



Topics in Numerical Analysis II

Computational Inverse Problems

Lecturer: Bangti Jin (b.jin@cuhk.edu.hk)

Chinese University of Hong Kong

October 16, 2022



The pros and cons of Tikhonov regularization

- relatively well developed theory (existence, convergence)
currently the most popular approach for inverse problems
- the choice of the penalty and regularization parameter is crucial
- requires **repeated** solution of optimization problem via optimizers
- ...



Iterative regularization

for the linear system

$$Ax = y,$$

with $\mathbb{R}^{m \times n}$, $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$

- there are a lot of iterative methods for solving linear systems: conjugate gradient, Krylov subspace (GMRES, MINRES, ...), ...
- an iterative solver attempts to solve the problem by finding successive approximations, starting from some x^0 , and then

$$x^{n+1} = f_n(x^n, x^{n-1}, \dots)$$

- Typically the update involves multiplications by A and its adjoint A^\top , but not explicit computation of the inverse.



why iterative methods

- sometimes the only feasible choice if the problem involves a large number of variables, making the direct methods (e.g., Gauss elimination) prohibitive
- especially practical if multiplications by A / A^* are cheap e.g., dedicated implementation on GPUs
- usually not designed for ill-posed equations, but often possesses **self-regularizing** properties: if the iterations are terminated before the solution starts to fit to noise (i.e., **early stopping**), one often obtains reasonable solutions for inverse problems.
- iteration index plays the role of *regularization parameter*



How to construct an iterative solver for

$$Ax = y?$$

minimize the residual

$$J(x) = \frac{1}{2} \|Ax - y\|^2$$

The global minimizer is given by $x^\dagger := A^\dagger y$, i.e., solving normal equation

$$A^*Ax = A^*y$$

and is orthogonal to $\ker(A)$. (this is not good for inverse problems.)



Instead we proceed iteratively: given x^0 , compute

$$\begin{aligned}x^{k+1} &= x^k - \beta \nabla J(x^k) \\ &= x^k - \beta A^*(Ax^k - y)\end{aligned}$$

Can one use it to construct approximation ?

- $\beta > 0$, step size (learning rate)
- This method is known as the Landweber method Landweber Amer. J. Math. 1951
- Landweber + early stopping is regularizing (later)
- optimal convergence rates (later)
- can be slow ...



Regularizing properties of Landweber method

The k th iteration of the Landweber iteration (with zero initial) can be written explicitly

$$\begin{aligned}x_{k+1} &= \sum_{j=0}^k (I - A^*A)^j A^* y = \sum_{j=0}^k (I - VS^*U^*USV^*)^j VS^*U^* y \\ &= \sum_{j=0}^k \sum_{i=1}^r (1 - s_i^2)^j s_i(y, u_i) v_i = \sum_{i=1}^r \sum_{j=0}^k (1 - s_i^2)^j s_i(y, u_i) v_i \\ &= \sum_{i=1}^r s_i^{-1} (1 - (1 - s_i^2)^{k+1}) (y, u_i) v_i\end{aligned}$$

frequency principle: low-frequencies error decays faster than high-freq. ones



$$x_k = \sum_{j=1}^r s_j^{-1} (1 - (1 - \beta s_j^2)^k) (y, u_j) v_j, \quad k = 0, 1, \dots$$

Since $|1 - \beta s_j^2| < 1$ by assumption,

$$(1 - \beta s_j^2)^k \rightarrow 0 \quad \text{as } k \rightarrow \infty.$$

which is expected since

$$x^\dagger = \sum_{j=1}^r s_j^{-1} (y, u_j) v_j.$$



However, while $k \in \mathbb{N}$ is finite, the coefficients of $(y, u_j)v_j$ satisfy

$$\begin{aligned} s_j^{-1}(1 - (1 - \beta s_j^2)^k) &= s_j^{-1} \left(1 - \sum_{\ell=0}^k \binom{k}{\ell} (-1)^\ell \beta^\ell s_j^{2\ell} \right) \\ &= \sum_{\ell=1}^k \binom{k}{\ell} (-1)^{\ell+1} \beta^\ell s_j^{2\ell-1} \end{aligned}$$

which converges to zero as $s_j \rightarrow 0$ (for a fixed k)

