


Image Processing for Physicists

Prof. Pierre Thibault

pthibault@units.it

Chosen topics
in computational imaging



Overview


- Computational imaging
- Optimization
 - Characterization of linear systems
 - Standard optimization methods (in 1-D and N-D)
- Complex-valued variables: Wirtinger calculus
-

Computational imaging

1D optimization

- Bracketing

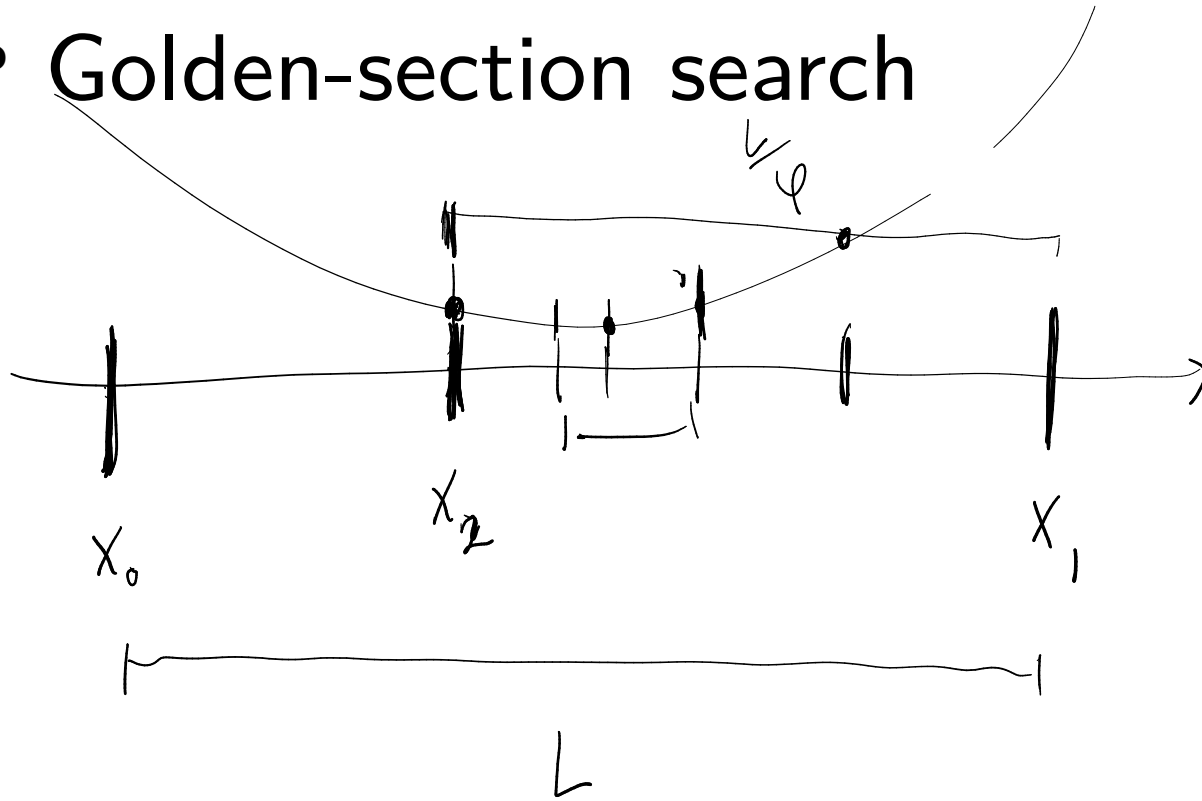
* commonly applied
on $f(x)$ costly to
evaluate, gradient
unknown or also costly



