Convergence and Non-Convergence of Algebraic Iterative Reconstruction Methods

Per Christian Hansen

DTU Compute, Technical University of Denmark

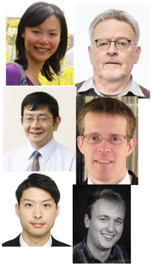

Joint work with

Tommy Elfving – Linköping University

Ken Hayami – NII, Tokyo

Michiel E. Hochstenbach – TU Eindhoven

Keiichi Morikuni – Univ. of Tsukuba

Yiqiu Dong & Nicolai A. B. Riis – DTU Compute

Overview of This Talk

Prelude

- X-ray CT model
- Reconstruction
- The algebraic approach

Fugue

- Stationary iterative reconstruction methods
- Their convergence
- Semi-convergence with noisy data
- Unmatched projectors
- Non-convergence and how to avoid it

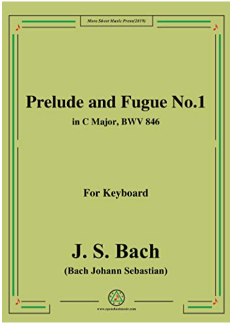

Prelude: X-Ray Computed Tomography (CT)

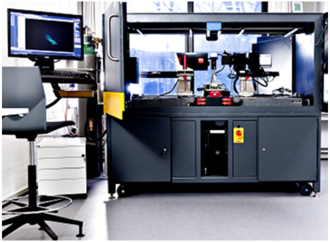

Lab scanner

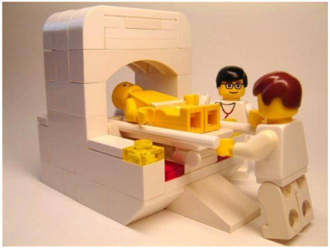

Medical scanner

Synchrotron

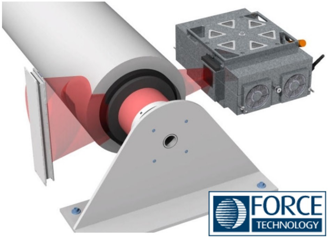

Industrial inspection

X-Ray Tomography and the Radon Transform

The Principle

Send X-rays through the object at different angles, and measure the attenuation.

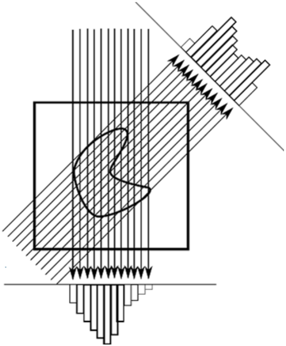

Lambert-Beer law $\rightarrow$ attenuation of X-ray through the object f is a line integral:

$$b_i = \int_{\mathrm{ray}_i} f(\xi_1, \xi_2)\, d\ell \ ,$$

$$f = \text{attenuation coef.}$$

A discrete version:

$$\boxed{A\,x = b}$$

$A \sim$ measurement geometry,

$x \sim$ reconstruction, $\quad b \sim$ data.

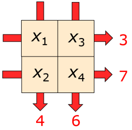

$$\begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} = \begin{pmatrix} 3 \\ 7 \\ 4 \\ 6 \end{pmatrix}$$

Modelling in CT: Forward and Back Projections

Forward projection $\mathcal{R}$, the Radon transform models the scanner physics via integration of the image f along lines $L_{\theta,s}$

$$\mathcal{R}[f](\theta, s) = \int_{L_{\theta,s}} f(\xi_1, \xi_2)\, d\ell = g(\theta, s) = \text{sinogram} \ .$$

Back projection $\mathcal{B}$ = adjoint($\mathcal{R}$), an abstraction, smears g back along $L_{\theta,s}$

$$\mathcal{B}[g](\xi_1, \xi_2) = \int_0^{2\pi} g(\theta, \xi_1 \cos\theta + \xi_2 \sin\theta)\, d\theta \ .$$

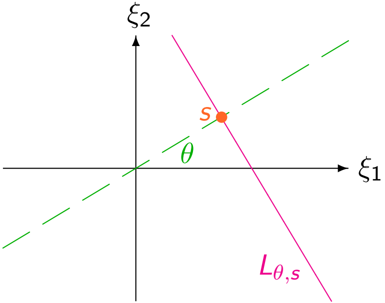

Reconstruction Algorithms

CT reconstruction is a mildly ill-posed inverse problem.

Lots of data + high resolution $\rightarrow$ large-scale computational problem.

Transform-based methods

Formulate the forward problem as a certain *transform*, then formulate a stable way to *invert* the transform.

Need to incorporate filtering in the inversion to obtain stability.

2D CT: Radon transform $\leftrightarrow$ filtered back projection (FBP).

