Unfolding Noisy measurements

Solving discrete linear inverse problems

Andrew Adare
University of Colorado

PHENIX HHJ meeting
November 29, 2012

Linear inverse problems

Continuous case

Inverse problem: obtain model parameters of a physical system, given a set of observed quantities.

Invert the process that transforms model parameters to measured data.

Solutions don't satisfy Hadamard's three conditions of existence, uniqueness, and stability required for classification as a well-posed problem.

Linear inverse problems are modeled by a Fredholm integral equation of the first kind:

$$\int_0^1 K(s,t)f(t)\,\mathrm{d}t = g(s), \qquad 0 \le s \le 1 \tag{1}$$

Riemann-Lebesgue lemma

(in loose terms)

For any "arbitrary" finite kernel, the integrated Fourier components of f are progressively diminished as their order increases ($+$ and - portions increasingly smoothed into cancellation).

The opposite of this damping occurs in computing f from g: the high-frequency components of f are amplified.

Although mathematically correct, the solution is of little physical use.

Regularization: cure by imposing smoothness requirement on result (hopefully without biasing it).

Discrete linear inverse problems

The discrete version of equation 1 is the linear system

$$\mathrm{A}\mathbf{x} = \mathbf{b} \tag{2}$$

which includes the response matrix $\mathrm{A} \in \mathbb{R}^{m \times n}$, the solution vector $\mathbf{x} \in \mathbb{R}^n$, and the measured vector $\mathbf{b} = \mathbf{b}^{exact} + \mathbf{e} \in \mathbb{R}^m$.

Three cases for measured (m) and true/solution bin dimensions (n):

- $m < n$: underdetermined system (not considered here)
- $m = n$: critically constrained system (unique solution)
- $m > n$: overdetermined system (least squares problem)

In the latter case, the problem is expressed as

$$\min_{\mathbf{x}} \|\mathrm{A}\mathbf{x} - \mathbf{b}\|_2^2 \tag{3}$$

SVD analysis of the unfolding problem

A powerful tool for linear problems

For any matrix $A \in \mathbb{R}^{m \times n}$ with $m \geq n$,

$$A = U \Sigma V^T = \sum_{i=1}^{n} \sigma_i \mathbf{u}_i \mathbf{v}_i^T . \tag{4}$$

$U \in \mathbb{R}^{m \times n}$ and $V \in \mathbb{R}^{n \times n}$ are orthogonal matrices whose columns form an orthonormal basis, and $\Sigma = \mathrm{diag}(\sigma_1, \ldots, \sigma_n)$ contains the SV spectrum.

Singular value spectrum

Example: Gaussian convolution ($K(s,t) = K(s-t)$)

The singular values decay steeply until $i \approx 20$, where they drop below $\sigma_1 \times$ machine precision.

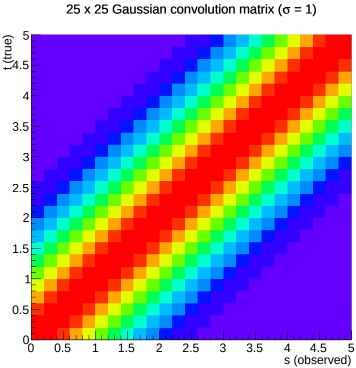

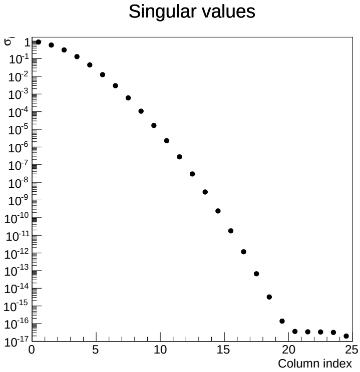

This behavior is typical for ill-posed problems, whose numerical rank (# nonzero σ_i values) is ill-determined.

SVD basis vectors

same example

The left singular vectors u_i are essentially Fourier components of A, and the singular values represent its power spectrum.