1D optimization

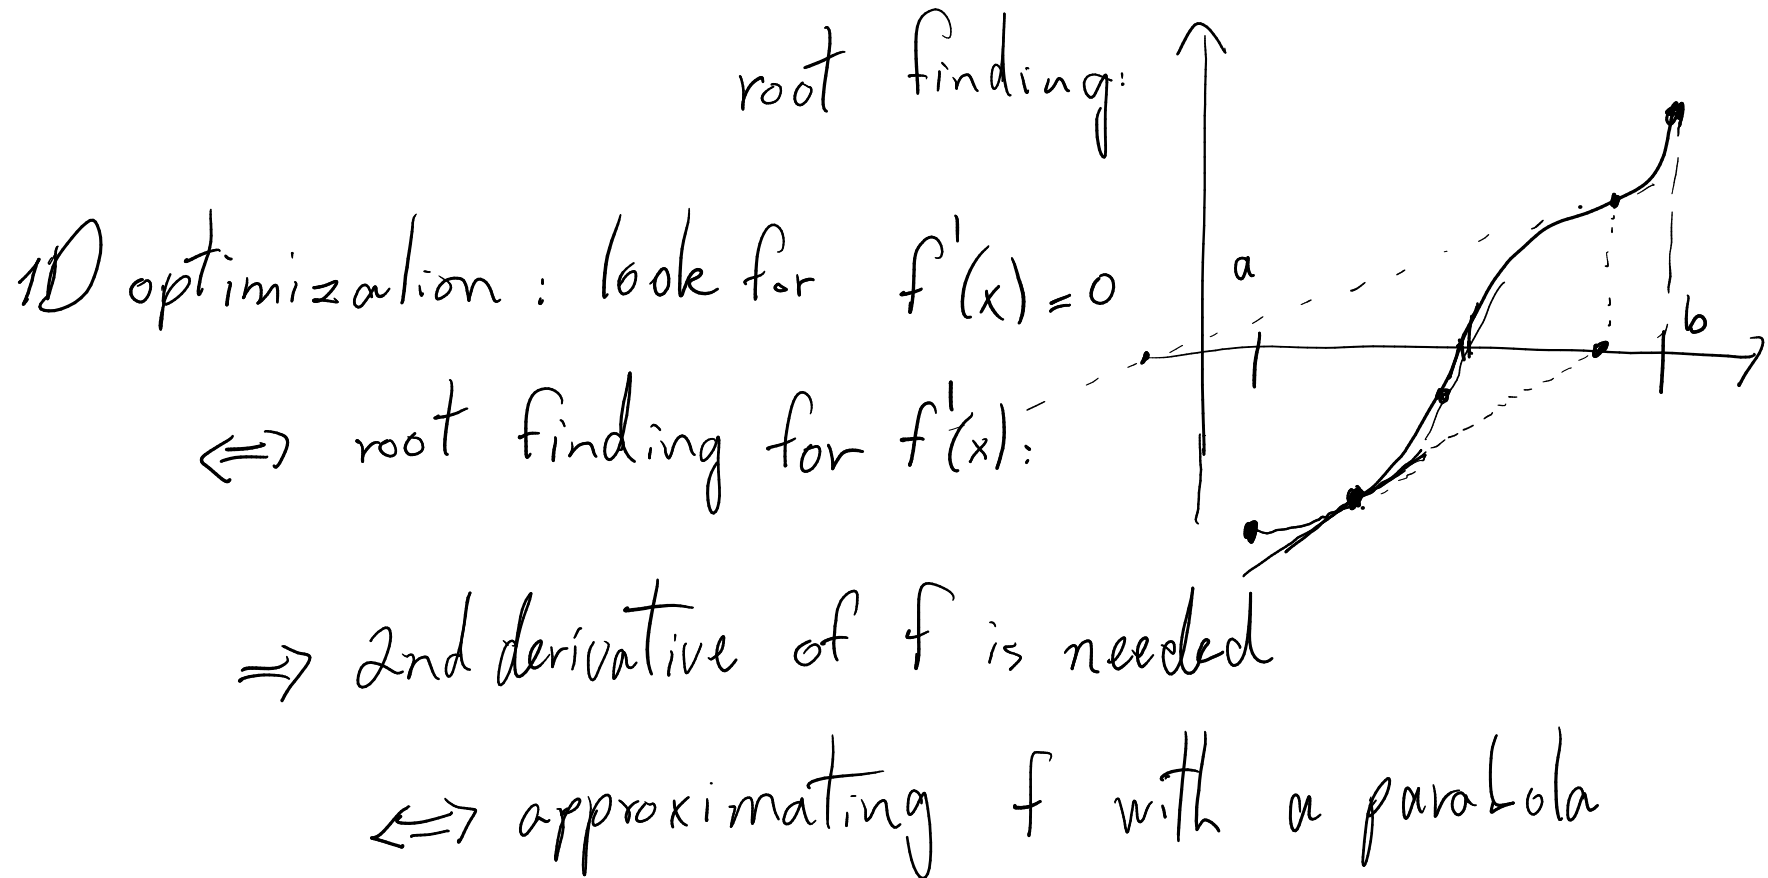
- Golden-section search

$$\varphi = \frac{\sqrt{5} + 1}{2} \approx 1.6181\dots$$




1D optimization

- Newton's method ("Newton-Raphson")