Algebraic iterative methods

Discretize the forward problem and solve the corresponding large-scale problem $A\,x = b$ by means of an *iterative method*.

Need to incorporate regularization in the iterative solver to obtain stability.

Filtered Back Projection Versus Algebraic Reconstruction

- FBP: fast, low memory, good results with sufficiently many good data.
- But *artifacts* appear with noisy and/or limited data.
- Difficult to incorporate constraints (e.g., nonnegativity).
- Algebraic iterative reconstruction methods are more flexible and adaptive – but require more computational work.

Example with 3% noise and an incomplete set of projection angles:

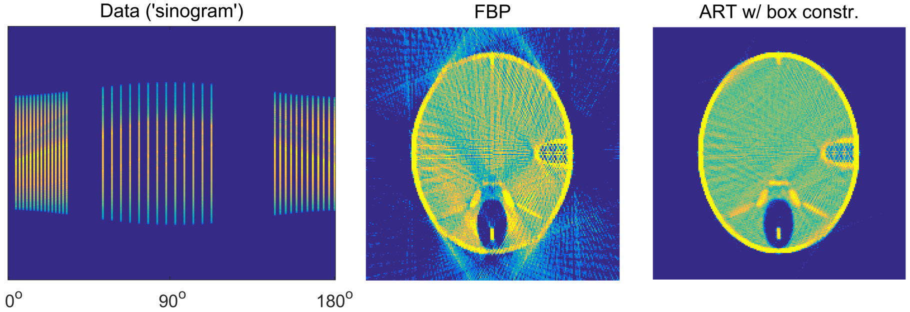

Storage Considerations

$N \times N$ image: each X-ray intersects at most $2N$ pixels $\rightarrow$ at most $2N$ nonzero elements in each row of A (at most $3N$ in 3D) $\rightarrow$ A is *sparse*.

Can still be problematic. 3D example: 1000 projection angles, 1000×1000 detector pixels, $1000 \times 1000 \times 1000$ voxels $\rightarrow$ number of non-zeros in A is of the order 10^{12} $\sim$ several Terabytes of memory.

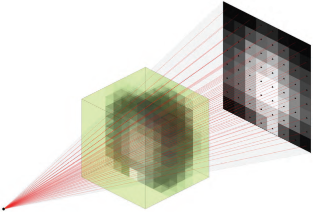

Alternative: use projection models to compute the matrix multiplications – the forward and back projections – "on the fly." We avoid the impossible task of storing A, at the price of having to recompute the matrix elements each time we need them.

More details here:

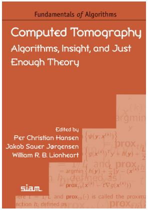

Fugue: Stationary Iterative Reconstruction Methods

A general class of iterative methods:

$$x^{k+1} = x^k + \omega\, D\, A^T M\, (b - A\, x^k), \qquad k = 0, 1, 2, \ldots$$

Diagonal matrices	D	M
Landweber	I	I
Gradient descent = steepest descent		
Cimmino	I	$\frac{1}{m}\,\mathrm{diag}\left(\frac{1}{\|a_i\|_2^2}\right)$
Landweber with row normalization		
CAV	I	$\mathrm{diag}\left(\frac{1}{\|a_i\|_S^2}\right)$
Component Averaging		$S = \mathrm{diag}(\mathtt{nnz}(\text{column } j))$
DROP	S^{-1}	$\mathrm{diag}\left(\frac{1}{\|a_i\|_2^2}\right)$
Diagonally relaxed orthogonal projection		
SIRT (aka SART)	$\mathrm{diag}\left(\frac{1}{\|a^j\|_1}\right)$	$\mathrm{diag}\left(\frac{1}{\|a_i\|_1}\right)$
Simultaneous iterative reconstruction technique		

Notation: $a_i = A(i,:) = \text{row}, \quad a^j = A(:,j) = \text{column.}$

Example of Convergence for Cimmino

Image size: 128×128.
Data: 360 projection angles in $[1^\circ, 360^\circ]$,
181 detector pixels.

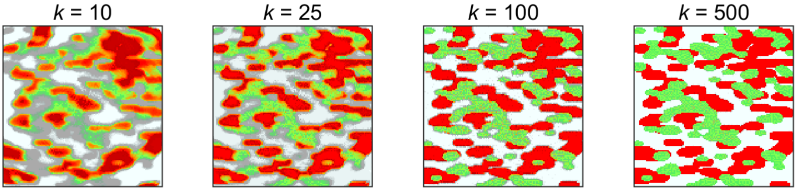

We must be concerned with three types of convergence:

1. **Convergence** of the iterative method.
2. **Semi-convergence** in the face of noisy data.
3. **Non-convergence** when forward and back projections don't match.

