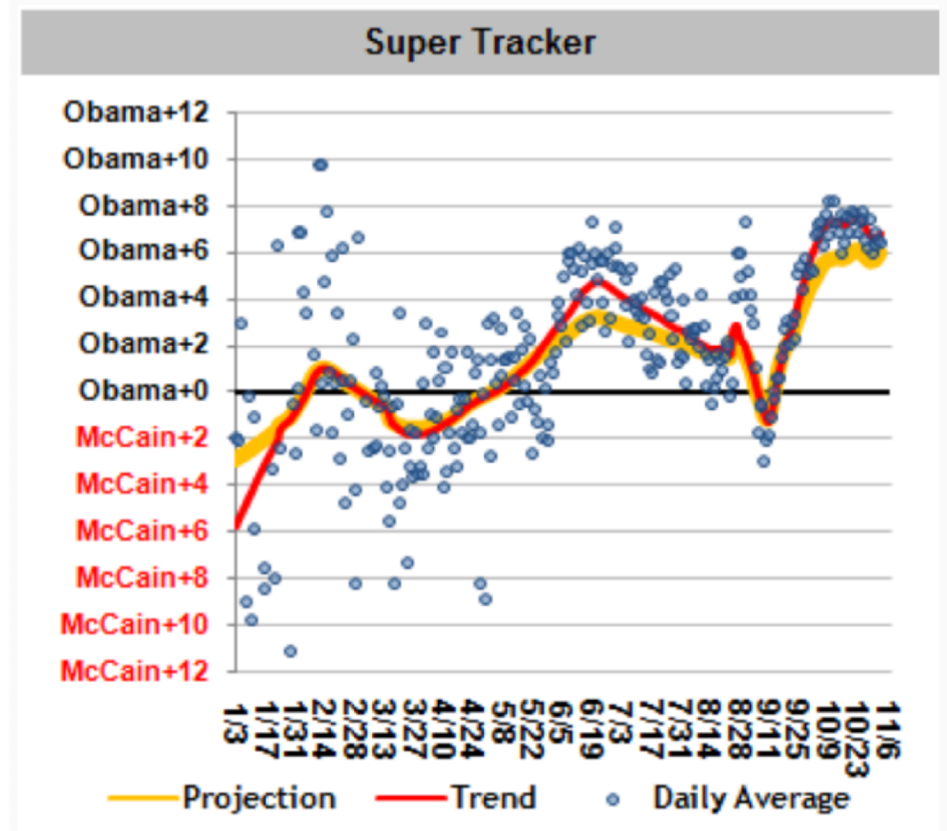
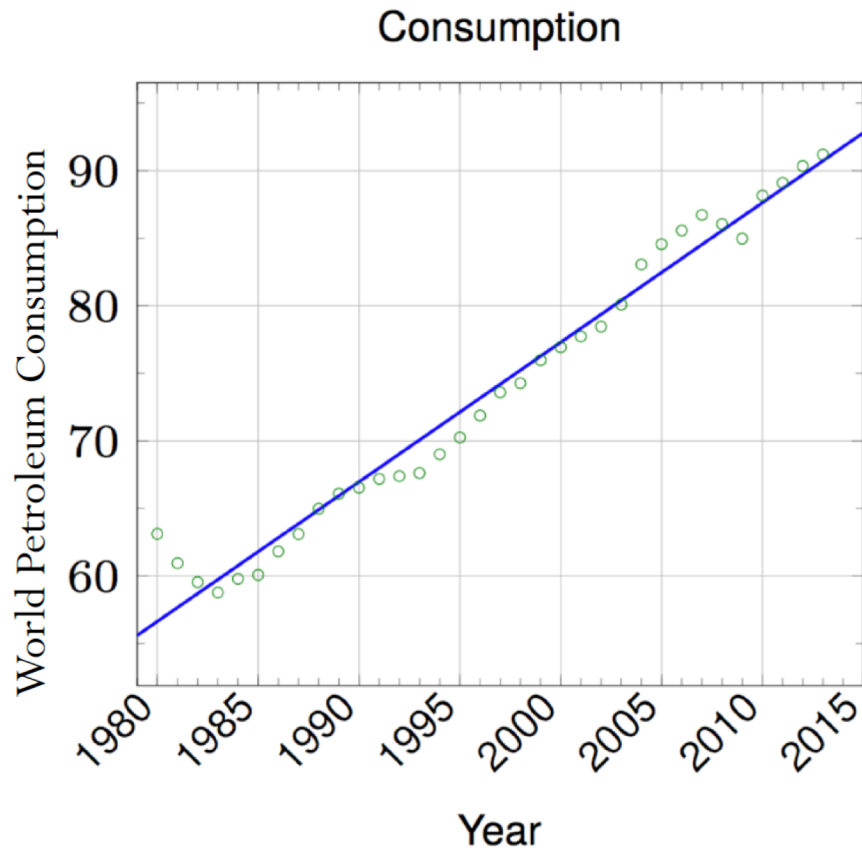