1D optimization

- Newton's method ("Newton-Raphson")



$$f(x) \approx f(x_0) + f'(x_0)(x-x_0) + \frac{1}{2} f''(x_0)(x-x_0)^2$$

$$f' = 0 \Leftrightarrow f'(x_0) + f''(x_0)(x-x_0) = 0$$

$$\Rightarrow x = x_0 - \frac{f'(x_0)}{f''(x_0)}$$

1D optimization

- Newton's method ("Newton-Raphson")

rule: iterate

$$x_{k+1} = x_k - \frac{f'(x_k)}{f''(x_k)}$$

Brent method: combines golden section,
Newton method

Optimization

(multidimensional)


- Gradient descent ("steepest descent")

take a step in the direction

$$-\nabla f$$

* how far?

↳ what is the appropriate
step size?



$$\vec{x}_{k+1} = \vec{x}_k - s_k \nabla f(\vec{x}_k)$$

"step size" \leftrightarrow "learning rate"

Optimization

- Steepest descent with exact line search

* compute $\nabla f(\vec{x}_k)$

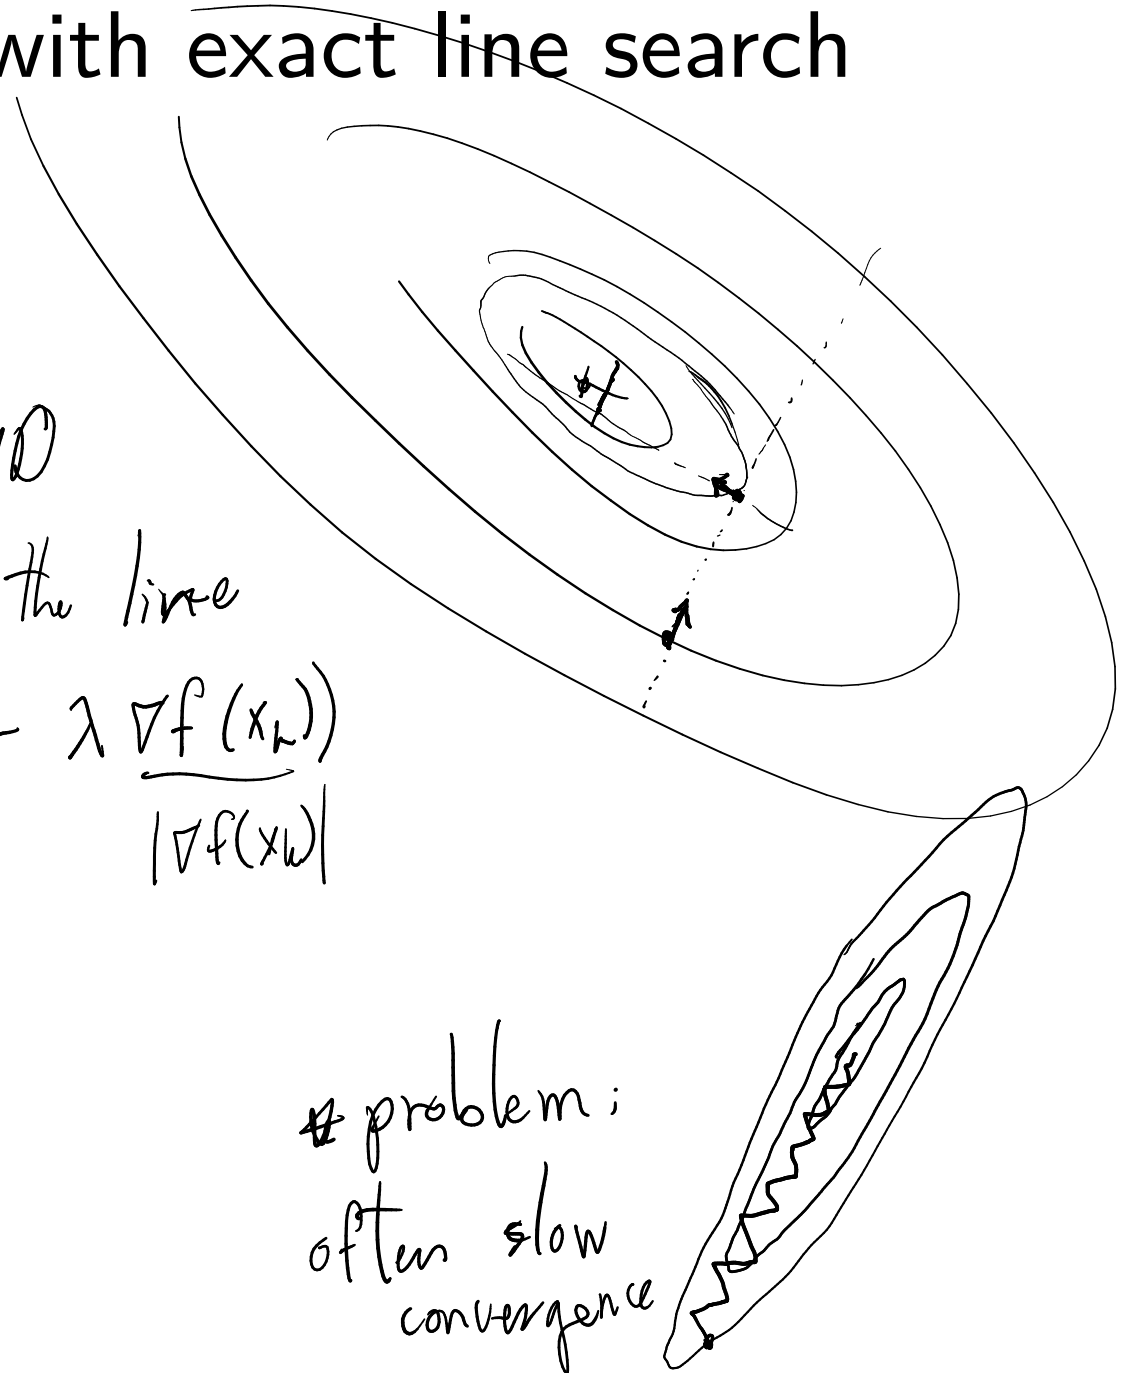
* carry out a full 1D
minimization along the line