Asymptotic Convergence for Cimmino

Follows from Nesterov (2004)

Assume that A is invertible and scaled such that $\|A\|_2^2 = m$.

$$\|x^k - \bar{x}\|_2^2 \leq \left(1 - \frac{2}{1+\kappa^2}\right)^k \|x^0 - \bar{x}\|_2^2 \ ,$$

where $\bar{x} = A^{-1}b$ and $\kappa = \|A\|_2 \, \|A^{-1}\|_2$. This is linear convergence.

When κ is large then we have the approximate upper bound

$$\|x^k - \bar{x}\|_2^2 \lesssim (1 - 2/\kappa^2)^k \, \|x^0 - \bar{x}\|_2^2 \ ,$$

showing that in each iteration the error is reduced by a factor $1 - 2/\kappa^2$.

Real Problems Have Noisy Data

A standard topic in numerical linear algebra: solve $A\,x = b$.

Don't do this for inverse problems with noisy data!

The right-hand side b (the data) is a sum of noise-free data $\bar{b} = A\,\bar{x}$ from the ground-truth image $\bar{x}$ plus a noise component e:

$$b = A\,\bar{x} + e, \qquad \bar{x} = \text{ground truth}, \qquad e = \text{noise}.$$

The naïve solution $x^{\text{naïve}} = A^{-1} b$ is undesired, because it has a large component coming from the noise in the data:

$$x^{\text{naïve}} = A^{-1} b = A^{-1}(A\,\bar{x} + e) = \bar{x} + A^{-1} e.$$

The component $A^{-1} e$ dominates over $\bar{x}$, because A is ill conditioned.

But something interesting happens during the iterations . . .

The Reconstruction Error With Noisy Data

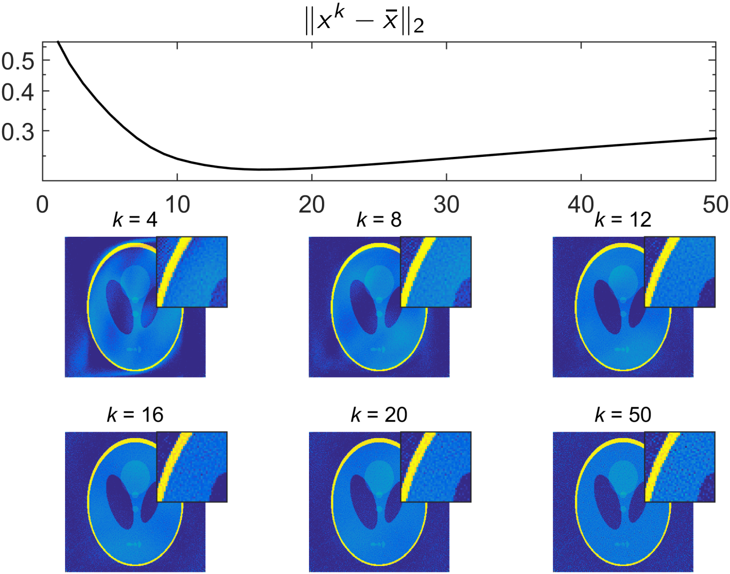

Semi-Convergence

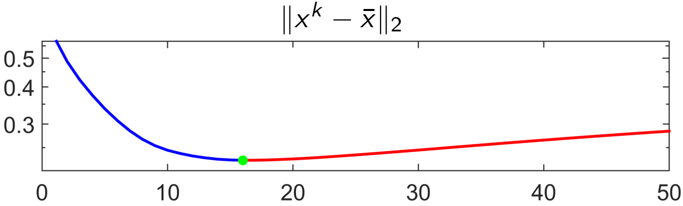

- In the initial iterations x^k approaches the unknown ground truth $\bar{x}$.
- During later iterations x^k converges to the undesired $x^{\text{naïve}} = A^{-1}b$.
- Stop the iterations when the convergence behavior changes.

Then we achieve a regularized solution: an approximation to the noise-free solution which is not too perturbed by the noise in the data.

- Today we explain *why* we have semi-convergence for noisy data.
- How to stop the iterations at the right time is a *different story*.

Convergence Analysis: Split the Error

Let $\bar{x}^k$ denote the iterates for a noise-free right-hand side. We consider:

$$\underbrace{x^k - \bar{x}}_{\text{total error}} = \underbrace{x^k - \bar{x}^k}_{\text{noise error}} + \underbrace{\bar{x}^k - \bar{x}}_{\text{iteration error}}$$

We expect the iteration error to decrease and the noise error to increase.