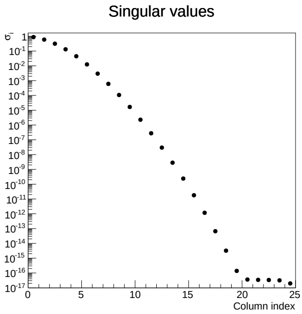

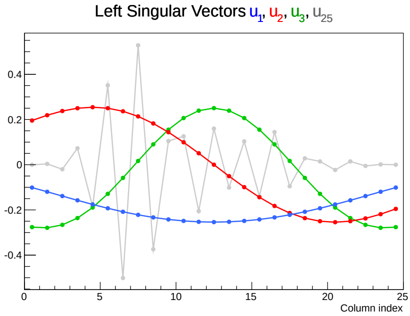

The lower-i components contain reliable information, while the highest components are noise-dominated.

Condition of the response matrix

large condition number ⇔ strong noise amplification

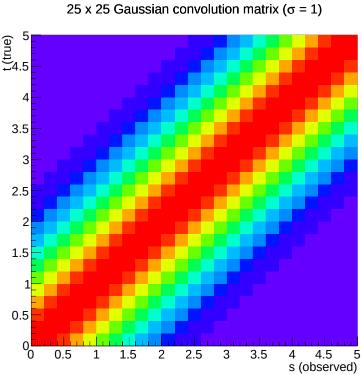

The 2-norm of A is
$||A||_2 = \max_{||x||_2=1} ||Ax||_2$.

The condition number of A is

$$\mathrm{cond}(A) = ||A||_2 ||A^{-1}||_2 = \sigma_1/\sigma_n.$$

Guideline: if the matrix condition number is 10^n, the accuracy in x is $15 - n$ digits for double precision.

This seems hopeless, but severely ill-conditioned matrices like this one (cond(A) $\sim 10^{17}$) can still allow reasonable solutions.

Regularization is critical!

Why direct inversion of A fails

Short answer: Ill-conditioned A + noise in $\mathbf{b}$

Direct algebraic solution: invert A to solve $A\mathbf{x} = \mathbf{b}$.

Since $A^{-1} = V\Sigma^{-1}U^T$, the direct (unregularized) solution is

$$\mathbf{x} = \mathrm{A}^{-1}\mathbf{b} = \sum_{i=1}^{n} \frac{\mathbf{u}_i^T \cdot \mathbf{b}}{\sigma_i} v_i. \tag{5}$$

What we have in practice[1] is

$$\mathbf{x} = \mathrm{A}^{-1}\mathbf{b}^{exact} + \mathrm{A}^{-1}\mathbf{e} = \sum_{i=1}^{n} \left(\frac{\mathbf{u}_i^T \cdot \mathbf{b}^{exact}}{\sigma_i} v_i + \frac{\mathbf{u}_i^T \cdot \mathbf{e}}{\sigma_i} v_i \right). \tag{6}$$

The final term is associated with high frequencies and small singular values $\rightarrow$ noise dominates at large i.

[1] Assume here the error on A $\ll$ measurement error.

A test problem

Gaussian smearing applied to $\mathbf{x}^{true} = e^{-t}$ via Monte Carlo sampling (10^4 events).

Simulates

- Poisson fluctuations in **b**
- Shape distortion
- Inefficiency

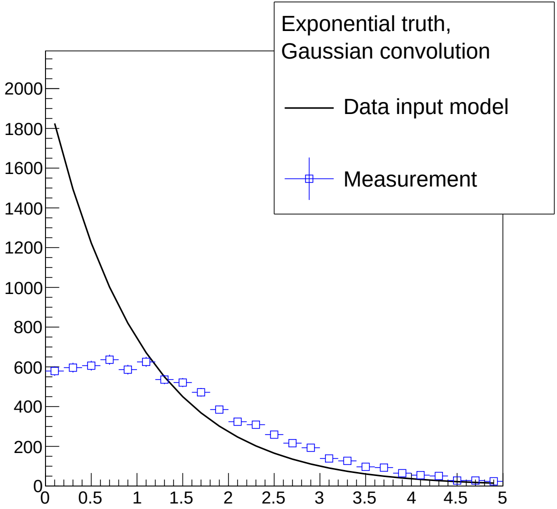

A test problem

Gaussian smearing applied to $\mathbf{x}^{true} = e^{-t}$ via Monte Carlo sampling (10^4 events).

