

Argonne
NATIONAL
LABORATORY
... for a brighter future

U.S. Department of Energy
THE UNIVERSITY OF CHICAGO
Office of Science
U.S. DEPARTMENT OF ENERGY


A U.S. Department of Energy Laboratory managed by The University of Chicago

A Casual Introduction to Least-squares Fitting: A [mostly] descriptive approach

Brian H. Toby

Outline

- Linear Algebra: a cheap intro
- Least-Squares Minimization
 - Linear
 - Non-linear
- Least-square's weakness: Correlation
- Uncertainty estimation for fitted parameters
- Resistance: one bad point can do you in



Linear Algebra: for solution of simultaneous equations

- Linear Algebra provides a compact way to deal with simultaneous equations:

$$A_{11}x_1 + A_{12}x_2 + A_{13}x_3 + \dots + A_{1m}x_m = b_1$$


$$A_{21}x_1 + A_{22}x_2 + A_{23}x_3 + \dots + A_{2m}x_m = b_2$$

$$A_{n1}x_1 + A_{n2}x_2 + A_{n3}x_3 + \dots + A_{nm}x_m = b_n$$

or equivalently with n equations, $\sum_j A_{ij}x_j = b_i$, where we want to find the x_j values knowing A_{ij} and b_i

can be written as $\mathbf{A} \mathbf{x} = \mathbf{b}$ where

 - \mathbf{A} is a (n by m) matrix;
 - \mathbf{b} is a column vector (or m by 1 matrix)
 - \mathbf{x} is a row vector (or 1 by n matrix)
- Solving for \mathbf{x} : $\mathbf{A}^{-1}\mathbf{A} \mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$ or $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$




Linear Algebra

- Matrix \mathbf{A} with m rows and n columns is composed of $n \times m$ elements A_{ij} :

$$\mathbf{A} = \begin{pmatrix} A_{11} & A_{21} & \dots & A_{m1} \\ A_{12} & A_{22} & \dots & A_{m2} \\ \dots & \dots & \dots & \dots \\ A_{1n} & A_{2n} & \dots & A_{mn} \end{pmatrix}$$
- Matrix multiplication: $\mathbf{C} = \mathbf{A} \mathbf{B}$, $C_{ij} = \sum_k A_{ik} B_{kj}$
 Note that in general, $\mathbf{A} \mathbf{B} \neq \mathbf{B} \mathbf{A}$,

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} = \begin{pmatrix} a\alpha + b\gamma & a\beta + b\delta \\ c\alpha + d\gamma & c\beta + d\delta \end{pmatrix}$$
- Matrix transpose, \mathbf{A}^T
 - if $\mathbf{B} = \mathbf{A}^T$ then $B_{ij} = A_{ji}$
$$\mathbf{A} = \begin{pmatrix} A_{11} & A_{21} & A_{31} \\ A_{12} & A_{22} & A_{32} \end{pmatrix} \quad \mathbf{A}^T = \mathbf{B} = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \\ A_{31} & A_{32} \end{pmatrix}$$



Matrix Inversion

- Identity Matrix:
 - diagonal elements = 1
 - off-diagonal elements = 0

$$\mathbf{1} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{pmatrix}$$

- Inverse of Matrix: $\mathbf{A}^{-1} \mathbf{A} = \mathbf{1}$

$$\mathbf{A}^{-1} \mathbf{A} = \begin{pmatrix} A_{11}^{-1} & A_{12}^{-1} & A_{13}^{-1} \\ A_{21}^{-1} & A_{22}^{-1} & A_{23}^{-1} \\ A_{31}^{-1} & A_{32}^{-1} & A_{33}^{-1} \end{pmatrix} \begin{pmatrix} A_{11} & A_{21} & A_{31} \\ A_{21} & A_{22} & A_{23} \\ A_{31} & A_{32} & A_{33} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