Then we have *semi-convergence* when the noise error starts to dominate:

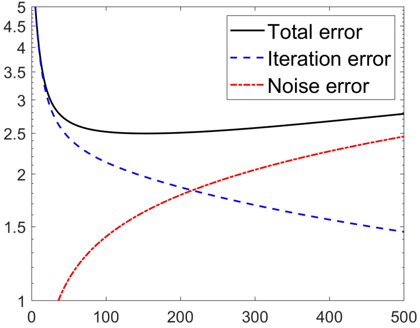

Analysis of Semi-Convergence for Cimmino

We use the SVD: $M^{\frac{1}{2}} A = \sum_{i=1}^{n} u_i \, \sigma_i \, v_i^T$

Van der Sluis & Van der Vorst (1990); Elfving, Nikazad, H (2010)

The iterate x^k is a **filtered SVD solution:**

$$x^k = \sum_{i=1}^{n} \varphi_i^{[k]} \, \frac{u_i^T (M^{\frac{1}{2}} b)}{\sigma_i} \, v_i , \qquad \varphi_i^{[k]} = 1 - \left(1 - \omega \, \sigma_i^2\right)^k .$$

Recall that we solve *noisy* systems $A\,x = b$ with $b = A\,\bar{x} + e$. Then:

$$x^k - \bar{x} = \underbrace{\sum_{i=1}^{n} \varphi_i^{[k]} \, \frac{u_i^T (M^{\frac{1}{2}} e)}{\sigma_i} \, v_i}_{\substack{\text{noise error} \\ \text{increases monotonically}}} - \underbrace{\sum_{i=1}^{n} (1 - \varphi_i^{[k]}) \, v_i^T \, \bar{x} \, v_i}_{\substack{\text{iteration error} \\ \text{decreases monotonically}}} .$$

Studies of Semi-Convergence

Semi-convergence has been analyzed by several authors:

- F. Natterer, *The Mathematics of Computerized Tomography* (1986)
- A. van der Sluis & H. van der Vorst, *SIRT- and CG-type methods for the iterative solution of sparse linear least-squares problems* (1990)
- M. Bertero & P. Boccacci, *Inverse Problems in Imaging* (1998)
- M. Kilmer & G. W. Stewart, *Iterative regularization and MINRES* (1999)
- H. W. Engl, M. Hanke & A. Neubauer, *Regularization of Inverse Problems* (2000)
- T. Elfving, H & T. Nikazad, *Semi-convergence properties of Kaczmarz's method* (2014)
- B. S. van Lith, H & M. E. Hochstenbach, *A twin error gauge for Kaczmarz's iterations* (2021)

So Far, So Good

At this time we have obtained an understanding of **convergence** and **semi-convergence** for algebraic iterative reconstruction methods.

I also promised to discuss **non-convergence** – to do that, I must briefly look at discretization methods for CT problems.

Projectors and Matrices

Multiplication with $A \leftrightsquigarrow$ action of forward projector $\mathcal{R}$.

Multiplication with $B \leftrightsquigarrow$ action of back projector $\mathcal{B} = \text{adjoint}(\mathcal{R})$.

When we can store A then we use A^T for back projection B, and our stationary iterative methods solve least squares problems associated with the normal equations $A^T A x = A^T b$.

When A is *too large to store*, we must use matrix-free multiplications of the forward projector and the back projector – cf. the Appendix.

HPC software: computational efficiency takes priority $\rightarrow B \neq A^T$.

We must study the influence of **unmatched** projector/backprojector pairs on the computed solutions and the convergence of the iterations.

Convergence Analysis for Unmatched Pairs

Substituting B for A^T in Landweber leads to the **BA Iteration**

$$\boxed{x^{k+1} = x^k + \omega\, B\,(b - A\,x^k)\ , \qquad \omega > 0.}$$

A *fixed-point iteration* that is not related to solving a minimization problem!

- Any fixed point x^* satisfies the **unmatched normal equations**

$$\boxed{B\,A\,x^* = B\,b.}$$

- If $B\,A$ is invertible then $x^* = (B\,A)^{-1} B\,b$.
- If $\mathcal{N}(B\,A) = \mathcal{N}(A)$ and $b \in \mathcal{R}(A)$ then $A\,x^* = b$.