Simulates

- Poisson fluctuations in **b**
- Shape distortion
- Inefficiency

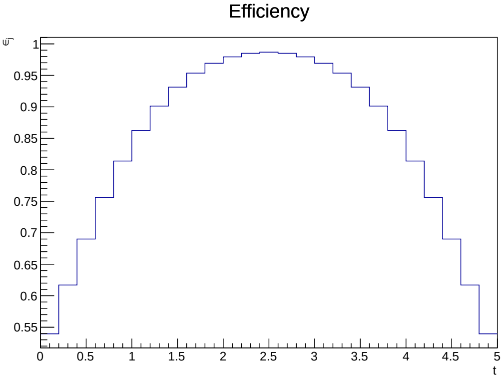

SVD analysis

for the Gaussian convolution example

$\mathbf{u}_i^T \cdot \mathbf{b}$ stops falling near $i = 6$, where $\frac{\mathbf{u}_i^T \cdot \mathbf{b}^{exact}}{\sigma_i} v_i \approx \frac{\mathbf{u}_i^T \cdot \mathbf{e}}{\sigma_i} v_i$.

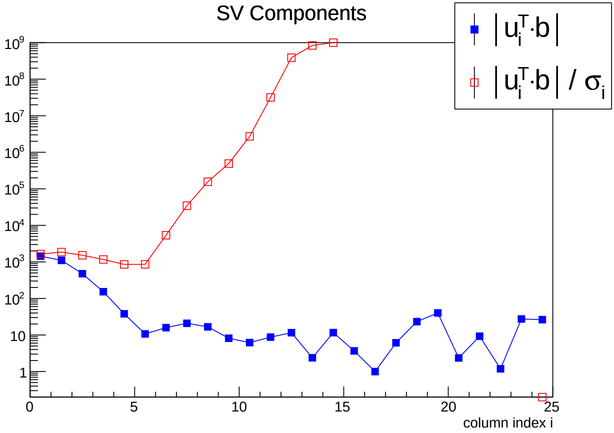

But the σ_i (shown earlier) continue falling steeply.
The noise-dominated $i > 6$ terms rapidly explode.

Including regularization

The general-form Tikhonov-Phillips problem

Add a term to the least squares problem (eq. 3) that includes a smoothing matrix L:

$$\min_{\mathbf{x}} \left\{ \|A\mathbf{x} - \mathbf{b}\|_2^2 + \lambda^2 \|L\mathbf{x}\|_2^2 \right\} \quad \textit{or} \quad \min_{\mathbf{x}} \left\| \begin{pmatrix} A \\ \lambda L \end{pmatrix} \mathbf{x} - \begin{pmatrix} \mathbf{b} \\ 0 \end{pmatrix} \right\|_2 \tag{7}$$

Examples of L include the $n \times n$ identity matrix and the second-order finite difference operator

$$L_2 = \begin{pmatrix} 1 & -2 & 1 & & \\ & \ddots & \ddots & \ddots & \\ & & 1 & -2 & 1 \end{pmatrix} \quad \in \mathbb{R}^{(n-2)\times n}. \tag{8}$$

The Generalized SVD

A joint decomposition of A and L

The GSVD of $A \in \mathbb{R}^{m\times n}$ and $L \in \mathbb{R}^{p\times n}$ is

$$A = UCX^{-1} \quad \textit{and} \quad L = VSX^{-1} \tag{9}$$

where $U \in \mathbb{R}^{m\times m}$, $V \in \mathbb{R}^{p\times p}$, and $X^{-1} \in \mathbb{R}^{n\times n}$ are orthogonal matrices, and[2]

$$C = \begin{pmatrix} I_{n-p} & 0 \\ 0 & \Sigma_p \\ 0 & 0 \end{pmatrix} \in \mathbb{R}^{m\times n} \quad \textit{and} \quad S = \begin{pmatrix} 0 & M_p \end{pmatrix} \in \mathbb{R}^{p\times n} \tag{10}$$

contain the generalized singular values

$$\begin{aligned} \Sigma_p &= diag(\alpha_i), \quad 1 \geq \alpha_1 \geq \alpha_2 \geq \ldots \alpha_p \geq 0 \\ M_p &= diag(\beta_i), \quad 0 \leq \beta_1 \leq \beta_2 \leq \ldots \beta_p \leq 1. \end{aligned}$$

[2] This ordering is reversed from some common conventions.

