positron range modeling for statistical PET image reconstruction. In *Proc. IEEE Nuclear Science Symp. Medical Imaging Conf.*, 2003.
- [13] M. S. Bazaraa, H. D. Sherali, and C. M. Shetty. *Nonlinear programming: Theory and algorithms*. Wiley, New York, 1993.
- [14] D. P. Bertsekas. A new class of incremental gradient methods for least squares problems. *SIAM J. Optim.*, 7(4):913–926, November 1997.
- [15] D. P. Bertsekas. *Nonlinear programming*. Athena Scientific, Belmont, 2 edition, 1999.

- [16] D. Blatt, A. Hero, and H. Gauchman. An incremental gradient method that converges with a constant step size. *SIAM J. Optim.*, 2004. Submitted.
- [17] C. Bouman and K. Sauer. A generalized Gaussian image model for edge-preserving MAP estimation. *IEEE Trans. Image Processing*, 2(3):296–310, July 1993.
- [18] J. A. Browne and A. R. De Pierro. A row-action alternative to the EM algorithm for maximizing likelihoods in emission tomography. *IEEE Trans. Med. Imag.*, 15(5):687–699, October 1996.
- [19] C. L. Byrne. Accelerating the EMLL algorithm and related iterative algorithms by rescaled block-iterative methods. *IEEE Trans. Image Processing*, 7(1):100–109, January 1998.
- [20] W. Byrne and A. Gunawardana. Comments on “Efficient training algorithms for HMMs using incremental estimation”. *IEEE Trans. Speech Audio Processing*, 8(6):751–754, November 2000.
- [21] R. E. Carson and K. Lange. The EM parametric image reconstruction algorithm. *J. Am. Stat. Ass.*, 80(389):20–22, March 1985.
- [22] M. E. Casey and E. J. Hoffman. Quantitation in positron emission computed tomography: 7 a technique to reduce noise in accidental coincidence measurements and coincidence efficiency calibration. *J. Comput. Assist. Tomogr.*, 10(5):845–850, 1986.
- [23] Y. Censor, P. P. B. Eggermont, and D. Gordon. Strong underrelaxation in Kaczmarz’s method for inconsistent systems. *Numerische Mathematik*, 41:83–92, 1983.
- [24] Y. Censor, D. Gordon, and R. Gordon. Component averaging: An efficient iterative parallel algorithm for large and sparse unstructured problems. *Parallel Computing*, 27(6):777–808, May 2001.
- [25] Y. Censor and S. A. Zenios. *Parallel optimization: Theory, algorithms, and applications*. Oxford University Press, New York, 1997.
- [26] P. C. Chiao, W. L. Rogers, N. H. Clinthorne, J. A. Fessler, and A. O. Hero. Model-based estimation for dynamic cardiac studies using ECT. *IEEE Trans. Med. Imag.*, 13(2):217–226, June 1994.
- [27] C. Cobelli, D. Foster, and G. Toffolo. *Tracer kinetics in biomedical research: from data to model*. Kluwer Academic, New York, 2000.
- [28] R. Correa and C. Lemarchal. Convergence of some algorithms for convex minimization. *Mathematical Programming*, 62:261–275, 1993.
- [29] A. R. De Pierro. On the relation between the ISRA and the EM algorithm for positron emission tomography. *IEEE Trans. Med. Imag.*, 12(2):328–333, June 1993.
- [30] A. R. De Pierro. A modified expectation maximization algorithm for penalized likelihood estimation in emission tomography. *IEEE Trans. Med. Imag.*, 14(1):132–137, March 1995.
- [31] A. R. De Pierro and M. E. B. Yamagishi. Fast EM-like methods for maximum ‘a posteriori’ estimates in emission tomography. *IEEE Trans. Med. Imag.*, 20(4):280–288, April 2001.
- [32] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Stat. Soc. Ser. B*, 39(1):1–38, 1977.
- [33] J. A. Fessler and H. Erdoğ̃an. A paraboloidal surrogates algorithm for convergent penalized-likelihood emission image reconstruction. In *Proc. IEEE Nuclear Science Symp. Medical Imaging Conf.*, volume 2, pages 1132–1135, 1998.
- [34] J. A. Fessler. Penalized weighted least-squares image reconstruction for positron emission tomography. *IEEE Trans. Med. Imag.*, 13(2):290–300, June 1994.

