

Statistical Methods for Image Reconstruction

Jeffrey A. Fessler

EECS Department
The University of Michigan

NSS-MIC

Nov. 12, 2002

0.0

These annotated slides were prepared by Jeff Fessler for attendees of the NSS-MIC short course on statistical image reconstruction methods.

The purpose of the annotation is to provide supplemental details, and particularly to provide extensive literature references for further study.

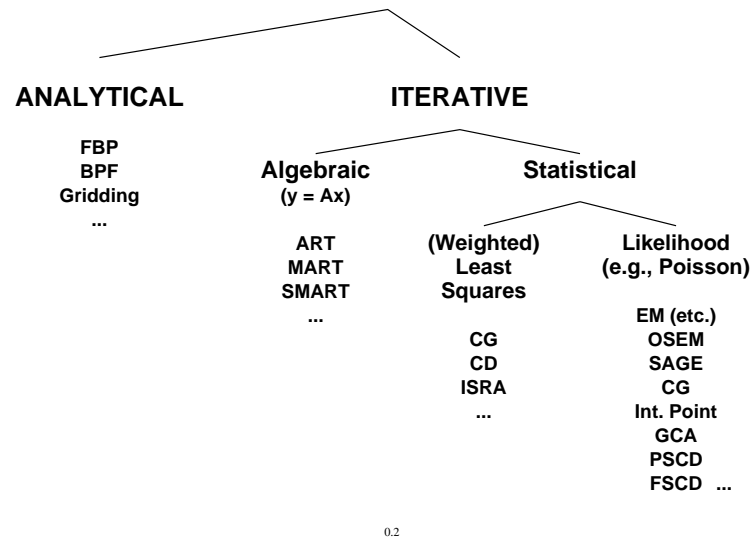
Image Reconstruction Methods (Simplified View)

Analytical
(FBP)

Iterative
(OSEM?)

0.1

Image Reconstruction Methods / Algorithms



0.2

Part of the goal is to bring order to this alphabet soup.

Outline

Part 0: Introduction / Overview / Examples

Part 1: From Physics to Statistics (Emission tomography)

- Assumptions underlying Poisson statistical model
- Emission reconstruction problem statement

Part 2: Four of Five Choices for Statistical Image Reconstruction

- Object parameterization
- System physical modeling
- Statistical modeling of measurements
- Cost functions and regularization

Part 3: Fifth Choice: Iterative algorithms

- Classical optimization methods
- Considerations: nonnegativity, convergence rate, ...
- Optimization transfer: EM etc.
- Ordered subsets / block iterative / incremental gradient methods

Part 4: Performance Analysis

- Spatial resolution properties
- Noise properties
- Detection performance

Part 5: Miscellaneous topics (?)

- ...

0.3

Emphasis on general principles rather than specific empirical results.

The journals (and conferences like NSS/MIC!) are replete with empirical comparisons.

Although the focus of examples in this course are PET / SPECT / CT, most of the principles apply equally well to other tomography problems like MR image reconstruction, optical / diffraction tomography, etc.

History

- Iterative method for X-ray CT (Hounsfield, 1968)
- ART for tomography (Gordon, Bender, Herman, JTB, 1970)
- Richardson/Lucy iteration for image restoration (1972, 1974)
- Weighted least squares for 3D SPECT (Goitein, NIM, 1972)
- Proposals to use Poisson likelihood for emission and transmission tomography
 - Emission: (Rockmore and Macovski, TNS, 1976)
 - Transmission: (Rockmore and Macovski, TNS, 1977)
- First expectation-maximization (EM) algorithms for Poisson model
 - Emission: (Shepp and Vardi, TMI, 1982)
 - Transmission: (Lange and Carson, JCAT, 1984)
- First regularized (aka Bayesian) Poisson emission reconstruction
 - Geman and McClure, ASA, 1985
- Ordered-subsets EM algorithm
 - Hudson and Larkin, TMI, 1994
- Commercial introduction of OSEM for PET scanners
 - circa 1997

0.4

X-ray CT patent: [2]

ART: [3–6]

Richardson/Lucy iteration for image restoration was not derived from ML considerations, but turns out to be the familiar ML-EM iteration [7, 8]

Emission: [9]

Transmission: [10]

General expectation-maximization (EM) algorithm (Dempster *et al.*, 1977) [11]

Emission EM algorithm: [12]

Transmission EM algorithm: [13]

Bayesian method for Poisson emission problem: [14]

OSEM [15]

Prior to the proposals for Poisson likelihood models, the Lawrence Berkeley Laboratory had proposed and investigated weighted least-squares (WLS) methods for SPECT (in 3D!) using iterative algorithms; see (Goitein, 1972) [16] and (Budinger and Gullberg, 1974) [17]. These methods became widely available in 1977 through the release of the Donner RECLBL package [18].

Of course there was lots of work ongoing based on “algebraic” reconstruction methods in the 1970s and before. But until WLS methods were proposed, this work was largely not “statistical.”

todo: According to Rob Lewitt: “On the history of iterative methods for tomographic reconstruction, David Kuhl was doing related work around 1966-72. I have a vague recollection that Bracewell’s 1956 paper mentions an iterative approach (as well as Fourier transforms).”

Why Statistical Methods?

- Object constraints (*e.g.*, nonnegativity, object support)
- Accurate physical models (less bias \Rightarrow improved quantitative accuracy)
 - improved spatial resolution?
 - (*e.g.*, nonuniform attenuation in SPECT)
- Appropriate statistical models (less variance \Rightarrow lower image noise)
 - (FBP treats all rays equally)
- Side information (*e.g.*, MRI or CT boundaries)
- Nonstandard geometries (“missing” data)

Disadvantages?

- Computation time
- Model complexity
- Software complexity

Analytical methods (a different short course!)


- Idealized mathematical model
 - Usually geometry only, greatly over-simplified physics
 - Continuum measurements
- No statistical model
- Easier analysis of properties (due to linearity)
 - e.g.*, Huesman (1984) FBP ROI variance for kinetic fitting

0.5

There is a continuum of physical system models that tradeoff accuracy and compute time. The “right” way to model the physics is usually too complicated, so one uses approximations. The sensitivity of statistical methods to those approximations needs more investigation.

FBP has its faults, but its properties (good and bad) are very well understood and hence predictable, due to its linearity. Spatial resolution, variance, ROI covariance (Huesman [19]), and autocorrelation have all been thoroughly analyzed (and empirical results agree with the analytical predictions). Only recently have such analyses been provided for *some* nonlinear reconstruction methods *e.g.*, [20–31].

What about Moore’s Law?



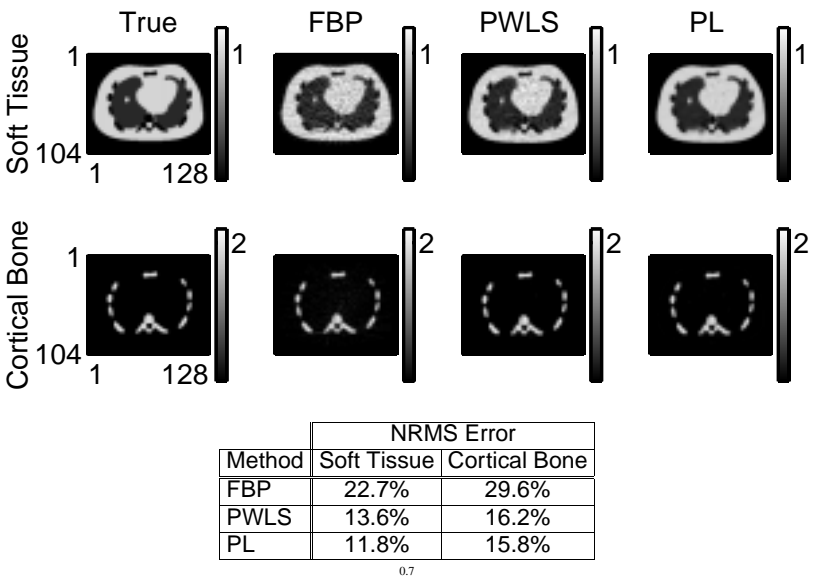
In this graph complexity is the number of lines of response (number of rays) acquired. The ECAT scanners can operate either in 2D mode (with septa in place) or 3D mode (with septa retracted) so those scanners have two points each.

I got this graph from Richard Leahy. I do not know where he got it, but presumably someone associated with Siemens/CTI since only CTI scanners (and their relatives) are represented.

There is considerable ongoing effort to reduce or minimize the compute time by more efficient algorithms.

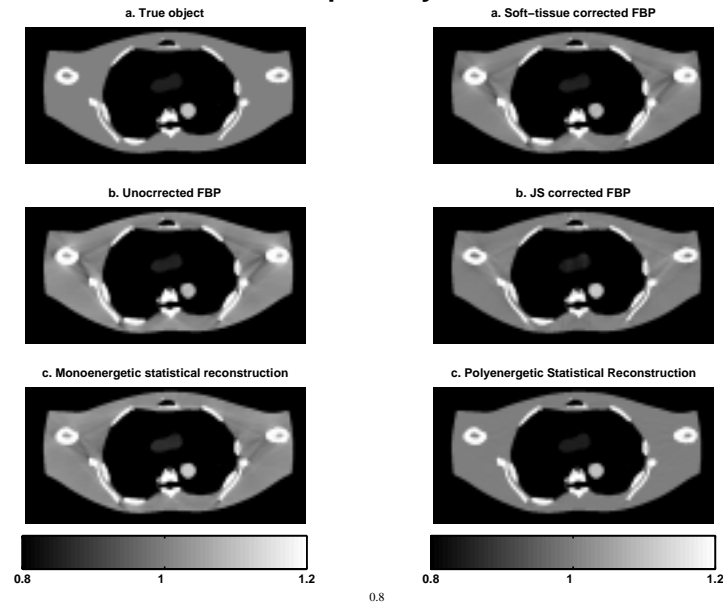
Moore's law for computing power increases will not alone solve all of the compute problems in image reconstruction. The problems increase in difficulty at nearly the same rate as the increase in compute power. (Consider the increased amount of data in 3D PET scanners relative to 2D.) (Or even the increased number of slices in 2D mode.) Or spiral CT, or fast dynamic MRI,... Therefore there is a need for further improvements in algorithms in addition to computer hardware advances.

Benefit Example: Statistical Models




Conventional FBP reconstruction of dual-energy X-ray CT data does not account for the noise properties of CT measurements and results in significant noise propagation into the soft tissue and cortical bone component images. Statistical reconstruction methods greatly reduces this noise, improving quantitative accuracy [32]. This is of potential importance for applications like bone density measurements.

Benefit Example: Physical Models




Conventional FBP ignores the polyenergetic X-ray source spectrum. Statistical/iterative reconstruction methods can build that spectrum into the model and nearly eliminate beam-hardening artifacts [33–35].

Benefit Example: Nonstandard Geometries



A SPECT transmission scan with 65cm distance between line source and standard Anger camera provides partially truncated sinogram views of most patients.

Truncated Fan-Beam SPECT Transmission Scan




0.10

The FBP reconstruction method is essentially ruined by the sinogram truncation.

Despite the partial truncation, each pixel is *partly* sampled by "line integrals" at some range of angles. With the benefit of spatial regularization, nonnegativity constraints, and statistical models, a statistical reconstruction method (PWLS in this case) can recover an attenuation map that is comparable to that obtained with an untruncated scan.

We have shown related benefits in PET with missing sinogram data [36].

One Final Advertisement: Iterative MR Reconstruction



0.11

MR signal equation:

$$s(t) = \int f(\vec{x}) \exp(-i\omega(\vec{x})t) \exp(-i2\pi k(\vec{x}) \cdot \vec{x}) d\vec{x}$$

- Due to field inhomogeneity, signal is *not* Fourier transform of object.
- Measure off-resonance field-map $\omega(\vec{x})$ using two displaced echos
- Penalized WLS cost function minimized by conjugate gradient
- System matrix A includes off-resonance effects
- Fast algorithm using NUFFT and time-segmentation

[37–39]

Hopefully that is enough motivation, so, on with the methodology!

Part 1: From Physics to Statistics

or
“What quantity is reconstructed?”
(in emission tomography)

Outline

- Decay phenomena and fundamental assumptions
- Idealized detectors
- Random phenomena
- Poisson measurement statistics
- State emission tomography reconstruction problem

1.1

The fact that “the measurements are Poisson” is well known, but the underlying assumptions for this fact are perhaps less so.

Since some systems in some modes of operation violate the assumptions (high deadtime), it may be worth revisiting the assumptions before launching into statistical reconstruction methods.

What Object is Reconstructed?

In *emission imaging*, our aim is to image the *radiotracer distribution*.

The what?

At time $t = 0$ we inject the patient with some *radiotracer*, containing a “large” number N of metastable atoms of some radionuclide.

Let $\vec{X}_k(t) \in \mathbb{R}^3$ denote the position of the k th *tracer atom* at time t . These positions are influenced by blood flow, patient physiology, and other unpredictable phenomena such as Brownian motion.

The ultimate imaging device would provide an exact list of the spatial locations $\vec{X}_1(t), \dots, \vec{X}_N(t)$ of all tracer atoms for the entire scan.

Would this be enough?

1.2

Here, “large” means in the statistical sense that for large N and small success probability, the binomial distribution closely approximates the Poisson distribution. However, in the biological sense N is tiny relative to the number of atoms in the body, hence a “tracer.”

Time-of-flight (TOF) PET with perfect timing resolution and perfect spatial resolution would approach this ultimate imaging device, but even then it would only provide the position of each tracer atom at the time it decays, rather than for all times.

Atom Positions or Statistical Distribution?

Repeating a scan would yield different tracer atom sample paths $\{\vec{x}_k(t)\}_{k=1}^N$.

\therefore statistical formulation

Assumption 1. The spatial locations of individual tracer atoms at any time $t \geq 0$ are *independent* random variables that are all *identically distributed* according to a common probability density function (pdf) $f_{\vec{x}(t)}(\vec{x})$.

This pdf is determined by patient physiology and tracer properties.

Larger values of $f_{\vec{x}(t)}(\vec{x})$ correspond to “hot spots” where the tracer atoms tend to be located at time t . Units: inverse volume, *e.g.*, atoms per cubic centimeter.

The *radiotracer distribution* $f_{\vec{x}(t)}(\vec{x})$ is the quantity of interest.

(Not $\{\vec{x}_k(t)\}_{k=1}^N$!)

1.3

This independence assumption should be very reasonable when *trace* quantities of radiotracer are injected. If a very large quantity of a radiotracer such as a neuroreceptor agent were injected, then the first wave of tracer atoms (and molecules) to reach the brain could occupy all or most available receptor sites, denying access to later arriving tracer atoms. This would lead to a statistical dependence between the tracer atom locations. Rarely are such large quantities injected, so our i.i.d. assumption is a reasonable starting point.

The identically distributed assumption should be reasonable if the injection duration is relatively short. If the injection is long enough that the patient's physiology has changed between the time the first tracer atoms were injected and the time the last tracer atoms were injected, then the distributions of those atoms' spatial locations will differ. However, in practice we can very rarely distinguish the first tracer atoms from the later tracer atoms, so the numbering from $k = 1$ to N is arbitrary, and could even be considered a random permutation of the indices, in which case the sample paths $\vec{x}_k(t)$ are again identically distributed according to an average radiotracer distribution. An exception would be multiple-injection studies with different radiotracers.


“density” would be more logical than “distribution”

Since real instruments have finite temporal resolution, we never really observe $f_{\vec{x}(t)}(\vec{x})$, but at best something like

$$\int_{t_1}^{t_2} f_{\vec{x}(t)}(\vec{x}) dt.$$

Except in list mode...

Example: Perfect Detector



True radiotracer distribution $f_{\vec{x}(t)}(\vec{x})$ at some time t .


A realization of $N = 2000$ i.i.d. atom positions (dots) recorded “exactly.”

Little similarity!

1.4

By “exactly” I mean there is no error in measuring the locations of the 2000 atoms. A non-ideal detector PSF would randomly relocate each point relative to its ideal location [40].

Binning/Histogram Density Estimator




Estimate of $f_{\vec{X}(t)}(\vec{x})$ formed by histogram binning of $N = 2000$ points. Ramp remains difficult to visualize.


1.5

Discretize object domain into small square pixels. Count number of atoms positioned within each square. Display gray-scale image with intensity proportional to number of atoms within each square.

Kernel Density Estimator



Gaussian kernel density estimator for $f_{\vec{X}(t)}(\vec{x})$ from $N = 2000$ points.



Horizontal profiles at $x_2 = 3$ through density estimates.

1.6

Given a collection $\vec{X}_1, \dots, \vec{X}_N$ of independently and identically distributed random variables drawn from a distribution with pdf $f(\vec{x})$, a kernel density estimator for $f(\vec{x})$ [41] is defined by

$$\hat{f}(\vec{x}) = \frac{1}{N} \sum_{k=1}^N g(\vec{x} - \vec{X}_k),$$

for some kernel function g , e.g., Gaussian. Basically each "point" \vec{X}_k is blurred out and these blobs are added up to make an image.

The maximum likelihood estimator for a pdf $f(\vec{x})$ is essentially a set of Dirac delta functions at each \vec{X}_k . This is generally not a good representation of $f(\vec{x})$.

Thus, even with a perfect imaging system, there is still a reconstruction problem!

Poisson Spatial Point Process

Assumption 2. The number of injected tracer atoms N has a Poisson distribution with some mean

$$\mu_N \triangleq E[N] = \sum_{n=0}^{\infty} nP[N=n].$$

Let $N(B)$ denote the number of tracer atoms that have spatial locations in any set $B \subset \mathbb{R}^3$ (VOI) at time t_0 after injection.

$N(\cdot)$ is called a *Poisson spatial point process*.

Fact. For any set B , $N(B)$ is Poisson distributed with mean:

$$E[N(B)] = E[N]P[\vec{X} \in B] = \mu_N \int_B f_{\vec{X}(t_0)}(\vec{x}) d\vec{x}.$$

Poisson N injected atoms + i.i.d. locations \Rightarrow Poisson point process

1.7

VOI = volume of interest

To be completely rigorous, “any set” really means “any Borel set” since we are assuming implicitly that $f_{\vec{X}(t)}(\vec{x})$ is a density with respect to Lebesgue measure.

Explicitly:


$$N(B) \triangleq \sum_{k=1}^N 1_{\{\vec{X}_k(t_0) \in B\}}.$$

One can also show the following. If B_1 and B_2 are two disjoint subsets of \mathbb{R}^3 , then $N(B_1)$ and $N(B_2)$ are independent Poisson random variables.

Thus, the spatial locations of tracer atoms are governed by Poisson statistics.

And we have not even mentioned radioactive decay yet!

Illustration of Point Process ($\mu_N = 200$)



1.8

Four realizations of tracer atom locations distributed according to the radiotracer distribution shown earlier. In this case $E[N] = 200$.

The number of points falling within the ROI in each realization is a random variable with a Poisson distribution.

Radionuclide Decay


Preceding quantities are all unobservable.
We “observe” a tracer atom only when it decays and emits photon(s).

The time that the k th tracer atom decays is a random variable T_k .

Assumption 3. The T_k 's are statistically *independent* random variables, and are independent of the (random) spatial location.

Assumption 4. Each T_k has an exponential distribution with mean $\mu_T = t_{1/2}/\ln 2$.

Cumulative distribution function: $P[T_k \leq t] = 1 - \exp(-t/\mu_T)$



We only observe *some* of the tracer atoms when they decay; most are never recorded.

The independence assumptions are reasonable physically except in cases of stimulated emissions, since the decay of a given nucleus is not “affected to any significant extent by events occurring outside the nucleus” [42, p. 22].

The exponential distribution is consistent with empirical observations [43]. Also, the exponential distribution is the unique distribution that is consistent with the assumption that “the probability of decay of an atom is independent of the age of that atom” [43, p. 470]. In statistical terms, this characteristic is called the *memoryless property*, and can be expressed mathematically as follows:

$$P[T_k \leq t | T_k \geq t_0] = P[T_k \leq t - t_0] \quad \text{for } t \geq t_0.$$

Decay half-life $t_{1/2}$ is the time for which $P[T_k \leq t_{1/2}] = 1/2$. Solving yields $t_{1/2} = \mu_T \ln 2$.

Statistics of an Ideal Decay Counter

Let $K(t, B)$ denote the number of tracer atoms that decay by time t , and that were located in the VOI $B \subset \mathbb{R}^3$ at the time of decay.

Fact. $K(t, B)$ is a *Poisson counting process* with mean

$$E[K(t, B)] = \int_0^t \int_B \lambda(\vec{x}, \tau) d\vec{x} d\tau,$$

where the (nonuniform) *emission rate density* is given by

$$\lambda(\vec{x}, t) \triangleq \mu_N \frac{e^{-t/\mu_T}}{\mu_T} \cdot f_{\vec{x}(t)}(\vec{x}).$$

Ingredients: “dose,” “decay,” “distribution”

Units: “counts” per unit time per unit volume, e.g., $\mu\text{Ci/cc}$.

“Photon emission is a Poisson process”

What about the actual measurement statistics?

Since we will analyze the statistical properties of “real” photon-counting detectors, it is natural to first examine the statistical properties of an “ideal” decay counter.

By definition $K(0, B) = 0$ and $K(t, B) = \sum_{k=1}^N \mathbf{1}_{\{T_k \leq t, \vec{x}_k(T_k) \in B\}}$.

The units of $\lambda(\vec{x}, t)$ are “counts” per unit time per unit volume. Note that $K(t, B)$ is an *inhomogenous Poisson process* since its rate varies with time (due to decay).

This counter is “ideal” because it omnisciently counts every decay, regardless of scatter or absorption, and regardless of the decay rate (no deadtime).

The emission rate density at a point \vec{x} is proportional to the (mean) number of administered photons (dose), corrected for decay, and to the (local) radiotracer density at \vec{x} .

Hot regions contain more tracer atoms, and hence produce more decays.

For t small, $\int_0^t \frac{e^{-s/\mu_T}}{\mu_T} ds = 1 - e^{-t/\mu_T} \approx t/\mu_T$, so the overall emission rate is proportional to μ_N/μ_T for small time intervals.

The above facts do not ensure that the *measurements* have Poisson distributions. That conclusion requires additional assumptions discussed below.

Idealized Detector Units

A nuclear imaging system consists of n_d conceptual *detector units*.

Assumption 5. Each decay of a tracer atom produces a recorded count in at most one detector unit.

Let $S_k \in \{0, 1, \dots, n_d\}$ denote the index of the incremented detector unit for decay of k th tracer atom. ($S_k = 0$ if decay is undetected.)

Assumption 6. The S_k 's satisfy the following conditional independence:

$$P(S_1, \dots, S_N | N, T_1, \dots, T_N, \vec{X}_1(\cdot), \dots, \vec{X}_N(\cdot)) = \prod_{k=1}^N P(S_k | \vec{X}_k(T_k)).$$

The recorded bin for the k th tracer atom's decay depends only on its position when it decays, and is independent of all other tracer atoms.

(No event pileup; no deadtime losses.)

1.11

Examples of a “detector unit” (perhaps better named “recorder unit”)

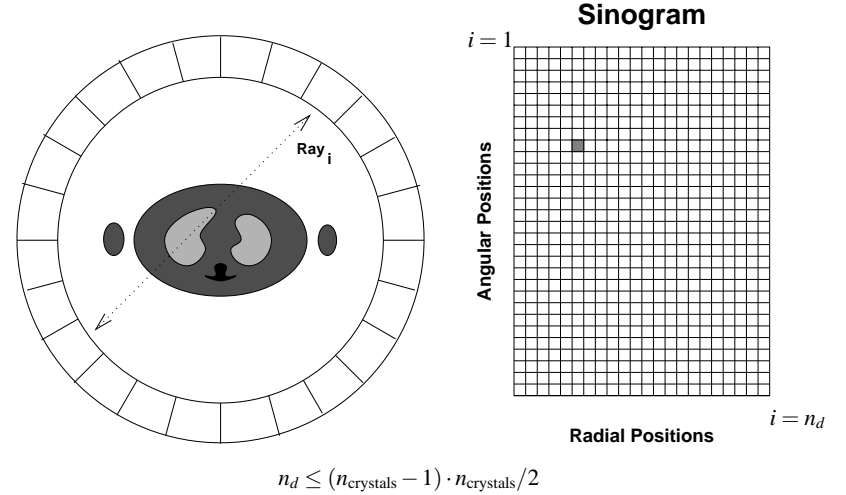
- pair of PET crystals (sinogram bin)
- Anger camera projection bin, at a particular projection angle (sinogram bin)
- Anger camera projection bin at some angle in some energy bin (for multiple energy window acquisitions)

In our terminology, detector units need not correspond to physical detectors. For example, in a 2D PET system consisting of a ring of n physical detectors, there could be as many as $n_d = n(n-1)$ detector units, each of which corresponds to a pair of physical detectors in electronic coincidence. In a SPECT system based on a single rotating gamma camera that collects $n_z \times n_x$ projection images at each of n_θ projection angles, there would be $n_d = n_z \cdot n_x \cdot n_\theta$ detector units. Systems that bin measurements into multiple energy windows would have even more detector units.

Assumption 5 applies to most nuclear imaging systems. If a system assigns fractions of an event to different detector units, or can assign events to more than one detector unit, then the measurements are probably non-Poisson and need more complicated analysis.

Assumption 6 is reasonable for moderate count rates. At high count rates, deadtime losses cause a decrease in the detection probabilities, *i.e.*, $P[S_k = 0]$ will depend on the T_j 's for $k \neq j$. Measurements affected by deadtime are not Poisson distributed, and need more complicated analysis [44–46].