Least Squares and Data Fitting

Data fitting

How do we best fit a set of data points?



Linear Least Squares – Fitting with a line

Given m data points $\{\{t_1, y_1\}, \dots, \{t_m, y_m\}\}$, we want to find the function

$$y = \alpha + \beta t$$

that best fit the data (or better, we want to find the parameters α, β).

Thinking geometrically, we can think "what is the line that most nearly passes through all the points?"

Linear Least Squares

$$\begin{array}{ccc} \begin{bmatrix} 1 & t_1 \\ \vdots & \vdots \\ 1 & t_m \end{bmatrix} & \begin{bmatrix} \alpha \\ \beta \end{bmatrix} & = & \begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix} & \mathbf{A} \mathbf{x} = \mathbf{b} \\ \mathbf{m \times n} & \mathbf{n \times 1} & & \mathbf{m \times 1} & \end{array}$$

- We want to find the appropriate linear combination of the columns of \mathbf{A} that makes up the vector \mathbf{b} .
- If a solution exists that satisfies $\mathbf{A} \mathbf{x} = \mathbf{b}$ then $\mathbf{b} \in \text{range}(\mathbf{A})$
- In most cases, $\mathbf{b} \notin \text{range}(\mathbf{A})$ and $\mathbf{A} \mathbf{x} = \mathbf{b}$ **does not have an exact solution!**
- Therefore, an overdetermined system is better expressed as

$$\mathbf{A} \mathbf{x} \cong \mathbf{b}$$

Linear Least Squares

- Find $\mathbf{y} = \mathbf{A} \mathbf{x}$ which is closest to the vector \mathbf{b}
- What is the vector $\mathbf{y} = \mathbf{A} \mathbf{x} \in \text{range}(\mathbf{A})$ that is closest to vector \mathbf{y} in the Euclidean norm?

Linear Least Squares

- **Least Squares:** find the solution \mathbf{x} that minimizes the residual

$$\mathbf{r} = \mathbf{b} - \mathbf{A} \mathbf{x}$$

- Let's define the function ϕ as the square of the 2-norm of the residual

$$\phi(\mathbf{x}) = \|\mathbf{b} - \mathbf{A} \mathbf{x}\|_2^2$$

- Then the least squares problem becomes

$$\min_{\mathbf{x}} \phi(\mathbf{x})$$

- Suppose $\phi: \mathcal{R}^m \rightarrow \mathcal{R}$ is a smooth function, then $\phi(\mathbf{x})$ reaches a (local) maximum or minimum at a point $\mathbf{x}^* \in \mathcal{R}^m$ only if

$$\nabla \phi(\mathbf{x}^*) = 0$$

How to find the minimizer?