$$g(\lambda) = f(\vec{x}_k - \lambda \frac{\nabla f(\vec{x}_k)}{|\nabla f(\vec{x}_k)|})$$

= "line search"


* iterate

* problem;
often slow
convergence



Optimization

- Steepest descent with exact line search



one possible solution:
use "preconditioner"
preconditioner: transform^(linear)
on \vec{x} to make the function
less elongated.

Optimization

- ~~Conjugate gradients~~ Newton's method

$$f(\vec{x}) \approx \underbrace{f(\vec{x}_0) + \nabla f(\vec{x}_0) \cdot (\vec{x} - \vec{x}_0)}_{\text{minimum is at}} + \frac{1}{2} (\vec{x} - \vec{x}_0)^T \underbrace{H}_{\text{Hessian matrix}} (\vec{x} - \vec{x}_0)$$

$$\vec{x} = \vec{x}_0 - H^{-1} \cdot \nabla f$$

powerful but not very useful
for large problems because Hessian
is too big (or expensive to compute)


Optimization

- Conjugate gradients ← one of many "quasi-Newton" methods

* compute the gradient $\vec{g}_k = \nabla f(x_k)$

* compute correction factor

$$\beta_k = \frac{\vec{g}_k \cdot (\vec{g}_k - \vec{g}_{k-1})}{\vec{g}_{k-1} \cdot \vec{g}_{k-1}}$$



* search direction:

$$\vec{h}_k = -(\vec{g}_k - \beta_k \vec{h}_{k-1}) \text{ direction}$$

to avoid "undoing" minimization that has already been done.


Optimization

- Nelder-Mead (downhill simplex method)

used if gradient unknown or difficult to compute
(like a multidimensional version of golden section)

* uses a $N+1$ simplex in N dimensions

↳ Nd version of a triangle (2d)
tetrahedron (3d)



* reflection




* expansion



* contraction



* shrink



Optimization

- Stochastic gradient descent

- * if cost function is large sum of terms (e.g. least squares with many data, or likelihood function)
- * idea: compute gradient using subset of the data
- * especially useful with very large and redundant datasets