PET Example



1.12

Here each “detector unit” is a pair of detectors in electronic coincidence.

SPECT Example



1.13

Here each “detector unit” is an (angle,bin) pair.

Or for multiple energy windows:

$$n_d = n_{\text{radial_bins}} \cdot n_{\text{angular_steps}} \cdot n_{\text{energy_windows}}$$

1.13

p1frame


Detector Unit Sensitivity Patterns

Spatial localization:

$s_i(\vec{x}) \triangleq$ probability that decay at \vec{x} is recorded by i th detector unit.

Idealized Example. Shift-invariant PSF: $s_i(\vec{x}) = h(\vec{k}_i \cdot \vec{x} - r_i)$

- r_i is the radial position of i th ray
- \vec{k}_i is the unit vector orthogonal to i th parallel ray
- $h(\cdot)$ is the shift-invariant radial PSF (e.g., Gaussian bell or rectangular function)




1.14

$h(r) = \delta(r)$ corresponds to the sensitivity pattern of the ideal line integral.

1.14

p1frame

Example: SPECT Detector-Unit Sensitivity Patterns



Two representative $s_i(\vec{x})$ functions for a collimated Anger camera.

1.15


Here each detector unit corresponds to a particular element of a SPECT sinogram.

In this model, the detector response is a Gaussian function whose FWHM increases monotonically as a function of the distance from the camera face to a point \vec{x} in object space.

Strictly speaking, with this Gaussian model, $s_i(\vec{x}) \neq 0$ everywhere. This would be impractical to compute and store. In practice we approximate $s_i(\vec{x})$ by a truncated version which is set to zero everywhere the Gaussian function is sufficiently small.

The s_i functions shown above are purely geometric (detector response) and do not include the effects of scatter or attenuation.

Example: PET Detector-Unit Sensitivity Patterns



1.16

These sensitivity patterns account for the parallax and crystal penetration effects in ring PET systems.

Detector Unit Sensitivity Patterns

$s_i(\vec{x})$ can include the effects of

- geometry / solid angle
- collimation
- scatter
- attenuation
- detector response / scan geometry
- duty cycle (dwell time at each angle)
- detector efficiency
- positron range, noncollinearity
- ...

System sensitivity pattern:


$$s(\vec{x}) \triangleq \sum_{i=1}^{n_d} s_i(\vec{x}) = 1 - s_0(\vec{x}) \leq 1$$

(probability that decay at location \vec{x} will be detected *at all* by system)

1.17

Subtle point: for a moving imaging system, like a rotating gamma camera, the detector-unit sensitivity patterns are time varying, *i.e.*, we should write $s_i(\vec{x}, t)$. If the radiotracer distribution is static, then even moving imaging systems fit within the framework described here. However, if the radiotracer distribution changes with time as the imaging system is moving, then more complicated image formation models and image reconstruction algorithms are needed to avoid artifacts, *e.g.*, [47–56].

(Overall) System Sensitivity Pattern: $s(\vec{x}) = \sum_{i=1}^{n_d} s_i(\vec{x})$




Example: collimated 180° SPECT system with uniform attenuation.

1.18

$s(\vec{x})$ = probability that a decay at location \vec{x} will be detected *at all* by system during the course of a scan.

Although we call this the “system” sensitivity pattern, it also depends on the object’s attenuation and scatter properties.

Detection Probabilities $s_i(\vec{x}_0)$ (vs det. unit index i)



1.19

This is a sinogram with 60 radial positions by 100 angles, so $n_d = 6000$ and i varies from 1 (upper left hand corner) to 6000 (lower right hand corner) in lexicographic ordering.

For a system with perfect detector response the right curve would be a thin sinusoid. The thickness of the above sinusoid comes from the finite detector response. The depth-dependence of this (SPECT) detector response is also evident.

Summary of Random Phenomena

- Number of tracer atoms injected N
- Spatial locations of tracer atoms $\{\vec{x}_k\}_{k=1}^N$
- Time of decay of tracer atoms $\{T_k\}_{k=1}^N$
- Detection of photon $[S_k \neq 0]$
- Recording detector unit $\{S_k\}_{i=1}^{n_d}$

1.20

We have made assumptions about the nature of the distributions of each of the above random variables.

Emission Scan

Record events in each detector unit for $t_1 \leq t \leq t_2$.

$Y_i \triangleq$ number of events recorded by i th detector unit during scan, for $i = 1, \dots, n_d$.

$$Y_i \triangleq \sum_{k=1}^N \mathbf{1}_{\{S_k=i, T_k \in [t_1, t_2]\}}.$$

The collection $\{Y_i : i = 1, \dots, n_d\}$ is our *sinogram*.

Note $0 \leq Y_i \leq N$.

Fact. Under Assumptions 1-6 above,

$$Y_i \sim \text{Poisson} \left\{ \int s_i(\vec{x}) \lambda(\vec{x}) d\vec{x} \right\} \quad (\text{cf "line integral"})$$

and Y_i 's are statistically independent random variables, where the *emission density* is given by

$$\lambda(\vec{x}) = \mu_N \int_{t_1}^{t_2} \frac{1}{\mu_T} e^{-t/\mu_T} f_{\vec{X}(t)}(\vec{x}) dt.$$

(Local number of decays per unit volume during scan.)

Ingredients:

- dose (injected)
- duration of scan
- decay of radionuclide
- distribution of radiotracer

1.21

The emission density $\lambda(\vec{x})$ is proportional to the decay-weighted time integral of the radiotracer distribution over the scan interval.

Poisson Statistical Model (Emission)

Actual measured counts = “foreground” counts + “background” counts.

Sources of background counts:

- cosmic radiation / room background
- random coincidences (PET)
- scatter not account for in $s_i(\vec{x})$
- “crosstalk” from transmission sources in simultaneous T/E scans
- anything else not accounted for by $\int s_i(\vec{x}) \lambda(\vec{x}) d\vec{x}$

Assumption 7.

The background counts also have independent Poisson distributions.

Statistical model (continuous to discrete)

$$Y_i \sim \text{Poisson} \left\{ \int s_i(\vec{x}) \lambda(\vec{x}) d\vec{x} + r_i \right\}, \quad i = 1, \dots, n_d$$

r_i : mean number of “background” counts recorded by i th detector unit.

1.22

The detector unit sensitivity pattern $s_i(\vec{x})$ in principle includes both direct (unscattered) and scattered photons, *i.e.*,

$$s_i(\vec{x}) = s_i^{\text{direct}}(\vec{x}) + s_i^{\text{scatter}}(\vec{x}).$$

For simplicity, we often only include the direct component $s_i^{\text{direct}}(\vec{x})$ in the integral model, *i.e.*, one assumes

$$Y_i \sim \text{Poisson} \left\{ \int s_i^{\text{direct}}(\vec{x}) \lambda(\vec{x}) d\vec{x} + r_i \right\},$$

in which case the remaining counts due to $\int s_i^{\text{scatter}}(\vec{x}) \lambda(\vec{x}) d\vec{x}$ should be included in the r_i 's (and determined separately by some method such as using multiple energy windows).

Emission Reconstruction Problem

Estimate the emission density $\lambda(\cdot)$ using (something like) this model:

$$Y_i \sim \text{Poisson} \left\{ \int s_i(\vec{x}) \lambda(\vec{x}) d\vec{x} + r_i \right\}, \quad i = 1, \dots, n_d.$$

Knowns:

- $\{Y_i = y_i\}_{i=1}^{n_d}$: observed counts from each detector unit
- $s_i(\vec{x})$ sensitivity patterns (determined by system models)
- r_i 's : background contributions (determined separately)

Unknown: $\lambda(\vec{x})$

1.23

How to determine the r_i 's is a broad topic that could occupy its own short course. It depends on what effects are relevant. Here are some standard methods.

- cosmic radiation / room background: determine from a "blank" scan
- transmission crosstalk: determine from same energy window in another head (for multiple-head SPECT systems with only one transmission head)
- random coincidences (PET): delayed-window method
- scatter: use calculated scatter contributions from attenuation map and initial estimate of $\lambda(\vec{x})$, or estimate from another energy window, or ...

And of course the s_i 's are not really known at the outset since they include the effects of patient attenuation which must be determined by some type of transmission scan. But that is a separate topic...

1.23

p1frame

List-mode acquisitions

Recall that conventional sinogram is temporally binned:

$$Y_i \triangleq \sum_{k=1}^N 1_{\{S_k=i, T_k \in [t_1, t_2]\}}.$$

This binning discards temporal information.

List-mode measurements: record all (detector,time) pairs in a list, *i.e.*,

$$\{(S_k, T_k) : k = 1, \dots, N\}.$$

List-mode dynamic reconstruction problem:

Estimate $\lambda(\vec{x}, t)$ given $\{(S_k, T_k)\}$.

1.24

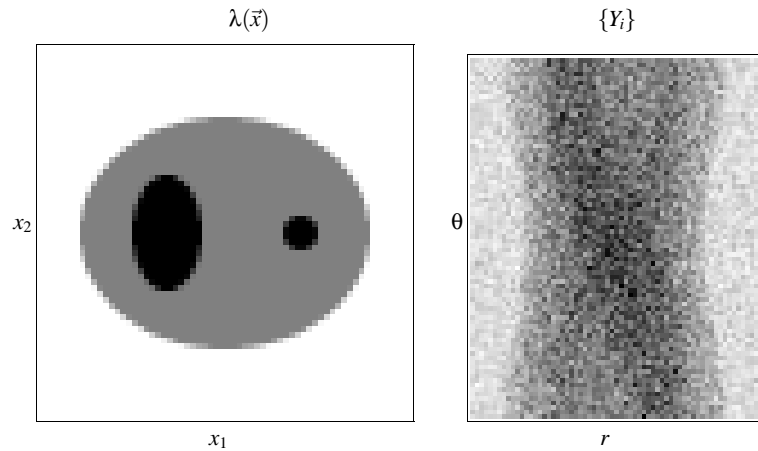
Recent treatment of dynamic PET reconstruction from list-mode measurements: [57].

To be precise, the list is actually $\{(S_k, T_k) : k = 1, \dots, N, S_k \neq 0\}$ since we only can record the times of recorded events.

1.24


p1frame

Emission Reconstruction Problem - Illustration



1.25

Example: MRI “Sensitivity Pattern”



Each “k-space sample” corresponds to a sinusoidal pattern weighted by:

- RF receive coil sensitivity pattern
- phase effects of field inhomogeneity
- spin relaxation effects.

$$y_i = \int f(\vec{x}) c_{\text{RF}}(\vec{x}) \exp(-i\omega(\vec{x})t_i) \exp(-t_i/T_2(\vec{x})) \exp(-i2\pi\vec{k}(t_i) \cdot \vec{x}) d\vec{x} + \epsilon_i$$

1.26

Of course the noise in MR is not Poisson, but this seemed like the best place for this slide...

Part 2: Five Categories of Choices

- Object parameterization: function $f(\vec{r})$ vs finite coefficient vector \mathbf{x}
- System physical model: $\{s_i(\vec{r})\}$
- Measurement statistical model $y_i \sim [?]$
- Cost function: data-mismatch and regularization
- Algorithm / initialization

No perfect choices - one can critique all approaches!

2.1

Often these choices are made implicitly rather than explicitly. Leaving the choices implicit fortifies the common belief among non-experts that there are basically two kinds of reconstruction algorithms, FBP and “iterative.”

In fact, the choices one makes in the above five categories can greatly affect the results.

In my opinion, every paper describing iterative image reconstruction methods (or results thereof) should make as explicit as possible what choices were made in each of the above categories.

Choice 1. Object Parameterization

Finite measurements: $\{y_i\}_{i=1}^{n_d}$. Continuous object: $f(\vec{r})$. Hopeless?

All models are wrong but some models are useful.

Linear *series expansion* approach. Replace $f(\vec{r})$ by $\mathbf{x} = (x_1, \dots, x_{n_p})$ where

$$f(\vec{r}) \approx \tilde{f}(\vec{r}) = \sum_{j=1}^{n_p} x_j b_j(\vec{r}) \leftarrow \text{“basis functions”}$$

Forward projection:

$$\begin{aligned} \int s_i(\vec{r}) f(\vec{r}) d\vec{r} &= \int s_i(\vec{r}) \left[\sum_{j=1}^{n_p} x_j b_j(\vec{r}) \right] d\vec{r} = \sum_{j=1}^{n_p} \left[\int s_i(\vec{r}) b_j(\vec{r}) d\vec{r} \right] x_j \\ &= \sum_{j=1}^{n_p} a_{ij} x_j = [\mathbf{A}\mathbf{x}]_i, \text{ where } a_{ij} \triangleq \int s_i(\vec{r}) b_j(\vec{r}) d\vec{r} \end{aligned}$$

- Projection integrals become finite summations.
- a_{ij} is contribution of j th basis function (*e.g.*, voxel) to i th detector unit.
- The units of a_{ij} and x_j depend on the user-selected units of $b_j(\vec{r})$.
- The $n_d \times n_p$ matrix $\mathbf{A} = \{a_{ij}\}$ is called the *system matrix*.

2.2

In principle it is not entirely hopeless to reconstruct a continuous $f(\vec{r})$ from a finite set of measurements. This is done routinely in the field of nonparametric regression [58] (the generalization of linear regression that allows for fitting smooth functions rather than just lines). But it is complicated in tomography...

Van De Walle, Barrett, *et al.* [59] have proposed pseudoinverse calculation method for MRI reconstruction from a continuous-object / discrete-data formulation, based on the general principles of Bertero *et al.* [60]. If the pseudo-inverse could truly be computed once-and-for-all then such an approach could be practically appealing. However, in practice there are object-dependent effects, such as nonuniform attenuation in SPECT and magnetic field inhomogeneity in MRI, and these preclude precomputation of the required SVDs. So pseudo-inverse approaches are impractical for typical realistic physical models.

(Linear) Basis Function Choices

- Fourier series (complex / not sparse)
- Circular harmonics (complex / not sparse)
- Wavelets (negative values / not sparse)
- Kaiser-Bessel window functions (blobs)
- Overlapping circles (disks) or spheres (balls)
- Polar grids, logarithmic polar grids
- “Natural pixels” $\{s_i(\vec{r})\}$
- B-splines (pyramids)
- Rectangular pixels / voxels (rect functions)
- Point masses / bed-of-nails / lattice of points / “comb” function
- Organ-based voxels (e.g., from CT), ...

Considerations

- Represent $f(\vec{r})$ “well” with moderate n_p
- Orthogonality? (not essential)
- “Easy” to compute a_{ij} ’s and/or $A\mathbf{x}$
- Rotational symmetry
- If stored, the system matrix A should be sparse (mostly zeros).
- Easy to represent nonnegative functions e.g., if $x_j \geq 0$, then $f(\vec{r}) \geq 0$.
A sufficient condition is $b_j(\vec{r}) \geq 0$.

2.3

“Well” \equiv approximation error less than estimation error

Many published “projector / backprojector pairs” are not based explicitly on any particular choice of basis.

Many bases have the desirable approximation property that one can form arbitrarily accurate approximations to $f(\vec{r})$ by taking n_p sufficiently large. (This is related to *completeness*.) Exceptions include “natural pixels” (a finite set) and the point-lattice “basis” (usually).

Some pixel-driven backprojectors could be interpreted implicitly as point-mass object models. This model works fine for FBP, but causes artifacts for iterative methods.

Mazur *et al.* [61] approximate the shadow of each pixel by a rect function, instead of by a trapezoid. “As the shapes of pixels are artifacts of our digitisation of continuous real-world images, consideration of alternative orientation or shapes for them seems reasonable.” However, they observe slightly worse results that worsen with iteration!

Classic series-expansion reference [62]

Organ-based voxel references include [51, 63–67]

Nonlinear Object Parameterizations

Estimation of intensity *and* shape (e.g., location, radius, etc.)

Surface-based (homogeneous) models

- Circles / spheres
- Ellipses / ellipsoids
- Superquadrics
- Polygons
- Bi-quadratic triangular Bezier patches, ...

Other models

- Generalized series $f(\vec{r}) = \sum_j x_j b_j(\vec{r}, \theta)$
- Deformable templates $f(\vec{r}) = b(T_\theta(\vec{r}))$
- ...

Considerations

- Can be considerably more parsimonious
- If correct, yield greatly reduced estimation error
- Particularly compelling in limited-data problems
- Often oversimplified (all models are wrong but...)
- Nonlinear dependence on location induces non-convex *cost functions*, complicating optimization

2.4

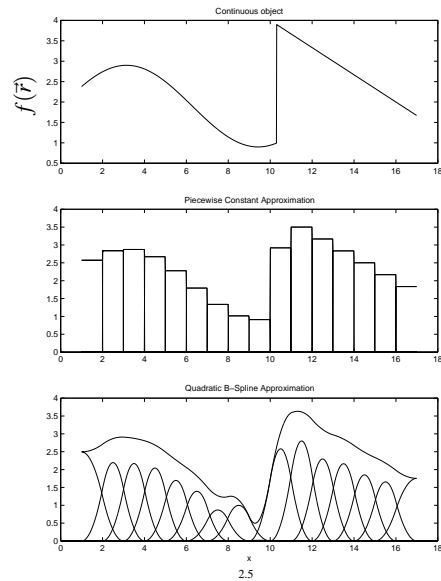
Disks [68, 69]

Polygons [70]

Generalized series [71]

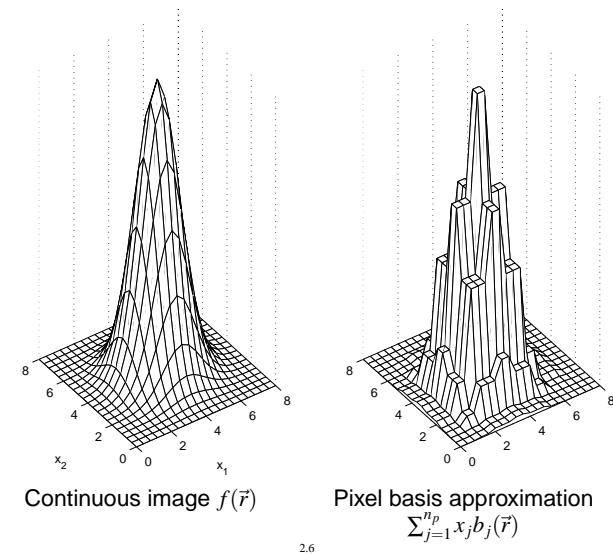
Bi-quadratic triangular Bezier patches [72]

Example Basis Functions - 1D



Neither the pixels nor the blobs are ideal in the above example, though both could reduce the average approximation error as low as needed by taking n_p sufficiently large.

Pixel Basis Functions - 2D



My tentative recommendation: use pixel / voxel basis.

- Simple
- Perfectly matched to digital displays
- Maximally sparse system matrix

Or use blobs (rotationally symmetric Kaiser-Bessel windows)

- Easy to compute projections "on the fly" due to rotational symmetry.
- Differentiable, nonnegative.
- Parsimony advantage using body-centered cubic packing

Discrete-Discrete Emission Reconstruction Problem

Having chosen a basis and *linearly* parameterized the emission density...

Estimate the emission density coefficient vector $\mathbf{x} = (x_1, \dots, x_{n_p})$ (aka “image”) using (something like) this statistical model:

$$y_i \sim \text{Poisson} \left\{ \sum_{j=1}^{n_p} a_{ij} x_j + r_i \right\}, \quad i = 1, \dots, n_d.$$

- $\{y_i\}_{i=1}^{n_d}$: observed counts from each detector unit
- $\mathbf{A} = \{a_{ij}\}$: system matrix (determined by system models)
- r_i 's : background contributions (determined separately)

Many image reconstruction problems are “find \mathbf{x} given \mathbf{y} ” where

$$y_i = g_i([\mathbf{A}\mathbf{x}]_i) + \varepsilon_i, \quad i = 1, \dots, n_d.$$

2.7

Called the “discrete-discrete” estimation problem since both the measurement vector and the image vector are “discretized” (finite dimensional).

In contrast, FBP is derived from the “continuous-continuous” Radon transform model.

Choice 2. System Model

System matrix elements: $a_{ij} = \int s_i(\vec{r}) b_j(\vec{r}) d\vec{r}$

- scan geometry
- collimator/detector response
- attenuation
- scatter (object, collimator, scintillator)
- duty cycle (dwell time at each angle)
- detector efficiency / dead-time losses
- positron range, noncollinearity, crystal penetration, ...
- ...

Considerations

- Improving system model can improve
 - Quantitative accuracy
 - Spatial resolution
 - Contrast, SNR, detectability
- Computation time (and storage vs compute-on-fly)
- Model uncertainties
(e.g., calculated scatter probabilities based on noisy attenuation map)
- Artifacts due to over-simplifications

2.8

For the pixel basis, a_{ij} is the probability that a decay in the j th pixel is recorded by the i th detector unit.

Attenuation enters a_{ij} differently in PET and SPECT.

Measured System Model?


Determine a_{ij} 's by scanning a voxel-sized cube source over the imaging volume and recording counts in all detector units (separately for each voxel).

- Avoids mathematical model approximations
- Scatter / attenuation added later (object dependent), approximately
- Small probabilities \Rightarrow long scan times
- Storage
- Repeat for every voxel size of interest
- Repeat if detectors change

2.9

Certainly worth doing occasionally (at least for some voxels) to check system models, particularly for complicated geometries.

“Line Length” System Model



2.10

Mathematically, the corresponding detector unit sensitivity pattern is

$$s_i(\vec{r}) = \delta(\vec{k}_i \cdot \vec{r} - \tau_i),$$

where δ denotes the Dirac impulse function.

This model is usually applied with the pixel basis, but can be applied to any basis.

Does not exactly preserve counts, *i.e.*, in general

$$\int f(\vec{r}) \, d\vec{r} \neq \sum_{i=1}^{n_d} \sum_{j=1}^{n_n} a_{ij} x_j$$


Leads to artifacts.

Units are wrong too. (Reconstructed x will have units inverse length.)

Perhaps reasonable for X-ray CT, but unnatural for emission tomography. (Line segment length is a probability?)

In short: I recommend using almost anything else!

“Strip Area” System Model



2.11

Accounts for finite detector width.

Mathematically, the corresponding detector unit sensitivity pattern is

$$s_i(\vec{r}) = \text{rect}\left(\frac{\vec{k}_i \cdot \vec{r} - \tau_i}{w}\right),$$

where w is the detector width.


Can exactly preserve counts, since all areas are preserved, provided that the width w is an integer multiple of the center-to-center ray spacing.

Most easily applied to the pixel basis, but in principle applies to any choice.


A little more work to compute than line-lengths, but worth the extra effort (particularly when pre-computed).

(Implicit) System Sensitivity Patterns

$$\sum_{i=1}^{n_d} a_{ij} \approx s(\vec{r}_j) = \sum_{i=1}^{n_d} s_i(\vec{r}_j)$$



Line Length



Strip Area

2.12

Backprojection of a uniform sinogram.

Explicitly:


$$\sum_{i=1}^{n_d} a_{ij} = \sum_{i=1}^{n_d} \int s_i(\vec{r}) b_j(\vec{r}) d\vec{r} = \int \left[\sum_{i=1}^{n_d} s_i(\vec{r}) \right] b_j(\vec{r}) d\vec{r} = \int s(\vec{r}) b_j(\vec{r}) d\vec{r} \approx s(\vec{r}_j)$$

where \vec{r}_j is center of j th basis function.

Shows probability for each pixel that an emission from that pixel will be detected somewhere.

These nonuniformities propagate into the reconstructed images, except when sinograms are simulated from the same model of course.

Point-Lattice Projector/Backprojector



2.13

This model is used implicitly in many pixel-driven backprojection subroutines.

Mathematically, the corresponding detector unit sensitivity pattern is


$$s_i(\vec{r}) = \text{tri} \left(\frac{\vec{k}_i \cdot \vec{r} - \tau_i}{\Delta_r} \right),$$

where Δ_r is the ray spacing.

This is a reasonable enough detector response model (more realistic than an impulse at least), but the problem arises when combining it with the bed-of-nails "basis" functions $b_j(\vec{r}) = \delta(\vec{r} - \vec{r}_j)$.

Point-Lattice Artifacts

Projections (sinograms) of uniform disk object:



2.14

A little radial smoothing can reduce these artifacts, but that just hides the underlying deficiency of the point-lattice model.

The disadvantage of the strip-area model is that it is more work to compute the areas of intersection than it is to compute line lengths or linear interpolations for point-lattice model. There may be less "symmetries" to exploit in the strip-area model too.

For 2D we precompute and store the a_{ij} 's, so the "extra work" is irrelevant. This becomes more cumbersome for 3D, but it has been done.

Forward- / Back-projector “Pairs”

Forward projection (image domain to projection domain):

$$\bar{y}_i = \int s_i(\vec{r}) f(\vec{r}) d\vec{r} = \sum_{j=1}^{n_p} a_{ij} x_j = [\mathbf{A}\mathbf{x}]_i, \text{ or } \bar{\mathbf{y}} = \mathbf{A}\mathbf{x}$$

Backprojection (projection domain to image domain):

$$\mathbf{A}'\mathbf{y} = \left\{ \sum_{i=1}^{n_d} a_{ij} y_i \right\}_{j=1}^{n_p}$$

Often $\mathbf{A}'\mathbf{y}$ is implemented as $\mathbf{B}\mathbf{y}$ for some “backprojector” $\mathbf{B} \neq \mathbf{A}'$

Least-squares solutions (for example):


$$\hat{\mathbf{x}} = [\mathbf{A}'\mathbf{A}]^{-1} \mathbf{A}'\mathbf{y} \neq [\mathbf{B}\mathbf{A}]^{-1} \mathbf{B}\mathbf{y}$$

2.15

Algorithms are generally derived using a single \mathbf{A} matrix, and usually the quantity $\mathbf{A}'\mathbf{y}$ appears somewhere in the derivation.