- To minimize the 2-norm of the residual vector

$$\min_{\mathbf{x}} \phi(\mathbf{x}) = \|\mathbf{b} - \mathbf{A} \mathbf{x}\|_2^2 \quad \text{or} \quad \phi(\mathbf{x}) = (\mathbf{b} - \mathbf{A} \mathbf{x})^T (\mathbf{b} - \mathbf{A} \mathbf{x})$$

Summary:

- \mathbf{A} is a $m \times n$ matrix, where $m > n$.
- m is the number of data pair points. n is the number of parameters of the “best fit” function.
- Linear Least Squares problem $\mathbf{A} \mathbf{x} \cong \mathbf{b}$ *always* has solution.
- The Linear Least Squares solution \mathbf{x} minimizes the square of the 2-norm of the residual:

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

- One method to solve the minimization problem is to solve the system of **Normal Equations**

$$\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}$$

- Let's see some examples and discuss the limitations of this method.


Example:

Demo: “Fit a line - Least Squares example”

$$\begin{bmatrix} 1 & t_1 \\ \vdots & \vdots \\ 1 & t_m \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \cong \begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix}$$

5 × 2 2 × 1 5 × 1

Solve: $\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}$




Data fitting - not always a line fit!

- Does not need to be a line! For example, here we are fitting the data using a quadratic curve.
- Linear Least Squares:

The problem is **linear in its coefficients!**

Which function is not suitable for linear least squares?

- A) $y = a + b x + c x^2 + d x^3$
- B) $y = x(a + b x + c x^2 + d x^3)$
- C) $y = a \sin(x) + b / \cos(x)$
- D) $y = a \sin(x) + x / \cos(bx)$
- E) $y = a e^{-2x} + b e^{2x}$




More examples

Demo “Make some noise”

We want to find the coefficients of the quadratic function that best fits the data points:

$$y = x_0 + x_1 t + x_2 t^2$$



The data points were generated by adding random noise to the function


$$f(t) = 0.8 - t + t^2$$

We would not want our “fit” curve to pass through the data points exactly as we are looking to model the general trend and not capture the noise.

Data fitting

$$\begin{bmatrix} 1 & t_1 & t_1^2 \\ \vdots & \vdots & \vdots \\ 1 & t_m & t_m^2 \end{bmatrix} \begin{bmatrix} x_0 \\ x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix}$$

Solve: $\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}$



Computational Cost





$$\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}$$

- Compute $\mathbf{A}^T \mathbf{A}$: $O(mn^2)$
- Factorize $\mathbf{A}^T \mathbf{A}$: LU $\rightarrow O\left(\frac{2}{3}n^3\right)$, Cholesky $\rightarrow O\left(\frac{1}{3}n^3\right)$
- Solve $O(n^2)$
- Since $m > n$ the overall cost is $O(mn^2)$

Short questions

Given the data in the table below, which of the plots shows the line of best fit in terms of least squares?

x	1	2	5
y	2	18	12



Short questions

Given the data in the table below, and the least squares model

$$y = c_1 + c_2 \sin(t\pi) + c_3 \sin(t\pi/2) + c_4 \sin(t\pi/4)$$

written in matrix form as

$$A \begin{bmatrix} c_1 \\ c_2 \\ c_3 \\ c_4 \end{bmatrix} \cong \mathbf{y}$$

Write the matrix A

t_i	y_i
0.5	0.72
1.0	0.79
1.5	0.72
2.0	0.97
2.5	1.03
3.0	0.96
3.5	1.00

Demo: “Ice example”

Condition number for Normal Equations

Finding the least square solution of $\mathbf{A} \mathbf{x} \cong \mathbf{b}$ (where \mathbf{A} is full rank matrix) using the Normal Equations

$$\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}$$

has some advantages, since we are solving a square system of linear equations with a symmetric matrix (and hence it is possible to use decompositions such as Cholesky Factorization)

However, the normal equations tend to worsen the conditioning of the matrix.

$$\text{cond}(\mathbf{A}^T \mathbf{A}) = (\text{cond}(\mathbf{A}))^2$$

How can we solve the least square problem without squaring the condition of the matrix?