- [35] J. A. Fessler. ASPIRE 3.0 user's guide: A sparse iterative reconstruction library. Technical Report 293, Comm. and Sign. Proc. Lab., Dept. of EECS, Univ. of Michigan, Ann Arbor, MI, 48109-2122, July 1995. Available from <http://www.eecs.umich.edu/~fessler>.
- [36] J. A. Fessler. Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): Applications to tomography. *IEEE Trans. Image Processing*, 5(3):493–506, March 1996.
- [37] J. A. Fessler. Grouped coordinate descent algorithms for robust edge-preserving image restoration. In *Proc. SPIE 3170, Im. Recon. and Restor. II*, pages 184–194, 1997.
- [38] J. A. Fessler. Statistical image reconstruction methods for transmission tomography. In M. Sonka and J. Michael Fitzpatrick, editors, *Handbook of Medical Imaging, Volume 2. Medical Image Processing and Analysis*, pages 1–70. SPIE, Bellingham, 2000.
- [39] J. A. Fessler and A. O. Hero. Space-alternating generalized expectation-maximization algorithm. *IEEE Tr. Sig. Proc.*, 42(10):2664–2677, October 1994.
- [40] J. A. Fessler and A. O. Hero. Penalized maximum-likelihood image reconstruction using space-alternating generalized EM algorithms. *IEEE Trans. Image Processing*, 4(10):1417–1429, October 1995.
- [41] J. A. Fessler and W. L. Rogers. Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs. *IEEE Trans. Image Processing*, 5(9):1346–1358, September 1996.
- [42] R. Gordon, R. Bender, and G. T. Herman. Algebraic reconstruction techniques (ART) for the three-dimensional electron microscopy and X-ray photography. *J. Theor. Biol.*, 29:471–481, 1970.
- [43] P. J. Green. Iteratively reweighted least squares for maximum likelihood estimation, and some robust and resistant alternatives. *J. Roy. Stat. Soc. Ser. B*, 46(2):149–92, 1984.
- [44] A. Gunawardana and W. Byrne. Convergence of EM variants. Technical Report CLSP Research Note No. 32, ECE Dept., Johns Hopkins University, February 1999.
- [45] A. J. R. Gunawardana. *The information geometry of EM variants for speech and image processing*. PhD thesis, Johns Hopkins University, Baltimore, MD., 2001.
- [46] C. Hamaker and D. C. Solmon. The angles between the null spaces of x rays. *J. Math. Anal. Applic.*, 62:1–22, 1978.
- [47] G. T. Herman and L. B. Meyer. Algebraic reconstruction techniques can be made computationally efficient. *IEEE Trans. Med. Imag.*, 12(3):600–609, September 1993.
- [48] E. J. Hoffman, S. C. Huang, M. E. Phelps, and D. E. Kuhl. Quantitation in positron emission computed tomography: 4 Effect of accidental coincidences. *J. Comput. Assist. Tomogr.*, 5(3):391–400, 1981.
- [49] D. Hogg, K. Thielemans, S. Mustafovic, and T. J. Spinks. A study of bias for various iterative reconstruction methods in PET. In *Proc. IEEE Nuclear Science Symp. Medical Imaging Conf.*, volume 3, pages 1519–1523, 2002.
- [50] I. Hsiao, A. Rangarajan, and G. Gindi. A new convergent MAP reconstruction algorithm for emission tomography using ordered subsets and separable surrogates. In *Proc. IEEE Intl. Symp. Biomedical Imaging*, pages 409–412, 2002.
- [51] I. T. Hsiao and G. Gindi. Noise propagation from attenuation correction into PET reconstructions. *IEEE Trans. Nucl. Sci.*, 49(1):90–97, February 2002.
- [52] I. T. Hsiao, A. Rangarajan, and G. Gindi. A provably convergent OS-EM like reconstruction algorithm for emission tomography. In *Proc. SPIE 4684, Medical Imaging 2002: Image Proc.*, pages 10–19, 2002.

- [53] I. T. Hsiao, A. Rangarajan, P. Khurd, and G. Gindi. An accelerated convergent ordered subsets algorithm for emission tomography. *Phys. Med. Biol.*, 49(11):2145–2156, June 2004.