Shi, Wei, Zhang (2011); Elfving, H (2018)

The **BA Iteration** converges to a solution of $B\,A\,x = B\,b$ if and only if

$$0 < \omega < \frac{2\,\mathrm{Re}(\lambda_j)}{|\lambda_j|^2} \quad \text{and} \quad \mathrm{Re}(\lambda_j) > 0, \qquad \{\lambda_j\} = \mathtt{eig}(B\,A)\ .$$

Iteration Error and Noise Error When $\mathrm{Re}(\lambda_j) > 0 \ \forall j$

Elfving, H (2018)

The *iteration error* is given by

$$\bar{x}^k - \bar{x}^* = T^k(\bar{x}^0 - \bar{x}) \ , \quad \bar{x}^0 = \text{initial vector} \ , \quad T = I - \omega\, B\, A \ ,$$

and it follows that we have **linear** convergence:

$$\|\bar{x}^k - \bar{x}\|_2 \ \le \ \|T^k\|_2\, \|\bar{x}^0 - \bar{x}\|_2 \ \le \ \|T\|_2^k\, \|\bar{x}^0 - \bar{x}\|_2 .$$

With $b = A\,\bar{x} + e$ the *noise error* satisfies

$$\boxed{\|x^k - \bar{x}^k\|_2 \ \le \ (\omega\, c \|B\|_2)\, k\, \|e\|_2}$$

where we define the constant c by: $\sup_j \|(I - \omega B A)^j\|_2 \le c$.

I.e., the upper bound grows linearly with the number of iterations k.

Numerical Example of Non-Convergence – **no Noise**

Parallel-beam CT, unmatched pair from *ASTRA*, 64×64 Shepp-Logan phantom, 90 proj. angles, 60 detector pixels, $\min \mathrm{Re}(\lambda_j) = -6.4 \cdot 10^{-8}$.

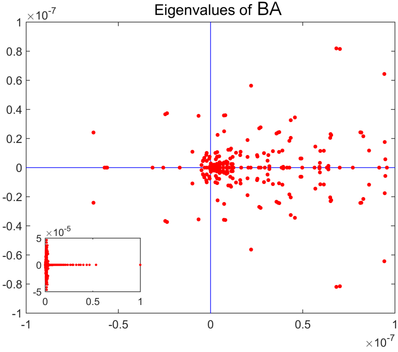

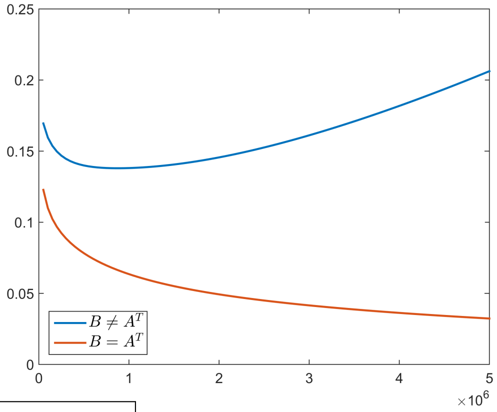

Non-convergence is the most common case.

What To Do?

1. Ask the software developers to change their implementation of forward projection and/or back projection?
   $\rightarrow$ Significant loss of computational efficiency.

2. Use mathematics to *fix* the nonconvergence.
   $\rightarrow$ What we do here.

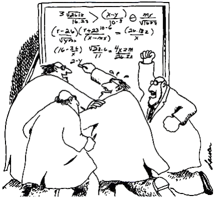

Take inspiration from the Tikhonov problem

$$\min_x \left\{ \|A\,x - b\|_2^2 + \alpha\,\|x\|_2^2 \right\} \ ,$$

for which a gradient step takes the form

$$\begin{aligned} x^{k+1} &= x^k - \omega\,(A^T(b - A\,x) + \alpha\,x^k) \\ &= (1 - \alpha\,\omega)\,x^k + \omega\,A^T(b - A\,x^k)\ . \end{aligned}$$

Note the factor $(1 - \alpha\,\omega)$.

The Shifted BA Iteration

Many thanks to Tommy Elfving for originally suggesting this.

We define the **shifted** version of the BA Iteration:

$$x^{k+1} = (1 - \alpha\,\omega)\,x^k + \omega\,B\,(b - A\,x^k)\;, \qquad \omega > 0$$

with just one extra factor $(1 - \alpha\,\omega)$; simple to implement.

This Shifted BA Iteration is equivalent to applying the BA Iteration with the substitutions

$$A \to \begin{bmatrix} A \\ \sqrt{\alpha}\,I \end{bmatrix}, \qquad B \to [\,B\,,\,\sqrt{\alpha}\,I\,]\;, \qquad b \to \begin{bmatrix} b \\ 0 \end{bmatrix}.$$

Hence it is "easy" to perform the convergence analysis ...

Convergence Results

Dong, H, Hochstenbach, Riis (2019)

Let λ_j denote those eigenvalues of BA that are different from $-\alpha$.

Then the Shifted BA Iteration converges to a fixed point if and only if α and ω satisfy

$$0 < \omega < 2 \, \frac{\operatorname{Re} \lambda_j + \alpha}{|\lambda_j|^2 + \alpha\,(\alpha + 2 \operatorname{Re} \lambda_j)} \qquad \text{and} \qquad \operatorname{Re} \lambda_j + \alpha > 0 \; .$$

The fixed point x_α^* satisfies

$$(BA + \alpha\, I)\, x_\alpha^* = B\, b \; .$$

This result tells us how to choose the shift parameter α:

Choose α just large enough that $\operatorname{Re} \lambda_j + \alpha > 0$ for all j.