Optimization

- Gauss-Newton method

Newton algorithm applied to non-linear least square cost function

$$f(\vec{\beta}) = \sum_i | \underbrace{M(\vec{\beta}; x_i)}_{\text{model}} - \underbrace{y_i}_{\text{measured data}} |^2$$

$$= \sum_i |r_i(\vec{\beta})|^2$$

$$\frac{\partial f}{\partial \beta_j} = 2 \sum_i r_i(\vec{\beta}) \underbrace{\frac{\partial r_i}{\partial \beta_j}}_{J_{ij}} = 2 J^T r$$

$$\frac{\partial^2 f}{\partial \beta_k \partial \beta_l} = 2 \underbrace{\sum_i \frac{\partial r_i}{\partial \beta_k} \frac{\partial r_i}{\partial \beta_l}}_{(J^T J)_{kl}} + 2 \sum_i r_i \underbrace{\frac{\partial^2 r_i}{\partial \beta_k \partial \beta_l}}_{\text{neglect this term}}$$

Newton: $\vec{\beta}^{(n+1)} = \vec{\beta}^{(n)} - (J^T J)^{-1} J^T r(\vec{\beta}^{(n)})$

Optimization

- Levenberg–Marquardt algorithm

$$f(\vec{\beta}) \approx f(\vec{\beta}_0) + \nabla f^T (\vec{\beta} - \vec{\beta}_0) + (\vec{\beta} - \vec{\beta}_0)^T H (\vec{\beta} - \vec{\beta}_0)$$

↓

$$\nabla f^T + H(\vec{\beta} - \vec{\beta}_0) = 0$$

$$\cancel{J}^T r + \cancel{J}^T \underbrace{J}_{\downarrow \text{replace with}} (\vec{\beta} - \vec{\beta}_0) = 0$$

$$(J^T J + \lambda \mathbb{1}) (\vec{\beta} - \vec{\beta}_0)$$

$$\Rightarrow \vec{\beta} - \vec{\beta}_0 = -(J^T J + \lambda \mathbb{1})^{-1} J^T r$$

$$\vec{\beta}^{(n+1)} = \vec{\beta}^{(n)} - (J^T J + \lambda \mathbb{1})^{-1} J^T r$$

Optimization

- Levenberg–Marquardt algorithm

$$\text{if } \lambda \rightarrow \infty \quad \vec{\beta}^{(n+1)} = \vec{\beta}^{(n)} - \frac{1}{2\lambda} \overbrace{2J^T r}^{\nabla f}$$

λ introduces a bias towards gradient descent

λ : Marquardt parameter - can be adjusted dynamically

Wirtinger derivatives

* Start with $f(z_1, z_2)$ where $z_1, z_2 \in \mathbb{C}$

* Introduce a change of variables:

$$z_1 = w_1 + w_2 \quad z_2 = w_1 - w_2 \quad w_1, w_2 \in \mathbb{C}$$

$$F(w_1, w_2) := f(z_1(w_1, w_2), z_2(w_1, w_2))$$

$$f(z_1, z_2) = F\left(\frac{z_1 + z_2}{2}, \frac{z_1 - z_2}{2}\right)$$

$$\frac{\partial f}{\partial z_1} = \frac{\partial F}{\partial w_1} \frac{\partial w_1}{\partial z_1} + \frac{\partial F}{\partial w_2} \frac{\partial w_2}{\partial z_1} = \frac{1}{2} \left(\frac{\partial F}{\partial w_1} + \frac{\partial F}{\partial w_2} \right)$$

$$\frac{\partial f}{\partial z_2} = \frac{1}{2} \left(\frac{\partial F}{\partial w_1} - \frac{\partial F}{\partial w_2} \right)$$

Wirtinger derivatives

Special case: $w_1 = x + i0$
 $w_2 = 0 + iy$

$$\Rightarrow z_1 = x + iy = z \quad z_2 = x - iy = z^*$$

$$\frac{\partial F}{\partial z} = \frac{1}{2} \left(\frac{\partial F}{\partial x} - i \frac{\partial F}{\partial y} \right) \quad \frac{\partial F}{\partial z^*} = \frac{1}{2} \left(\frac{\partial F}{\partial x} + i \frac{\partial F}{\partial y} \right)$$

special case: if $f \in \mathbb{R}$

then,

$$\left(\frac{\partial f}{\partial z} \right)^* = \frac{1}{2} \left(\frac{\partial f}{\partial x} + i \frac{\partial f}{\partial y} \right) = \frac{\partial f}{\partial z^*}$$

Wirtinger derivatives

* If f is a cost function, it is real (and probably non-negative)

$$\text{then, } \frac{\partial f}{\partial z} = 0 \iff \frac{\partial f}{\partial z^*} = 0 \quad \text{because } \left(\frac{\partial f}{\partial z}\right)^* = \frac{\partial f}{\partial z^*}$$