Rank of a matrix

Suppose \mathbf{A} is a $m \times n$ rectangular matrix where $m > n$:

$$\mathbf{A} = \begin{pmatrix} \vdots & \dots & \vdots & \dots & \vdots \\ \mathbf{u}_1 & \dots & \mathbf{u}_n & \dots & \mathbf{u}_m \\ \vdots & \dots & \vdots & \dots & \vdots \end{pmatrix} \begin{pmatrix} \sigma_1 & & & & \\ & \ddots & & & \\ & & \sigma_n & & \\ & & 0 & & \\ & & \vdots & & \\ & & 0 & & \end{pmatrix} \begin{pmatrix} \dots & \mathbf{v}_1^T & \dots \\ \vdots & \vdots & \vdots \\ \dots & \mathbf{v}_n^T & \dots \end{pmatrix}$$

$$\mathbf{A} = \begin{pmatrix} \vdots & \dots & \vdots \\ \mathbf{u}_1 & \dots & \mathbf{u}_n \\ \vdots & \dots & \vdots \end{pmatrix} \begin{pmatrix} \dots & \sigma_1 \mathbf{v}_1^T & \dots \\ \vdots & \vdots & \vdots \\ \dots & \sigma_n \mathbf{v}_n^T & \dots \end{pmatrix} = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T + \sigma_2 \mathbf{u}_2 \mathbf{v}_2^T + \dots + \sigma_n \mathbf{u}_n \mathbf{v}_n^T$$

Rank of a matrix

- The rank of \mathbf{A} equals the number of non-zero singular values which is the same as the number of non-zero diagonal elements in $\mathbf{\Sigma}$.
- Rounding errors may lead to small but non-zero singular values in a rank deficient matrix, hence the rank of a matrix determined by the number of non-zero singular values is sometimes called “effective rank”.
- The right-singular vectors (columns of \mathbf{V}) corresponding to vanishing singular values span the null space of \mathbf{A} .
- The left-singular vectors (columns of \mathbf{U}) corresponding to the non-zero singular values of \mathbf{A} span the range of \mathbf{A} .

Back to least squares...

Normal Equations: $\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}$

SVD to solve linear least squares problems

\mathbf{A} is a $m \times n$ rectangular matrix where $m > n$, and hence the SVD decomposition is given by:

$$\mathbf{A} = \begin{pmatrix} \vdots & \dots & \vdots \\ \mathbf{u}_1 & \dots & \mathbf{u}_m \\ \vdots & \dots & \vdots \end{pmatrix} \begin{pmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_n \\ & & 0 \\ & & \vdots \\ & & 0 \end{pmatrix} \begin{pmatrix} \dots & \mathbf{v}_1^T & \dots \\ \vdots & \vdots & \vdots \\ \dots & \mathbf{v}_n^T & \dots \end{pmatrix}$$

We want to find the least square solution of $\mathbf{A} \mathbf{x} \cong \mathbf{b}$, where $\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$


or better expressed in reduced form: $\mathbf{A} = \mathbf{U}_R \mathbf{\Sigma}_R \mathbf{V}^T$

Recall Reduced SVD $m > n$

$$A = U_R \Sigma_R V^T$$

$m \times n$ $m \times n$ $n \times n$ $n \times n$

$n \times n$



Shapes of the Reduced SVD

Suppose you compute a reduced SVD $A = U\Sigma V^T$ of a 10×14 matrix A . What will the shapes of U , Σ , and V be?

Hint: Remember the transpose on V !

The shape of U will be \times .

The shape of Σ will be \times .


The shape of V will be \times .


SVD to solve linear least squares problems

$$\mathbf{A} = \mathbf{U}_R \mathbf{\Sigma}_R \mathbf{V}^T$$

$$\mathbf{A} = \begin{pmatrix} \vdots & \dots & \vdots \\ \mathbf{u}_1 & \dots & \mathbf{u}_n \\ \vdots & \dots & \vdots \end{pmatrix} \begin{pmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_n \end{pmatrix} \begin{pmatrix} \dots & \mathbf{v}_1^T & \dots \\ \vdots & \vdots & \vdots \\ \dots & \mathbf{v}_n^T & \dots \end{pmatrix}$$