While, k is small enough, no coefficients of $(y, u_j)v_j$ is so large that the component of the measurement noise in the direction u_j is not amplified in an uncontrolled manner.



discrepancy principle

Let y be a noisy version of some underlying exact data y^\dagger and

$$\|y - y^\dagger\| = \delta > 0.$$



The Morozov discrepancy principle works for the Landweber iteration is similar to the truncated SVD and Tikhonov regularization: choose the smallest $k \geq 0$ s.t.

$$\|y^\delta - Ax_k^\delta\| \leq \delta.$$



Example: heat conduction



- w : the simulated heat distribution at $T = 0.1$
- f^\dagger : wedge function (initial data)
- $A = e^{TB}$ forward operator, $\|A\| = 1$, $\beta = 1$
- a small amount of noise to the measurement, discrepancy principle



exact solution v.s. least-squares solution




solution for exact data







convergence for exact data







results for noisy data (1% noise)






results for noisy data



discrepancy principle stops at $k = 6$



results for noisy data





convergence issue revisited

When the problem is ill-posed, $I - \beta A^* A$ is not a contraction !

problem setting:

$$Ax = y^\delta,$$

with $\|y^\delta - y^\dagger\| \leq \delta$, construct an approximation $x_{k(\delta)}^\delta$ by Landweber method, i.e.,

$$x_k^\delta = x_{k-1}^\delta - A^*(Ax_{k-1}^\delta - y^\delta), \quad k = 1, \dots,$$

(with $x_0^\delta = 0$, $\|A\| \leq 1$, by rescaling) s.t.

$$\lim_{\delta \rightarrow 0^+} x_{k(\delta)}^\delta = x^\dagger?$$



Let A be injective with dense range $\mathcal{R}(A) \subset Y$. If $y^\dagger \in D(A^\dagger)$, then $x_k \rightarrow A^\dagger y^\dagger$ as $k \rightarrow \infty$. If $y \notin D(A^\dagger)$, then $\|x_k\| \rightarrow \infty$ as $k \rightarrow \infty$.



important fact: $\|(I - AA^*)^k A^*\| \leq (k + 1)^{-1/2}$

If $y \in \mathcal{R}(A)$, then $y = Kx^\dagger$ for some x^\dagger , and it follows that

$$\begin{aligned}x^\dagger - x_k &= x^\dagger - x_{k-1} - A^*A(x^\dagger - x_{k-1}) \\ &= (I - A^*A)(x^\dagger - x_{k-1}) = (I - A^*A)^k x^\dagger\end{aligned}$$

For $\|A\| \leq 1$, the operator $(I - A^*A)^k$, $k \in \mathbb{N}$ are uniformly bounded by one, and for every $x^\dagger \in \mathcal{R}(A^*)$,

$$(I - A^*A)^k x^\dagger = (I - A^*A)^k A^* w \Rightarrow 0$$

Meanwhile, if A is injective and has dense range, $\mathcal{R}(A^*)$ is dense in X . Thus $(I - A^*A)^k$ converges pointwise to zero on a dense subset of X and Banach-Steinhaus theorem implies that the error converges to zero for each $x^\dagger \in \mathcal{X}$ as $k \rightarrow \infty$.



The iterate x_k is given by

$$x_{k+1} = \sum_{j=0}^k (I - A^*A)^j A^* y^\dagger$$

with $y^\dagger \in D(A^\dagger)$, then $A^* y^\dagger = A^* A x^\dagger$ for $x^\dagger = A^\dagger y^\dagger$, and