- Inverse of 3x3 matrix (from wikipedia)

$$\mathbf{A}^{-1} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix}^{-1} = \frac{1}{|\mathbf{A}|} \begin{bmatrix} ei - fh & ch - bi & bf - ce \\ fg - di & ai - cg & cd - af \\ dh - eg & bg - ah & ae - bd \end{bmatrix}$$

$$|\mathbf{A}| = a(ei - fh) - b(di - fg) + c(dh - eg)$$

Summary: Part 1

- You have now had a very brief introduction to linear algebra and should understand the concept of a matrix

Singular Matrices

- If any column (or row) in a matrix is repeated, the matrix cannot be inverted. The same is true if a column (or row) is repeated multiplied by a constant
- A matrix that cannot be inverted is called **singular**

$$\mathbf{A} = \begin{pmatrix} a & b & na \\ d & e & nd \\ g & h & ng \end{pmatrix}^{-1} \quad |\mathbf{A}| = \begin{vmatrix} a & b & na \\ d & e & nd \\ g & h & ng \end{vmatrix} = a(cng - ndh) - b(dng - ndg) + na(dh - eg) = 0$$

$$\mathbf{A}^{-1} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix}^{-1} = \frac{1}{|\mathbf{A}|} \begin{bmatrix} ei - fh & ch - bi & bf - ce \\ fg - di & ai - cg & cd - af \\ dh - eg & bg - ah & ae - bd \end{bmatrix}$$

$$|\mathbf{A}| = a(ei - fh) - b(di - fg) + c(dh - eg)$$

Nearly Singular Matrices

- When columns are nearly equivalent, we start subtracting numbers that are almost equal from each other.
 - This is a very bad thing in computer math as it causes round-off errors to be increased.

Round-off error example

$$\infty\text{-precision arithmetic: } 64 \left(\frac{1}{8}\right) \left(\frac{1}{8}\right) - 1 = 0$$

$$\text{Repeat with two significant-figures: } 64(0.13)(0.13) - 1 = 64(0.017) - 1 = 1.1 - 1 = 0.1$$

$$\mathbf{A} = \begin{pmatrix} a & b & na \\ d & e & nd \\ g & h & n(g + \delta) \end{pmatrix}^{-1} \quad |\mathbf{A}| = a(en[g + \delta] - ndh) - b(dn[g + \delta] - ndg) + na(dh - eg)$$

Summary: Part 2

- You should now understand that a singular matrix is one that cannot be inverted
- A matrix that is nearly singular in theory can be inverted, but in practice inversion will be highly inaccurate due to round-off errors

Terminology of Least-Squares

Data: n observations, y_i , measured at independent variable setting x_i

Model: a function that predicts the observations: $Y(x_i, \mathbf{p})$

- Linear Model: $Y(x_i, \mathbf{p}) = p_1 f_1(x_i) + p_2 f_2(x_i) + \dots$
- Non-linear Model: $Y(x_i, \mathbf{p}) = f(x_i, p_1, p_2, \dots)$

Parameters: m terms $p_1, p_2, p_3, \dots, p_m$ that determine the values that are computed from the model

Refine: Find values for parameters, \mathbf{p} , to yield the **best fit** between the model $Y(x_i, \mathbf{p})$ and observations y_i

Best fit: Means the finding the minimum for $\sum w_i [y_i - Y(x_i, \mathbf{p})]^2$ where $w_i = [1 / \sigma(y_i)]^2$ (Note: σ is standard uncertainty on y_i)

Linear Least-Squares

Linear Model: $Y(x_i, \mathbf{p}) = p_1 f_1(x_i) + p_2 f_2(x_i) + \dots = \sum_k p_k f_k(x_i)$

- Goal: Find $p_1, p_2, p_3, \dots, p_m$ that minimize $\sum_i w_i [y_i - Y(x_i, \mathbf{p})]^2$

set derivative w/r each parameter to zero: $\partial / \partial p_j \sum_i w_i [y_i - Y(x_i, \mathbf{p})]^2 = 0$

Gives m coupled equations: $\sum_i w_i y_i \partial Y / \partial p_j = \sum_i w_i Y(x_i, \mathbf{p}) \partial Y / \partial p_j$

