

1 Vector & Matrix Norms

A **vector norm** on \mathbb{R}^n is a real map $\|\cdot\| : \mathbb{R}^n \rightarrow \mathbb{R}$ satisfying the following:

1. $\forall \vec{x} \in \mathbb{R}^n. [\vec{x} \neq \vec{0} \implies \|\vec{x}\| > 0]$.
2. $\forall \lambda \in \mathbb{R}. \forall \vec{x} \in \mathbb{R}^n. [\|\lambda \vec{x}\| = |\lambda| \|\vec{x}\|]$.
3. $\forall \vec{x}, \vec{y} \in \mathbb{R}^n. [\|\vec{x} + \vec{y}\| \leq \|\vec{x}\| + \|\vec{y}\|]$.

For $p > 0$ the ℓ_p -norm of \vec{x} is defined as:

$$\|\vec{x}\|_p = \left(\sum_{i=1}^n |x_i|^p \right)^{\frac{1}{p}}$$

$$\|\vec{x}\|_1 = \sum_{i=1}^n |x_i|$$

$$\|\vec{x}\|_2 = \sqrt{\sum_{i=1}^n x_i^2}$$

$$\|\vec{x}\|_\infty = \max_{1 \leq i \leq n} |x_i|$$

Any vector norms $\|\cdot\|_a, \|\cdot\|_b$ in \mathbb{R}^n are **equivalent**: $\exists r, s > 0$ s.t. $\forall \vec{x} \in \mathbb{R}^n. [r\|\vec{x}\|_a \leq \|\vec{x}\|_b \leq s\|\vec{x}\|_a]$.

A **matrix norm** on $\mathbb{R}^{m \times n}$ is a real map $\|\cdot\| : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ satisfying the following:

1. $\forall A \in \mathbb{R}^{m \times n}. [A \neq 0 \implies \|A\| > 0]$.
2. $\forall \lambda \in \mathbb{R}. \forall A \in \mathbb{R}^{m \times n}. [\|\lambda A\| = |\lambda| \|A\|]$.
3. $\forall A, B \in \mathbb{R}^{m \times n}. [\|A + B\| \leq \|A\| + \|B\|]$.

In a **sub-multiplicative** matrix norm, $\forall A \in \mathbb{R}^{m \times n}, \forall B \in \mathbb{R}^{n \times p}. [\|AB\| \leq \|A\| \|B\|]$. We define ℓ_p -norms of a matrix as:

$$\|A\|_1 = \max_j \sum_{i=1}^m |a_{ij}|$$

$$\|A\|_2 = \sigma_1(A)$$

$$\|A\|_\infty = \|A^T\|_1$$

$$\|A\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2}$$

A matrix norm on $\mathbb{R}^{m \times n}$ is **consistent** on vector norms $\|\cdot\|_a$ on \mathbb{R}^n and $\|\cdot\|_b$ on \mathbb{R}^m iff $\forall A \in \mathbb{R}^{m \times n}, \forall \vec{x} \in \mathbb{R}^n. [\|A\vec{x}\|_b \leq \|A\| \|\vec{x}\|_a]$. If $a = b$ then they are **compatible**. For a vector norm, the **subordinate** matrix norm is $\forall A \in \mathbb{R}^{m \times n}, \|A\| = \max\{\|A\vec{x}\| : \vec{x} \in \mathbb{R}^n, \|\vec{x}\| = 1\}$. A vector norm is compatible with its subordinate matrix norm.

2 Linear Maps on \mathbb{C}^n

The **standard inner product** is defined as $\forall \vec{u}, \vec{v} \in \mathbb{C}^n. \langle \vec{u}, \vec{v} \rangle = \vec{u}^T \vec{v} = \sum_{i=1}^n \bar{u}_i v_i$.

3 Least Square Method

The **orthogonal projection** $\pi_U : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is $\vec{v} \mapsto \pi_U(\vec{v}) = U(U^T U)^{-1} U^T \vec{v}$. $\text{im } A \perp \ker A^T$, hence $\forall \vec{x} \in \mathbb{R}^m$, there exist unique $\vec{x}_i \in \text{im } A, \vec{x}_k \in \ker A$ s.t. $\vec{x} = \vec{x}_i + \vec{x}_k$.