Spectral filtering in the GSVD basis

For Tikhonov's problem

Using the GSVD, the regularized solution is

$$x_{reg} = XF \begin{pmatrix} I_{n-p} & 0 \\ 0 & \Sigma^{\dagger} \end{pmatrix} U^T \mathbf{b} \tag{11}$$

where $\Sigma^{\dagger}$ is the Moore-Penrose pseudoinverse of Σ and F is a diagonal matrix containing 1's ($i = 1 \ldots n - p$) and the Tikhonov filter factors

$$f_i = \frac{\gamma_i^2}{\gamma_i^2 + \lambda^2}, \quad \gamma_i = \frac{\alpha_i}{\beta_i}, \quad i = n - p + 1 \ldots n. \tag{12}$$

In terms of the column vectors of U^T and X,

$$x_{reg} = \sum_{i=1}^{n-p} (\mathbf{u}_i^T \cdot \mathbf{b})\mathbf{x}_i + \sum_{i=n-p+1}^{n} f_i \frac{\mathbf{u}_i^T \cdot \mathbf{b}}{\alpha_i} \mathbf{x}_i \tag{13}$$

Analysis of the regularized problem

This time, expanded in the Generalized SVD basis ($\lambda = 0.38$)

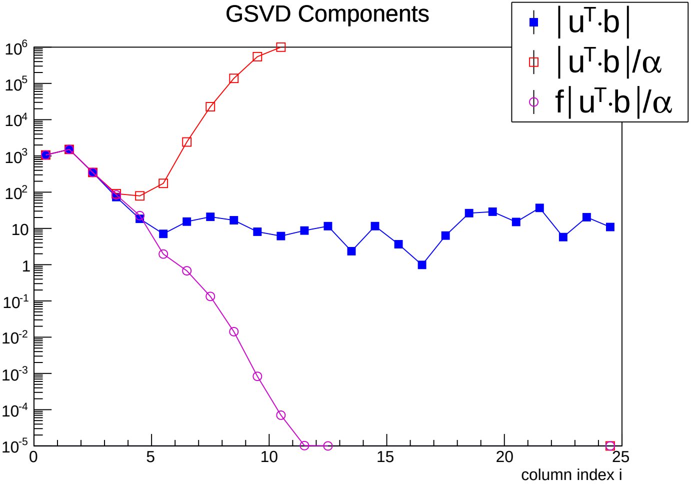

Applying Tikhonov filter factors f suppresses the noise-dominated coefficients.

GSVD solution

$\lambda = 0.38$, shown without uncertainties

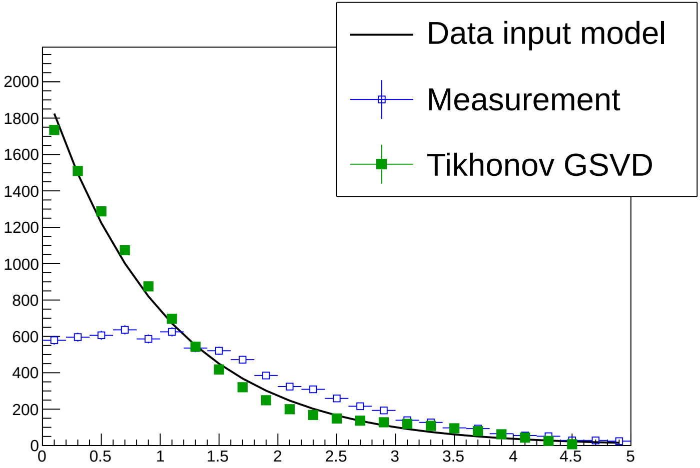

Not too bad, but is this the best parameter choice? How is λ selected?