If the product $\mathbf{A}'\mathbf{y}$ is implemented by some $\mathbf{B}\mathbf{y}$ for $\mathbf{B} \neq \mathbf{A}'$, then all convergence properties, statistical properties, etc. of the theoretical algorithm may be lost by the implemented algorithm.

Mismatched Backprojector $\mathbf{B} \neq \mathbf{A}'$



2.16

Noiseless 3D PET data, images are $n_x \times n_y \times n_z = 64 \times 64 \times 4$, with $n_u \times n_v \times n_\theta \times n_\phi = 62 \times 10 \times 60 \times 3$ projections. 15 iterations of PWLS-CG, initialized with the true image. True object values range from 0 to 2. Display windowed to [0.7, 1.3] to highlight artifacts.

In this case mismatch arises from a ray-driven forward projector but a pixel-driven back projector.

Another case where mismatch can arise is in “rotate and sum” projection / backprojection methods, if implemented carelessly.


The problem with mismatched backprojectors arises in iterative reconstruction because multiple iterations are generally needed, so discrepancies between \mathbf{B} and \mathbf{A}' can accumulate.

Such discrepancies may matter more for regularized methods where convergence is desired, then for unregularized methods where one stops well before convergence, but this is merely speculation.

The deliberate use of mismatched projectors/backprojectors has been called the “dual matrix” approach.

todo: add unregularized comparison with mismatched

Horizontal Profiles



2.17

This was from noiseless simulated data!

System Model Tricks

- Factorize (*e.g.*, PET Gaussian detector response)

$$\mathbf{A} \approx \mathbf{S}\mathbf{G}$$

(geometric projection followed by Gaussian smoothing)

- Symmetry
- Rotate and Sum
- Gaussian diffusion for SPECT Gaussian detector response
- Correlated Monte Carlo (Beekman *et al.*)

In all cases, consistency of backprojector with \mathbf{A}' requires care.

2.18


Separability [73]
 Symmetry [74–77]
 Rotators [78]
 Gaussian diffusion [79, 80]
 Sparse storage [81]
 Forward projector tricks [82–95]
 Correlated Monte Carlo [96]

My recommendation: in 2D PET and SPECT, precompute and store as accurate of a system matrix as you can, strip-area model at least.

This topic could be a talk in itself...

Other issues: include scatter in \mathbf{A} or not?

SPECT System Model



Complications: nonuniform attenuation, depth-dependent PSF, Compton scatter

2.19

Choice 3. Statistical Models

After modeling the system physics, we have a deterministic “model:”

$$y_i \approx g_i([Ax]_i)$$

for some functions g_i , *e.g.*, $g_i(l) = l + r_i$ for emission tomography.

Statistical modeling is concerned with the “ \approx ” aspect.

Considerations

- More accurate models:
 - can lead to lower variance images,
 - may incur additional computation,
 - may involve additional algorithm complexity
(*e.g.*, proper transmission Poisson model has nonconcave log-likelihood)
- Statistical model errors (*e.g.*, deadtime)
- Incorrect models (*e.g.*, log-processed transmission data)

2.20

“Complexity” can just mean “inconvenience.” It would certainly be more convenient to precorrect the sinogram data for effects such as randoms, attenuation, scatter, detector efficiency, etc., since that would save having to store those factors for repeated use during the iterations. But such pre-corrections destroy the Poisson statistics and lead to suboptimal performance (higher variance).

More accurate statistical models may also yield lower bias, but bias is often dominated by approximations in the system model (neglected scatter, etc.) and by resolution effects induced by regularization.

Statistical Model Choices for Emission Tomography

- “None.” Assume $\mathbf{y} - \mathbf{r} = \mathbf{Ax}$. “Solve algebraically” to find \mathbf{x} .
- White Gaussian noise. Ordinary least squares: minimize $\|\mathbf{y} - \mathbf{Ax}\|^2$
- Non-white Gaussian noise. Weighted least squares: minimize

$$\|\mathbf{y} - \mathbf{Ax}\|_W^2 = \sum_{i=1}^{n_d} w_i (y_i - [\mathbf{Ax}]_i)^2, \text{ where } [\mathbf{Ax}]_i \triangleq \sum_{j=1}^{n_p} a_{ij}x_j$$

(e.g., for Fourier rebinned (FORE) PET data)

- Ordinary Poisson model (ignoring or precorrecting for background)

$$y_i \sim \text{Poisson}\{[\mathbf{Ax}]_i\}$$

- Poisson model

$$y_i \sim \text{Poisson}\{[\mathbf{Ax}]_i + r_i\}$$

- Shifted Poisson model (for randoms precorrected PET)

$$y_i = y_i^{\text{prompt}} - y_i^{\text{delay}} \sim \text{Poisson}\{[\mathbf{Ax}]_i + 2r_i\} - 2r_i$$

2.21

These are all for the emission case.

GE uses WLS for FORE data [97].

The shifted-Poisson model for randoms-precorrected PET is described in [98–101].

Snyder *et al.* used similar models for CCD imaging [102, 103].

Missing from the above list: deadtime model [44].

My recommendations.

- If the data is uncorrected, then use Poisson model above.
- If the data was corrected for random coincidences, use shifted Poisson model.
- If the data has been corrected for other stuff, consider using WLS, e.g. [104, 105].
- Try not to correct the data so that the first choice can be used!

Classic reason for WLS over Poisson was compute time. This has been obviated by recent algorithm advances. Now the choice should be made statistically.

Preprocessing: randoms subtraction, Fourier or multislice rebinning (3d to 2d), attenuation, scatter, det. eff., etc.

Shifted Poisson model for PET

Precorrected random coincidences: $y_i = y_i^{\text{prompt}} - y_i^{\text{delay}}$

$$y_i^{\text{prompt}} \sim \text{Poisson}\{[\mathbf{Ax}]_i + r_i\}$$

$$y_i^{\text{delay}} \sim \text{Poisson}\{r_i\}$$

$$E[y_i] = [\mathbf{Ax}]_i$$

$$\text{Var}\{y_i\} = [\mathbf{Ax}]_i + 2r_i \quad \text{Mean} \neq \text{Variance} \Rightarrow \text{not Poisson!}$$

Statistical model choices

- Ordinary Poisson model: ignore randoms

$$[y_i]_+ \sim \text{Poisson}\{[\mathbf{Ax}]_i\}$$

Causes bias due to truncated negatives

- Data-weighted least-squares (Gaussian model):

$$y_i \sim N([\mathbf{Ax}]_i, \hat{\sigma}_i^2), \quad \hat{\sigma}_i^2 = \max(y_i + 2\hat{r}_i, \sigma_{\min}^2)$$

Causes bias due to data-weighting

- Shifted Poisson model (matches 2 moments):

$$[y_i + 2\hat{r}_i]_+ \sim \text{Poisson}\{[\mathbf{Ax}]_i + 2\hat{r}_i\}$$

Insensitive to inaccuracies in \hat{r}_i .

2.22

Ideally, \hat{r}_i would be calculated based on block singles rates and detector efficiency information.

Practically, using $\hat{r}_i \approx 2$ or so works fine for usual whole-body FDG PET scans [98–101].

If \hat{r}_i is too small, then there will be many remaining negatives that will be truncated by the $[]_+$ operation, leading to biased images.

If \hat{r}_i is too large, the statistical properties approach those of unweighted least squares, yielding unnecessarily noisy images and slower convergence of EM-type algorithms.

Shifted-Poisson Model for X-ray CT

Model with both photon variability and readout noise:

$$y_i \sim \text{Poisson}\{\bar{y}_i(\boldsymbol{\mu})\} + N(0, \sigma^2)$$

Shifted Poisson approximation

$$y_i + \sigma^2 \sim \text{Poisson}\{\bar{y}_i(\boldsymbol{\mu}) + \sigma^2\}$$

or just use WLS...

Complications:

- Intractability of likelihood for Poisson+Gaussian
- Compound Poisson distribution due to photon-energy-dependent detector signal.

X-ray statistical models is a current research area in several groups!

2.23

For Poisson+Gaussian, see [102, 103].

For compound Poisson distribution, see [106, 107].

Choice 4. Cost Functions

Components:

- *Data-mismatch* term
- *Regularization* term (and regularization parameter β)
- Constraints (*e.g.*, nonnegativity)

$$\Psi(\mathbf{x}) = \text{DataMismatch}(\mathbf{y}, \mathbf{Ax}) + \beta \cdot \text{Roughness}(\mathbf{x})$$

$$\hat{\mathbf{x}} \triangleq \arg \min_{\mathbf{x} \geq 0} \Psi(\mathbf{x})$$

Actually *several* sub-choices to make for Choice 4 ...

Distinguishes “statistical methods” from “algebraic methods” for “ $\mathbf{y} = \mathbf{Ax}$.”

2.24

β sometimes called *hyperparameter*

Why Cost Functions?

(vs “procedure” e.g., adaptive neural net with wavelet denoising)

Theoretical reasons

ML is based on minimizing a *cost function*: the negative log-likelihood

- ML is asymptotically consistent
- ML is asymptotically unbiased
- ML is asymptotically efficient (under true statistical model...)
- Estimation: Penalized-likelihood achieves uniform CR bound asymptotically
- Detection: Qi and Huesman showed analytically that MAP reconstruction outperforms FBP for SKE/BKE lesion detection (T-MI, Aug. 2001)

Practical reasons

- Stability of estimates (if Ψ and algorithm chosen properly)
- Predictability of properties (despite nonlinearities)
- Empirical evidence (?)

2.25

Stability means that running “too many iterations” will not compromise image quality.

Asymptotically efficient means that the variance of ML estimator approaches that given by the Cramer-Rao lower bound, which is a bound on the variance of unbiased estimators.

But nuclear imaging is not asymptotic (too few counts), and system models are always approximate, and we regularize which introduces bias anyway.

Uniform CR bound generalizes CR bound to biased case [108]

Bottom line: haven't found anything better, seen plenty that are worse (LS vs ML in low count)

OSEM vs MAP [73, 109]

Qi and Huesman [31]

2.25

p2stat

Bayesian Framework

Given a prior distribution $p(x)$ for image vectors x , by Bayes' rule:

$$\text{posterior: } p(x|y) = p(y|x)p(x)/p(y)$$

so

$$\log p(x|y) = \log p(y|x) + \log p(x) - \log p(y)$$

- $-\log p(y|x)$ corresponds to data mismatch term
- $-\log p(x)$ corresponds to regularizing penalty function

Maximum a posteriori (MAP) estimator:

$$\hat{x} = \arg \max_x \log p(x|y)$$

- Has certain optimality properties (provided $p(y|x)$ and $p(x)$ are correct).
- Same form as Ψ

2.26

I avoid the Bayesian terminology because

- Images drawn from the “prior” distributions almost never look like real objects
 - The risk function associated with MAP estimation seems less natural to me than a quadratic risk function. The quadratic choice corresponds to conditional mean estimation $\hat{x} = E[x|y]$ which is very rarely used by those who describe Bayesian methods for image formation.
 - I often use penalty functions $R(x)$ that depend on the data y , which can hardly be called “priors.”
-

2.26

© J. Fessler, November 6, 2002

p2stat

Choice 4.1: Data-Mismatch Term

Options (for emission tomography):

- Negative log-likelihood of statistical model. Poisson *emission* case:

$$-L(\mathbf{x}; \mathbf{y}) = -\log p(\mathbf{y}|\mathbf{x}) = \sum_{i=1}^{n_d} ([\mathbf{Ax}]_i + r_i) - y_i \log([\mathbf{Ax}]_i + r_i) + \log y_i!$$

- Ordinary (unweighted) least squares: $\sum_{i=1}^{n_d} \frac{1}{2} (y_i - \hat{r}_i - [\mathbf{Ax}]_i)^2$
- Data-weighted least squares: $\sum_{i=1}^{n_d} \frac{1}{2} (y_i - \hat{r}_i - [\mathbf{Ax}]_i)^2 / \hat{\sigma}_i^2$, $\hat{\sigma}_i^2 = \max(y_i + \hat{r}_i, \sigma_{\min}^2)$, (causes bias due to data-weighting).
- Reweighted least-squares: $\hat{\sigma}_i^2 = [\mathbf{Ax}]_i + \hat{r}_i$
- Model-weighted least-squares (nonquadratic, but convex!)

$$\sum_{i=1}^{n_d} \frac{1}{2} (y_i - \hat{r}_i - [\mathbf{Ax}]_i)^2 / ([\mathbf{Ax}]_i + \hat{r}_i)$$

- Nonquadratic cost-functions that are robust to outliers
- ...

Considerations

- Faithfulness to statistical model vs computation
- Ease of optimization (convex?, quadratic?)
- Effect of statistical modeling errors

2.27

Poisson probability mass function (PMF):

$$p(\mathbf{y}|\mathbf{x}) = \prod_{i=1}^{n_d} e^{-\hat{y}_i} \hat{y}_i^{y_i} / y_i! \quad \text{where } \hat{\mathbf{y}} \triangleq \mathbf{Ax} + \mathbf{r}$$

Reweighted least-squares [110]

Model-weighted least-squares [111, 112]

$$f(l) = \frac{1}{2} (y - r - l)^2 / (l + r) \quad \check{f}(l) = y^2 / (l + r)^3 > 0$$

Robust norms [113, 114]

Generally the data-mismatch term and the statistical model go hand-in-hand.

Choice 4.2: Regularization

Forcing too much “data fit” gives noisy images

Ill-conditioned problems: small data noise causes large image noise

Solutions:

- Noise-reduction methods
- True regularization methods

Noise-reduction methods

- Modify the *data*
 - Prefilter or “denoise” the sinogram measurements
 - Extrapolate missing (*e.g.*, truncated) data
- Modify an *algorithm* derived for an ill-conditioned problem
 - Stop algorithm before convergence
 - Run to convergence, post-filter
 - Toss in a filtering step every iteration or couple iterations
 - Modify update to “dampen” high-spatial frequencies [115]

2.28

FBP with an apodized ramp filter belongs in the “modify the algorithm” category. The FBP method is derived based on a highly idealized system model. The solution so derived includes a ramp filter, which causes noise amplification if used unmodified. Throwing in apodization of the ramp filter attempts to “fix” this problem with the FBP “algorithm.”

The fault is not with the *algorithm* but with the problem definition and *cost function*. Thus the fix should be to the latter, not to the algorithm.

The estimate-maximize smooth (EMS) method [116] uses filtering every iteration.

The continuous image $f(\vec{r})$ - discrete data problem is *ill-posed*.

If the discrete-discrete problem has a full column rank system matrix \mathbf{A} , then that problem is well-posed, but still probably ill-conditioned.

Noise-Reduction vs True Regularization

Advantages of noise-reduction methods

- Simplicity (?)
- Familiarity
- Appear less subjective than using penalty functions or priors
- Only fiddle factors are # of iterations, amount of smoothing
- Resolution/noise tradeoff usually varies with iteration (stop when image looks good - in principle)
- Changing post-smoothing does not require re-iterating

Advantages of true regularization methods

- Stability
- Faster convergence
- Predictability
- Resolution can be made object independent
- Controlled resolution (*e.g.*, spatially uniform, edge preserving)
- Start with decent image (*e.g.*, FBP) \Rightarrow reach solution faster.

2.29

Running many iterations followed by post-filtering seems preferable to aborting early by stopping rules [117, 118].

Lalush *et al.* reported small differences between post-filtering and MAP reconstructions with an entropy prior [119].

Slijpen and Beekman conclude that post-filtering slightly more accurate than “oracle” filtering between iterations for SPECT reconstruction [120].

True Regularization Methods

Redefine the *problem* to eliminate ill-conditioning, rather than patching the data or algorithm!

- Use bigger pixels (fewer basis functions)
 - Visually unappealing
 - Can only preserve edges coincident with pixel edges
 - Results become even less invariant to translations
- Method of sieves (constrain image roughness)
 - Condition number for “pre-emission space” can be even worse
 - Lots of iterations
 - Commutability condition rarely holds exactly in practice
 - Degenerates to post-filtering in some cases
- Change *cost function* by adding a roughness penalty / prior
 - Disadvantage: apparently subjective choice of penalty
 - Apparent difficulty in choosing penalty parameters (cf apodizing filter / cutoff frequency in FBP)

2.30

Big pixels [121]

Sieves [122, 123]

Lots of iterations for convergence [117, 124]

Penalty Function Considerations

- Computation
- Algorithm complexity
- Uniqueness of minimizer of $\Psi(x)$
- Resolution properties (edge preserving?)
- # of adjustable parameters
- Predictability of properties (resolution and noise)

Choices

- separable vs nonseparable
- quadratic vs nonquadratic
- convex vs nonconvex

2.31

There is a huge literature on different regularization methods. Of the many proposed methods, and many anecdotal results illustrating properties of such methods, only the “lowly” quadratic regularization method has been shown *analytically* to yield detection results that are superior to FBP [31].

Penalty Functions: Separable vs Nonseparable

Separable

- Identity norm: $R(x) = \frac{1}{2}x'Ix = \sum_{j=1}^{n_p} x_j^2/2$
penalizes large values of x , but causes “squashing bias”
- Entropy: $R(x) = \sum_{j=1}^{n_p} x_j \log x_j$
- Gaussian prior with mean μ_j , variance σ_j^2 : $R(x) = \sum_{j=1}^{n_p} \frac{(x_j - \mu_j)^2}{2\sigma_j^2}$
- Gamma prior $R(x) = \sum_{j=1}^{n_p} p(x_j, \mu_j, \sigma_j)$ where $p(x, \mu, \sigma)$ is Gamma pdf

The first two basically keep pixel values from “blowing up.”

The last two encourage pixels values to be close to prior means μ_j .

$$\text{General separable form: } R(x) = \sum_{j=1}^{n_p} f_j(x_j)$$

Simple, but these do not explicitly enforce smoothness.

2.32

The identity norm penalty is a form of Tikhinov-Miller regularization [125].

The Gaussian and Gamma bias the results towards the prior image. This can be good or bad depending on whether the prior image is correct or not! If the prior image comes from a normal database, but the patient is abnormal, such biases would be undesirable.

For arguments favoring maximum entropy, see [126]. For critiques of maximum entropy regularization, see [127–129].

Penalty Functions: Separable vs Nonseparable

Nonseparable (partially couple pixel values) to penalize *roughness*

x_1	x_2	x_3
x_4	x_5	

Example

$$R(\mathbf{x}) = (x_2 - x_1)^2 + (x_3 - x_2)^2 + (x_5 - x_4)^2 + (x_4 - x_1)^2 + (x_5 - x_2)^2$$

2	2	2
2	1	

$$R(\mathbf{x}) = 1$$

3	3	1
2	2	

$$R(\mathbf{x}) = 6$$

1	3	1
2	2	

$$R(\mathbf{x}) = 10$$

Rougher images \Rightarrow greater $R(\mathbf{x})$

2.33

If diagonal neighbors were included there would be 3 more terms in this example.

Roughness Penalty Functions

First-order neighborhood and pairwise pixel differences:

$$R(\mathbf{x}) = \sum_{j=1}^{n_p} \frac{1}{2} \sum_{k \in N_j} \psi(x_j - x_k)$$

$N_j \triangleq$ *neighborhood* of j th pixel (e.g., left, right, up, down)
 ψ called the *potential function*

Finite-difference approximation to continuous roughness measure:

$$R(f(\cdot)) = \int \|\nabla f(\vec{r})\|^2 d\vec{r} = \int \left| \frac{\partial}{\partial x} f(\vec{r}) \right|^2 + \left| \frac{\partial}{\partial y} f(\vec{r}) \right|^2 + \left| \frac{\partial}{\partial z} f(\vec{r}) \right|^2 d\vec{r}.$$

Second derivatives also useful:
 (More choices!)

$$\left. \frac{\partial^2}{\partial x^2} f(\vec{r}) \right|_{\vec{r}=\vec{r}_j} \approx f(\vec{r}_{j+1}) - 2f(\vec{r}_j) + f(\vec{r}_{j-1})$$

$$R(\mathbf{x}) = \sum_{j=1}^{n_p} \psi(x_{j+1} - 2x_j + x_{j-1}) + \dots$$

2.34

For differentiable basis functions (e.g., B-splines), one can find $\int \|\nabla f(\vec{r})\|^2 d\vec{r}$ exactly in terms of coefficients, e.g., [130].

See Gindi *et al.* [131, 132] for comparisons of first and second order penalties.

Penalty Functions: General Form

$$R(\mathbf{x}) = \sum_k \psi_k([C\mathbf{x}]_k) \quad \text{where} \quad [C\mathbf{x}]_k = \sum_{j=1}^{n_p} c_{kj}x_j$$

Example:

x_1	x_2	x_3
x_4	x_5	

$$C\mathbf{x} = \begin{bmatrix} -1 & 1 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & -1 & 1 \\ -1 & 0 & 0 & 1 & 0 \\ 0 & -1 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} = \begin{bmatrix} x_2 - x_1 \\ x_3 - x_2 \\ x_5 - x_4 \\ x_4 - x_1 \\ x_5 - x_2 \end{bmatrix}$$

$$R(\mathbf{x}) = \sum_{k=1}^5 \psi_k([C\mathbf{x}]_k) = \psi_1(x_2 - x_1) + \psi_2(x_3 - x_2) + \psi_3(x_5 - x_4) + \psi_4(x_4 - x_1) + \psi_5(x_5 - x_2)$$

2.35

This form is general enough to cover nearly all the penalty functions that have been used in tomography. Exceptions include priors based on nonseparable line-site models [133–136]. and the median root prior [137, 138], both of which are nonconvex.

It is just coincidence that C is square in this example. In general, for a $n_x \times n_y$ image, there are $n_x(n_y - 1)$ horizontal pairs and $n_y(n_x - 1)$ vertical pairs, so C will be a $(2n_xn_y - n_x - n_y) \times (n_xn_x)$ very sparse matrix (for a first-order neighborhood consisting of horizontal and vertical cliques).

Concretely, for a $n_x \times n_y$ image ordered lexicographically, for a first-order neighborhood we use

$$C = \begin{bmatrix} I_{n_y} \otimes D_{n_x} \\ D_{n_y} \otimes I_{n_x} \end{bmatrix}$$

where \otimes denotes the Kronecker product and D_n denotes the following $(n - 1) \times n$ matrix:

$$D_n \triangleq \begin{bmatrix} -1 & 1 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & \ddots & \ddots & 0 \\ 0 & 0 & 0 & -1 & 1 \end{bmatrix}.$$

Penalty Functions: Quadratic vs Nonquadratic

$$R(\mathbf{x}) = \sum_k \psi_k([C\mathbf{x}]_k)$$

Quadratic ψ_k

If $\psi_k(t) = t^2/2$, then $R(\mathbf{x}) = \frac{1}{2}\mathbf{x}'C'C\mathbf{x}$, a quadratic form.


- Simpler optimization
- Global smoothing

Nonquadratic ψ_k

- Edge preserving
- More complicated optimization. (This is essentially solved in convex case.)
- Unusual noise properties
- Analysis/prediction of resolution and noise properties is difficult
- More adjustable parameters (e.g., δ)

Example: Huber function. $\psi(t) \triangleq \begin{cases} t^2/2, & |t| \leq \delta \\ \delta|t| - \delta^2/2, & |t| > \delta \end{cases}$


2.36



Lower cost for large differences \Rightarrow edge preservation

2.37

Edge-Preserving Reconstruction Example



A transmission example would be preferable...

2.38

In terms of ROI quantification, a nonquadratic penalty may outperform quadratic penalties for certain types of objects (especially phantom-like piecewise smooth objects). But the benefits of nonquadratic penalties for visual tasks is largely unknown.

The smaller δ is in the Huber penalty, the stronger the degree of edge preservation, and the more unusual the noise effects. In this case I used $\delta = 0.4$, for a phantom that is 0 in background, 1 in white matter, 4 in graymatter. Thus δ is one tenth the maximum value, as has been recommended by some authors.

Penalty Functions: Convex vs Nonconvex

Convex

- Easier to optimize
- Guaranteed unique minimizer of Ψ (for convex negative log-likelihood)

Nonconvex

- Greater degree of edge preservation
- Nice images for piecewise-constant phantoms!
- Even more unusual noise properties
- Multiple extrema
- More complicated optimization (simulated / deterministic annealing)
- Estimator \hat{x} becomes a discontinuous function of data Y

Nonconvex examples

- “broken parabola”

$$\psi(t) = \min(t^2, t_{\max}^2)$$

- true median root prior:

$$R(\mathbf{x}) = \sum_{j=1}^{n_p} \frac{(x_j - \text{median}_j(\mathbf{x}))^2}{\text{median}_j(\mathbf{x})} \quad \text{where } \text{median}_j(\mathbf{x}) \text{ is local median}$$

Exception: orthonormal wavelet threshold *denoising* via nonconvex potentials!

2.39

The above form is not exactly what has been called the median root prior by Alenius *et al.* [139]. They have used $\text{median}_j(x^{(n)})$ which is not a true prior since it depends on the previous iteration. Hsiao, Rangarajan, and Ginda have developed a very interesting prior that is similar to the “medial root prior” but is convex [140].

For nice analysis of nonconvex problems, see the papers by Mila Nikolova [141].