$$x^\dagger - x_{k+1} = x^\dagger - A^* A \sum_{j=0}^k (I - A^*A)^j x^\dagger$$

since $A^* A \sum_{j=0}^k (I - A^*A)^j = I - (I - A^*A)^{k+1}$, i.e.,

$$x^\dagger - x_{k+1} = (I - A^*A)^{k+1} x^\dagger$$

by the singular system (s_j, u_j, v_j) of A

$$\|x_{k+1} - x^\dagger\|^2 = \sum_j (1 - s_j^2)^{2(k+1)} (x^\dagger, v_j)^2$$

uniformly bounded + Lebesgue dominated convergence theorem



convergence rate under the source condition:

$$x^\dagger = A^* w$$

\Rightarrow

$$x^\dagger - x_{k+1} = (I - A^* A)^{k+1} x^\dagger = (I - A^* A)^{k+1} A^* w$$

by means of spectral decomposition:

$$\|x_{k+1} - x^\dagger\|^2 = \sum_j s_j^2 (1 - s_j^2)^{2(k+1)} (w, u_j)^2$$

Note that $\sup_{\lambda \in [0,1]} \lambda(1-\lambda)^k \leq (k+1)^{-1}$. The error decay

$$\|x_{k+1} - x^\dagger\|^2 \leq (2k+3)^{-1} \sum_j (w, u_j)^2 = (2k+3)^{-1} \|w\|^2$$

i.e.,

$$\|x_{k+1} - x^\dagger\| \leq (2k+2)^{-1/2} \|w\|$$



Let y^\dagger, y^δ be a pair of data with $\|y^\delta - y^\dagger\| \leq \delta$. Then

$$\|x_k - x_k^\dagger\| \leq \sqrt{k}\delta, \quad k \geq 0.$$

$$x_k - x_k^\delta = \sum_{j=0}^{k-1} (I - A^*A)^j A^* (y^\dagger - y^\delta) := R_k (y^\dagger - y^\delta)$$

and

$$\|R_k\|^2 = \|R_k R_k^*\| = \left\| \sum_{j=0}^{k-1} (I - A^*A)^j (I - (I - A^*A)^k) \right\| \leq \left\| \sum_{j=0}^{k-1} (I - A^*A)^j \right\| \leq k.$$



error analysis:

$$x^\dagger - x_k^\delta = x^\dagger - x_k + x_k - x_k^\delta$$

- approximation error: $x^\dagger - x_k$ converges to zero as $k \rightarrow \infty$
- data error $x_k - x_k^\delta$, of order $k^{\frac{1}{2}} \delta$
- convergence: $k(\delta) \rightarrow \infty$ and $\delta^2 k(\delta) \rightarrow 0$ as $\delta \rightarrow 0 \Rightarrow x_{k(\delta)}^\delta \rightarrow x^\dagger$
- optimal convergence (under $x^\dagger = A^* w$ + a priori choice of $k(\delta)$)

$$\|x_k^\delta - x^\dagger\| \leq c\delta^{\frac{1}{2}}, \quad \text{with } k(\delta) = \delta^{-1}.$$

- **semiconvergence: the regularizing property of iterative methods (for ill-posed) problems ultimately depend on reliable stopping**



bad news: the divergence point is not easy to determine by monitoring the residual $\|y^\delta - Ax_k^\delta\|!$

If $\|A\| \leq 1$ and if $\mathcal{R}(A)$ is dense in Y , then the norm of $r_k = y^\delta - Af_k$ decreases monotonically to zero.

$$\|r_k\|^2 \leq \langle r_{k-1}, r_k \rangle \leq \|r_{k-1}\|^2, \quad k = 1, 2, \dots,$$

and both inequalities are strict unless y^δ is zero or an eigenfunction of AA^* for the eigenvalue $\lambda = 1$



■ basic identity

$$r_k = (I - AA^*)r_{k-1} = \dots = (I - AA^*)^k y^\delta$$

and by Banach-Steinhaus theorem, $Ax_k \rightarrow y^\delta$ for any $y^\delta \in Y$, because $\|A\| \leq 1$ and $\mathcal{R}(A)$ is dense in Y .