- [54] P. J. Huber. *Robust statistics*. Wiley, New York, 1981.
- [55] H. M. Hudson and R. S. Larkin. Accelerated image reconstruction using ordered subsets of projection data. *IEEE Trans. Med. Imag.*, 13(4):601–609, December 1994.
- [56] R. H. Huesman and B. M. Mazoyer. Kinetic data analysis with a noisy input function. *Phys. Med. Biol.*, 32(12):1569–1579, December 1987.
- [57] D. R. Hunter and K. Lange. Rejoinder to discussion of “Optimization transfer using surrogate objective functions”. *J. Computational and Graphical Stat.*, 9(1):53–59, March 2000.
- [58] M. W. Jacobson and J. A. Fessler. Properties of optimization transfer algorithms on convex feasible sets. *SIAM J. Optim.*, 2003. Submitted.
- [59] P. M. Joseph and R. D. Spital. A method for correcting bone induced artifacts in computed tomography scanners. *J. Comput. Assist. Tomogr.*, 2:100–108, January 1978.
- [60] P. Khurd, I. T. Hsiao, A. Rangarajan, and G. Gindi. A globally convergent regularized ordered-subset EM algorithm for list-mode reconstruction. *IEEE Trans. Nucl. Sci.* To appear.
- [61] V. M. Kibardin. Decomposition into functions in the minimization problem. *Automat. Remote Control*, 40:1311–1323, 1980. Translation of *Avtomatika i Telemekhanika*, vol. 9, pp. 66–79, Sept. 1979.
- [62] P. E. Kinahan, J. A. Fessler, and J. S. Karp. Statistical image reconstruction in PET with compensation for missing data. *IEEE Trans. Nucl. Sci.*, 44(4):1552–7, August 1997.
- [63] K. C. Kiwiel. Convergence of approximate and incremental subgradient methods for convex optimization. *SIAM J. Optim.*, 14(3):807–840, 2004.
- [64] R. Koeppe. *Data analysis and image processing*, 2002.
- [65] H. Kudo, H. Nakazawa, and T. Saito. Convergent block-iterative method for general convex cost functions. In *Proc. of the 1999 Int. Mtg. Fully 3D Im. Recon. in Rad. Nuc. Med.*, pages 247–250, 1999.
- [66] H. Kudo, H. Nakazawa, and T. Saito. Block-gradient method for image reconstruction in emission tomography. *Trans. IEICE*, J83-D-II(1):63–73, January 2000. In Japanese.
- [67] K. Lange. Convergence of EM image reconstruction algorithms with Gibbs smoothing. *IEEE Trans. Med. Imag.*, 9(4):439–446, December 1990. Corrections, T-MI, 10:2(288), June 1991.
- [68] K. Lange and R. Carson. EM reconstruction algorithms for emission and transmission tomography. *J. Comput. Assist. Tomogr.*, 8(2):306–316, April 1984.
- [69] K. Lange and J. A. Fessler. Globally convergent algorithms for maximum a posteriori transmission tomography. *IEEE Trans. Image Processing*, 4(10):1430–1438, October 1995.
- [70] K. Lange, D. R. Hunter, and I. Yang. Optimization transfer using surrogate objective functions. *J. Computational and Graphical Stat.*, 9(1):1–20, March 2000.
- [71] H Lantéri, M. Roche, and C. Aime. Penalized maximum likelihood image restoration with positivity constraints: multiplicative algorithms. *Inverse Prob.*, 18(5):1397–1419, October 2002.
- [72] H Lantéri, M. Roche, O. Cuevas, and C. Aime. A general method to devise maximum-likelihood signal restoration multiplicative algorithms with non-negativity constraints. *Signal Processing*, 81(5):945–974, May 2001.

- [73] R. M. Leahy and J. Qi. Statistical approaches in quantitative positron emission tomography. *Statistics and Computing*, 10(2):147–165, April 2000.
- [74] E. L. Lehmann and G. Casella. *Theory of point estimation*. Springer-Verlag, New York, 1998.
- [75] C. S. Levin and E. J. Hoffman. Calculation of positron range and its effect on the fundamental limit of positron emission tomography. *Phys. Med. Biol.*, 44(3):781–799, March 1999.