\Rightarrow to minimize a cost function with complex-valued arguments (z_1, z_2, z_3, \dots) we only need to

$$\text{set } \frac{\partial f}{\partial z_1} = \frac{\partial f}{\partial z_2} = \dots = 0 \quad \text{or} \quad \frac{\partial f}{\partial z_1^*} = \frac{\partial f}{\partial z_2^*} = \dots = 0$$

(frequently used with Fourier space cost functions)

Singular Value Decomposition

Any matrix can be decomposed as

$$A = U \Sigma V^t$$

where U and V are square unitary matrices and Σ has only non-negative diagonal entries.

$$\begin{array}{c} | \\ | \\ m \\ | \\ | \end{array} \begin{array}{c} \text{--- } n \text{ ---} \\ \left[\begin{array}{c} A \end{array} \right] \end{array} = \begin{array}{c} | \\ | \\ m \\ | \\ | \end{array} \begin{array}{c} \text{--- } m \text{ ---} \\ \left[\begin{array}{c} U \end{array} \right] \end{array} \begin{array}{c} \text{--- } n \text{ ---} \\ \left[\begin{array}{c} \sigma_1 \quad 0 \\ \sigma_2 \quad \quad \\ \sigma_3 \quad \quad \\ \vdots \quad \quad \\ 0 \end{array} \right] \end{array} \begin{array}{c} \left[\begin{array}{c} V^t \\ \text{--- } n \text{ ---} \end{array} \right] \begin{array}{c} | \\ | \\ n \\ | \\ | \end{array} \end{array}$$

Singular Value Decomposition

Relationship with eigen-decomposition:

$A^+ A$ is square. What are the eigen-vectors and eigenvalues of $A^+ A$?

$$A^+ A w = \lambda w$$

$$\begin{bmatrix} | & | & | & \dots \\ w_1 & w_2 & w_3 & \dots \\ | & | & | & \dots \end{bmatrix}$$

$$A^+ A W = \Lambda W$$

$$\begin{bmatrix} \lambda_1 & & & 0 \\ & \lambda_2 & & \\ & & \lambda_3 & \\ 0 & & & \dots \end{bmatrix}$$

$$W^+ A^+ A W = \Lambda$$

$$A^+ A = V \Sigma^T \underbrace{U^+ U}_I \Sigma V^+ = V \Sigma^T \Sigma V^+$$

Singular Value Decomposition

$$W^T \Lambda W = W^T A^T A W = W^T \underbrace{V^T \Sigma \Sigma^T V}_{\uparrow S} V^T W$$
$$S = \begin{bmatrix} \sigma_1^2 & & & 0 \\ & \sigma_2^2 & & \\ & & \sigma_3^2 & \\ 0 & & & \ddots \end{bmatrix}$$

$$\Rightarrow V^T = W, \quad \Lambda = S$$

conclusion: * the eigenvalues of $A^T A$ are $\lambda_1 = \sigma_1^2$,
 $\lambda_2 = \sigma_2^2$,

⋮

* the columns of V^T are the eigenvectors of $A^T A$

\Rightarrow convenient in some instances to compute σ_i and V

Singular Value Decomposition

rank is at most $\min(m, n)$

rank = number of non-zero singular values.

pseudo-inverse:

definition: $(A^T A)^{-1} A^T$

$$(V \Sigma^T U^T U \Sigma V^T)^{-1} V \Sigma U^T$$

$$= (V S V^T)^{-1} V \Sigma U^T$$

$$= V S^{-1} \underbrace{V^T V}_{I} \Sigma U^T$$

$$= V S^{-1} \Sigma U^T$$

maybe doesn't exist

$$= V \begin{bmatrix} \frac{1}{\sigma_1} & & \\ & \frac{1}{\sigma_2} & \\ & & \ddots \end{bmatrix} u^T$$

* if $\sigma_i \rightarrow 0$, inversion becomes unstable
 \rightarrow characterized by the ratio $\frac{\sigma_{\max}}{\sigma_{\min}}$

$$\frac{\sigma_{\max}}{\sigma_{\min}} = \text{conditioning number } N$$

$N = 1$: well-conditioned, invertible ...

$N \rightarrow \infty$: ill-conditioned system

Dimensionality reduction

approximate Σ with fewer non-zero elements

(truncate the list of singular values, removing the smallest)