For orthonormal wavelet denoising, the cost functions [142] usually have the form

$$\Psi(\mathbf{x}) = \|\mathbf{y} - \mathbf{Ax}\|^2 + \sum_{j=1}^{n_p} \psi(x_j)$$


where \mathbf{A} is an orthonormal. When \mathbf{A} is orthonormal we can write: $\|\mathbf{y} - \mathbf{Ax}\|^2 = \|\mathbf{A}'\mathbf{y} - \mathbf{x}\|^2$, so

$$\Psi(\mathbf{x}) = \sum_{j=1}^{n_p} (x_j - [\mathbf{A}'\mathbf{y}]_j)^2 + \psi(x_j)$$

which completely separates into n_p 1-D minimization problems, each of which has a unique minimizer for all useful potential functions.

2.39


p2reg



2.40

2.40

Local Extrema and Discontinuous Estimators



Small change in data \Rightarrow large change in minimizer \hat{x} .
Using convex penalty functions obviates this problem.

2.41

[114] discuss discontinuity

Augmented Regularization Functions

Replace roughness penalty $R(x)$ with $R(x|b) + \alpha R(b)$,
where the elements of b (often binary) indicate boundary locations.

- Line-site methods
- Level-set methods

Joint estimation problem:

$$(\hat{x}, \hat{b}) = \arg \min_{x, b} \Psi(x, b), \quad \Psi(x, b) = -L(x; y) + \beta R(x|b) + \alpha R(b).$$

Example: b_{jk} indicates the presence of edge between pixels j and k :

$$R(x|b) = \sum_{j=1}^{n_p} \sum_{k \in N_j} (1 - b_{jk}) \frac{1}{2} (x_j - x_k)^2$$

Penalty to discourage too many edges (e.g.):

$$R(b) = \sum_{jk} b_{jk}.$$

- Can encourage local edge continuity
- Require annealing methods for minimization

2.42

Line-site methods: [133–136].

Level-set methods: [143–145].

For the simple *non-interacting* line-site penalty function $R(b)$ given above, one can perform the minimization over b analytically, yielding an equivalent regularization method of the form $R(x)$ with a broken parabola potential function [146].

More sophisticated line-site methods use neighborhoods of line-site variables to encourage local boundary continuity [133–136].

Modified Penalty Functions

$$R(\mathbf{x}) = \sum_{j=1}^{n_p} \frac{1}{2} \sum_{k \in N_j} w_{jk} \Psi(x_j - x_k)$$

Adjust weights $\{w_{jk}\}$ to

- Control resolution properties
- Incorporate anatomical side information (MR/CT)
(avoid smoothing across anatomical boundaries)

Recommendations

- Emission tomography:
 - begin with quadratic (nonseparable) penalty functions
 - Consider modified penalty for resolution control and choice of β
 - Use modest regularization and post-filter more if desired
- Transmission tomography (attenuation maps), X-ray CT
 - consider convex nonquadratic (*e.g.*, Huber) penalty functions
 - choose δ based on attenuation map units (water, bone, etc.)
 - choice of regularization parameter β remains nontrivial,
learn appropriate values by experience for given study type

2.43

Resolution properties [25, 147, 148]

Side information (a very incomplete list) [149–160].

Choice 4.3: Constraints

- Nonnegativity
- Known support
- Count preserving
- Upper bounds on values
e.g., maximum μ of attenuation map in transmission case

Considerations

- Algorithm complexity
- Computation
- Convergence rate
- Bias (in low-count regions)
- ...

2.44

Sometimes it is stated that the ML-EM algorithm “preserves counts.” This only holds when $r_i = 0$ in the statistical model. The count-preserving property originates from the likelihood, not the algorithm. The ML estimate, under the Poisson model, happens to preserve counts. It is fine that ML-EM does so every iteration, but that does not mean that it is superior to other algorithms that get to the optimum \hat{x} faster without necessarily preserving counts along the way.

I do not recommend artificially renormalizing each iteration to try to “preserve counts.”

Open Problems

Modeling

- Noise in a_{ij} 's (system model errors)
- Noise in \hat{r}_i 's (estimates of scatter / randoms)
- Statistics of corrected measurements
- Statistics of measurements with deadtime losses

Cost functions

- Performance prediction for nonquadratic penalties
- Effect of nonquadratic penalties on detection tasks
- Choice of regularization parameters for nonquadratic regularization

2.45

Deadtime statistics are analyzed in [45, 46]. Bottom line: in most SPECT and PET systems with paralyzable deadtime, the measurements are non-Poisson, but the mean and variance are nearly identical. So presumably the Poisson statistical model is adequate, provided the deadtime losses are included in the system matrix A .

Summary

- 1. Object parameterization: function $f(\vec{r})$ vs vector x
- 2. System physical model: $s_i(x)$
- 3. Measurement statistical model $Y_i \sim \boxed{?}$
- 4. Cost function: data-mismatch / regularization / constraints

Reconstruction Method = Cost Function + Algorithm

Naming convention:

- ML-EM, MAP-OSL, PL-SAGE, PWLS+SOR, PWLS-CG, ...

2.46

Part 3. Algorithms

Method = Cost Function + Algorithm

Outline

- Ideal algorithm
- Classical general-purpose algorithms
- Considerations:
 - nonnegativity
 - parallelization
 - convergence rate
 - monotonicity
- Algorithms tailored to *cost functions* for imaging
 - Optimization transfer
 - EM-type methods
 - Poisson emission problem
 - Poisson transmission problem
- Ordered-subsets / block-iterative algorithms

3.1

Choosing a *cost function* is an important part of imaging science.

Choosing an algorithm should be mostly a matter of numerical methods.

Nevertheless, it gets a lot of attention by imaging scientists since our *cost functions* have forms that can be exploited to get faster convergence than general-purpose methods.

Why iterative algorithms?

- For nonquadratic Ψ , no closed-form solution for minimizer.
- For quadratic Ψ with nonnegativity constraints, no closed-form solution.
- For quadratic Ψ without constraints, closed-form solutions:

$$\text{PWLS: } \hat{\mathbf{x}} = [\mathbf{A}'\mathbf{W}\mathbf{A} + \mathbf{R}]^{-1}\mathbf{A}'\mathbf{W}\mathbf{y}$$

$$\text{OLS: } \hat{\mathbf{x}} = [\mathbf{A}'\mathbf{A}]^{-1}\mathbf{A}'\mathbf{y}$$

Impractical (memory and computation) for realistic problem sizes.
 \mathbf{A} is sparse, but $\mathbf{A}'\mathbf{A}$ is not.

All algorithms are imperfect. No single best solution.

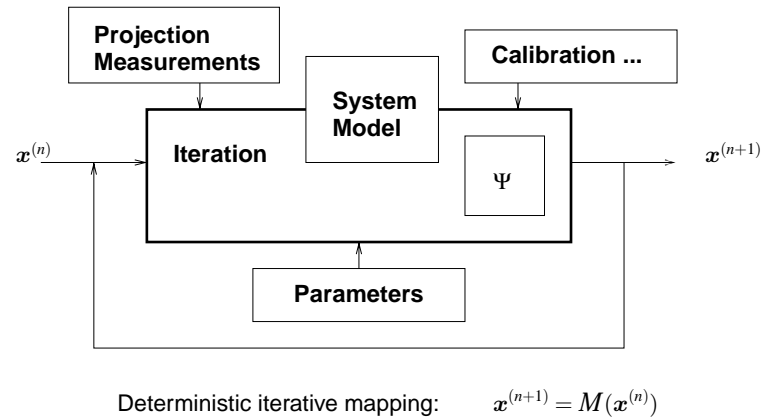
3.2

Singular value decomposition (SVD) techniques have been proposed for the OLS *cost function* as a method for reducing the computation problem, *e.g.*, [161–170].

The idea is that one could precompute the pseudo-inverse of \mathbf{A} “once and for all.” However \mathbf{A} includes physical effects like attenuation, which change for every patient. And for data-weighted least squares, \mathbf{W} changes for each scan too.

Image reconstruction never requires the matrix inverse $[\mathbf{A}'\mathbf{A}]^{-1}$; all that is required is a solution to the normal equations $[\mathbf{A}'\mathbf{A}]\hat{\mathbf{x}} = \mathbf{A}'\mathbf{y}$ which is easier, but still nontrivial.

General Iteration



3.3

There are also stochastic iterative algorithms, such as simulated annealing [133] and the stochastic EM algorithm [171].

Ideal Algorithm

$$x^* \triangleq \arg \min_{x \geq 0} \Psi(x) \quad (\text{global minimizer})$$

Properties

stable and convergent
converges quickly
globally convergent
fast
robust
user friendly

$\{x^{(n)}\}$ converges to x^* if run indefinitely
 $\{x^{(n)}\}$ gets "close" to x^* in just a few iterations
 $\lim_n x^{(n)}$ independent of starting image $x^{(0)}$
requires minimal computation per iteration
insensitive to finite numerical precision
nothing to adjust (e.g., acceleration factors)

parallelizable
simple
flexible
(matrix stored by row or column or projector/backprojector)

(when necessary)
easy to program and debug
accommodates any type of system model

Choices: forgo one or more of the above

3.4

One might argue that the "ideal algorithm" would be the algorithm that produces x^{true} . In the framework presented here, it is the job of the *cost function* to try to make $x^* \approx x^{\text{true}}$, and the job of the algorithm to find x^* by minimizing Ψ .

In fact, *nothing* in the above list really has to do with image quality. In the statistical framework, image quality is determined by Ψ , not by the algorithm.

Note on terminology: "algorithms" do not really converge, it is the *sequence* of estimates $\{x^{(n)}\}$ that converges, but everyone abuses this all the time, so I will too.

Classic Algorithms

Non-gradient based

- Exhaustive search
- Nelder-Mead simplex (amoeba)

Converge very slowly, but work with nondifferentiable *cost functions*.

Gradient based

- Gradient descent

$$\mathbf{x}^{(n+1)} \triangleq \mathbf{x}^{(n)} - \alpha \nabla \Psi(\mathbf{x}^{(n)})$$

Choosing α to ensure convergence is nontrivial.

- Steepest descent

$$\mathbf{x}^{(n+1)} \triangleq \mathbf{x}^{(n)} - \alpha_n \nabla \Psi(\mathbf{x}^{(n)}) \quad \text{where} \quad \alpha_n \triangleq \arg \min_{\alpha} \Psi(\mathbf{x}^{(n)} - \alpha \nabla \Psi(\mathbf{x}^{(n)}))$$

Computing α_n can be expensive.

Limitations

- Converge slowly.
- Do not easily accommodate nonnegativity constraint.

3.5

Nice discussion of optimization algorithms in [172].

Row and column gradients:

$$\nabla \Psi(\mathbf{x}) = \left[\frac{\partial}{\partial x_1} \Psi, \frac{\partial}{\partial x_2} \Psi, \dots, \frac{\partial}{\partial x_{n_p}} \Psi \right], \quad \mathbf{\nabla} = \nabla'$$

Using gradients excludes nondifferentiable penalty functions such as the Laplacian prior which involves $|x_j - x_k|$. See [173–175] for solutions to this problem.

3.5

p3alg

Gradients & Nonnegativity - A Mixed Blessing

Unconstrained optimization of differentiable *cost functions*:

$$\nabla \Psi(\mathbf{x}) = 0 \quad \text{when} \quad \mathbf{x} = \mathbf{x}^*$$

- A necessary condition always.
- A sufficient condition for strictly convex *cost functions*.
- Iterations search for zero of gradient.

Nonnegativity-constrained minimization:

Karush-Kuhn-Tucker conditions


$$\left. \frac{\partial}{\partial x_j} \Psi(\mathbf{x}) \right|_{\mathbf{x}=\mathbf{x}^*} \quad \text{is} \quad \begin{cases} = 0, & x_j^* > 0 \\ \geq 0, & x_j^* = 0 \end{cases}$$

- A necessary condition always.
- A sufficient condition for strictly convex *cost functions*.
- Iterations search for ???
- $0 = x_j^* \frac{\partial}{\partial x_j} \Psi(\mathbf{x}^*)$ is a necessary condition, but never sufficient condition.

3.6

3.6


Karush-Kuhn-Tucker Illustrated



3.7

The usual condition $\frac{\partial}{\partial x_j} \Psi(x) = 0$ only applies for pixels where the nonnegativity constraint is inactive.

Why Not Clip Negatives?



Newton-Raphson with negatives set to zero each iteration. Fixed-point of iteration is not the constrained minimizer!

3.8

By clipped negatives, I mean you start with some nominal algorithm $M_0(x)$ and modify it to be: $x^{(n+1)} = M(x^{(n)})$ where $M(x) = [M_0(x)]_+$ and the j th element of $[x]_+$ is x_j if $x_j > 0$ or 0 if $x_j \leq 0$. Basically, you run your favorite iteration and then set any negatives to zero before proceeding to the next iteration.

Simple 2D quadratic problem. Curves show contours of equal value of the *cost function* Ψ .

Same problem arises with upper bounds too.

The above problem applies generally to simultaneous update iterative methods. For sequential update methods, such as coordinate descent, clipping works fine.

Newton-Raphson Algorithm

$$\mathbf{x}^{(n+1)} = \mathbf{x}^{(n)} - [\nabla^2 \Psi(\mathbf{x}^{(n)})]^{-1} \nabla \Psi(\mathbf{x}^{(n)})$$

Advantage:

- Super-linear convergence rate (if convergent)

Disadvantages:

- Requires twice-differentiable Ψ
- Not guaranteed to converge
- Not guaranteed to monotonically decrease Ψ
- Does not enforce nonnegativity constraint
- Impractical for image recovery due to matrix inverse

General purpose remedy: bound-constrained Quasi-Newton algorithms

3.9

$\nabla^2 \Psi(\mathbf{x})$ is called the *Hessian matrix*. It is a $n_p \times n_p$ matrix (where n_p is the dimension of \mathbf{x}). The j, k th element of it is $\frac{\partial^2}{\partial x_j \partial x_k} \Psi(\mathbf{x})$.

A “matrix inverse” actually is not necessary. One can rewrite the above iteration as $\mathbf{x}^{(n+1)} = \mathbf{x}^{(n)} - \mathbf{d}^{(n)}$ where $\mathbf{d}^{(n)}$ is the solution to the system of equations: $\nabla^2 \Psi(\mathbf{x}^{(n)}) \mathbf{d}^{(n)} = \nabla \Psi(\mathbf{x}^{(n)})$. Unfortunately, this is a non-sparse $n_p \times n_p$ system of equations, requiring $O(n_p^3)$ flops to solve, which is expensive. Instead of solving the system exactly one could use approximate iterative techniques, but then it should probably be considered a preconditioned gradient method rather than Newton-Raphson.

Quasi-Newton algorithms [176–179] [180, p. 136] [181, p. 77] [182, p. 63].

bound-constrained Quasi-Newton algorithms (LBFGS) [178, 183–186].

3.9

p3alg

Newton’s Quadratic Approximation

2nd-order Taylor series:

$$\Psi(\mathbf{x}) \approx \phi(\mathbf{x}; \mathbf{x}^{(n)}) \triangleq \Psi(\mathbf{x}^{(n)}) + \nabla \Psi(\mathbf{x}^{(n)})^T (\mathbf{x} - \mathbf{x}^{(n)}) + \frac{1}{2} (\mathbf{x} - \mathbf{x}^{(n)})^T \nabla^2 \Psi(\mathbf{x}^{(n)}) (\mathbf{x} - \mathbf{x}^{(n)})$$

Set $\mathbf{x}^{(n+1)}$ to the (“easily” found) minimizer of this quadratic approximation:

$$\begin{aligned} \mathbf{x}^{(n+1)} &\triangleq \arg \min_{\mathbf{x}} \phi(\mathbf{x}; \mathbf{x}^{(n)}) \\ &= \mathbf{x}^{(n)} - [\nabla^2 \Psi(\mathbf{x}^{(n)})]^{-1} \nabla \Psi(\mathbf{x}^{(n)}) \end{aligned}$$


Can be nonmonotone for Poisson emission tomography log-likelihood, even for a single pixel and single ray:

$$\Psi(x) = (x + r) - y \log(x + r)$$

3.10

3.10

Nonmonotonicity of Newton-Raphson



3.11

Consideration: Monotonicity

An algorithm is monotonic if

$$\Psi(\mathbf{x}^{(n+1)}) \leq \Psi(\mathbf{x}^{(n)}), \quad \forall \mathbf{x}^{(n)}.$$

Three categories of algorithms:

- Nonmonotonic (or unknown)
- Forced monotonic (e.g., by line search)
- Intrinsically monotonic (by design, simplest to implement)

Forced monotonicity

Most nonmonotonic algorithms can be converted to forced monotonic algorithms by adding a line-search step:

$$\mathbf{x}^{\text{temp}} \triangleq M(\mathbf{x}^{(n)}), \quad \mathbf{d} = \mathbf{x}^{\text{temp}} - \mathbf{x}^{(n)}$$

$$\mathbf{x}^{(n+1)} \triangleq \mathbf{x}^{(n)} - \alpha_n \mathbf{d}^{(n)} \quad \text{where} \quad \alpha_n \triangleq \arg \min_{\alpha} \Psi(\mathbf{x}^{(n)} - \alpha \mathbf{d}^{(n)})$$

Inconvenient, sometimes expensive, nonnegativity problematic.

3.12

Although monotonicity is not a necessary condition for an algorithm to converge globally to \mathbf{x}^* , it is often the case that global convergence and monotonicity go hand in hand. In fact, for strictly convex Ψ , algorithms that monotonically decrease Ψ each iteration are guaranteed to converge under reasonable regularity conditions [187].

Any algorithm containing a line search step will have difficulties with nonnegativity. In principle one can address these problems using a “bent-line” search [188], but this can add considerable computation per iteration.

Conjugate Gradient Algorithm

Advantages:

- Fast converging (if suitably preconditioned) (in unconstrained case)
- Monotonic (forced by line search in nonquadratic case)
- Global convergence (unconstrained case)
- Flexible use of system matrix A and tricks
- Easy to implement in unconstrained quadratic case
- Highly parallelizable

Disadvantages:

- Nonnegativity constraint awkward (slows convergence?)
- Line-search awkward in nonquadratic cases

Highly recommended for unconstrained quadratic problems (*e.g.*, PWLS without nonnegativity). Useful (but perhaps not ideal) for Poisson case too.

3.13

CG is like steepest descent, but the search direction is modified each iteration to be conjugate to the previous search direction.

Preconditioners [189, 190]

Poisson case [109, 191, 192].

Consideration: Parallelization

Simultaneous (fully parallelizable)

update all pixels simultaneously using all data
EM, Conjugate gradient, ISRA, OSL, SIRT, MART, ...

Block iterative (ordered subsets)

update (nearly) all pixels using one subset of the data at a time
OSEM, RBBI, ...

Row action

update many pixels using a single ray at a time
ART, RAMLA

Pixel grouped (multiple column action)

update some (but not all) pixels simultaneously a time, using all data
Grouped coordinate descent, multi-pixel SAGE
(Perhaps the most nontrivial to implement)

Sequential (column action)

update one pixel at a time, using all (relevant) data
Coordinate descent, SAGE

3.14

Sequential algorithms are the least parallelizable since one cannot update the second pixel until the first pixel has been updated (to preserve monotonicity and convergence properties).

SAGE [193, 194]

Grouped coordinate descent [195]

Multi-pixel SAGE [196]

RAMLA [197]

OSEM [15]

RBBI [198–200]

ISRA [201–203]

OSL [204, 205]

Coordinate Descent Algorithm

aka Gauss-Siedel, successive over-relaxation (SOR), iterated conditional modes (ICM)

Update one pixel at a time, holding others fixed to their most recent values:

$$x_j^{\text{new}} = \arg \min_{x_j \geq 0} \Psi(x_1^{\text{new}}, \dots, x_{j-1}^{\text{new}}, x_j, x_{j+1}^{\text{old}}, \dots, x_{n_p}^{\text{old}}), \quad j = 1, \dots, n_p$$

Advantages:

- Intrinsically monotonic
- Fast converging (from good initial image)
- Global convergence
- Nonnegativity constraint trivial

Disadvantages:

- Requires column access of system matrix A
- Cannot exploit some “tricks” for A
- Expensive “arg min” for nonquadratic problems
- Poorly parallelizable

3.15

Fast convergence shown by Sauer and Bouman with clever frequency-domain analysis [206].


Any ordering can be used. Convergence rate may vary with ordering.

Global convergence even with negatives clipped [207].

One can replace the “arg min” with a one-dimensional Newton-Raphson step [195, 208–210]. However, this change then loses the guarantee of monotonicity for nonquadratic Ψ . Also, evaluating the second partial derivatives of Ψ with respect to x_j is expensive (costs an extra modified backprojection per iteration) [195].

The paraboloidal surrogates coordinate descent (PSCD) algorithm circumvents these problems [211].


Constrained Coordinate Descent Illustrated



3.16

In this particular case, the nonnegativity constraint led to exact convergence in 1.5 iterations.

Coordinate Descent - Unconstrained



3.17

In general coordinate descent converges at a linear rate [104, 206].

Interestingly, for this particular problem the nonnegativity constraint accelerated convergence.

Coordinate-Descent Algorithm Summary

Recommended when all of the following apply:

- quadratic or nearly-quadratic convex *cost function*
- nonnegativity constraint desired
- precomputed and stored system matrix A with column access
- parallelization not needed (standard workstation)

Cautions:

- Good initialization (e.g., properly scaled FBP) essential.
(Uniform image or zero image cause slow initial convergence.)
- Must be programmed carefully to be efficient.
(Standard Gauss-Siedel implementation is suboptimal.)
- Updates high-frequencies fastest \Rightarrow poorly suited to unregularized case

Used daily in UM clinic for 2D SPECT / PWLS / nonuniform attenuation

3.18

In saying “not good for the unregularized case” I am assuming one does not really wish to find the minimizer of Ψ in that case. If you really want the minimizer of Ψ in the unregularized case, then coordinate descent may still be useful.

Summary of General-Purpose Algorithms

Gradient-based

- Fully parallelizable
- Inconvenient line-searches for nonquadratic *cost functions*
- Fast converging in unconstrained case
- Nonnegativity constraint inconvenient

Coordinate-descent

- Very fast converging
- Nonnegativity constraint trivial
- Poorly parallelizable
- Requires precomputed/stored system matrix

CD is well-suited to moderate-sized 2D problem (e.g., 2D PET), but poorly suited to large 2D problems (X-ray CT) and fully 3D problems

Neither is ideal.

\therefore need *special-purpose algorithms* for image reconstruction!

3.19

Interior-point methods for general-purpose constrained optimization have recently been applied to image reconstruction [212] and deserve further examination.

Data-Mismatch Functions Revisited

For fast converging, intrinsically monotone algorithms, consider the form of Ψ .

WLS:

$$-L(\mathbf{x}) = \sum_{i=1}^{n_d} \frac{1}{2} w_i (y_i - [\mathbf{Ax}]_i)^2 = \sum_{i=1}^{n_d} h_i([\mathbf{Ax}]_i), \quad \text{where } h_i(l) \triangleq \frac{1}{2} w_i (y_i - l)^2.$$

Emission Poisson log-likelihood:

$$-L(\mathbf{x}) = \sum_{i=1}^{n_d} ([\mathbf{Ax}]_i + r_i) - y_i \log([\mathbf{Ax}]_i + r_i) = \sum_{i=1}^{n_d} h_i([\mathbf{Ax}]_i) \\ \text{where } h_i(l) \triangleq (l + r_i) - y_i \log(l + r_i).$$

Transmission Poisson log-likelihood:

$$-L(\mathbf{x}) = \sum_{i=1}^{n_d} \left(b_i e^{-[\mathbf{Ax}]_i} + r_i \right) - y_i \log \left(b_i e^{-[\mathbf{Ax}]_i} + r_i \right) = \sum_{i=1}^{n_d} h_i([\mathbf{Ax}]_i) \\ \text{where } h_i(l) \triangleq (b_i e^{-l} + r_i) - y_i \log(b_i e^{-l} + r_i).$$

MRI, polyenergetic X-ray CT, confocal microscopy, image restoration, ...
All have same *partially separable* form.

3.20

All the algorithms discussed this far are generic; they can be applied to any differentiable Ψ .

General Imaging Cost Function

General form for data-mismatch function:

$$-L(\mathbf{x}) = \sum_{i=1}^{n_d} h_i([\mathbf{A}\mathbf{x}]_i)$$

General form for regularizing penalty function:

$$R(\mathbf{x}) = \sum_k \psi_k([\mathbf{C}\mathbf{x}]_k)$$

General form for *cost function*:

$$\Psi(\mathbf{x}) = -L(\mathbf{x}) + \beta R(\mathbf{x}) = \sum_{i=1}^{n_d} h_i([\mathbf{A}\mathbf{x}]_i) + \beta \sum_k \psi_k([\mathbf{C}\mathbf{x}]_k)$$


Properties of Ψ we can exploit:

- summation form (due to independence of measurements)
- convexity of h_i functions (usually)
- summation argument (inner product of \mathbf{x} with i th row of \mathbf{A})

Most methods that use these properties are forms of *optimization transfer*.

3.21

Optimization Transfer Illustrated



3.22

This figure does not do justice to the problem. A one-dimensional Ψ is usually easy to minimize. The problem is in multiple dimensions.

Optimization Transfer

General iteration:

$$\mathbf{x}^{(n+1)} = \arg \min_{\mathbf{x} \geq 0} \phi(\mathbf{x}; \mathbf{x}^{(n)})$$

Monotonicity conditions (Ψ decreases provided these hold):


- $\phi(\mathbf{x}^{(n)}; \mathbf{x}^{(n)}) = \Psi(\mathbf{x}^{(n)})$ (matched current value)
- $\nabla_{\mathbf{x}} \phi(\mathbf{x}; \mathbf{x}^{(n)}) \Big|_{\mathbf{x}=\mathbf{x}^{(n)}} = \nabla \Psi(\mathbf{x}) \Big|_{\mathbf{x}=\mathbf{x}^{(n)}}$ (matched gradient)
- $\phi(\mathbf{x}; \mathbf{x}^{(n)}) \geq \Psi(\mathbf{x}) \quad \forall \mathbf{x} \geq 0$ (lies above)

These 3 (sufficient) conditions are satisfied by the Q function of the EM algorithm (and SAGE).