- [76] Q. Li, E. Asma, and R. M. Leahy. A fast fully 4D incremental gradient reconstruction algorithm for list mode PET data. In *Proc. IEEE Intl. Symp. Biomedical Imaging*, pages 555–558, 2004.
- [77] D. G. Luenberger. *Optimization by vector space methods*. Wiley, New York, 1969.
- [78] D. G. Luenberger. *Linear and nonlinear programming*. Addison-Wesley, Massachusetts, 2 edition, 1984.
- [79] A. Macovski. *Medical imaging systems*. Prentice-Hall, New Jersey, 1983.
- [80] S. Matej and R. M. Lewitt. Practical considerations for 3-D image reconstruction using spherically symmetric volume elements. *IEEE Trans. Med. Imag.*, 15(1):68–78, February 1996.
- [81] B. M. Mazoyer, R. H. Huesman, T. F. Budinger, and B. L. Knittel. Dynamic PET data study. *J. Comput. Assist. Tomogr.*, 10(4):645–653, 1986.
- [82] R. R. Meyer. Sufficient conditions for the convergence of monotonic mathematical programming algorithms. *J. Comput. System. Sci.*, 12(1):108–121, 1976.
- [83] C. Michel, X. Liu, S. Sanabria, M. Lonnew, M. Sibomana, A. Bol, C. Comtat, P. E. Kinahan, D. W. Townsend, and M. Defrise. Weighted schemes applied to 3D-OSEM reconstruction in PET. In *Proc. IEEE Nuclear Science Symp. Medical Imaging Conf.*, volume 3, pages 1152–1157, 1999.
- [84] C. Michel, M. Sibomana, A. Bol, X. Bernard, M. Lonnew, M. Defrise, C. Comtat, P. E. Kinahan, and D. W. Townsend. Preserving Poisson characteristics of PET data with weighted OSEM reconstruction. In *Proc. IEEE Nuclear Science Symp. Medical Imaging Conf.*, volume 2, pages 1323–1329, 1998.
- [85] E. Ü. Mumcuoğlu, R. Leahy, S. R. Cherry, and Z. Zhou. Fast gradient-based methods for Bayesian reconstruction of transmission and emission PET images. *IEEE Trans. Med. Imag.*, 13(4):687–701, December 1994.
- [86] E. Ü. Mumcuoğlu, R. M. Leahy, and S. R. Cherry. Bayesian reconstruction of PET images: methodology and performance analysis. *Phys. Med. Biol.*, 41(9):1777–1807, September 1996.
- [87] R. Neal and G. E. Hinton. A view of the EM algorithm that justifies incremental, sparse and other variants. In M. I. Jordan, editor, *Learning in Graphical Models*, pages 255–268. Kluwer, Dordrecht, 1998.
- [88] A. Nedić and D. Bertsekas. Convergence rate of incremental subgradient algorithms. In S. Uryasev and P. M. Pardalos, editors, *Stochastic Optimization: Algorithms and Applications*, pages 263–304. Kluwer, New York, 2000.
- [89] A. Nedić and D. P. Bertsekas. Incremental subgradient methods for nondifferentiable optimization. *SIAM J. Optim.*, 12(1):109–138, 2001.
- [90] T. E. Nichols, J. Qi, E. Asma, and R. M. Leahy. Spatiotemporal reconstruction of list mode PET data. *IEEE Trans. Med. Imag.*, 21(4):396–404, April 2002.
- [91] M. Nikolova and M. Ng. Fast image reconstruction algorithms combining half-quadratic regularization and preconditioning. In *Proc. IEEE Intl. Conf. on Image Processing*, volume 1, pages 277–80, 2001.

- [92] J. Nocedal. Large scale unconstrained optimization. In I. S. Duff and G. A. Watson, editors, *The State of the Art in Numerical Analysis*, pages 311–338. Clarendon Press, Oxford, 1997.
- [93] J. Nuyts, S. Stroobants, P. Dupont, S. Vleugels, P. Flamen, and L. Mortelmans. Reducing loss of image quality due to the attenuation artifact in uncorrected PET whole body images. *J. Nuc. Med.*, 43(8):1054–1062, August 2002.
- [94] J. M. Ollinger. Model-based scatter correction for fully 3D PET. *Phys. Med. Biol.*, 41(1):153–76, January 1996.
- [95] J. M. Ollinger and J. A. Fessler. Positron emission tomography. *IEEE Signal Processing Mag.*, 14(1):43–55, January 1997.