■ second inequality

$$\langle r_{k-1}, r_k \rangle = \langle r_{k-1}, (I - AA^*)r_{k-1} \rangle = \|r_{k-1}\|^2 - \|A^*r_{k-1}\|^2 \leq \|r_{k-1}\|^2$$

This inequality is strict, unless $A^*r_{k-1} = 0$, i.e., $r_{k-1} = 0$ ($\mathcal{R}(A)$ is dense in Y). When $r_{k-1} = 0$, we distinguish two cases.

- For $k = 2$, this gives $r_1 = y^\delta - Ax_1 = 0$.
- for $k \geq 3$, then $r_{k-1} = 0$ implies $r_{k-2} \in \mathcal{N}(I - AA^*)$, and also $r_{k-2} \in \mathcal{R}(I - AA^*)$. Thus $r_j = 0$ for all $j = 1, \dots, k - 1$



■ first inequality

$$\|r_k\|^2 = \langle r_k, r_k \rangle = \langle (I - AA^*)r_{k-1}, r_k \rangle = \langle r_{k-1}, r_k \rangle - \langle A^*r_{k-1}, A^*r_k \rangle$$

the first inequality follows if $\langle A^*r_{k-1}, A^*r_k \rangle \geq 0$:

$$\begin{aligned} \langle A^*r_{k-1}, A^*r_k \rangle &= \langle A^*r_{k-1}, A^*(I - AA^*)r_{k-1} \rangle \\ &= \langle A^*r_{k-1}, (I - A^*A)A^*r_{k-1} \rangle \geq 0 \end{aligned}$$

It remains to study when the first inequality fails to be strict. This occurs when $A^*r_{k-1} \in \mathcal{N}(I - A^*A)$ for $k \geq 2$. Meanwhile, $A^*r_{k-1} = (I - A^*A)A^*r_{k-2} \in \mathcal{R}(I - A^*A) \perp \mathcal{N}(I - A^*A)$, hence $A^*r_{k-1} \in \mathcal{N}(I - A^*A)$ can only happen $A^*r_{k-1} = 0$, and as above $Ax_1 = g^\delta$. For $k = 1$, it follows that $0 = (I - A^*A)A^*r_0 = A^*r_1$, which also gives $r_1 = 0$.



discrepancy principle: choose $k(\delta)$ s.t.

$$\|y^\delta - Ax_{k(\delta)}^\delta\| \leq \tau\delta \leq \|y^\delta - Ax_k^\delta, \quad \forall 0 \leq k < k(\delta),$$

with $\tau > 1$ fixed

motivation

- x_k^δ approximates x^\dagger better than x_{k-1}^δ , for all $k \leq k(\delta) - 1$.
That is, the discrepancy principle ensures the monotone decreasing of the iteration error
- The sequence $\|y^\delta - Ax_k^\delta\|$ is monotonically decreasing.



If $k(\delta) \leq 1$, the assertion is trivial. For $k(\delta) \geq 2$, $1 \leq k \leq k(\delta) - 1$:

$$\begin{aligned}
 & \|x_k^\delta - x^\dagger\|^2 - \|x_{k-1}^\delta - x^\dagger\|^2 \\
 &= \|x_{k-1}^\delta - x^\dagger + A^* r_{k-1}\|^2 - \|x_{k-1}^\delta - x^\dagger\|^2 \\
 &= 2\langle x_{k-1}^\delta - x^\dagger, A^* r_{k-1} \rangle + \|A^* r_{k-1}\|^2 \\
 &= 2\langle Ax_{k-1}^\delta - y^\dagger, r_{k-1} \rangle + \langle AA^* r_{k-1}, r_{k-1} \rangle \\
 &= 2\langle y^\delta - y^\dagger, r_{k-1} \rangle + 2\langle Ax_{k-1}^\delta - y^\delta, r_{k-1} \rangle + \langle AA^* r_{k-1}, r_{k-1} \rangle \\
 &\leq 2\|y^\delta - y^\dagger\| \|r_{k-1}\| - \|r_{k-1}\|^2 - \langle (I - AA^*) r_{k-1}, r_{k-1} \rangle \\
 &= 2\|y^\delta - y^\dagger\| \|r_{k-1}\| - \|r_{k-1}\|^2 - \langle r_k, r_{k-1} \rangle \\
 &\leq 2\|y^\delta - y^\dagger\| \|r_{k-1}\| - \|r_{k-1}\|^2 - \|r_k\|^2 \\
 &\leq 2(\delta - \|r_k\|) \|r_{k-1}\| - (\|r_{k-1}\| - \|r_k\|)^2 \leq 2(\delta - \|r_k\|) \|r_{k-1}\|
 \end{aligned}$$