"Perturbation" Result

How much do we perturb the *solution* $\bar{x}^*_\alpha$ – the fixed point – when we introduce $\alpha > 0$?

Dong, H, Hochstenbach, Riis (2019)

Assume that $BA + \alpha I$ is nonsingular and the right-hand side is noise-free with $b = \bar{b} = A\bar{x}$. Then the corresponding fixed point $\bar{x}^*_\alpha$ satisfies

$$\bar{x} - \bar{x}^*_\alpha = \alpha\,(BA + \alpha\,I)^{-1}\bar{x}\ .$$

Notice the factor α.

With a small α – just large enough to ensure convergence – we compute a slightly perturbed solution (instead of computing nothing).

Alternative: Solve the Unmatched Normal Equations

Instead of "fixing" a stationary method designed for solving another problem, just solve the unmatched normal equations in one of the forms

$$\boxed{B\,A\,x = B\,b} \qquad \text{or} \qquad \boxed{A\,B\,y = b\,, \quad x = B\,y}$$

The left- or right-preconditioned **GMRES** method for (A, b) immediately presents itself as a good choice with B as the preconditioner.

<u>BA-GMRES</u> solves $B\,A\,x = B\,b$ with B as a left preconditioner.

<u>AB-GMRES</u> solves $A\,B\,y = b\,,\; x = B\,y$ with B as a right preconditioner.

Advantages:

- these methods always converge,
- no need for relaxation parameter or shift parameter,
- we have semi-convergence, cf. Calvetti, Lewis, Reichel (2002) and Gazzola, Novati (2016).

Solving the Unmatched Normal Equations

Hayami, Yin, Ito (2010)

AB-GMRES solves $\min_y \|A\,B\,y - b\|_2$, $x = B\,y$ (B = right precond.)

▷ $\min_x \|A\,x - b\|_2 = \min_z \|A\,B\,z - b\|_2$ holds for all b if an only if $\text{range}(A\,B) = \text{range}(A)$, e.g., if $\text{range}(B) = \text{range}(A^T)$.

BA-GMRES solves $\min_x \|B\,A\,x - B\,b\|_2$ (B = left preconditioner)

▷ the problems $\min_x \|A\,x - b\|_2$ and $\min_x \|B\,A\,x - B\,b\|_2$ are equivalent for all b if and only if $\text{range}(B^T B\,A) = \text{range}(A)$, e.g., if $\text{range}(B^T) = \text{range}(A)$.

Both methods use the same Krylov subspace $\mathcal{K}_k(B\,A, B\,b)$ for the solution, but they use different objective functions.
They are identical to LSQR/LSMR when $B = A^T$.

Conditions are impossible to check in a given problem, but it works ↬

Image has 420×420 pixels, 600 projection angles, 420 detector pixels.

AB-GMRES $\|x_k - \bar{x}\|_2 / \|\bar{x}\|_2$

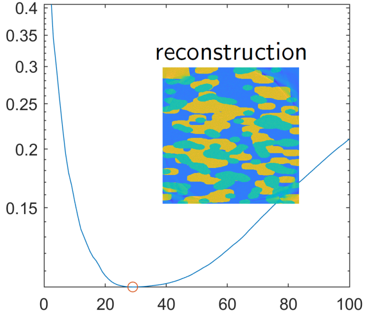

BA-GMRES $\|x_k - \bar{x}\|_2 / \|\bar{x}\|_2$

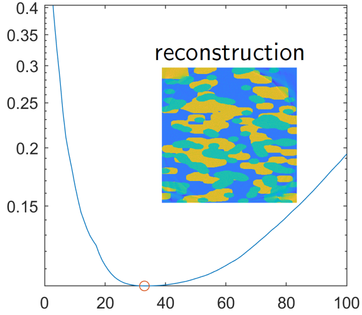

- Semi-convergence is evident (SVD analysis in Appendix).
- Same minimum reconstruction error $\|x_k - \bar{x}\|_2 / \|\bar{x}\|_2 \approx 0.10$ for both.
- Discrepancy principle and NCP-criterion stopping rules work well.
- Slightly fewer iterations for AB-GMRES in this example.

Coda

Facts

- The algebraic approach is very flexible; calls for iterative methods.
- Need to use matrix-free implementations for large-scale problems.

Convergence

- Good understanding of convergence for noise-free data.
- Emerging: good understanding of semi-convergence for noisy data.
- Non-convergence is caused by unmatched forward and back projectors.
- We avoid it with the right choice of algorithms.