Note that $\partial Y / \partial p_j = f_j(x_i)$ so the m coupled equations become:

$$\sum_i w_i y_i f_j(x_i) = \sum_i w_i [\sum_k p_k f_k(x_i)] f_j(x_i) = \sum_k p_k \sum_i w_i f_k(x_i) f_j(x_i)$$

Define: $A_{ij} = f_j(x_i) / \sigma(y_i)$; $b_i = y_i / \sigma(y_i)$

This gives m coupled equations: $\sum_i b_i A_{ij} = \sum_k p_k \sum_i A_{ij} A_{ik}$

Recast using linear algebra: $\mathbf{A}^T \mathbf{b} = \mathbf{A}^T \mathbf{A} \mathbf{p}$ or solving for \mathbf{p} :

$$(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b} = \mathbf{p} \quad \text{This allows the } \mathbf{p} \text{ values to be determined directly}$$

Non-Linear Least-Squares (Gauss-Newton)

- With a non-linear model, $Y(x_i, \mathbf{p}) = f(x_i, p_1, p_2, \dots)$, it is not possible to solve for \mathbf{p}

- Remembering the Taylor expansion:

$$f(x_i, \mathbf{p} + \delta) = f(x_i, \mathbf{p}) + \delta (\partial f / \partial \mathbf{p}) + \delta^2 (\partial^2 f / \partial \mathbf{p}^2) / 2 + \dots$$

- Multi-parameter Taylor expansion around approximate values for \mathbf{p} :

$$Y(x_i, p_1 + \delta_1, p_2 + \delta_2, \dots) = Y(x_i, p_1, p_2, \dots) + \sum_k \delta_k (\partial Y / \partial p_k) + \sum_k \delta_k^2 (\partial^2 Y / \partial p_k^2) / 2 + \dots$$

- as before, set $\partial / \partial p_j \sum_i w_i [y_i - Y(x_i, \mathbf{p})]^2 = 0$; solve for δ_k

$$m \text{ coupled equations: } \sum_i w_i [y_i - Y(x_i, \mathbf{p})] (\partial Y / \partial p_j) = \sum_k \delta_k \sum_i w_i (\partial Y / \partial p_k) (\partial Y / \partial p_j)$$

Define: $A_{ij} = (\partial Y(x_i, \mathbf{p}) / \partial p_j) / \sigma(y_i)$; $b_i = [y_i - Y(x_i, \mathbf{p})] / \sigma(y_i)$

This gives m coupled equations: $\sum_i b_i A_{ij} = \sum_k \delta_k \sum_i A_{ij} A_{ik}$

Recast using linear algebra: $\mathbf{A}^T \mathbf{b} = \mathbf{A}^T \mathbf{A} \delta$ or $(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b} = \delta$

Refinement is iterative process, starting from approximate \mathbf{p} values

More on Least-squares

- A is called the **Design Matrix**: $A_{ij} = \partial Y / \partial p_j / \sigma(y_i)$
- $H = A^T A$ is called the **Hessian Matrix**
- The inverse of the Hessian, $H^{-1} = (A^T A)^{-1}$, is called the **Covariance Matrix** (Einstein called it the **Variance-Covariance Matrix**)
- The Hessian measures, evaluated for all data points, how the model responds to changes in parameters :
 - $H_{ij} = \sum_k (\partial Y / \partial p_j) (\partial Y / \partial p_i) / \sigma(y_k)$

Summary: Part 3

You should now understand

- the difference between linear and non-linear least squares
- why non-linear least squares is iterative and requires starting with approximate values for parameters
- how LS refinement uses weights, differences and depends of the derivatives of w/r to parameters
- commonly used terms: covariance matrix, Hessian matrix

Correlation:

- If two (or more) parameters have the same effect on the model, the derivatives are the same and the Hessian is singular
- If two (or more) parameters have very similar effects, the derivatives are nearly the same and the Hessian is nearly singular -- round-off dominates!

When parameters have similar effects on the fit they are said to be correlated

Correlation: the Achilles' Heel of Least-Squares


- Least-squares works best with parameters that have very different effects on the model
- If parameters have exactly the same effect on the model (are completely correlated), the p values cannot be determined
- Least-squares performs poorly when p_i values have similar effects (are correlated.)