$(r_k = y^\delta - Af_k + \text{monotonicity lemma})$



If A is injective with dense range \mathcal{A} and $\|A\| \leq 1$. The DP is consistent.



case (i): finite termination

If $k_j = k(\delta_{n_j}) = k$ is always the same for some subsequence (δ_{n_j}) .
Then the approximations $x_{k_j}^\delta$ satisfy

$$x_{k_j}^\delta = R_k y^{\delta_{n_j}} \rightarrow R_k y^\dagger, \quad j \rightarrow \infty$$

Meanwhile

$$y^{\delta_{n_j}} - Ax_{k_j}^\delta \rightarrow y^\dagger - AR_k y^\dagger, \quad j \rightarrow \infty$$

The discrepancy principle implies $\|y^{\delta_{n_j}} - Ax_{k_j}^\delta\| \leq \delta_{n_j}$ so
 $R_k y^\dagger = A^{-1} y^\dagger$ and hence $x_{k_j}^\delta \rightarrow A^{-1} y^\dagger = x^\dagger$



case (ii) infinite termination

for any subseq. (δ_{n_j}) with $k_{n_j} \rightarrow \infty$, let $\tilde{x}_{k_{n_j}}^{\delta_{n_j}} = R_{k_{n_j}-1} y^{\delta_{n_j}}$, the next to last iterate. Then for $y^\dagger \in \mathcal{R}(A)$, there holds $R_k y^\dagger \rightarrow \mathcal{A}^{-1} y^\dagger \Rightarrow$ there exists k_ϵ s.t. $\|R_{k_\epsilon} y^\dagger - x^\dagger\| \leq \epsilon$. Since $k_{n_j} \rightarrow \infty$, there is j_ϵ such that $k_{n_j} > k_\epsilon$ for all $j \geq j_\epsilon$, and the monotonicity lemma \Rightarrow

$$\begin{aligned} \|\tilde{x}_{k_{n_j}}^{\delta_{n_j}} - x^\dagger\| &\leq \|R_{k_\epsilon} y^{\delta_{n_j}} - x^\dagger\| \leq \|R_{k_\epsilon} y^{\delta_{n_j}} - R_{k_\epsilon} y^\dagger\| + \|R_{k_\epsilon} y^\dagger - x^\dagger\| \\ &\leq \|R_{k_\epsilon} (y^{\delta_{n_j}} - y^\dagger)\| + \epsilon \end{aligned}$$

for all $j \geq j_\epsilon$. Since $y^{\delta_{n_j}} \rightarrow y^\dagger$, $\|R_{k_\epsilon} (y^{\delta_{n_j}} - y^\dagger)\| \rightarrow 0$ as $j \rightarrow \infty$, i.e.,

$$\limsup_{j \rightarrow \infty} \|\tilde{x}_{k_{n_j}}^{\delta_{n_j}} - x^\dagger\| \leq \epsilon$$

Since ϵ is arbitrarily chosen, it follows $\tilde{x}_{k_{n_j}}^{\delta_{n_j}} \rightarrow x^\dagger$. Since

$$x_{k_{n_j}}^{\delta_{n_j}} = \tilde{x}_{k_{n_j}}^{\delta_{n_j}} + A^*(y^{\delta_{n_j}} - A\tilde{x}_{k_{n_j}}^{\delta_{n_j}}) \rightarrow x^\dagger$$



The Landweber iteration with discrepancy principle is order-optimal.