- [96] J. M. Ortega and W. C. Rheinboldt. *Iterative solution of nonlinear equations in several variables*. Academic, New York, 1970.
- [97] A. M. Ostrowski. *Solution of equations in Euclidean and Banach spaces*. Academic, New York, 3 edition, 1973.
- [98] M. E. Phelps, J. C. Mazziotta, and H. R. Schelbert (eds.). *Positron emission tomography and autoradiography*. Ravel Press, New York: NY, 1986.
- [99] D. G. Politte and D. L. Snyder. Corrections for accidental coincidences and attenuation in maximum-likelihood image reconstruction for positron-emission tomography. *IEEE Trans. Med. Imag.*, 10(1):82–89, March 1991.
- [100] B. T. Polyak. *Introduction to optimization*. Optimization Software Inc, New York, 1987.
- [101] J. Qi and R. H. Huesman. Theoretical study of lesion detectability of MAP reconstruction using computer observers. *IEEE Trans. Med. Imag.*, 20(8):815–22, August 2001.
- [102] J. Qi and R. M. Leahy. Resolution and noise properties MAP reconstruction for fully 3D PET. *IEEE Trans. Med. Imag.*, 19(5):493–506, May 2000.
- [103] J. Qi, R. M. Leahy, C. Hsu, T. H. Farquhar, and S. R. Cherry. Fully 3D Bayesian image reconstruction for the ECAT EXACT HR+. *IEEE Trans. Nucl. Sci.*, 45(3):1096–1103, June 1998.
- [104] A. Rahmim, M. Lenox, A. J. Reader, C. Michel, Z. Burbar, T. J. Ruth, and V. Sossi. Statistical list-mode image reconstruction for the high resolution research tomograph. *Phys. Med. Biol.*, 49:4239–4258, August 2004.
- [105] C. R. Rao. *Linear statistical inference and its applications*. Wiley, New York, 1973.
- [106] B. W. Reutter, G. T. Gullberg, and R. H. Huesman. Direct least squares estimation of spatiotemporal distributions from dynamic SPECT projections using a spatial segmentation and temporal B-splines. *IEEE Trans. Med. Imag.*, 19(5):434–450, May 2000.
- [107] B. W. Reutter, G. T. Gullberg, and R. H. Huesman. Effects of temporal modelling on the statistical uncertainty of spatiotemporal distributions estimated directly from dynamic SPECT projections. *Phys. Med. Biol.*, 47:2673–2683, 2002.
- [108] S. O. Rice. Uniform asymptotic expansions for saddle point integrals—Application to a probability distribution occurring in noise theory. *Bell Syst. Tech. J.*, 47:1971–2013, November 1968.
- [109] A. J. Rockmore and A. Macovski. A maximum likelihood approach to emission image reconstruction from projections. *IEEE Trans. Nucl. Sci.*, 23:1428–1432, 1976.
- [110] K. A. Ross. *Elementary analysis: the theory of calculus*. Springer-Verlag, New York, 1980.
- [111] K. Sauer and C. Bouman. A local update strategy for iterative reconstruction from projections. *IEEE Tr. Sig. Proc.*, 41(2):534–548, February 1993.

- [112] L. A. Shepp and Y. Vardi. Maximum likelihood reconstruction for emission tomography. *IEEE Trans. Med. Imag.*, 1(2):113–122, October 1982.
- [113] L. A. Shepp, Y. Vardi, J. B. Ra, S. K. Hilal, and Z. H. Cho. Maximum likelihood PET with real data. *IEEE Trans. Nucl. Sci.*, 31(2):910–913, April 1984.
- [114] D. L. Snyder. Statistical analysis of dynamic tracer data. *IEEE Tr. Biomed. Engin.*, 20(1):11–20, January 1973.
- [115] D. L. Snyder. *Random point processes*. Wiley, New York, 1975.
- [116] D. L. Snyder, C. W. Helstrom, A. D. Lanterman, M. Faisal, and R. L. White. Compensation for readout noise in CCD images. *J. Opt. Soc. Amer. A*, 12(2):272–283, February 1995.
- [117] J. A. Sorenson and M. E. Phelps. *Physics in nuclear medicine*. Saunders, Philadelphia, 2 edition, 1987.