We want to find the least square solution of $\mathbf{A} \mathbf{x} \cong \mathbf{b}$, where $\mathbf{A} = \mathbf{U}_R \mathbf{\Sigma}_R \mathbf{V}^T$








Solving Least Squares Problem with SVD (summary)

Cost of SVD:

- Find \mathbf{x} that satisfies $\min_{\mathbf{x}} \|\mathbf{b} - \mathbf{A} \mathbf{x}\|_2^2$
- Find \mathbf{y} that satisfies $\min_{\mathbf{y}} \|\boldsymbol{\Sigma}_R \mathbf{y} - \mathbf{U}_R^T \mathbf{b}\|_2^2$
- Propose \mathbf{y} that is solution of $\boldsymbol{\Sigma}_R \mathbf{y} = \mathbf{U}_R^T \mathbf{b}$

Cost:

- Evaluate: $\mathbf{z} = \mathbf{U}_R^T \mathbf{b}$ 
- Set: $y_i = \begin{cases} \frac{z_i}{\sigma_i}, & \text{if } \sigma_i \neq 0 \\ 0, & \text{otherwise} \end{cases} \quad i = 1, \dots, n$ 
- Then compute $\mathbf{x} = \mathbf{V} \mathbf{y}$ 

Solving Least Squares Problem with SVD (summary)

- If $\sigma_i \neq 0$ for $\forall i = 1, \dots, n$, then the solution $\mathbf{y} = \mathbf{V} (\boldsymbol{\Sigma}_R)^{-1} \mathbf{U}_R^T \mathbf{b}$ is unique (and not a “choice”).
- If at least one of the singular values is zero, then the proposed solution \mathbf{y} is the one with the smallest 2-norm ($\|\mathbf{y}\|_2$ is minimal) that minimizes the 2-norm of the residual $\|\boldsymbol{\Sigma}_R \mathbf{y} - \mathbf{U}_R^T \mathbf{b}\|_2$
- Since $\|\mathbf{x}\|_2 = \|\mathbf{V} \mathbf{y}\|_2 = \|\mathbf{y}\|_2$, then the solution \mathbf{x} is also the one with the smallest 2-norm ($\|\mathbf{x}\|_2$ is minimal) for all possible \mathbf{x} for which $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ is minimal.

Pseudo-Inverse

- **Problem:** Σ may not be invertible
- **How to fix it:** Define the Pseudo Inverse
- **Pseudo-Inverse of a diagonal matrix:**

$$(\Sigma^+)_i = \begin{cases} \frac{1}{\sigma_i}, & \text{if } \sigma_i \neq 0 \\ 0, & \text{if } \sigma_i = 0 \end{cases}$$

- **Pseudo-Inverse of a matrix A :**

$$A^+ = V\Sigma^+U^T$$

Solving Least Squares Problem with SVD (summary)

Solve $\mathbf{A} \mathbf{x} \cong \mathbf{b}$ or $\mathbf{U}_R \boldsymbol{\Sigma}_R \mathbf{V}^T \mathbf{x} \cong \mathbf{b}$

$$\mathbf{x} \cong \mathbf{V} (\boldsymbol{\Sigma}_R)^+ \mathbf{U}_R^T \mathbf{b}$$

Example:

Consider solving the least squares problem $\mathbf{A} \mathbf{x} \cong \mathbf{b}$, where the singular value decomposition of the matrix $\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T \mathbf{x}$ is:

$$\begin{bmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} & 0 & 0 \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 14 & 0 & 0 \\ 0 & 14 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \mathbf{x} \cong \begin{bmatrix} 12 \\ 9 \\ 9 \\ 10 \end{bmatrix}$$

Determine $\|\mathbf{b} - \mathbf{A} \mathbf{x}\|_2$

Clicker question

Suppose you have $\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$ calculated. What is the cost of solving

$$\min_{\mathbf{x}} \|\mathbf{b} - \mathbf{A} \mathbf{x}\|_2^2 ?$$

- A) $O(n)$
- B) $O(n^2)$
- C) $O(mn)$
- D) $O(m)$
- E) $O(m^2)$