If A is injective with dense range, and $y^\dagger = Ax^\dagger$ with $x^\dagger = A^*w$, with $\|w\| \leq \rho$. Then

$$\|x_{k\delta}^\delta - x^\dagger\| \leq 3\sqrt{\rho}\sqrt{\delta} + 2\delta.$$



Let k'_δ be the first iteration for which $\|y^\delta - Af_{k'_\delta}^\delta\| \leq 2\delta$. Then the assumption on x^\dagger implies

$$r_k = y^\delta - Ax_k^\delta = (I - AA^*)^k y^\delta = (I - AA^*)^k (y^\delta - y^\dagger) + (I - AA^*)^k AA^* w$$

for every $k \in \mathbb{N}_0$. Hence,

$$2\delta < \|y^\delta - Ax_k^\delta\| \leq \delta + (k + 1)^{-1} \|w\|, \quad k = 0, \dots, k'_\delta$$

Setting $k = k'_\delta - 1$ yields $k'_\delta \leq \rho/\delta$. Then the error decomposition \Rightarrow

$$\|x_{k'_\delta}^\delta - x^\dagger\| \leq \|(I - R_{k'_\delta} K)K^* w\| + \sqrt{\rho\delta}$$

$$\begin{aligned} \|(I - R_k A)A^* w\|^2 &= \langle (I - AA^*)w, A(I - A^* A)^k A^* w \rangle \\ &= \langle (I - AA^*)^k w, (I - AA^*)^k y^\dagger \rangle \end{aligned}$$

with

$$\begin{aligned} \|(I - AA^*)^k y^\dagger\| &\leq \|(I - AA^*)^k (y^\delta - y)\| + \|(I - AA^*)^k y^\delta\| \\ &\leq \|I - AA^*\| \delta + \|r_k\| \end{aligned}$$



Thus, for $k = k'_\delta$,

$$\|(I - R_{k'_\delta} A)A^* w\|^2 \leq \|(I - AA^*)\| \|w\| \|(I - AA^*)^k y^\dagger\| \leq 3\rho\delta$$

Thus

$$\|x_{k'_\delta}^\delta - x^\dagger\| \leq (\sqrt{3} + 1)\sqrt{\rho\delta} \leq 3\sqrt{\rho\delta}$$

Then by the monotonicity, we have

$$\|x_{k_\delta-1}^\delta - x^\dagger\| \leq 3\sqrt{\rho\delta}$$

Then

$$\begin{aligned} \|x_{k_\delta}^\delta - x^\dagger\| &\leq \|x_{k_\delta-1}^\delta - x^\dagger\| + \|A^*\| \|y^\delta - Ax_{k_\delta-1}^\delta\| \\ &\leq 3\sqrt{\rho\delta} + \|y^\delta - Ax_{k'_\delta}^\delta\| \leq 3\sqrt{\rho\delta} + 2\delta. \end{aligned}$$



- The Landweber method can be slow ...
- How to accelerate the computation ...
 - simple: Anderson acceleration
 - simple: Kaczmarz / stochastic gradient descent
 - complex: conjugate gradient, MINRES
 - ...



Anderson acceleration for fixed point equation: $x_{n+1} = T(x_n)$



- let $g(x) = T(x) - x$, $g_k = g(x_k)$
- set x_0 and $m \geq 1$ (memory parameter)

D. G. Anderson. Iterative Procedures for Nonlinear Integral Equations. J. the ACM. 1965; 12 (4): 547–560

```
 $x_1 = T(x_0)$   
for  $k = 1, 2, \dots$  do  
   $m_k = \min(m, k)$   
   $G_k = [g_{k-m_k} \ \dots \ g_k]$   
   $\alpha_k = \arg \min_{\sum_{i=0}^{m_k} \alpha_i = 1} \|G_k \alpha\|$   
   $x_{k+1} = \sum_{i=0}^{m_k} (\alpha_k)_i f(x_{k-m_k+i})$   
end for
```



Landweber iteration v.s. Anderson acceleration for gravity





Anderson acceleration

- Landweber: slow convergence vs. slow divergence
- Anderson: fast convergence v.s. fast divergence
- no analysis of the regularizing property of AA ! (open)
-