The 3rd condition is *not* satisfied by the Newton-Raphson quadratic approximation, which leads to its nonmonotonicity.

3.23

Optimization Transfer in 2d



3.24

Optimization Transfer of EM Algorithm

E-step: choose surrogate function $\phi(x; x^{(n)})$

M-step: minimize surrogate function

$$x^{(n+1)} = \arg \min_{x \geq 0} \phi(x; x^{(n)})$$

Designing surrogate functions

- Easy to “compute”
- Easy to minimize
- Fast convergence rate

Often mutually incompatible goals \therefore compromises


3.25

From the point of view of “per iteration convergence rate,” the optimal “surrogate function” would be just Ψ itself. However, then the M-step is very difficult (in fact it is the original optimization problem). Such an “algorithm” would converge in one very expensive “iteration.”

3.25

p3x

Convergence Rate: Slow




3.26

© J. Fessler, November 6, 2002

3.26

p3x


Convergence Rate: Fast



3.27

Tradeoff between curvature and ease of M-step... Can we beat this tradeoff?

Tool: Convexity Inequality



g convex $\Rightarrow g(\alpha x_1 + (1 - \alpha)x_2) \leq \alpha g(x_1) + (1 - \alpha)g(x_2)$ for $\alpha \in [0, 1]$

More generally: $\alpha_k \geq 0$ and $\sum_k \alpha_k = 1 \Rightarrow g(\sum_k \alpha_k x_k) \leq \sum_k \alpha_k g(x_k)$. Sum outside!

3.28

The emission Poisson ray log-likelihood h_i is strictly convex on $(-r_i, \infty)$. This turns out to be adequate for the derivation.

Example 1: Classical ML-EM Algorithm

Negative Poisson log-likelihood *cost function* (unregularized):

$$\Psi(\mathbf{x}) = \sum_{i=1}^{n_d} h_i([\mathbf{Ax}]_i), \quad h_i(l) = (l + r_i) - y_i \log(l + r_i).$$

Intractable to minimize directly due to summation within logarithm.

Clever trick due to De Pierro (let $\bar{y}_i^{(n)} = [\mathbf{Ax}^{(n)}]_i + r_i$):

$$[\mathbf{Ax}]_i = \sum_{j=1}^{n_p} a_{ij} x_j = \sum_{j=1}^{n_p} \left[\frac{a_{ij} x_j^{(n)}}{\bar{y}_i^{(n)}} \right] \left(\frac{x_j}{x_j^{(n)}} \bar{y}_i^{(n)} \right).$$

Since the h_i 's are *convex* in Poisson emission model:

$$h_i([\mathbf{Ax}]_i) = h_i \left(\sum_{j=1}^{n_p} \left[\frac{a_{ij} x_j^{(n)}}{\bar{y}_i^{(n)}} \right] \left(\frac{x_j}{x_j^{(n)}} \bar{y}_i^{(n)} \right) \right) \leq \sum_{j=1}^{n_p} \left[\frac{a_{ij} x_j^{(n)}}{\bar{y}_i^{(n)}} \right] h_i \left(\frac{x_j}{x_j^{(n)}} \bar{y}_i^{(n)} \right)$$

$$\Psi(\mathbf{x}) = \sum_{i=1}^{n_d} h_i([\mathbf{Ax}]_i) \leq \phi(\mathbf{x}; \mathbf{x}^{(n)}) \triangleq \sum_{i=1}^{n_d} \sum_{j=1}^{n_p} \left[\frac{a_{ij} x_j^{(n)}}{\bar{y}_i^{(n)}} \right] h_i \left(\frac{x_j}{x_j^{(n)}} \bar{y}_i^{(n)} \right)$$

Replace convex *cost function* $\Psi(\mathbf{x})$ with *separable* surrogate function $\phi(\mathbf{x}; \mathbf{x}^{(n)})$.

3.29

The clever (multiplicative) trick in the first equation is due to Alvaro De Pierro [203].

Note that the bracketed terms sum over j to unity.

I believe that this is the shortest and simplest possible derivation of the ML-EM algorithm, out of five distinct derivations I have seen.

This derivation is complete only for the case $r_i = 0$. It is easily generalized to $r_i \neq 0$.

“ML-EM Algorithm” M-step

E-step gave separable surrogate function:

$$\phi(\mathbf{x}; \mathbf{x}^{(n)}) = \sum_{j=1}^{n_p} \phi_j(x_j; \mathbf{x}^{(n)}), \quad \text{where } \phi_j(x_j; \mathbf{x}^{(n)}) \triangleq \sum_{i=1}^{n_d} \left[\frac{a_{ij} x_j^{(n)}}{\bar{y}_i^{(n)}} \right] h_i \left(\frac{x_j}{x_j^{(n)}} \bar{y}_i^{(n)} \right).$$

M-step separates:

$$\mathbf{x}^{(n+1)} = \arg \min_{\mathbf{x} \geq 0} \phi(\mathbf{x}; \mathbf{x}^{(n)}) \Rightarrow x_j^{(n+1)} = \arg \min_{x_j \geq 0} \phi_j(x_j; \mathbf{x}^{(n)}), \quad j = 1, \dots, n_p$$

Minimizing:

$$\frac{\partial}{\partial x_j} \phi_j(x_j; \mathbf{x}^{(n)}) = \sum_{i=1}^{n_d} a_{ij} h_i \left(\frac{x_j}{x_j^{(n)}} \bar{y}_i^{(n)} \right) = \sum_{i=1}^{n_d} a_{ij} \left[1 - \frac{y_i}{\bar{y}_i^{(n)} x_j / x_j^{(n)}} \right] \Big|_{x_j = x_j^{(n+1)}} = 0.$$

Solving (in case $r_i = 0$):

$$x_j^{(n+1)} = x_j^{(n)} \left[\sum_{i=1}^{n_d} a_{ij} \frac{y_i}{[\mathbf{Ax}^{(n)}]_i} \right] / \left(\sum_{i=1}^{n_d} a_{ij} \right), \quad j = 1, \dots, n_p$$

- Derived without any statistical considerations, unlike classical EM formulation.
- Uses only convexity and algebra.
- Guaranteed monotonic: surrogate function ϕ satisfies the 3 required properties.
- M-step trivial due to *separable surrogate*.

3.30

When $r_i = 0$, $\dot{h}_i(l) \triangleq \frac{d}{dl} h_i(l) = 1 - y_i/l$.

Case where $r_i \neq 0$ can also be handled with more algebra. Just replace final $[\mathbf{Ax}^{(n)}]_i$ with $\bar{y}_i^{(n)} = [\mathbf{Ax}^{(n)}]_i + r_i$.

To be rigorous, we should check that the Karush-Kuhn-Tucker condition holds for our minimizer of $\phi_j(\cdot; \mathbf{x}^{(n)})$. It does, provided $\mathbf{x}^{(n)} \geq 0$.

I prefer this derivation over the statistical EM derivation, even though we are doing statistical image reconstruction. Statistics greatly affect the design of Ψ , but minimizing Ψ is really just a numerical problem, and statistics need not have any role in that.

ML-EM is Scaled Gradient Descent

$$\begin{aligned} x_j^{(n+1)} &= x_j^{(n)} \left[\sum_{i=1}^{n_d} a_{ij} \frac{y_i}{\bar{y}_i^{(n)}} \right] / \left(\sum_{i=1}^{n_d} a_{ij} \right) \\ &= x_j^{(n)} + x_j^{(n)} \left[\sum_{i=1}^{n_d} a_{ij} \left(\frac{y_i}{\bar{y}_i^{(n)}} - 1 \right) \right] / \left(\sum_{i=1}^{n_d} a_{ij} \right) \\ &= \boxed{x_j^{(n)} - \left(\frac{x_j^{(n)}}{\sum_{i=1}^{n_d} a_{ij}} \right) \frac{\partial}{\partial x_j} \Psi(\mathbf{x}^{(n)})}, \quad j = 1, \dots, n_p \end{aligned}$$


$$\mathbf{x}^{(n+1)} = \mathbf{x}^{(n)} + D(\mathbf{x}^{(n)}) \nabla \Psi(\mathbf{x}^{(n)})$$

This particular diagonal scaling matrix remarkably

- ensures monotonicity,
- ensures nonnegativity.

3.31

Consideration: Separable vs Nonseparable



Contour plots: loci of equal function values.

Uncoupled vs coupled minimization.

3.32

To find the minimizer of a separable function, one can minimize separately with respect to each argument. To find the minimizer of a nonseparable function, one must consider the variables together. In this sense the minimization problem "couples" together the unknown parameters.

Separable Surrogate Functions (Easy M-step)

The preceding EM derivation structure applies to *any cost function* of the form

$$\Psi(\mathbf{x}) = \sum_{i=1}^{n_d} h_i([A\mathbf{x}]_i).$$

of ISRA (for nonnegative LS), “convex algorithm” for transmission reconstruction

Derivation yields a separable surrogate function

$$\Psi(\mathbf{x}) \leq \phi(\mathbf{x}; \mathbf{x}^{(n)}), \text{ where } \phi(\mathbf{x}; \mathbf{x}^{(n)}) = \sum_{j=1}^{n_p} \phi_j(x_j; \mathbf{x}^{(n)})$$

M-step separates into 1D minimization problems (fully parallelizable):

$$\mathbf{x}^{(n+1)} = \arg \min_{\mathbf{x} \geq 0} \phi(\mathbf{x}; \mathbf{x}^{(n)}) \Rightarrow x_j^{(n+1)} = \arg \min_{x_j \geq 0} \phi_j(x_j; \mathbf{x}^{(n)}), \quad j = 1, \dots, n_p$$

Why do EM / ISRA / convex-algorithm / etc. converge so slowly?

3.33

Unfortunately, choosing additively separable surrogate functions generally leads to very high curvature surrogates, which gives very slow convergence rates. EM is the classic example.

The classic EM algorithm is simple to implement precisely because it uses separable surrogate functions.


The derivation of the “convex algorithm” for the Poisson transmission problem [213] and the convergence proof of the ISRA algorithm [203] use a very similar derivation.

Clarify: the self-similar surrogate function is easy to minimize because it is separable. So even though L and Q are composed of the same ray-log likelihood functions, the latter is easier to minimize because it is separable.

3.33

p3x


Separable vs Nonseparable



Separable surrogates (e.g., EM) have high curvature \therefore slow convergence. Nonseparable surrogates can have lower curvature \therefore faster convergence. Harder to minimize? Use paraboloids (quadratic surrogates).

3.34

High Curvature of EM Surrogate



3.35

Sublinear convergence rate of EM shown in [214].


1D Parabola Surrogate Function

Find parabola $q_i^{(n)}(l)$ of the form:

$$q_i^{(n)}(l) = h_i(\ell_i^{(n)}) + h_i'(\ell_i^{(n)})(l - \ell_i^{(n)}) + c_i^{(n)} \frac{1}{2} (l - \ell_i^{(n)})^2, \text{ where } \ell_i^{(n)} \triangleq [Ax^{(n)}]_i$$

Satisfies tangent condition. Choose curvature to ensure “lies above” condition:

$$c_i^{(n)} \triangleq \min \left\{ c \geq 0 : q_i^{(n)}(l) \geq h_i(l), \quad \forall l \geq 0 \right\}.$$



Lower curvature!

Paraboloidal Surrogate

Combining 1D parabola surrogates yields *paraboloidal surrogate*:

$$\Psi(\mathbf{x}) = \sum_{i=1}^{n_d} h_i([\mathbf{A}\mathbf{x}]_i) \leq \phi(\mathbf{x}; \mathbf{x}^{(n)}) = \sum_{i=1}^{n_d} q_i^{(n)}([\mathbf{A}\mathbf{x}]_i)$$

$$\text{Rewriting: } \phi(\delta + \mathbf{x}^{(n)}; \mathbf{x}^{(n)}) = \Psi(\mathbf{x}^{(n)}) + \nabla \Psi(\mathbf{x}^{(n)})\delta + \frac{1}{2}\delta' \mathbf{A}' \text{diag}\{c_i^{(n)}\} \mathbf{A}\delta$$

Advantages

- Surrogate $\phi(\mathbf{x}; \mathbf{x}^{(n)})$ is *quadratic*, unlike Poisson log-likelihood
 \Rightarrow easier to minimize
- Not separable (unlike EM surrogate)
- Not self-similar (unlike EM surrogate)
- Small curvatures \Rightarrow fast convergence
- Intrinsically monotone global convergence
- Fairly simple to derive / implement

Quadratic minimization

- Coordinate descent
 - + fast converging
 - + Nonnegativity easy
 - precomputed column-stored system matrix
- Gradient-based quadratic minimization methods
 - Nonnegativity inconvenient

3.37

Instead of coordinate descent, one could also apply nonnegativity-constrained conjugate gradient.


PSCD recommended for 2D emission Poisson likelihood when system matrix precomputed and stored by columns.

Example: PSCD for PET Transmission Scans

FBP

PL-OSTR-16
4 iterations

PL-PSCD
10 iterations



- square-pixel basis
- strip-integral system model
- shifted-Poisson statistical model
- edge-preserving convex regularization (Huber)
- nonnegativity constraint
- inscribed circle support constraint
- paraboloidal surrogate coordinate descent (PSCD) algorithm

3.38

Separable Paraboloidal Surrogate

To derive a parallelizable algorithm apply another De Pierro trick:

$$[\mathbf{Ax}]_i = \sum_{j=1}^{n_p} \pi_{ij} \left[\frac{a_{ij}}{\pi_{ij}} (x_j - x_j^{(n)}) + \ell_i^{(n)} \right], \quad \ell_i^{(n)} = [\mathbf{Ax}^{(n)}]_i.$$

Provided $\pi_{ij} \geq 0$ and $\sum_{j=1}^{n_p} \pi_{ij} = 1$, since parabola q_i is convex:

$$q_i^{(n)}([\mathbf{Ax}]_i) = q_i^{(n)} \left(\sum_{j=1}^{n_p} \pi_{ij} \left[\frac{a_{ij}}{\pi_{ij}} (x_j - x_j^{(n)}) + \ell_i^{(n)} \right] \right) \leq \sum_{j=1}^{n_p} \pi_{ij} q_i^{(n)} \left(\frac{a_{ij}}{\pi_{ij}} (x_j - x_j^{(n)}) + \ell_i^{(n)} \right)$$

$$\therefore \phi(\mathbf{x}; \mathbf{x}^{(n)}) = \sum_{i=1}^{n_d} q_i^{(n)}([\mathbf{Ax}]_i) \leq \tilde{\phi}(\mathbf{x}; \mathbf{x}^{(n)}) \triangleq \sum_{i=1}^{n_d} \sum_{j=1}^{n_p} \pi_{ij} q_i^{(n)} \left(\frac{a_{ij}}{\pi_{ij}} (x_j - x_j^{(n)}) + \ell_i^{(n)} \right)$$

Separable Paraboloidal Surrogate:

$$\tilde{\phi}(\mathbf{x}; \mathbf{x}^{(n)}) = \sum_{j=1}^{n_p} \phi_j(x_j; \mathbf{x}^{(n)}), \quad \phi_j(x_j; \mathbf{x}^{(n)}) \triangleq \sum_{i=1}^{n_d} \pi_{ij} q_i^{(n)} \left(\frac{a_{ij}}{\pi_{ij}} (x_j - x_j^{(n)}) + \ell_i^{(n)} \right)$$

Parallelizable M-step (cf gradient descent!):

$$x_j^{(n+1)} = \arg \min_{x_j \geq 0} \phi_j(x_j; \mathbf{x}^{(n)}) = \left[x_j^{(n)} - \frac{1}{d_j^{(n)}} \frac{\partial}{\partial x_j} \Psi(\mathbf{x}^{(n)}) \right]_+, \quad d_j^{(n)} = \sum_{i=1}^{n_d} \frac{a_{ij}^2}{\pi_{ij}} c_i^{(n)}$$

Natural choice is $\pi_{ij} = |a_{ij}|/|a|_i$, $|a|_i = \sum_{j=1}^{n_p} |a_{ij}|$
3.39

De Pierro's "additive trick" in [215].

For the natural choice $\pi_{ij} = |a_{ij}|/|a|_i$, we have

$$d_j^{(n)} = \sum_{i=1}^{n_d} |a_{ij}| |a|_i c_i^{(n)}$$

Example: Poisson ML Transmission Problem

Transmission negative log-likelihood (for i th ray):

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

Optimal (smallest) parabola surrogate curvature (Erdoğan, T-MI, Sep. 1999):

$$c_i^{(n)} = c(\ell_i^{(n)}, h_i), \quad c(l, h) = \begin{cases} \left[2 \frac{h(0) - h(l) + \dot{h}(l)l}{l^2} \right]_+, & l > 0 \\ [\ddot{h}(l)]_+, & l = 0. \end{cases}$$

Separable Paraboloidal Surrogate Algorithm:

Precompute $|a|_i = \sum_{j=1}^{n_p} a_{ij}$, $i = 1, \dots, n_d$

$$\ell_i^{(n)} = [\mathbf{Ax}^{(n)}]_i, \quad (\text{forward projection})$$

$$\bar{y}_i^{(n)} = b_i e^{-\ell_i^{(n)}} + r_i \quad (\text{predicted means})$$

$$\bar{h}_i^{(n)} = 1 - y_i / \bar{y}_i^{(n)} \quad (\text{slopes})$$

$$c_i^{(n)} = c(\ell_i^{(n)}, h_i) \quad (\text{curvatures})$$

$$x_j^{(n+1)} = \left[x_j^{(n)} - \frac{1}{d_j^{(n)}} \frac{\partial}{\partial x_j} \Psi(\mathbf{x}^{(n)}) \right]_+ = \left[x_j^{(n)} - \frac{\sum_{i=1}^{n_d} a_{ij} \bar{h}_i^{(n)}}{\sum_{i=1}^{n_d} |a_{ij}| |a|_i c_i^{(n)}} \right]_+, \quad j = 1, \dots, n_p$$

Monotonically decreases *cost function* each iteration.
3.40

No logarithm!

Note that this algorithm never takes the logarithm of the transmission data, since it is based directly on a statistical model for the raw measurements. This is a significant part of the reason why it works well for low-count measurements.

Optimal parabola surrogate curvature for transmission problem [211]. Emission problem [216].

A Matlab m-file for this algorithm is available from
<http://www.eecs.umich.edu/~fessler/code>
 as transmission/tml_sps.m

Related m-files also of interest include transmission/tpl_osps.m

The MAP-EM M-step “Problem”

Add a penalty function to our surrogate for the negative log-likelihood:

$$\begin{aligned}\Psi(\mathbf{x}) &= -L(\mathbf{x}) + \beta R(\mathbf{x}) \\ \phi(\mathbf{x}; \mathbf{x}^{(n)}) &= \sum_{j=1}^{n_p} \phi_j(x_j; \mathbf{x}^{(n)}) + \beta R(\mathbf{x})\end{aligned}$$

$$\text{M-step: } \mathbf{x}^{(n+1)} = \arg \min_{\mathbf{x} \geq 0} \phi(\mathbf{x}; \mathbf{x}^{(n)}) = \arg \min_{\mathbf{x} \geq 0} \sum_{j=1}^{n_p} \phi_j(x_j; \mathbf{x}^{(n)}) + \beta R(\mathbf{x}) = ?$$

For nonseparable penalty functions, the M-step is coupled \therefore difficult.

Suboptimal solutions

- Generalized EM (GEM) algorithm (coordinate descent on ϕ)
Monotonic, but inherits slow convergence of EM.
- One-step late (OSL) algorithm (use outdated gradients) (Green, T-MI, 1990)

$$\frac{\partial}{\partial x_j} \phi(\mathbf{x}; \mathbf{x}^{(n)}) = \frac{\partial}{\partial x_j} \phi_j(x_j; \mathbf{x}^{(n)}) + \beta \frac{\partial}{\partial x_j} R(\mathbf{x}) \stackrel{?}{\approx} \frac{\partial}{\partial x_j} \phi_j(x_j; \mathbf{x}^{(n)}) + \beta \frac{\partial}{\partial x_j} R(\mathbf{x}^{(n)})$$

Nonmonotonic. Known to diverge, depending on β .

Temptingly simple, but *avoid*!

Contemporary solution

- Use separable surrogate for penalty function too (De Pierro, T-MI, Dec. 1995)
Ensures monotonicity. Obviates all reasons for using OSL!

OSL [204, 205]

GEM [217–219]

De Pierro's separable penalty derived in [215].

3.41

p3x

De Pierro's MAP-EM Algorithm

Apply separable paraboloidal surrogates to penalty function:

$$R(\mathbf{x}) \leq R_{\text{SPS}}(\mathbf{x}; \mathbf{x}^{(n)}) = \sum_{j=1}^{n_p} R_j(x_j; \mathbf{x}^{(n)})$$

$$\text{Overall separable surrogate: } \phi(\mathbf{x}; \mathbf{x}^{(n)}) = \sum_{j=1}^{n_p} \phi_j(x_j; \mathbf{x}^{(n)}) + \beta \sum_{j=1}^{n_p} R_j(x_j; \mathbf{x}^{(n)})$$

The M-step becomes fully parallelizable:

$$x_j^{(n+1)} = \arg \min_{x_j \geq 0} \phi_j(x_j; \mathbf{x}^{(n)}) - \beta R_j(x_j; \mathbf{x}^{(n)}), \quad j = 1, \dots, n_p.$$

Consider quadratic penalty $R(\mathbf{x}) = \sum_k \psi([C\mathbf{x}]_k)$, where $\psi(t) = t^2/2$.

If $\gamma_{kj} \geq 0$ and $\sum_{j=1}^{n_p} \gamma_{kj} = 1$ then

$$[C\mathbf{x}]_k = \sum_{j=1}^{n_p} \gamma_{kj} \left[\frac{c_{kj}}{\gamma_{kj}} (x_j - x_j^{(n)}) + [C\mathbf{x}^{(n)}]_k \right].$$

Since ψ is convex:

$$\begin{aligned}\psi([C\mathbf{x}]_k) &= \psi \left(\sum_{j=1}^{n_p} \gamma_{kj} \left[\frac{c_{kj}}{\gamma_{kj}} (x_j - x_j^{(n)}) + [C\mathbf{x}^{(n)}]_k \right] \right) \\ &\leq \sum_{j=1}^{n_p} \gamma_{kj} \psi \left(\frac{c_{kj}}{\gamma_{kj}} (x_j - x_j^{(n)}) + [C\mathbf{x}^{(n)}]_k \right)\end{aligned}$$

Often we just choose

$$\gamma_{kj} = \begin{cases} \frac{1}{\text{number of nonzero } c_{kj}\text{'s in } k\text{th row of } C}, & c_{kj} \neq 0 \\ 0, & \text{otherwise,} \end{cases}$$

which satisfies the two conditions $\gamma_{kj} \geq 0$ and $\sum_{j=1}^{n_p} \gamma_{kj} = 1$, e.g.

$$C = \begin{bmatrix} -1 & 1 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & -1 & 1 \\ -1 & 0 & 0 & 1 & 0 \\ 0 & -1 & 0 & 0 & 1 \end{bmatrix}, \quad \{\gamma_{kj}\} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & 0 & 0 & \frac{1}{2} \end{bmatrix}.$$

Alternatively we use the choice

$$\gamma_{kj} = \frac{|c_{kj}|}{\sum_{j=1}^{n_p} |c_{kj}|},$$

which happens to yield the same result when the elements of C are just ± 1 as in the above example. For non-unity c_{kj} 's, the latter ratio seems to be preferable in terms of convergence rate.

3.42

p3x

De Piero's Algorithm Continued

So $R(\mathbf{x}) \leq R(\mathbf{x}; \mathbf{x}^{(n)}) \triangleq \sum_{j=1}^{n_p} R_j(x_j; \mathbf{x}^{(n)})$ where

$$R_j(x_j; \mathbf{x}^{(n)}) \triangleq \sum_k \gamma_{kj} \Psi \left(\frac{c_{kj}}{\gamma_{kj}} (x_j - x_j^{(n)}) + [\mathbf{C}\mathbf{x}^{(n)}]_k \right)$$

M-step: Minimizing $\phi_j(x_j; \mathbf{x}^{(n)}) + \beta R_j(x_j; \mathbf{x}^{(n)})$ yields the iteration:

$$x_j^{(n+1)} = \frac{x_j^{(n)} \sum_{i=1}^{n_d} a_{ij} y_i / \bar{y}_i^{(n)}}{B_j + \sqrt{B_j^2 + \left(x_j^{(n)} \sum_{i=1}^{n_d} a_{ij} y_i / \bar{y}_i^{(n)} \right) \left(\beta \sum_k c_{kj}^2 / \gamma_{kj} \right)}}$$

$$\text{where } B_j \triangleq \frac{1}{2} \left[\sum_{i=1}^{n_d} a_{ij} + \beta \sum_k \left(c_{kj} [\mathbf{C}\mathbf{x}^{(n)}]_k - \frac{c_{kj}^2}{\gamma_{kj}} x_j^{(n)} \right) \right], \quad j = 1, \dots, n_p$$

and $\bar{y}_i^{(n)} = [\mathbf{A}\mathbf{x}^{(n)}]_i + r_i$.