- [118] S. Sotthivirat and J. A. Fessler. Statistical reconstruction for digital holography. *J. Opt. Soc. Amer. A*. To appear.
- [119] S. Sotthivirat and J. A. Fessler. Relaxed ordered-subsets algorithm for penalized-likelihood image restoration. *J. Opt. Soc. Amer. A*, 20(3):439–449, March 2003.
- [120] T. J. Spinks, T. Jones, M. C. Gilardi, and J. D. Heather. Physical performance of the latest generation of commercial positron scanner. *IEEE Trans. Nucl. Sci.*, 35(1):721–725, February 1988.
- [121] H. Stark and J. W. Woods. *Probability, random processes, and estimation theory for engineers*. Prentice-Hall, Englewood Cliffs, NJ, 2 edition, 1994.
- [122] J. W. Stayman and J. A. Fessler. Compensation for nonuniform resolution using penalized-likelihood reconstruction in space-variant imaging systems. *IEEE Trans. Med. Imag.* To appear.
- [123] J. W. Stayman and J. A. Fessler. Regularization for uniform spatial resolution properties in penalized-likelihood image reconstruction. *IEEE Trans. Med. Imag.*, 19(6):601–615, June 2000.
- [124] J. W. Stayman and J. A. Fessler. Compensation for nonuniform resolution using penalized-likelihood reconstruction in space-variant imaging systems. *IEEE Tr. Med. Im.*, 23(3):269–284, March 2004.
- [125] J. Stoer and R. Bulirsch. *Introduction to numerical analysis*. Springer-Verlag, New York, 1980.
- [126] G. Strang. *Linear algebra and its applications*. Academic Press, New York, 1980.
- [127] C. C. Watson, D. Newport, M. E. Casey, R. A. De Kemp, R. S. Beanlands, and M. Schmand. Evaluation of simulation-based scatter correction for 3-D PET cardiac imaging. *IEEE Trans. Nucl. Sci.*, 44(1):90–97, February 1997.
- [128] K. Wienhard, L. Eriksson, S. Grootoink, M. Casey, U. Pietrzyk, and W. D. Heiss. Performance evaluation of a new generation positron scanner ECAT EXACT. *J. Comput. Assist. Tomogr.*, 16(5):804–813, September 1992.
- [129] M. Yavuz. *Statistical tomographic image reconstruction methods for randoms-precorrected PET measurements*. PhD thesis, Univ. of Michigan, Ann Arbor, MI, 48109-2122, 1999.
- [130] M. Yavuz and J. A. Fessler. New statistical models for randoms-precorrected PET scans. In J. Duncan and G. Gindi, editors, *Information Processing in Medical Im.*, volume 1230 of *Lecture Notes in Computer Science*, pages 190–203. Springer-Verlag, Berlin, 1997.
- [131] M. Yavuz and J. A. Fessler. Statistical image reconstruction methods for randoms-precorrected PET scans. *Med. Imag. Anal.*, 2(4):369–378, December 1998.

- [132] M. Yavuz and J. A. Fessler. Penalized-likelihood estimators and noise analysis for randoms-precorrected PET transmission scans. *IEEE Trans. Med. Imag.*, 18(8):665–674, August 1999.
- [133] M. Yavuz and J. A. Fessler. Maximum likelihood emission image reconstruction for randoms-precorrected PET scans. In *Proc. IEEE Nuclear Science Symp. Medical Imaging Conf.*, volume 2, pages 15,229–15,233, 2000.
- [134] A. Yendiki and J. A. Fessler. A comparison of rotation- and blob-based system models for 3D SPECT with depth-dependent detector response. *Phys. Med. Biol.*, 49(11):2157–2168, June 2004.
- [135] D. F. Yu and J. A. Fessler. Mean and variance of coincidence photon counting with deadtime. *Nucl. Instr. Meth. Phys. Res. A.*, 488(1-2):362–374, August 2002.
- [136] W. Zangwill. *Nonlinear programming, a unified approach*. Prentice-Hall, NJ, 1969.
- [137] G. Zubal, G. Gindi, M. Lee, C. Harrell, and E. Smith. High resolution anthropomorphic phantom for Monte Carlo analysis of internal radiation sources. In *IEEE Symposium on Computer-Based Medical Systems*, pages 540–547, 1990.