Parameter scan

$\lambda = 0.01, 0.02, \ldots, 1.0$

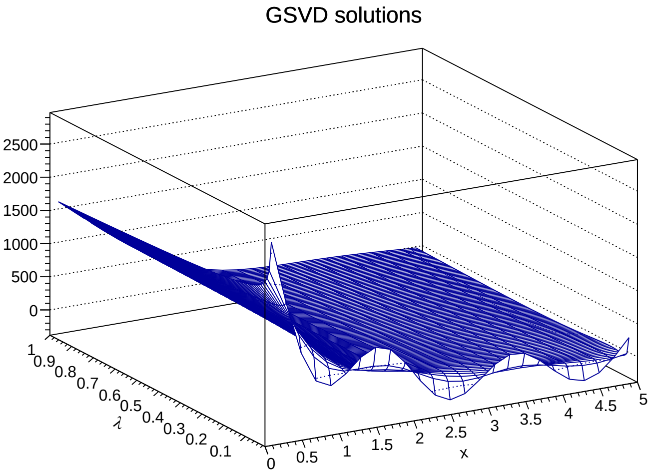

Solution evolves from under-smoothed to over-smoothed.

Generalized cross-validation

Tries to find λ such that $A\mathbf{x}_\lambda$ predicts $\mathbf{b}^{exact}$ as well as possible

Skipping the derivation, λ_{GCV} is the minimum of

$$G(\lambda) = \frac{\|A\mathbf{x} - \mathbf{b}\|_2^2}{(m - \sum_{i=1}^{n} f_i)^2} \tag{14}$$

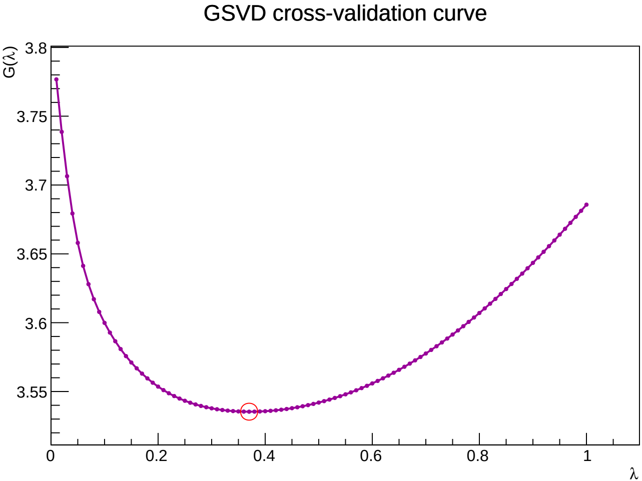

The L-Curve criterion

Parametric curve of curvature vs. agreement

Smoothing norm $\|L\mathbf{x}_\lambda\|_2$ vs. residual norm $\|A\mathbf{x}_\lambda - \mathbf{b}\|_2$.

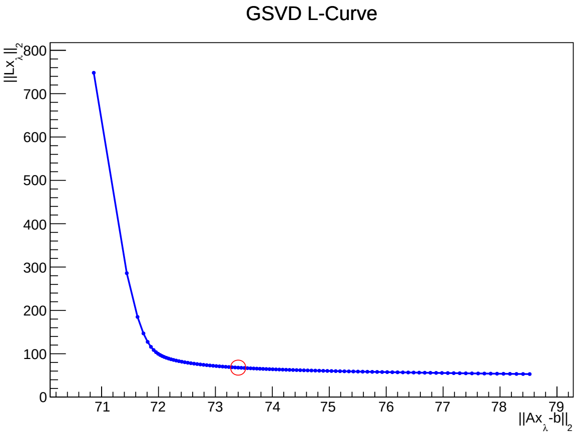

Moving up the vertical axis $\Leftrightarrow$ less regularization (smaller λ).
Moving across the horizontal axis $\Leftrightarrow$ more regularization (larger λ).
Optimal value occurs near kink. λ_{GCV} circled for reference.

More to discuss. . .

out of time

There are many more topics:

- Weighting **x** to improve regularization (often critical for spectra)
- Error propagation and covariance
- Iterative unfolding methods
- Uncertainty scaling (prewhitening)
- Solutions with discontinuities
- Handling boundary conditions
- Dealing with background

If interest, these can be presented in a future meeting.
In the meantime, the algorithms and examples are available:

```
git clone https://github.com/andrewadare/utils.git
git clone https://github.com/andrewadare/unfolding.git
```