Advantages: Intrinsically monotone, nonnegativity, fully parallelizable.
Requires only a couple % more computation per iteration than ML-EM

Disadvantages: Slow convergence (like EM) due to separable surrogate

3.43

As a concrete example, consider $R(\mathbf{x}) = \sum_{j=1}^{n_p} \frac{1}{2} \sum_{k \in N_j} \frac{1}{2} (x_j - x_k)^2$ with N_j corresponding to the $|N_j|$ nearest neighbors to the j th pixel. For this penalty with the choice $\gamma_{kj} = |c_{kj}|/c_k$ where $c_k = \sum_{j=1}^{n_p} |c_{kj}| = |N_j|$, the separable surrogate is [215]:

$$R_j(x_j; \mathbf{x}^{(n)}) = \sum_{k \in N_j} \frac{1}{2} \frac{1}{|N_j|} \left(|N_j| (x_j - x_j^{(n)}) + x_j^{(n)} - x_k^{(n)} \right)^2.$$

Matlab m-file available from [http://www.eecs.umich.edu/~fessler/code/as emission/eql_emdp.m](http://www.eecs.umich.edu/~fessler/code/as%20emission/eql_emdp.m)

Caution: use stable quadratic roots [172] (slightly more complicated than above).

One can make an ordered-subsets version of De Piero's MAP-EM easily. Such an approach is preferable to the OSL version of OS-EM mentioned by Hudson and Larkin [15].

One can do multiple M-step subiterations for minimal additional computation with some improvement in convergence rate.

For a tomography problem with a 64×64 image, 64×80 sinogram, and strip-area system matrix, De Piero's MAP-EM algorithm requires 4% more flops per iteration than classic ML-EM.

3.43

p3x

Ordered Subsets Algorithms

aka *block iterative* or *incremental gradient* algorithms

The gradient appears in essentially every algorithm:

$$\frac{\partial}{\partial x_j} \Psi(\mathbf{x}) = \sum_{i=1}^{n_d} a_{ij} \dot{h}_i([\mathbf{A}\mathbf{x}]_i).$$

This is a *backprojection* of a sinogram of the derivatives $\{\dot{h}_i([\mathbf{A}\mathbf{x}]_i)\}$.

Intuition: with half the angular sampling, this backprojection would be fairly similar

$$\frac{1}{n_d} \sum_{i=1}^{n_d} a_{ij} \dot{h}_i(\cdot) \approx \frac{1}{|S|} \sum_{i \in S} a_{ij} \dot{h}_i(\cdot),$$

where S is a subset of the rays.

To "OS-ize" an algorithm, replace all backprojections with partial sums.

Recall typical iteration:

$$\mathbf{x}^{(n+1)} = \mathbf{x}^{(n)} - \mathbf{D}(\mathbf{x}^{(n)}) \nabla \Psi(\mathbf{x}^{(n)}).$$

3.44

The dramatic improvements in apparent "convergence rate" of OSEM over classic ML-EM are due largely to the fact that the latter converges so slowly.

Modern, faster converging algorithms may benefit much less from OS modifications.

Richard Larkin (personal communication) has described the development of OSEM as something of a fortuitous programming "accident." In the course of developing software to implement the E-ML-EM algorithm, he first implemented a version that updated the image immediately after the reprojection of each view. Later he implemented the classical E-ML-EM algorithm but found it to give worse images (in the early iterations). (Due of course to its slow convergence.) The "immediate update" version turns out to be OSEM with 1 view per subset.


Several publications hinted at the use of subsets of projection views for acceleration, *e.g.*, [220–223], and D. Polite's 1983 dissertation. But it was the paper by Larking and Hudson that incited widespread use of OSEM [15].

In the general optimization literature, such algorithms are called *incremental gradient* methods [224–228], and they date back to the 1970's [229].

3.44

p3x

Geometric View of Ordered Subsets



Two subset case: $\Psi(x) = f_1(x) + f_2(x)$ (e.g., odd and even projection views).


For $x^{(n)}$ far from x^* , even partial gradients should point roughly towards x^* .

For $x^{(n)}$ near x^* , however, $\nabla \Psi(x) \approx 0$, so $\nabla f_1(x) \approx -\nabla f_2(x) \Rightarrow$ cycles!

Issues. "Subset gradient balance": $\nabla \Psi(x) \approx M \nabla f_k(x)$. Choice of ordering.

3.45


Incremental Gradients (WLS, 2 Subsets)



3.46

Here the initial image $x^{(0)}$ is far from the solution so the incremental gradients, i.e., the gradients computed from just the even or odd angles, agree well with the full gradient.

Subset Gradient Imbalance



3.47

Here the first subset was angles 0-90°, and the second subset was angles 90-180°, roughly speaking. Now the incremental gradients do not agree as well with the full gradient. (Of course the *sum* of the two incremental gradients would still equal the full gradient.) This imbalance is expected to slow “convergence.”

Problems with OS-EM


- Non-monotone
- Does not converge (may cycle)
- Byrne’s RBBI approach only converges for consistent (noiseless) data
- ∴ unpredictable
 - What resolution after n iterations?
Object-dependent, spatially nonuniform
 - What variance after n iterations?
 - ROI variance? (*e.g.*, for Huesman’s WLS kinetics)

3.48

Soares and Glick *et al.* [30] [29] have extended the work of Barrett *et al.* [22] to the OSEM case.

Wang *et al.* have extended it to the penalized case, for the OSL algorithm [28].

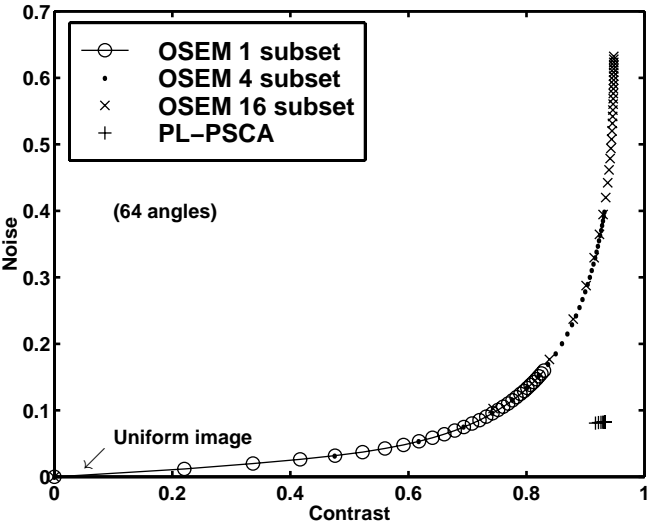
OSEM vs Penalized Likelihood




- 64×62 image
- 66×60 sinogram
- 10^6 counts
- 15% randoms/scatter
- uniform attenuation
- contrast in cold region
- within-region σ opposite side

3.49

Contrast-Noise Results



3.50



3.51

An Open Problem

Still no algorithm with all of the following properties:

- Nonnegativity easy
- Fast converging
- Intrinsically monotone global convergence
- Accepts any type of system matrix
- Parallelizable

Relaxed block-iterative methods

$$\Psi(\mathbf{x}) = \sum_{k=1}^K \Psi_k(\mathbf{x})$$

$$\mathbf{x}^{(n+(k+1)/K)} = \mathbf{x}^{(n+k/K)} - \alpha_n D(\mathbf{x}^{(n+k/K)}) \nabla \Psi_k(\mathbf{x}^{(n+k/K)}), \quad k = 0, \dots, K-1$$

Relaxation of step sizes:

$$\alpha_n \rightarrow 0 \text{ as } n \rightarrow \infty, \quad \sum_n \alpha_n = \infty, \quad \sum_n \alpha_n^2 < \infty$$

- ART
- RAMLA, BSREM (De Pierro, T-MI, 1997, 2001)
- Ahn and Fessler, NSS/MIC 2001, T-MI (to appear)

Proper relaxation can induce convergence, *but* still lacks monotonicity.

Choice of relaxation schedule requires experimentation.

3.52


Until such an algorithm is developed, OSEM will probably remain very popular...

RAMLA [197] (for ML only)

Kudo [230] does not give convergence proof in English...

BSREM [231] convergence proof requires some "a posteriori" assumptions. These have been eliminated in [232].


Relaxed OS-SPS



3.53

[233]

OSTR aka Transmission OS-SPS



Ordered subsets version of separable paraboloidal surrogates for PET transmission problem with nonquadratic convex *regularization*

Matlab m-file <http://www.eecs.umich.edu/~fessler/code/transmission/tpl-osps.m>

3.54

Ordered subsets for PET/SPECT transmission scans [234], and for X-ray CT [34].

Precomputed curvatures for OS-SPS

Separable Paraboloidal Surrogate (SPS) Algorithm:

$$x_j^{(n+1)} = \left[x_j^{(n)} - \frac{\sum_{i=1}^{n_d} a_{ij} \dot{h}_i([\mathbf{A}\mathbf{x}^{(n)}]_i)}{\sum_{i=1}^{n_d} a_{ij} |a_i| c_i^{(n)}} \right]_+, \quad j = 1, \dots, n_p$$

Ordered-subsets abandons monotonicity, so why use optimal curvatures $c_i^{(n)}$?

Precomputed curvature:

$$c_i = \ddot{h}_i(\hat{l}_i), \quad \hat{l}_i = \arg \min_l h_i(l)$$

Precomputed denominator (saves one backprojection each iteration!):

$$d_j = \sum_{i=1}^{n_d} a_{ij} |a_i| c_i, \quad j = 1, \dots, n_p.$$

OS-SPS algorithm with M subsets:

$$x_j^{(n+1)} = \left[x_j^{(n)} - \frac{\sum_{i \in S^{(n)}} a_{ij} \dot{h}_i([\mathbf{A}\mathbf{x}^{(n)}]_i)}{d_j/M} \right]_+, \quad j = 1, \dots, n_p$$

3.55

Precomputed parabola surrogate curvature for transmission problem and ordered subsets [211, 234].

For emission problem, $c_i \approx 1/y_i$.

For transmission problem, $c_i \approx y_i$.

Precomputed curvatures combined with suitable relaxation yields guaranteed convergence for convex problems [232].

Summary of Algorithms

- General-purpose optimization algorithms
- Optimization transfer for image reconstruction algorithms
- Separable surrogates \Rightarrow high curvatures \Rightarrow slow convergence
- Ordered subsets accelerate *initial* convergence
require relaxation for true convergence
- Principles apply to emission and transmission reconstruction
- Still work to be done...

3.56

Part 4. Performance Characteristics

- Spatial resolution properties
- Noise properties
- Detection properties

4.1

Spatial Resolution Properties

Choosing β can be painful, so ...

For true minimization methods:

$$\hat{x} = \arg \min_x \Psi(x)$$

the *local impulse response* is approximately (Fessler and Rogers, T-MI, Sep. 1996):

$$l_j(x) = \lim_{\delta \rightarrow 0} \frac{E[\hat{x}|x + \delta e_j] - E[\hat{x}|x]}{\delta} \approx [-\nabla^2 \Psi]^{-1} \nabla \Psi \frac{\partial}{\partial x_j} \bar{y}(x).$$

Depends only on chosen *cost function* and statistical model.
Independent of optimization algorithm.

- Enables prediction of resolution properties (provided Ψ is minimized)
- Useful for designing regularization penalty functions with desired resolution properties

$$l_j(x) \approx [A'WA + \beta R]^{-1} A'W A x^{\text{true}}.$$


- Helps choose β for desired spatial resolution

4.2

[25, 147, 235]

A commonly cited disadvantage of regularized methods is the need to select the regularization parameter β . One must also select the cutoff frequency for FBP, but at least that value is intuitive and works the same (resolution-wise) for all patients. Not so for stopping rules. The analysis in [25, 147, 235] brings some of the consistency of FBP-like resolution selection to statistical methods.

Modified Penalty Example, PET




- a) filtered backprojection
- b) Penalized unweighted least-squares
- c) PWLS with conventional regularization
- d) PWLS with certainty-based penalty [25]
- e) PWLS with modified penalty [147]

4.3

Figure from [147].

Modified Penalty Example, SPECT - Noiseless



4.4

Figure from [148].

Modified Penalty Example, SPECT - Noisy




Figure from [148].

These were preliminary results, and they initially casted doubt on the claim sometimes made that post-filtered EM (or OSEM) is equivalent to truly regularized image reconstruction.

However, it turns out we had not matched spatial resolution as carefully as needed...

Regularized vs Post-filtered, with Matched PSF

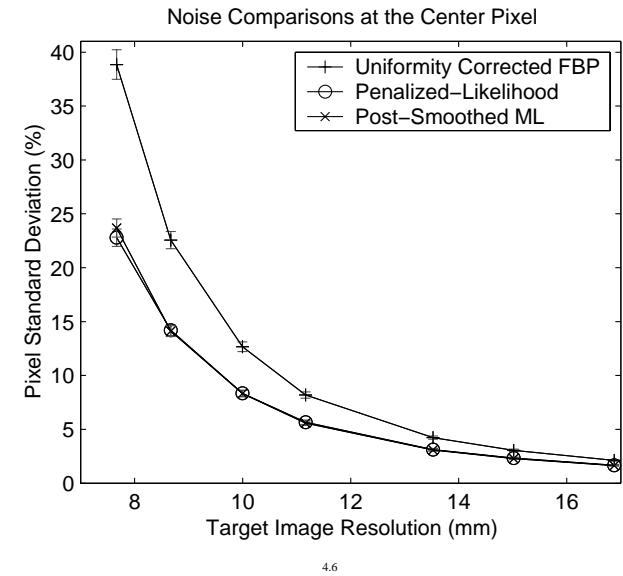


Figure from [236].

Reconstruction Noise Properties

For unconstrained (converged) minimization methods, the estimator is *implicit*:

$$\hat{x} = \hat{x}(y) = \arg \min_x \Psi(x, y).$$

What is $\text{Cov}\{\hat{x}\}$?

New simpler derivation.

Denote the column gradient by $g(x, y) \triangleq \nabla_x \Psi(x, y)$.

Ignoring constraints, the gradient is zero at the minimizer: $g(\hat{x}(y), y) = 0$.

First-order Taylor series expansion:

$$\begin{aligned} g(\hat{x}, y) &\approx g(x^{\text{true}}, y) + \nabla_x g(x^{\text{true}}, y)(\hat{x} - x^{\text{true}}) \\ &= g(x^{\text{true}}, y) + \nabla_x^2 \Psi(x^{\text{true}}, y)(\hat{x} - x^{\text{true}}). \end{aligned}$$

Equating to zero:

$$\hat{x} \approx x^{\text{true}} - [\nabla_x^2 \Psi(x^{\text{true}}, y)]^{-1} \nabla_x \Psi(x^{\text{true}}, y).$$

If the Hessian $\nabla^2 \Psi$ is weakly dependent on y , then

$$\boxed{\text{Cov}\{\hat{x}\} \approx [\nabla_x^2 \Psi(x^{\text{true}}, \bar{y})]^{-1} \text{Cov}\{\nabla_x \Psi(x^{\text{true}}, y)\} [\nabla_x^2 \Psi(x^{\text{true}}, \bar{y})]^{-1}.$$

If we further linearize w.r.t. the data: $g(x, y) \approx g(x, \bar{y}) + \nabla_y g(x, \bar{y})(y - \bar{y})$, then

$$\text{Cov}\{\hat{x}\} \approx [\nabla_x^2 \Psi]^{-1} (\nabla_x \nabla_y \Psi) \text{Cov}\{y\} (\nabla_x \nabla_y \Psi)' [\nabla_x^2 \Psi]^{-1}.$$

4.7

The latter approximation was derived in [24].

Covariance Continued

Covariance approximation:

$$\text{Cov}\{\hat{x}\} \approx [\nabla_x^2 \Psi(x^{\text{true}}, \bar{y})]^{-1} \text{Cov}\{\nabla_x \Psi(x^{\text{true}}, y)\} [\nabla_x^2 \Psi(x^{\text{true}}, \bar{y})]^{-1}$$

Depends only on chosen cost function and statistical model.

Independent of optimization algorithm.

- Enables prediction of noise properties
- Can make variance images
- Useful for computing ROI variance (e.g., for weighted kinetic fitting)
- Good variance prediction for quadratic regularization in nonzero regions
- Inaccurate for nonquadratic penalties, or in nearly-zero regions

4.8

Qi has developed an approximation that may help with the nonnegativity constraint [237].

Qi and Huesman's Detection Analysis

SNR of MAP reconstruction > SNR of FBP reconstruction (T-MI, Aug. 2001)

quadratic regularization
SKE/BKE task
prewhitened observer
non-prewhitened observer

Open issues

Choice of regularizer to optimize detectability?
(2002 MIC poster by Fessler & Yendiki.)

4.9

[31].


Part 5. Miscellaneous Topics

(Pet peeves and more-or-less recent favorites)

- Short transmission scans
- 3D PET options
- OSEM of transmission data (ugh!)
- Precorrected PET data
- Transmission scan problems
- List-mode EM
- List of other topics I wish I had time to cover...

5.1

PET Attenuation Correction (J. Nuyts)



5.2

Iterative reconstruction for 3D PET

- Fully 3D iterative reconstruction
- Rebinning / 2.5D iterative reconstruction
- Rebinning / 2D iterative reconstruction
 - PWLS
 - OSEM with attenuation weighting
- 3D FBP
- Rebinning / FBP

5.3

[105, 238, 239]

OSEM of Transmission Data?

Bai and Kinahan *et al.* “Post-injection single photon transmission tomography with ordered-subset algorithms for whole-body PET imaging”

- 3D penalty better than 2D penalty
- OSTR with 3D penalty better than FBP and OSEM
- standard deviation from a single realization to estimate noise can be misleading

Using OSEM for transmission data requires taking logarithm, whereas OSTR does not.

5.4

[240]

Precorrected PET data

C. Michel examined shifted-Poisson model, “weighted OSEM” of various flavors.

concluded attenuation weighting matters especially

5.5

[241, 242]

Transmission Scan Challenges

- Overlapping-beam transmission scans
- Polyenergetic X-ray CT scans
- Sourceless attenuation correction

All can be tackled with optimization transfer methods.

List-mode EM

$$\begin{aligned} x_j^{(n+1)} &= x_j^{(n)} \left[\sum_{i=1}^{n_d} a_{ij} \frac{y_i}{\bar{y}_i^{(n)}} \right] / \left(\sum_{i=1}^{n_d} a_{ij} \right) \\ &= \frac{x_j^{(n)}}{\sum_{i=1}^{n_d} a_{ij}} \sum_{i: y_i \neq 0} a_{ij} \frac{y_i}{\bar{y}_i^{(n)}} \end{aligned}$$

- Useful when $\sum_{i=1}^{n_d} y_i \leq \sum_{i=1}^{n_d} 1$
- Attenuation and scatter non-trivial
- Computing a_{ij} on-the-fly
- Computing sensitivity $\sum_{i=1}^{n_d} a_{ij}$ still painful
- List-mode ordered-subsets is naturally balanced

5.6

Overlapping beams [243].

Polyenergetic X-ray CT [34, 244].

5.7

Misc

- 4D regularization (reconstruction of dynamic image sequences)
- “Sourceless” attenuation-map estimation
- Post-injection transmission/emission reconstruction
- μ -value priors for transmission reconstruction
- Local errors in $\hat{\mu}$ propagate into emission image (PET and SPECT)

5.8

Dynamic reconstruction

- nonlinear models [47, 48, 50, 51, 54, 65, 70, 245–253]
- linear models [49, 52, 53, 55, 57, 254–258]
- KL-based approaches [56, 259–264]
- Motion/gating [265]

Sourceless attenuation [266–272]

Summary

- Predictability of resolution / noise and controlling spatial resolution argues for regularized *cost function*
- todo: Still work to be done...