So why use Least-Squares?

- If $\sigma(y_i)$ accurately describes the estimated error (standard uncertainty) in y_i and the model produces an ideal fit to the data ($\chi^2 \approx 1$) then the diagonal elements of the covariance matrix give the standard uncertainty in the parameters: $(A^T A)^{-1}_{jj} = \sigma(p_j)$
- Least-squares makes optimum use of data -- giving the result with the smallest possible statistically uncertainty.

Correlation: Example 1

Fit $y = \underline{n}\sin(x) + \underline{m}x + \underline{b}$ (all $x \ll \pi/2$)



$y = \underline{m}x + \underline{b}$



$y = \underline{n}\sin(x) + \underline{b}$

Note that \underline{m} and \underline{n} have almost exactly the same effect
The least-squares refinement will be prone to diverge!



In order to fit both \underline{m} and \underline{n} well: data over a wide range in x or extremely precise data are needed

Correlation: Example 2

Occupancies and displacement parameters correlate.

Why? Decreasing occupancy has a similar effect to increasing U_{iso}



- Example: Simulate change in NaCl X-ray diffraction due to 20% Na vacancies or due to a $\times 2.8$ increase in Na U_{iso} .




Note: with a smaller Q range, these changes would be even harder to discern

Correlation: Neutrons vs. X-rays

Same example: NaCl with 20% Na vacancies or $\times 2.8$ increase in Na U_{iso} but now simulated with neutron diffraction (can discern much better)



Correlation: Example 3



Fit $y = \underline{m}x + \underline{b} + \underline{c}$

Note: \underline{b} & \underline{c} are completely equivalent

Since the effect of changing \underline{b} & \underline{c} is exactly the same, the Hessian matrix cannot be inverted (is singular).

GSAS treats a singular Hessian by ignoring one of the two identical variables.

Exact correlation: Crystallographic Examples

- When symmetry is lowered, there will be complete correlation between:
 - [Formerly] equivalent unit cell constants
 - Sets of [formerly] equivalent atoms
 - Either manually change the parameters to break the equivalence or vary only one of the set to start.
- Vacancies are equivalent to partial substitution by a “lighter” atom
- Refining all atom positions in space groups with only translational symmetry (arbitrary origin)
- Complete correlation occurs any time two (or more) parameters have exactly the same effect on the fit.

Summary: Part 4

You should now understand

- how exact correlation in parameters leads to a singular Hessian
- why highly correlated parameters leads to a very inaccurate inversion of the Hessian – possibly causing a refinement to fail.

Uncertainty estimation for derived parameters


- Statistical error estimates are computed using covariance matrix
 - If $\sigma(y_i)$ accurately describes the estimated error (standard uncertainty) in y_i , and the model produces an ideal fit to the data ($\chi^2=1$) then the diagonal elements of the covariance matrix give the standard uncertainty in the parameters: $\mathbf{H}^{-1}_{ij} = [\sigma(p_i)]^2$
- For functions of fitted parameters, uncertainty also computed:
 - If $s = \mathbf{f} \mathbf{p}$ then $\sigma(s) = (\mathbf{H}^{-1})^T \mathbf{f} \mathbf{H}^{-1}$
 - Used for bond distances & angles (DISAGL)
 - Can also be used for total composition (from refined occupancies, implemented in GEOMETRY)

Least-Squares is not Resistant

Least-squares weighting assumes:

- uncertainty estimates on data are accurate
- model is accurate (no systematic errors)

“Bad” points skew refinements



Robust-Resistant algorithms limit the maximum leverage a poorly fitting data point may have (for example, by changing weighting.)

Final Summary

You have now seen

- the strength of least-squares: error estimates for refined parameters
- that weights need to reflect the actual uncertainty on a observation or “bad data” can yield a bad fit.

In conclusion:

- Linear algebra simplifies least squares fitting
- Understand how least squares fitting works
- Understand the strengths and weaknesses of least squares