Future

- Need more theory about semi-convergence for GMRES.
- Ready-to-use implementations for the CT community.

Appendix: Examples of Discretization Models

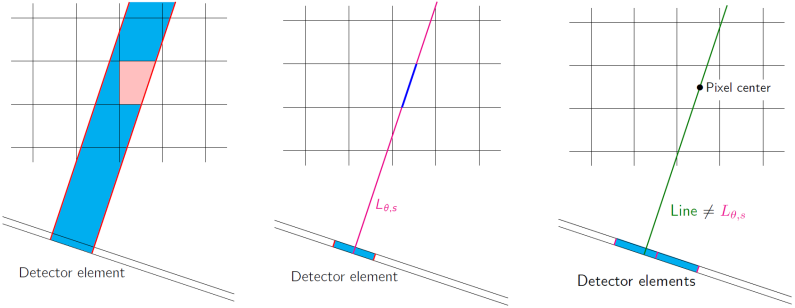

Forward strip model	Forward line model	Back projection model
Reflects the physics	Ray driven	Pixel driven
Not suited for GPUs	Suited for GPUs	Suited for GPUs

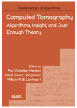

Forward line model: start from detector element centers.

Back projection model: start from image pixel centers and interpolate detector element values.

Appendix: Divergence and Convergence

Parallel-beam CT, 128×128 Shepp-Logan phantom, 90 projection angles in $[0^\circ, 180^\circ]$, 80 detector pixels; $m = 7\,200$ and $n = 16\,384$.
Both A and B are generated with the GPU-version of the ASTRA toolbox.

$\rho(BA) = 1.76 \cdot 10^4$
$\alpha = 1.85$

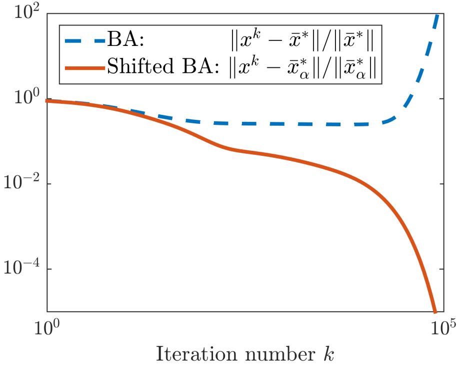

The BA Iteration diverges from $\bar{x}^* = (BA)^{-1} B\,\bar{b}$.

The Shifted BA Iteration converges to fixed point $\bar{x}^*_\alpha = (BA + \alpha\, I)^{-1} B\,\bar{b}$.

BA-GMRES: SVD Analysis, Small Matrix 23 040 × 16 384

Right-hand side: $\bar{b} = A\,\bar{x}$

$\bar{x}$ = ground truth

no noise

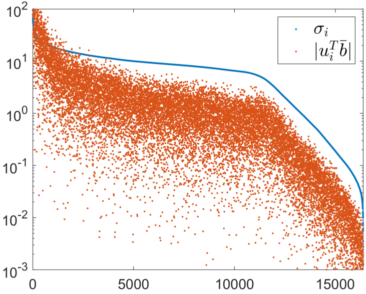

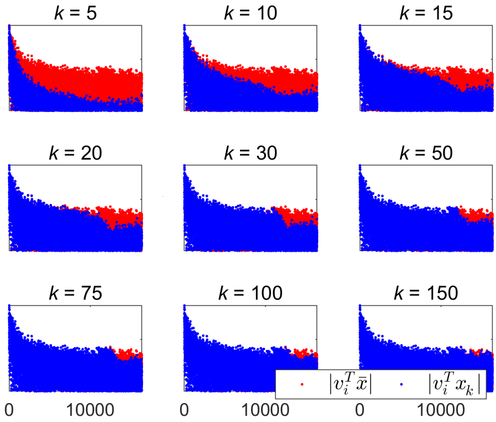

- Left plot is typical for X-ray CT problems; no rank deficiency.
- As k increases we capture more SVD components in x_k.
- At $k = 30$ we already capture the first 11 000 exact SVD components.