5.9

References

- [1] S. Webb. *From the watching of shadows: the origins of radiological tomography*. A. Hilger, Bristol, 1990.
- [2] G. Hounsfield. A method of apparatus for examination of a body by radiation such as x-ray or gamma radiation, 1972. US Patent 1283915. British patent 1283915. London, issued to EMI Ltd. Filed Aug. 1968. See [1, Ch. 5].
- [3] 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.*, 29:477–81, 1970.
- [4] R. Gordon and G. T. Herman. Reconstruction of pictures from their projections. *Comm. ACM*, 14:759–68, 1971.
- [5] G. T. Herman, A. Lent, and S. W. Rowland. ART: mathematics and applications (a report on the mathematical foundations and on the applicability to real data of the algebraic reconstruction techniques). *J. Theor. Biol.*, 42:1–32, 1973.
- [6] R. Gordon. A tutorial on ART (algebraic reconstruction techniques). *IEEE Tr. Nuc. Sci.*, 21:78–93, 1974.
- [7] R. Richardson. Bayesian-based iterative method of image restoration. *J. Opt. Soc. Am.*, 62(1):55–9, January 1972.
- [8] L. Lucy. An iterative technique for the rectification of observed distributions. *The Astronomical J.*, 79(6):745–54, June 1974.
- [9] A. J. Rockmore and A. Macovski. A maximum likelihood approach to emission image reconstruction from projections. *IEEE Tr. Nuc. Sci.*, 23:1428–32, 1976.
- [10] A. J. Rockmore and A. Macovski. A maximum likelihood approach to transmission image reconstruction from projections. *IEEE Tr. Nuc. Sci.*, 24(3):1023–35, June 1977.
- [11] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Stat. Soc. Ser. B*, 39(1):1–38, 1977.
- [12] L. A. Shepp and Y. Vardi. Maximum likelihood reconstruction for emission tomography. *IEEE Tr. Med. Im.*, 1(2):113–22, October 1982.
- [13] K. Lange and R. Carson. EM reconstruction algorithms for emission and transmission tomography. *J. Comp. Assisted Tomo.*, 8(2):306–16, April 1984.
- [14] S. Geman and D. E. McClure. Bayesian image analysis: an application to single photon emission tomography. In *Proc. of Stat. Comp. Sect. of Amer. Stat. Assoc.*, pages 12–8, 1985.
- [15] H. M. Hudson and R. S. Larkin. Accelerated image reconstruction using ordered subsets of projection data. *IEEE Tr. Med. Im.*, 13(4):601–9, December 1994.
- [16] M. Golin. Three-dimensional density reconstruction from a series of two-dimensional projections. *Nucl. Instr. Meth.*, 101(15):509–18, June 1972.
- [17] T. F. Budinger and G. T. Gulberg. Three dimensional reconstruction in nuclear medicine emission imaging. *IEEE Tr. Nuc. Sci.*, 21(3):2–20, 1974.
- [18] R. H. Huesman, G. T. Gulberg, W. L. Greenberg, and T. F. Budinger. *RECLBL library users manual*. Lawrence Berkeley Laboratory, Berkeley, CA, 1977.
- B.1
- [19] R. H. Huesman. A new fast algorithm for the evaluation of regions of interest and statistical uncertainty in computed tomography. *Phys. Med. Biol.*, 29(5):543–52, 1984.
- [20] D. W. Wilson and B. M. W. Tsui. Noise properties of filtered-backprojection and ML-EM reconstructed emission tomographic images. *IEEE Tr. Nuc. Sci.*, 40(4):1189–1203, August 1993.
- [21] D. W. Wilson and B. M. W. Tsui. Spatial resolution properties of FB and ML-EM reconstruction methods. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pages 1189–1193, 1993.
- [22] H. H. Barrett, D. W. Wilson, and B. M. W. Tsui. Noise properties of the EM algorithm: I. Theory. *Phys. Med. Biol.*, 39(5):833–46, May 1994.
- [23] D. W. Wilson, B. M. W. Tsui, and H. H. Barrett. Noise properties of the EM algorithm: II. Monte Carlo simulations. *Phys. Med. Biol.*, 39(5):847–72, May 1994.
- [24] J. A. Fessler. Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): Applications to tomography. *IEEE Tr. Im. Proc.*, 5(3):493–506, March 1996.
- [25] J. A. Fessler and W. L. Rogers. Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs. *IEEE Tr. Im. Proc.*, 5(9):1346–58, September 1996.
- [26] W. Wang and G. Gindi. Noise analysis of regularized EM SPECT reconstruction. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 3, pages 1933–7, 1996.
- [27] C. K. Abbey, E. Clarkson, H. H. Barrett, S. P. Mueller, and F. J. Rybicki. Approximate distributions for maximum likelihood and maximum a posteriori estimates under a Gaussian noise model. In J. Duncan and G. Gindi, editors, *Information Processing in Medical Im.*, pages 167–75. Springer-Verlag, Berlin, 1997.
- [28] W. Wang and G. Gindi. Noise analysis of MAP-EM algorithms for emission tomography. *Phys. Med. Biol.*, 42(11):2215–32, November 1997.
- [29] S. J. Glick and E. J. Soares. Noise characteristics of SPECT iterative reconstruction with a mis-matched projector-backprojector pair. *IEEE Tr. Nuc. Sci.*, 45(4):2183–8, August 1998.
- [30] E. J. Soares, C. L. Byrne, T. S. Pan, S. J. Glick, and M. A. King. Modeling the population covariance matrices of block-iterative expectation-maximization reconstructed images. In *Proc. SPIE 3034, Med. Im. 1997: Im. Proc.*, volume 1, pages 415–25, 1997.
- [31] J. Qi and R. H. Huesman. Theoretical study of lesion detectability of MAP reconstruction using computer observers. *IEEE Tr. Med. Im.*, 20(8):815–22, August 2001.
- [32] J. A. Fessler, I. Elbakri, P. Sukovic, and N. H. Clinthorne. Maximum-likelihood dual-energy tomographic image reconstruction. In *Proc. SPIE 4684, Medical Imaging 2002: Image Proc.*, pages 38–49, 2002.
- [33] Bruno De Man, J. Nuyts, P. Dupont, G. Marchal, and P. Suetens. An iterative maximum-likelihood polychromatic algorithm for CT. *IEEE Tr. Med. Im.*, 20(10):989–1008, October 2001.
- [34] I. A. Elbakri and J. A. Fessler. Statistical image reconstruction for polynuclear X-ray computed tomography. *IEEE Tr. Med. Im.*, 21(2):89–99, February 2002.
- B.2
- [35] I. A. Elbakri and J. A. Fessler. Segmentation-free statistical image reconstruction for polynuclear X-ray computed tomography. In *Proc. IEEE Intl. Symp. Biomedical Imaging*, pages 828–31, 2002.
- [36] P. E. Knaflitz, J. A. Fessler, and J. S. Karp. Statistical image reconstruction in PET with compensation for missing data. *IEEE Tr. Nuc. Sci.*, 44(4):1552–7, August 1997.
- [37] J. A. Fessler and B. P. Sutton. Tomographic image reconstruction using the nonuniform FFT. In *SIAM Conf. Imaging Sci., Abstract Book*, page 21, 2002. Invited presentation.
- [38] B. P. Sutton, J. A. Fessler, and D. Noll. A min-max approach to the nonuniform N-D FFT for rapid iterative reconstruction of MR images. In *Proc. Intl. Soc. Mag. Res. Med.*, page 763, 2001.
- [39] B. P. Sutton, D. C. Noll, and J. A. Fessler. Fast, iterative, field-corrected image reconstruction for MRI. In *Proc. IEEE Intl. Symp. Biomedical Imaging*, pages 489–92, 2002.
- [40] J. A. Fessler. Spatial resolution and noise tradeoffs in pinhole imaging system design: A density estimation approach. *Optics Express*, 2(6):237–53, March 1998.
- [41] B. W. Silverman. *Density estimation for statistics and data analysis*. Chapman and Hall, New York, 1986.
- [42] J. A. Sorenson and M. E. Phelps. *Physics in nuclear medicine*. Saunders, Philadelphia, 2 edition, 1987.
- [43] R. D. Evans. *The atomic nucleus*. McGraw-Hill, New York, 1985.
- [44] U. Engeland, T. Strker, and H. Luijck. Count-rate statistics of the gamma camera. *Phys. Med. Biol.*, 43(10):2339–47, October 1998.
- [45] D. F. Yu and J. A. Fessler. Mean and variance of singles photon counting with deadline. *Phys. Med. Biol.*, 45(7):2043–56, July 2000.
- [46] D. F. Yu and J. A. Fessler. Mean and variance of coincidence photon counting with deadline. *Nucl. Instr. Meth. Phys. Res. A*, 488(1-2):362–74, August 2002.
- [47] M. A. Limber, M. N. Limber, A. Celler, J. S. Barney, and J. M. Borwein. Direct reconstruction of functional parameters for dynamic SPECT. *IEEE Tr. Nuc. Sci.*, 42(4):1249–56, August 1995.
- [48] G. L. Zeng, G. T. Gulberg, and R. H. Huesman. Using linear time-invariant system theory to estimate kinetic parameters directly from projection measurements. *IEEE Tr. Nuc. Sci.*, 42(6-2):2339–46, December 1995.
- [49] E. Heibler, D. Oldenburg, T. Farncombe, and A. Celler. Direct estimation of dynamic parameters in SPECT tomography. *IEEE Tr. Nuc. Sci.*, 44(6-2):2425–30, December 1997.
- [50] R. H. Huesman, B. W. Reutter, G. L. Zeng, and G. T. Gulberg. Kinetic parameter estimation from SPECT cone-beam projection measurements. *Phys. Med. Biol.*, 43(4):373–82, April 1998.
- [51] B. W. Reutter, G. T. Gulberg, and R. H. Huesman. Kinetic parameter estimation from attenuated SPECT projection measurements. *IEEE Tr. Nuc. Sci.*, 45(6):3007–13, December 1998.
- [52] H. H. Bauschke, D. Noll, A. Celler, and J. M. Borwein. An EM algorithm for dynamic SPECT. *IEEE Tr. Med. Im.*, 18(3):252–61, March 1999.
- [53] T. Farncombe, A. Celler, D. Noll, J. Maeght, and R. Harrop. Dynamic SPECT imaging using a single camera rotation (dSPECT). *IEEE Tr. Nuc. Sci.*, 46(2):1055–61, August 1999.
- [54] J. S. Maltz. Direct recovery of regional tracer kinetics from temporally inconsistent dynamic ECT projections using dimension-reduced line-activity basis. *Phys. Med. Biol.*, 45(11):3413–29, November 2000.
- [55] B. W. Reutter, G. T. Gulberg, and R. H. Huesman. Direct least squares estimation of spatiotemporal distributions from dynamic SPECT projections using a spatial segmentation and temporal B-splines. *IEEE Tr. Med. Im.*, 19(5):434–50, May 2000.
- [56] J. S. Maltz. Optimal time-activity basis selection for exponential spectral analysis: application to the solution of large dynamic emission tomographic reconstruction problems. *IEEE Tr. Nuc. Sci.*, 48(4-2):1452–64, August 2001.
- [57] T. E. Nichols, J. Qi, E. Asma, and R. M. Leahy. Spatiotemporal reconstruction of list mode PET data. *IEEE Tr. Med. Im.*, 21(4):396–404, April 2002.
- [58] B. W. Silverman. Some aspects of the spline smoothing approach to non-parametric regression curve fitting. *J. Royal Stat. Soc. Ser. B*, 47(1):1–52, 1985.
- [59] R. Van de Walle, H. H. Barrett, K. J. Myers, M. I. Altbach, B. Desplanques, A. F. Gmitro, J. Cornelis, and I. Lemahieu. Reconstruction of MR images from data acquired on a general non-regular grid by pseudoinverse calculation. *IEEE Tr. Med. Im.*, 19(12):1180–7, December 2000.
- [60] M. Bertero, C. De Mol, and E. R. Pike. Linear inverse problems with discrete data, I. General formulation and singular system analysis. *Inverse Probl.*, 1(4):301–30, November 1985.
- [61] E. J. Mazur and R. Gordon. Interpolative algebraic reconstruction techniques without beam partitioning for computed tomography. *Med. Biol. Eng. Comput.*, 33(1):182–6, January 1995.
- [62] Y. Censor. Finite series expansion reconstruction methods. *Proc. IEEE*, 71(3):409–419, March 1983.
- [63] D. L. Snyder. Utilizing side information in emission tomography. *IEEE Tr. Nuc. Sci.*, 31(1):533–7, February 1984.
- [64] R. E. Carson, M. V. Green, and S. M. Larson. A maximum likelihood method for calculation of tomographic region-of-interest (ROI) values. *J. Nuc. Med.*, 26:P20, 1985.
- [65] R. E. Carson and K. Lange. The EM parametric image reconstruction algorithm. *J. Am. Stat. Ass.*, 80(389):20–2, March 1985.
- [66] R. E. Carson. A maximum likelihood method for region-of-interest evaluation in emission tomography. *J. Comp. Assisted Tomo.*, 10(4):654–63, July 1986.
- [67] A. R. Formicola. Least squares algorithm for region-of-interest evaluation in emission tomography. *IEEE Tr. Med. Im.*, 12:90–100, 1993.
- [68] D. J. Rossi and A. S. Wilsky. Reconstruction from projections based on detection and estimation of objects—Parts I & II: Performance analysis and robustness analysis. *IEEE Tr. Acoust. Sp. Sig. Proc.*, 32(4):886–906, August 1984.
- [69] S. P. Muller, M. F. Kijowski, S. C. Moore, and B. L. Holman. Maximum-likelihood estimation: a mathematical model for quantitation in nuclear medicine. *J. Nuc. Med.*, 31(10):1685–701, October 1990.
- B.4

- [70] P. C. Ohia, W. L. Rogers, N. H. Clinthorne, J. A. Fessler, and A. O. Hero. Model-based estimation for dynamic cardiac studies using ECT. *IEEE Tr. Med. Im.*, 13(2):217–26, June 1994.
- [71] Z. P. Liang, F. E. Boada, R. T. Constable, E. M. Haacke, P. C. Lauterbur, and M. R. Smith. Constrained reconstruction methods in MR imaging. *Reviews of Magnetic Resonance in Medicine*, 4:87–185, 1992.
- [72] G. S. Cunningham and A. Leinhardt. 4D reconstructions from low-count SPECT data using deformable models with smooth interior intensity variations. In *Proc. SPIE 3979: Medical Imaging 2000: Image Proc.*, 2000.
- [73] J. Qi, R. M. Leahy, E. U. Mumcuoglu, S. R. Cherry, A. Chatziananou, and T. H. Farguhar. High resolution 3D Bayesian image reconstruction for microPET. In *Conf. Record, 1997 Int'l Meeting on Fully 3D Image Recon. in Rad. and Nuc. Med.*, 1997.
- [74] T. Hebert, R. Leahy, and M. Singh. Fast MLE for SPECT using and intermediate polar representation and a stopping criterion. *IEEE Tr. Nuc. Sci.*, 35(1):615–9, 1988.
- [75] T. J. Hebert and R. Leahy. Fast methods for including attenuation in the EM algorithm. *IEEE Tr. Nuc. Sci.*, 37(2):754–758, April 1990.
- [76] C. A. Johnson, Y. Yen, R. E. Carson, R. L. Martino, and M. E. Daube-Witherspoon. A system for the 3D reconstruction of retacted-septa PET data using the EM algorithm. *IEEE Tr. Nuc. Sci.*, 42(4):1223–7, August 1995.
- [77] J. M. Ollinger and A. Goggin. Maximum likelihood reconstruction in fully 3D PET via the SAGE algorithm. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 3, pages 1594–8, 1996.
- [78] E. V. R. Di Bella, A. B. Barclay, R. L. Elser, and R. W. Schaler. Comparison of rotation-based methods for iterative reconstruction algorithms. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pages 1146–50, 1995.
- [79] M. I. Miller and B. Roysam. Bayesian image reconstruction for emission tomography incorporating Good's roughness prior on massively parallel processors. *Proc. Natl. Acad. Sci.*, 88:3223–3227, April 1991.
- [80] A. W. McCarthy and M. I. Miller. Maximum likelihood SPECT in clinical computation times using mesh-connected parallel processors. *IEEE Tr. Med. Im.*, 10(3):428–436, September 1991.
- [81] T. R. Miller and J. W. Wallis. Fast maximum-likelihood reconstruction. *J. Nuc. Med.*, 33(9):1710–11, September 1992.
- [82] B. E. Openheim. More accurate algorithms for iterative 3-dimensional reconstruction. *IEEE Tr. Nuc. Sci.*, 21:72–77, June 1974.
- [83] T. M. Peters. Algorithm for fast back- and reprojection in computed tomography. *IEEE Tr. Nuc. Sci.*, 28(4):3641–3647, August 1981.
- [84] P. M. Joseph. An improved algorithm for projecting rays through pixel images. *IEEE Tr. Med. Im.*, 1(3):192–6, November 1982.
- [85] R. L. Siddon. Fast calculation of the exact radiological path for a three-dimensional CT array. *Med. Phys.*, 12(2):252–255, March 1985.
- [86] R. Schiringer, S. Cool, and M. King. Area weighted convolutional interpolation for data reprojection in single photon emission computed tomography. *Med. Phys.*, 13(3):350–355, May 1986.
- [87] S. C. B. Lo. Strip and line pixel integrals with a square pixel matrix: A unified theory for computational CT projections. *IEEE Tr. Med. Im.*, 7(4):355–363, December 1988.
- B.5
- [88] Z. H. Cho, C. M. Chen, and S. Y. Lee. Incremental algorithm—a new fast backprojection scheme for parallel geometries. *IEEE Tr. Med. Im.*, 9(2):207–17, June 1990.
- [89] K. Zienons, H. Herzog, P. Boesli, and L. E. Feinendegen. Iterative image reconstruction with weighted pixel contributions to projection elements. *Eur. J. Nuc. Med.*, 19:387, 1992.
- [90] Y. J. He, A. Cai, and J. A. Sun. Incremental backprojection algorithm: modification of the searching flow scheme and utilization of the relationship among projection views. *IEEE Tr. Med. Im.*, 12(3):555–559, September 1993.
- [91] B. Sahiner and A. E. Yagle. A fast algorithm for backprojection with linear interpolation. *IEEE Tr. Im. Proc.*, 2(4):547–549, October 1993.
- [92] D. C. Yu and S. C. Huang. Study of reprojection methods in terms of their resolution loss and sampling errors. *IEEE Tr. Nuc. Sci.*, 40(4):1174–1178, August 1993.
- [93] P. Schmidlin. Improved iterative image reconstruction using variable projection binning and abbreviated convolution. *Eur. J. Nuc. Med.*, 21(9):930–6, September 1994.
- [94] W. Zhuang, S. S. Gopal, and T. J. Hebert. Numerical evaluation of methods for computing tomographic projections. *IEEE Tr. Nuc. Sci.*, 41(4):1660–1665, August 1994.
- [95] E. L. Johnson, H. Wang, J. W. McCormick, K. L. Greer, R. E. Coleman, and R. J. Jaszcak. Pixel driven implementation of filtered backprojection for reconstruction of fan beam SPECT data using a position dependent effective projection bin length. *Phys. Med. Biol.*, 41(8):1439–52, August 1996.
- [96] H. W. A. de Jong, E. T. P. Slippen, and F. J. Beekman. Acceleration of Monte Carlo SPECT simulation using convolution-based forced detection. *IEEE Tr. Nuc. Sci.*, 48(1):58–64, February 2001.
- [97] C. W. Stearns and J. A. Fessler. 3D PET Reconstruction with FORE and WLS-OS-EM. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, 2002. To appear.
- [98] M. Yavuz and J. A. Fessler. Objective functions for tomographic reconstruction from randoms-precorrected PET scans. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pages 1067–71, 1996.
- [99] M. Yavuz and J. A. Fessler. New statistical models for randoms-precorrected PET scans. In J. Durcan and G. Gindi, editors, *Information Processing in Medical Im.*, volume 1230 of *Lecture Notes in Computer Science*, pages 190–203. Springer Verlag, Berlin, 1997.
- [100] M. Yavuz and J. A. Fessler. Statistical image reconstruction methods for randoms-precorrected PET scans. *Med. Im. Anal.*, 2(4):369–78, December 1998.
- [101] M. Yavuz and J. A. Fessler. Penalized-likelihood estimators and noise analysis for randoms-precorrected PET transmission scans. *IEEE Tr. Med. Im.*, 18(6):665–74, August 1999.
- [102] D. L. Snyder, A. M. Hammoud, and R. L. White. Image recovery from data acquired with a charge-couple-device camera. *J. Opt. Soc. Am. A*, 10(5):1014–23, May 1993.
- [103] 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. Am. A*, 12(2):272–83, February 1995.
- B.6

- [104] J. A. Fessler. Penalized weighted least-squares image reconstruction for positron emission tomography. *IEEE Tr. Med. Im.*, 13(2):290–300, June 1994.
- [105] C. Comtat, P. E. Kinahan, M. Defrise, C. Michel, and D. W. Townsend. Fast reconstruction of 3D PET data with accurate statistical modeling. *IEEE Tr. Nuc. Sci.*, 45(3):1083–9, June 1998.
- [106] B. R. W. Signal statistics in x-ray computed tomography. In *Proc. SPIE 4882: Medical Imaging 2002: Med. Phys.*, 2002.
- [107] B. R. Whiting, L. J. Montagnino, and D. G. Polite. Modeling X-ray computed tomography sinograms, 2001. submitted to mp.
- [108] A. O. Hero, J. A. Fessler, and M. Umanan. Exploring estimator bias-variance tradeoffs using the uniform CR bound. *IEEE Tr. Sig. Proc.*, 44(8):2028–41, August 1996.
- [109] E. Mumcuoglu, R. Leahy, S. Cherry, and E. Hoffman. Accurate geometric and physical response modeling for statistical image reconstruction in high resolution PET scanners. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 3, pages 1569–73, 1996.
- [110] P. J. Green. Iteratively reweighted least squares for maximum likelihood estimation, and some robust and resistant alternatives. *J. Royal Stat. Soc. Ser. B*, 46(2):145–92, 1984.
- [111] J. M. Anderson, B. A. Mair, M. Rao, and C. H. Wu. A weighted least-squares method for PET. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pages 1292–6, 1995.
- [112] J. M. Anderson, B. A. Mair, M. Rao, and C-H. Wu. Weighted least-squares reconstruction methods for positron emission tomography. *IEEE Tr. Med. Im.*, 16(2):159–65, April 1997.
- [113] P. J. Huber. *Robust statistics*. Wiley, New York, 1981.
- [114] C. Bounian and K. Sauer. A generalized Gaussian image model for edge-preserving MAP estimation. *IEEE Tr. Im. Proc.*, 2(3):296–310, July 1993.
- [115] E. Tanaka. Improved iterative image reconstruction with automatic noise artifact suppression. *IEEE Tr. Med. Im.*, 11(1):21–7, March 1992.
- [116] B. W. Silverman, M. C. Jones, J. D. Wilson, and D. W. Nychka. A smoothed EM approach to indirect estimation problems, with particular reference to stereology and emission tomography. *J. Royal Stat. Soc. Ser. B*, 52(2):271–324, 1990.
- [117] T. R. Miller, J. W. Wallis, C. S. Butler, M. I. Miller, and D. L. Snyder. Improved brain SPECT by maximum-likelihood reconstruction. *J. Nuc. Med. (Abstr. Book)*, 33(5):964, May 1992.
- [118] F. J. Beekman, E. T. P. Slippen, and W. J. Nissen. Selection of task-dependent diffusion filters for the post-processing of SPECT images. *Phys. Med. Biol.*, 43(6):1713–30, June 1998.
- [119] D. S. Lalush and B. M. Tsui. Performance of ordered-subset reconstruction algorithms under conditions of extreme attenuation and truncation in myocardial SPECT. *J. Nuc. Med.*, 41(4):737–44, April 2000.
- [120] E. T. P. Slippen and F. J. Beekman. Comparison of post-filtering and filtering between iterations for SPECT reconstruction. *IEEE Tr. Nuc. Sci.*, 46(6):2233–8, December 1999.
- B.7
- [121] R. H. Huesman. The effects of a finite number of projection angles and finite lateral sampling of projections on the propagation of statistical errors in transverse section reconstruction. *Phys. Med. Biol.*, 22(3):511–521, 1977.
- [122] D. L. Snyder and M. I. Miller. The use of sieves to stabilize images produced with the EM algorithm for emission tomography. *IEEE Tr. Nuc. Sci.*, 32(9):3684–71, October 1985.
- [123] D. L. Snyder, M. I. Miller, L. J. Thomas, and D. G. Polite. Noise and edge artifacts in maximum-likelihood reconstructions for emission tomography. *IEEE Tr. Med. Im.*, 6(3):228–38, September 1987.
- [124] T. R. Miller and J. W. Wallis. Clinically important characteristics of maximum-likelihood reconstruction. *J. Nuc. Med.*, 33(9):1678–84, September 1992.
- [125] A. Tikhonov and V. Arsenin. *Solution of ill-posed problems*. Wiley, New York, 1977.
- [126] I. Csizsár. Why least squares and maximum entropy? An axiomatic approach to inference for linear inverse problems. *Ann. Stat.*, 19(4):2032–86, 1991.
- [127] D. L. Donoho, I. M. Johnstone, A. S. Stern, and J. C. Hoch. Does the maximum entropy method improve sensitivity. *Proc. Natl. Acad. Sci.*, 87(13):5066–8, July 1990.
- [128] D. L. Donoho, I. M. Johnstone, J. C. Hoch, and A. S. Stern. Maximum entropy and the nearly black object. *J. Royal Stat. Soc. Ser. B*, 54(1):41–81, 1992.
- [129] R. T. Constable and R. M. Henkelman. Why MEM does not work in MR image reconstruction. *Magnetic Resonance in Medicine*, 14(1):12–25, April 1990.
- [130] A. H. Delany and Y. Bresler. A fast and accurate Fourier algorithm for iterative parallel-beam tomography. *IEEE Tr. Im. Proc.*, 5(5):740–53, May 1996.
- [131] S. J. Lee, A. Rangarajan, and G. Gindi. A comparative study of the effects of using higher order mechanical priors in SPECT reconstruction. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 4, pages 1696–1700, 1994.
- [132] S. J. Lee, A. Rangarajan, and G. Gindi. Bayesian image reconstruction in SPECT using higher order mechanical models as priors. *IEEE Tr. Med. Im.*, 14(4):668–80, December 1995.
- [133] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and Bayesian restoration of images. *IEEE Tr. Patt. Anal. Mach. Int.*, 6(6):721–41, November 1984.
- [134] B. W. Silverman, C. Jennison, J. Stander, and T. C. Brown. The specification of edge penalties for regular and irregular pixel images. *IEEE Tr. Patt. Anal. Mach. Int.*, 12(10):1017–24, October 1990.
- [135] V. E. Johnson, W. H. Wong, X. Hu, and C. T. Chen. Image restoration using Gibbs priors: Boundary modeling, treatment of blurring, and selection of hyperparameter. *IEEE Tr. Patt. Anal. Mach. Int.*, 13(5):413–25, May 1991.
- [136] V. E. Johnson. A model for segmentation and analysis of noisy images. *J. Am. Stat. Ass.*, 89(425):230–41, March 1984.
- B.8

- [137] S. Aienius and U. Ruotsalainen. Bayesian image reconstruction for emission tomography based on median root prior. *Eur. J. Nuc. Med.*, 24(3):258–65, 1997.
- [138] S. Aienius, U. Ruotsalainen, and J. Asola. Using local median as the location of the prior distribution in iterative emission tomography image reconstruction. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, page 1726, 1997.
- [139] S. Aienius, U. Ruotsalainen, and J. Asola. Using local median as the location of the prior distribution in iterative emission tomography image reconstruction. *IEEE Tr. Nuc. Sci.*, 45(6):3097–104, December 1998.
- [140] I.-T. Hsiao, A. Rangarajan, and G. Gindi. A new convex edge-preserving median prior with applications to tomography, 2002.
- [141] M. Nikolova. Thresholding implied by truncated quadratic regularization. *IEEE Tr. Sig. Proc.*, 48(12):3437–50, December 2000.
- [142] A. Antoniadis and J. Fan. Regularization and wavelet approximations. *J. Am. Stat. Ass.*, 96(465):939–55, September 2001.
- [143] D. F. Yu and J. A. Fessler. Edge-preserving tomographic reconstruction with nonlocal regularization. In *Proc. IEEE Intl. Conf. on Image Processing*, volume 1, pages 29–33, 1998.
- [144] D. F. Yu and J. A. Fessler. Edge-preserving tomographic reconstruction with nonlocal regularization. *IEEE Tr. Med. Im.*, 21(2):159–73, February 2002.
- [145] J. Ye, Y. Bresler, and P. Moulin. A self-referencing level-set method for image reconstruction from sparse fourier samples. In *Proc. IEEE Intl. Conf. on Image Processing*, volume 2, pages 33–6, 2001.
- [146] M. J. Black and A. Rangarajan. On the unification of line processes, outlier rejection, and robust statistics with applications in early vision. *Intl. J. Comp. Vison*, 19(1):57–91, July 1996.
- [147] J. W. Stayman and J. A. Fessler. Regularization for uniform spatial resolution properties in penalized-likelihood image reconstruction. *IEEE Tr. Med. Im.*, 19(6):601–15, June 2000.
- [148] J. W. Stayman and J. A. Fessler. Nonnegative definite quadratic penalty design for penalized-likelihood reconstruction. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, pages M1–1, 2001.
- [149] C. T. Chen, X. Ouyang, W. H. Wong, and X. Hu. Improvement of PET image reconstruction using high-resolution anatomic images. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 3, page 2082, 1991. (Abstract).
- [150] R. Leahy and X. H. Yan. Statistical models and methods for PET image reconstruction. In *Proc. of Stat. Comp. Sect. of Amer. Stat. Assoc.*, pages 1–10, 1991.
- [151] J. A. Fessler, N. H. Clinthorne, and W. L. Rogers. Regularized emission image reconstruction using imperfect side information. *IEEE Tr. Nuc. Sci.*, 39(5):1464–71, October 1992.
- [152] I. G. Zubal, M. Lee, A. Rangarajan, C. R. Harrell, and G. Gindi. Bayesian reconstruction of SPECT images using registered anatomical images as priors. *J. Nuc. Med. (Abs. Book)*, 33(5):983, May 1992.
- [153] G. Gindi, M. Lee, A. Rangarajan, and I. G. Zubal. Bayesian reconstruction of functional images using anatomical information as priors. *IEEE Tr. Med. Im.*, 12(4):570–680, December 1993.
- B.9
- [154] X. Ouyang, W. H. Wong, V. E. Johnson, X. Hu, and C.-T. Chen. Incorporation of correlated structural images in PET image reconstruction. *IEEE Tr. Med. Im.*, 13(4):927, December 1994.
- [155] S.-J. Lee, G. B. Ghah, I. G. Zubal, and A. Rangarajan. Using ground-truth data to design priors in Bayesian SPECT reconstruction. In Y. Bizzi, C. Baillet, and R. D. Padoa, editors, *Information Processing in Medical Im.*, Kluwer, 1, 1995.
- [156] J. E. Bowsher, V. E. Johnson, T. G. Turkington, T. G. Jaszczak, C. E. Floyd, and R. E. Coleman. Bayesian reconstruction and use of anatomical a priori information for emission tomography. *IEEE Tr. Med. Im.*, 15(5):673–86, October 1996.
- [157] S. Sastry and R. E. Carson. Multimodality Bayesian algorithm for image reconstruction in positron emission tomography: a tissue composition model. *IEEE Tr. Med. Im.*, 16(6):750–61, December 1997.
- [158] R. Pirmuthu and A. O. Hero. Side information averaging method for PML emission tomography. In *Proc. IEEE Intl. Conf. on Image Processing*, 1998.
- [159] C. Comtat, P. E. Kinahan, J. A. Fessler, T. Beyer, D. W. Townsend, M. Defrise, and C. Michel. Reconstruction of 3d whole-body PET data using blurred anatomical labels. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 3, pages 1651–5, 1998.
- [160] A. O. Hero, R. Pirmuthu, S. R. Tius, and J. A. Fessler. Minimax emission computed tomography using high resolution anatomical side information and B-spline models. *IEEE Tr. Info. Theory*, 45(3):320–38, April 1999.
- [161] Y. S. Shin and Z. H. Cho. SVD pseudoinversion image reconstruction. *IEEE Tr. Acoust. Sp. Sig. Proc.*, 29(4):904–908, August 1981.
- [162] U. Ralf, D. N. Stroud, and W. R. Hendee. Improvement of lesion detection in scintigraphic images by SVD techniques for resolution recovery. *IEEE Tr. Med. Im.*, 5(1):35–44, March 1986.
- [163] D. A. Fish, J. Grochmalicki, and E. R. Pike. Scanning SVD method for restoration of images with space-variant blur. *J. Opt. Soc. Am. A*, 13(3):464–9, March 1996.
- [164] A. Caponnetto and M. Bertero. Tomography with a finite set of projections: singular value decomposition and resolution. *Inverse Prob.*, 13(5):1191–1205, October 1997.
- [165] A. K. Louis. Incomplete data problems in x-ray computerized tomography. I. Singular value decomposition of the limited angle transform. *Numerische Mathematik*, 48(3):251–62, 1986.
- [166] F. Natterer. Numerical treatment of ill-posed problems. In G. Talenti, editor, *Inverse Prob.*, volume 1225, pages 142–67. Berlin, Springer, 1986. Lecture Notes in Math.
- [167] R. C. Luu and L. D. Brown. Nonexistence of informative unbiased estimators in singular problems. *Ann. Stat.*, 21(1):1–13, March 1993.
- [168] J. Ouy and R. G. Pratt. Are our parameter estimators biased? The significance of finite-different regularization operators. *Inverse Prob.*, 11(2):397–424, April 1995.
- [169] I. M. Johnstone. On singular value decompositions for the Radon Transform and smoothness classes of functions. Technical Report 310, Dept. of Statistics, Stanford, January 1989.

- [170] M. F. Smith, C. E. Floyd, R. J. Jaszczak, and R. E. Coleman. Reconstruction of SPECT images using generalized matrix inverses. *IEEE Tr. Med. Im.*, 11(2), June 1992.
- [171] M. Lavielle. A stochastic algorithm for parametric and non-parametric estimation in the case of incomplete data. *Signal Processing*, 42(1):3–17, 1995.
- [172] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical recipes in C*. Cambridge Univ. Press, New York, 1988.
- [173] K. Sauer and C. Bouman. Bayesian estimation of transmission tomograms using local optimization operations. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 3, pages 2089–93, 1991.
- [174] K. Sauer and C. Bouman. Bayesian estimation of transmission tomograms using segmentation based optimization. *IEEE Tr. Nuc. Sci.*, 39(4):1144–52, August 1992.
- [175] D. P. Bertsekas and S. K. Mitter. A descent numerical method for optimization problems with nondifferentiable cost functionals. *SIAM J. Control*, 1837–92, 1973.
- [176] W. C. Davidson. Variable metric methods for minimization. Technical Report ANL-5990, AEC Research and Development Report, Argonne National Laboratory, USA, 1959.
- [177] H. F. Khalil, R. H. Byrd, and R. B. Schnabel. A theoretical and experimental study of the symmetric rank-one update. *SIAM J. Optim.*, 3(1):1–24, 1993.
- [178] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal. Algorithm 778: LBFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Tr. Math. Software*, 23(4):550–60, December 1997.
- [179] T. G. Kolda, D. P. O’Leary, and L. Nazareth. BFGS with update skipping and varying memory. *SIAM J. Optim.*, 8(4):1060–83, 1998.
- [180] K. Lange. *Numerical analysis for statisticians*. Springer-Verlag, New York, 1999.
- [181] B. T. Polyak. *Introduction to optimization*. Optimization Software Inc, New York, 1987.
- [182] D. P. Bertsekas. *Constrained optimization and Lagrange multiplier methods*. Academic-Press, New York, 1982.
- [183] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comp.*, 16:1190–1208, 1995.
- [184] L. Kaufman. Reduced storage, Quasi-Newton trust region approaches to function optimization. *SIAM J. Optim.*, 10(1):56–69, 1999.
- [185] M. Hanke, J. G. Nagy, and C. Vogel. Quasi-Newton approach to nonnegative image restorations. *Linear Algebra and its Applications*, 316(1):223–36, September 2000.
- [186] J. L. Morales and J. Nocedal. Automatic preconditioning by limited memory Quasi-Newton updating. *SIAM J. Optim.*, 10(4):1079–96, 2000.
- [187] R. R. Meyer. Sufficient conditions for the convergence of monotonic mathematical programming algorithms. *J. Comput. System. Sci.*, 12(1):109–21, 1976.
- B.11
- [188] L. Kaufman. Implementing and accelerating the EM algorithm for positron emission tomography. *IEEE Tr. Med. Im.*, 6(1):37–51, March 1987.
- [189] N. H. Clinthorne, T. S. Pan, P. C. Chiao, W. L. Rogers, and J. A. Stamos. Preconditioning methods for improved convergence rates in iterative reconstructions. *IEEE Tr. Med. Im.*, 12(1):79–83, March 1993.
- [190] J. A. Fessler and S. D. Booth. Conjugate-gradient preconditioning methods for shift-variant PET image reconstruction. *IEEE Tr. Im. Proc.*, 8(5):688–99, May 1999.
- [191] E. U. Mumcuoglu, R. Leahy, S. R. Cherry, and Z. Zhou. Fast gradient-based methods for Bayesian reconstruction of transmission and emission PET images. *IEEE Tr. Med. Im.*, 13(5):337–701, December 1994.
- [192] E. U. Mumcuoglu, 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.
- [193] J. A. Fessler and A. O. Hero. Space-alternating generalized expectation-maximization algorithm. *IEEE Tr. Sig. Proc.*, 42(10):2664–77, October 1994.
- [194] J. A. Fessler and A. O. Hero. Penalized maximum-likelihood image reconstruction using space-alternating generalized EM algorithms. *IEEE Tr. Im. Proc.*, 4(10):1417–29, October 1995.
- [195] J. A. Fessler, E. P. Ficaro, N. H. Clinthorne, and K. Lange. Grouped-coordinate ascent algorithms for penalized-likelihood transmission image reconstruction. *IEEE Tr. Med. Im.*, 16(2):166–75, April 1997.
- [196] J. A. Fessler and A. O. Hero. Space-alternating generalized EM algorithms for penalized maximum-likelihood image reconstruction. Technical Report 286, Comm. and Sig. Proc. Lab., Dept. of EECS, Univ. of Michigan, Ann Arbor, MI, 48109-2122, February 1994.
- [197] J. A. Browne and A. R. D. Pierro. A row-action alternative to the EM algorithm for maximizing likelihoods in emission tomography. *IEEE Tr. Med. Im.*, 15(5):687–99, October 1996.
- [198] C. L. Byrne. Block-iterative methods for image reconstruction from projections. *IEEE Tr. Im. Proc.*, 5(5):792–3, May 1996.
- [199] C. L. Byrne. Convergent block-iterative algorithms for image reconstruction from inconsistent data. *IEEE Tr. Im. Proc.*, 6(9):1296–304, September 1997.
- [200] C. L. Byrne. Accelerating the EMML algorithm and related iterative algorithms by rescaled block-iterative methods. *IEEE Tr. Im. Proc.*, 7(1):100–9, January 1998.
- [201] M. E. Daube-Witherspoon and G. Muehlelehner. An iterative image space reconstruction algorithm suitable for volume ECT. *IEEE Tr. Med. Im.*, 5(2):161–66, June 1986.
- [202] J. M. Ollinger. Iterative reconstruction-reprojection and the expectation-maximization algorithm. *IEEE Tr. Med. Im.*, 9(1):94–8, March 1990.
- [203] A. R. De Pierro. On the relation between the ISRA and the EM algorithm for positron emission tomography. *IEEE Tr. Med. Im.*, 12(2):328–33, June 1993.
- [204] P. J. Green. Bayesian reconstructions from emission tomography data using a modified EM algorithm. *IEEE Tr. Med. Im.*, 9(1):94–93, March 1990.

- [205] P. J. Green. On use of the EM algorithm for penalized likelihood estimation. *J. Royal Stat. Soc. Ser. B* 52(3):443–452, 1990.
- [206] K. Sauer and C. Bouman. A local update strategy for iterative reconstruction from projections. *IEEE Tr. Sig. Proc.*, 41(2):534–48, February 1993.
- [207] K. G. Murty. *Linear complementarity, linear and nonlinear programming*. Helderman Verlag, Berlin, 1988.
- [208] C. A. Bouman, K. Sauer, and S. S. Saquib. Tractable models and efficient algorithms for Bayesian tomography. In *Proc. IEEE Conf. Acoust. Speech Sig. Proc.*, volume 5, pages 2397–10, 1995.
- [209] C. A. Bouman and K. Sauer. A unified approach to statistical tomography using coordinate descent optimization. *IEEE Tr. Im. Proc.*, 5(3):480–92, March 1996.
- [210] J. A. Fessler. Hybrid Poisson/polynomial objective functions for tomographic image reconstruction from transmission scans. *IEEE Tr. Im. Proc.*, 4(10):1439–50, October 1995.
- [211] H. Erdoghan and J. A. Fessler. Monotonic algorithms for transmission tomography. *IEEE Tr. Med. Im.*, 18(9):801–14, September 1999.
- [212] C. A. Johnson, J. Seidel, and A. Seder. Interior point methodology for 3-D PET reconstruction. *IEEE Tr. Med. Im.*, 19(4):271–85, April 2000.
- [213] K. Lange and J. A. Fessler. Globally convergent algorithms for maximum a posteriori transmission tomography. *IEEE Tr. Im. Proc.*, 4(10):1430–8, October 1995.
- [214] J. A. Fessler, N. H. Clinthorne, and W. L. Rogers. On complete data spaces for PET reconstruction algorithms. *IEEE Tr. Nuc. Sci.*, 40(4):1055–61, August 1993.
- [215] A. De Piero. A modified expectation maximization algorithm for penalized likelihood estimation in emission tomography. *IEEE Tr. Med. Im.*, 14(1):132–137, March 1995.
- [216] J. A. Fessler and H. Erdoghan. A paraboloidal surrogates algorithm for convergent penalized-likelihood emission image reconstruction. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pages 1132–5, 1998.
- [217] T. Hebert and R. Leahy. A Bayesian reconstruction algorithm for emission tomography using a Markov random field prior. In *Proc. SPIE 1092, Med. Im. III: Im. Proc.*, pages 459–462, 1989.
- [218] T. Hebert and R. Leahy. A generalized EM algorithm for 3-D Bayesian reconstruction from Poisson data using Gibbs priors. *IEEE Tr. Med. Im.*, 8(2):194–202, June 1989.
- [219] T. J. Hebert and R. Leahy. Statistic-based MAP image reconstruction from Poisson data using Gibbs priors. *IEEE Tr. Sig. Proc.*, 40(9):2290–303, September 1992.
- [220] E. Tanaka. Utilization of non-negativity constraints in reconstruction of emission tomograms. In S. L. Bacharach, editor, *Information Processing in Medical Im.*, pages 379–83. Martinus-Nijhoff, Boston, 1985.
- [221] R. M. Lewitt and G. Muellerheiner. Accelerated iterative reconstruction for positron emission tomography based on the EM algorithm for maximum likelihood estimation. *IEEE Tr. Med. Im.*, 5(1):16–22, March 1986.
- B.13
- [222] T. Hebert, R. Leahy, and M. Singh. Three-dimensional maximum-likelihood reconstruction for an electronically collimated single-photon-emission imaging system. *J. Opt. Soc. Am. A*, 7(7):1305–13, July 1990.
- [223] S. Hella, P. Schmidlin, A. Lindén, G. Rosenqvist, and L. Eriksson. Iterative image reconstruction for emission tomography: A study of convergence and quantitation problems. *IEEE Tr. Nuc. Sci.*, 37(2):629–635, April 1990.
- [224] D. P. Bertsekas. A new class of incremental gradient methods for least squares problems. *SIAM J. Optim.*, 7(4):913–26, November 1997.
- [225] 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 255–68. Kluwer, Dordrecht, 1998.
- [226] A. Nedic and D. Bertsekas. Convergence rate of incremental subgradient algorithms. In S. Uryasev and P. M. Pardalos, editors, *Stochastic Optimization: Algorithms and Applications*, pages 283–304. Kluwer, New York, 2000.
- [227] A. Nedic, D. Bertsekas, and V. Borker. Distributed asynchronous incremental subgradient methods. In D. Buniaini, Y. Censor, S. Reich, editor, *Inherently Parallel Algorithms in Feasibility and Optimization and Their Applications*. Elsevier, Amsterdam, 2000.
- [228] A. Nedic and D. P. Bertsekas. Incremental subgradient methods for nondifferentiable optimization. *SIAM J. Optim.*, 12(1):109–38, 2001.
- [229] V. M. Kibardin. Decomposition into functions in the minimization problem. *Akronika / Tekhnicheskaya*, 9:66–79, September 1979. Translation: p. 1311–23 in Plenum Publishing Co. Adaptive Systems.
- [230] H. Kudo, H. Nakazawa, and T. Saito. Convergent block-iterative method for general convex cost functions. In *Proc. of the 1999 Int. Mtg. on Fully 3D Im. Recon. in Rad. Nuc. Med.*, pages 247–250, 1999.
- [231] A. De Piero and M. E. B. Yamagishi. Fast EM-like methods for maximum a posteriori estimates in emission tomography. *IEEE Tr. Med. Im.*, 20(4):280–8, April 2001.
- [232] S. Ann and J. A. Fessler. Globally convergent image reconstruction for emission tomography using relaxed ordered subsets algorithms. *IEEE Tr. Med. Im.*, 7, 2002. Submitted.
- [233] S. Ann and J. A. Fessler. Globally convergent ordered subsets algorithms: Application to tomography. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, pages M1–2, 2001.
- [234] H. Erdoghan and J. A. Fessler. Ordered subsets algorithms for transmission tomography. *Phys. Med. Biol.*, 44(11):2835–51, November 1999.
- [235] J. W. Stayman and J. A. Fessler. Spatially-variant roughness penalty design for uniform resolution in penalized-likelihood image reconstruction. In *Proc. IEEE Int. Conf. on Image Processing*, volume 2, pages 685–9, 1998.
- [236] J. W. Stayman and J. A. Fessler. Compensation for nonuniform resolution using penalized-likelihood reconstruction in space-variant imaging systems. *IEEE Tr. Med. Im.*, 7, 2002. Submitted.
- [237] J. Qi and P. M. Leahy. Resolution and noise properties of MAP reconstruction for fully 3D PET. In *Proc. of the 1999 Int. Mtg. on Fully 3D Im. Recon. in Rad. Nuc. Med.*, pages 35–9, 1999.
- [238] X. Liu, C. Comtat, C. Michel, P. Kinahan, M. Defrise, and D. Townsend. Comparison of 3D reconstruction with OSEM and FORE-OSEM for PET. In *Proc. of the 1999 Int. Mtg. on Fully 3D Im. Recon. in Rad. Nuc. Med.*, pages 39–42, 1999.
- [239] X. Liu, C. Comtat, C. Michel, P. Kinahan, M. Defrise, and D. Townsend. Comparison of 3D reconstruction with OSEM and FORE-OSEM for PET. *IEEE Tr. Med. Im.*, 20(8):304–13, August 2001.
- [240] C. Bai, E. E. Kinahan, D. Brasse, C. Comtat, and D. Townsend. Post-injection single photon transmission tomography with ordered-subset algorithms for whole-body PET imaging. *IEEE Tr. Nuc. Sci.*, 7, 2001.
- [241] C. Michel, M. Sbornina, A. Bol, X. Bernard, M. Lonneux, M. Defrise, C. Comtat, P. E. Kinahan, and D. W. Townsend. Preserving Poisson characteristics of PET data with weighted OSEM reconstruction. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, pages 1323–9, 1998.
- [242] C. Michel, X. Liu, S. Sanabria, M. Lonneux, M. Sbornina, 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 Nuc. Sci. Symp. Med. Im. Conf.*, volume 3, pages 1152–7, 1999.
- [243] D. F. Yu, J. A. Fessler, and E. P. Ficaro. Maximum likelihood transmission image reconstruction for overlapping transmission beams. *IEEE Tr. Med. Im.*, 19(11):1094–1105, November 2000.
- [244] I. Elbakri and J. A. Fessler. Ordered subsets transmission reconstruction with beam hardening correction for x-ray CT. In *Proc. SPIE 4322, Medical Imaging 2001: Image Proc.*, volume 1, pages 1–12, 2001.
- [245] D. L. Snyder. Parameter estimation for dynamic studies in emission-tomography systems having list-mode data. *IEEE Tr. Nuc. Sci.*, 31(2):325–31, April 1984.
- [246] J. M. Ollinger and D. L. Snyder. A preliminary evaluation of the use of the EM algorithm for estimating parameters in dynamic tracer-studies. *IEEE Tr. Nuc. Sci.*, 32(1):948–54, February 1985.
- [247] J. M. Ollinger and D. L. Snyder. An evaluation of an improved method for computing histograms in dynamic tracer studies using positron-emission tomography. *IEEE Tr. Nuc. Sci.*, 33(1):435–8, February 1986.
- [248] J. M. Ollinger. Estimation algorithms for dynamic tracer studies using positron-emission tomography. *IEEE Tr. Med. Im.*, 6(2):115–25, June 1987.
- [249] J. M. Ollinger. An evaluation of a practical algorithm for estimating histograms in dynamic tracer studies using positron-emission tomography. *IEEE Tr. Nuc. Sci.*, 34(1):348–53, February 1987.
- [250] F. O’Sullivan. Imaging radiotracer model parameters in PET: a mixture analysis approach. *IEEE Tr. Med. Im.*, 12(3):399–412, September 1993.
- [251] P. C. Chiao, W. L. Rogers, J. A. Fessler, N. H. Clinthorne, and A. O. Hero. Model-based estimation with boundary side information or boundary regularization. *IEEE Tr. Med. Im.*, 13(2):227–34, June 1994.
- [252] J. M. Bornewin and W. Sun. The stability analysis of dynamic SPECT systems. *Numerische Mathematik*, 77(3):283–98, 1997.
- [253] J. S. Maltz, E. Polak, and T. F. Budinger. Multiscale optimization algorithm for joint spatial and kinetic parameter estimation from dynamic ECT projection data. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 3, pages 1567–73, 1988.
- [254] D. S. Lalush and B. M. W. Tsui. Block-iterative techniques for fast 4D reconstruction using a priori motion models in gated cardiac SPECT. *Phys. Med. Biol.*, 43(4):875–86, April 1998.
- [255] T. E. Nichols, J. Qi, and R. M. Leahy. Continuous time dynamic PET imaging using list mode data. In A. Todd-Pokropek, A. Kuba, M. Sma, editor, *Information Processing in Medical Im.*, pages 98–111. Springer, Berlin, 1998.
- [256] E. Asma, T. E. Nichols, J. Qi, and R. M. Leahy. 4D PET image reconstruction from list mode data. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pages 1557–65, 2000.
- [257] U. Schmitt and A. K. Louis. Efficient algorithms for the regularization of dynamic inverse problems: I. Theory. *Inverse Prob.*, 18(3):645–68, June 2002.
- [258] U. Schmitt, A. K. Louis, C. Wolters, and M. Vauhkonen. Efficient algorithms for the regularization of dynamic inverse problems: II. Applications. *Inverse Prob.*, 18(3):659–76, June 2002.
- [259] C.-M. Kuo, J. T. Yip, J. Mukherjee, and M. N. Wernick. Image reconstruction for dynamic PET based on low-order approximation and restoration of the sinogram. *IEEE Tr. Med. Im.*, 16(6):727–37, December 1997.
- [260] J. Matthews, D. Bailey, P. Price, and V. Cunningham. The direct calculation of parametric images from dynamic PET data using maximum-likelihood iterative reconstruction. *Phys. Med. Biol.*, 42(6):1155–73, June 1997.
- [261] S. R. Meikle, J. C. Matthews, V. J. Cunningham, D. L. Bailey, L. Liveratos, T. Jones, and P. Price. Parametric image reconstruction using spectral analysis of PET projection data. *Phys. Med. Biol.*, 43(3):651–66, March 1998.
- [262] M. V. Narayanan, M. A. King, E. J. Soares, C. L. Byrne, P. H. Pretorius, and M. N. Wernick. Application of the Karhunen-Loeve transform to 4D reconstruction of cardiac gated SPECT images. *IEEE Tr. Nuc. Sci.*, 46(4):1001–8, August 1999.
- [263] M. N. Wernick, E. J. Infusino, and M. Milosevic. Fast spatio-temporal image reconstruction for dynamic PET. *IEEE Tr. Med. Im.*, 18(3):185–95, March 1999.
- [264] M. V. Narayanan, M. A. King, M. N. Wernick, C. L. Byrne, E. J. Soares, and P. H. Pretorius. Improved image quality and computation reduction in 4-D reconstruction of cardiac-gated SPECT images. *IEEE Tr. Med. Im.*, 19(5):423–33, May 2000.
- [265] J. E. Koss, D. L. Kirch, E. P. Little, T. K. Johnson, and P. P. Steele. Advantages of list-mode acquisition of dynamic cardiac data. *IEEE Tr. Nuc. Sci.*, 44(6):2431–8, December 1997.
- [266] Y. Censor, D. E. Gustafson, A. Lent, and H. Tuy. A new approach to the emission computerized tomography problem: simultaneous calculation of attenuation and activity coefficients. *IEEE Tr. Nuc. Sci.*, 26(2), January 1979.
- [267] P. R. Nunes. *Estimation algorithms for medical imaging including joint attenuation and emission reconstruction*. PhD thesis, Stanford Univ., Stanford, CA, June 1980.
- [268] M. S. Kaplan, D. R. Haynor, and H. Vijla. A differential attenuation method for simultaneous estimation of SPECT activity and attenuation distributions. *IEEE Tr. Nuc. Sci.*, 46(3):533–41, June 1999.
- [269] C. M. Laymon and T. G. Turkington. Calculation of attenuation factors from combined singles and coincidence emission projections. *IEEE Tr. Med. Im.*, 18(12):1194–200, December 1999.
- [270] R. Remeau, R. Clackdoyle, F. Noo, and G. Bal. Accurate attenuation correction in SPECT imaging using optimization of bilinear functions and assuming an unknown spatially-varying attenuation distribution. *Zeitschrift für Angewandte Mathematik und Mechanik*, 80(9):613–21, 2000.
- [271] A. Krol, J. E. Bowsher, S. H. Manglos, D. H. Feglin, M. P. Torreal, and F. D. Thomas. An EM algorithm for estimating SPECT emission and transmission parameters from emission data only. *IEEE Tr. Med. Im.*, 20(3):218–32, March 2001.

[272] J. Nuyts, P. Dupont, S. Stroobants, R. Bennink, L. Morelmans, and P. Suetens. Simultaneous maximum a-posteriori reconstruction of attenuation and activity distributions from emission sinograms. *IEEE Tr. Med. Im.*, 18(5):399–403, May 1999.

The literature on image reconstruction is enormous and growing. Many valuable publications are not included in this list, which is not intended to be comprehensive.

Slides and lecture notes available from: <http://www.eecs.umich.edu/~fessler/B.17>

Glossary of Symbols

\mathbf{R}^3	3-dimensional Euclidian space
\vec{r}	position in 3-dimensional Euclidian space
\mathcal{B}	subset of \mathbf{R}^3
$\lambda(\vec{r})$	emission density
$\mathbf{\Lambda}$	vector of emission density expansion coefficients
$\hat{\mathbf{\Lambda}}$	estimate of $\mathbf{\Lambda}$
$\vec{x}_k(t)$	position of k th tracer atom vs time t
$f_{\text{RTO}}^{(k)}(\vec{r})$	radotracer distribution function (pdf of $\vec{x}_k(t)$)
T_k	decay time of k th tracer atom
$t_{1/2}$	half-life of radionuclide
S_k	detector unit that records k th decay
μ_N	mean number of injected tracer atoms
μ_T	mean of T_k
$\mathbf{1}_{\{t\}}$	indicator function
n_p	number of pixels
n_d	number of detector units
s_j	sensitivity pattern of j th detector unit
$s(\vec{r})$	system sensitivity pattern
\mathbf{A}	matrix transpose
Ψ	cost function: $\Psi: \mathbf{R}^n \rightarrow \mathbf{R}$
\dot{h}_i	derivative of h_i
Y_i	counts recorded by i th detector unit during scan
y_i	random realization of Y_i
\tilde{y}_i	mean of y_i
r_i	mean number of background counts for i th detector unit
a_{ij}	system matrix element
λ_{ij}	j th element of $\mathbf{\Lambda}$
$[\mathbf{A}\mathbf{x}]_i$	shorthand for $\sum_{j=1}^{n_p} a_{ij}x_j$
$R(x)$	penalty function
\mathcal{W}	potential function
\mathcal{N}_j	neighborhood of j th pixel
$c_{\mathbf{N}_j}$	element of \mathcal{C}
\mathbf{x}	generic parameter vector to be estimated
$\mathbf{x}^{(n)}$	value of \mathbf{x} at n th iteration