BA-GMRES: SVD Analysis – Now With Noisy Data

Right-hand side: $b = A\,\bar{x} + e$

$\bar{x}$ = ground truth

$\|e\|_2 / \|\bar{b}\|_2 = 0.003$ Gaussian

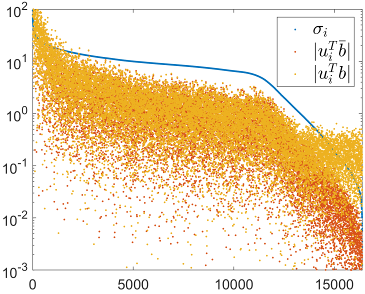

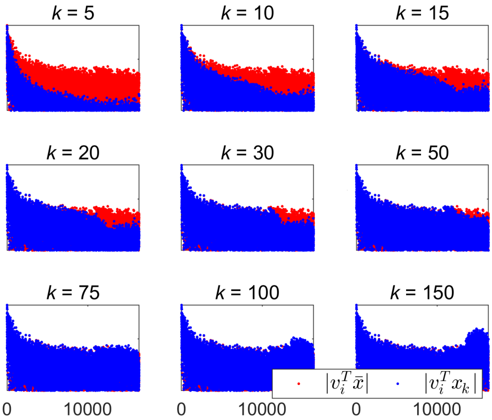

- Left plot is typical for X-ray CT problems; no rank deficiency.
- As k increases we capture more SVD components in x_k.
- At $k = 30$ we already capture the first 11 000 exact SVD components.
- Eventually we include noisy SVD components = **semi-convergence**.
- We obtain the best reconstruction after $k \approx 50$ iterations.

BA-GMRES: SVD Analysis – With More Noise

Right-hand side: $b = A\,\bar{x} + e$

$\bar{x}$ = ground truth

$\|e\|_2 / \|\bar{b}\|_2 = 0.03$ Gaussian

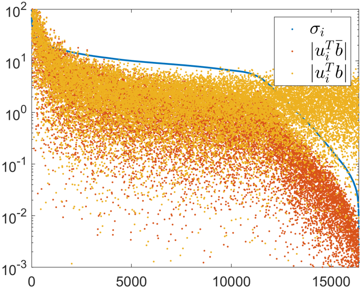

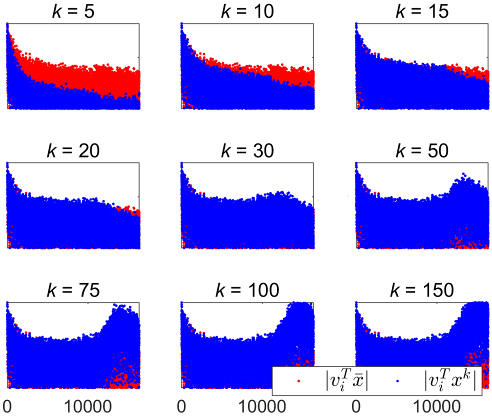

- Left plot: the "noise floor" increases..
- As k increases we capture more SVD components in x_k.
- At $k = 30$ we already capture the first 11 000 exact SVD components.
- Eventually we include noisy SVD components = **semi-convergence**.
- Now we obtain the best reconstruction after $k \approx 20$ iterations.

Appendix: Stopping Rules

We must terminate the iterations at the point of semi-convergence.

- **Discrepancy principle (DP):** terminates the iterations as soon as the residual norm is smaller than the noise level:

$$k_{\mathrm{DP}} = \text{the smallest } k \text{ for which } \|b - A\,x_k\|_2 \leq \tau\,\|e\|_2$$

  where $\tau \geq 1 =$ safety factor when we have a rough estimate of $\|e\|_2$.

- **NCP criterion:** uses the normalized cumulative periodogram to perform a spectral analysis of the residual vector $b - A\,x_k$ and identifies when the residual is close to being white noise – which indicates that all available information has been extracted from the noisy data.

For those who are curious: the L-curve criterion does not work, and we cannot implement generalized cross validation (GCV) efficiently.

Stopping Rules: Tests With 2 Different Back Projectors

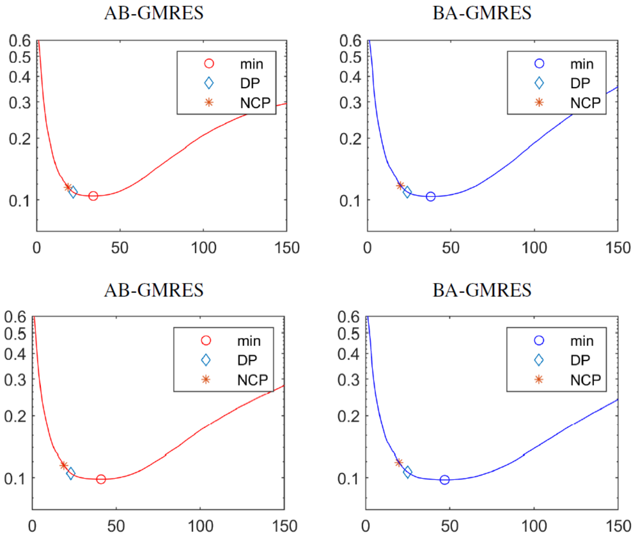

Both DP and NCP stop a bit too early – better than stopping too late.