Suppose $\forall \vec{x} \in \mathbb{R}^n$, $A\vec{x} = \vec{b}$ has no solution (i.e. $\vec{b} \notin \text{im } A$). The **Least Square Method** finds $x \in \mathbb{R}^n$ s.t. $\|A\vec{x} - \vec{b}\|_2$ is **minimized**.

This happens iff $\|A\vec{x} - \vec{b}\|_2 = 0$, or $A\vec{x} = \vec{b}$.

$A^T A\vec{x} = A^T \vec{b}$ is the **normal equation**, whose solution also solves the least square problem. To find the best **affine map** (straight line) $Y = mX + c$ we solve

$$A^T A \begin{bmatrix} c \\ m \end{bmatrix} = A^T Y \text{ where } A = \begin{bmatrix} 1 & X_1 \\ \vdots & \vdots \\ 1 & X_n \end{bmatrix}.$$

3.1 Linear Regression

Suppose we have a set of points (y_i, \vec{a}_i) where $y_i \in \mathbb{R}$, $\vec{a}_i \in \mathbb{R}^n$. We want to find a **model of best fit** with params $s_0 \in \mathbb{R}$, $\vec{s} \in \mathbb{R}^n$ s.t. the sum of errors squared $\sum_{i=1}^m (s_0 + \vec{s} \cdot \vec{a}_i - y_i)^2$ is minimized: $\forall i. [s_0 + \vec{s} \cdot \vec{a}_i \approx y_i]$. To solve with least squares, let:

$$A = \begin{bmatrix} 1 & a_{11} & \dots & a_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & a_{m1} & \dots & a_{mn} \end{bmatrix}, \vec{z} = \begin{bmatrix} s_0 \\ s_1 \\ \vdots \\ s_n \end{bmatrix} \in \mathbb{R}^{n+1}$$

$$\sum_{i=1}^m (s_0 + \vec{s} \cdot \vec{a}_i - y_i)^2 = \|A\vec{z} - \vec{y}\|_2^2$$

4 Spectral Decomposition

A matrix $Q \in \mathbb{R}^{n \times n}$ is **orthogonal** iff it is invertible and $Q^{-1} = Q^T$. This implies $|\det Q| = 1$, & \forall eigenvalues λ of $Q, |\lambda| = 1$.

A matrix $A \in \mathbb{R}^{n \times n}$ is **symmetric** iff $A^T = A$. Then, all eigenvalues of A are **real**, and their **algebraic and geometric multiplicities are equal**. Also, any eigenvectors with different eigenvalues are orthogonal.

Any **symmetric** matrix $A \in \mathbb{R}^{n \times n}$ can be diagonalized as $A = QDQ^T = QDQ^{-1}$, where Q is **orthogonal** and D is **diagonal** matrix of not necessarily distinct eigenvalues. To find the SD of A :

1. Find characteristic polynomial of A , solve to find eigenvalues λ_i of A .
2. For each distinct λ_i , find corresponding eigenspace E_{λ_i} .
3. For each E_{λ_i} , find orthonormal basis $(\vec{v}_{\lambda_1}, \vec{v}_{\lambda_2}, \dots, \vec{v}_{\lambda_i} \mid \dim E_{\lambda_i})$.
4. The bases \vec{v}_{λ_i} are columns of Q and eigenvectors of A . Done!

5 Singular Value Decomposition

A **symmetric** matrix $A \in \mathbb{R}^{n \times n}$ is **positive semi-definite** (*non-negative eigenvalues*) iff $\forall \vec{x} \in \mathbb{R}^n. \vec{x}^T A \vec{x} \geq 0$, and **positive definite** (*positive eigenvalues*) iff $\forall \vec{x} \in \mathbb{R}^n - \{0\}. \vec{x}^T A \vec{x} > 0$.

For any $A \in \mathbb{R}^{m \times n}$, $A^T A \in \mathbb{R}^{n \times n}$ and $AA^T \in \mathbb{R}^{m \times m}$ are **symmetric & positive semi-definite**.

For any $A \in \mathbb{R}^{m \times n}$, the **Singular Value Decomposition** of A is $A = USV^T$, where $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ are **orthogonal**, and $S = \text{diag}(\sigma_1, \dots, \sigma_p)$ and $p = \min(m, n)$ where $\sigma_1 \geq \dots \geq \sigma_p \geq 0$ are the **singular values** of A . This gives rise to properties:

- The rank $r = \text{rk}(A)$ is equal to the number of positive singular values in S : $\sigma_1 \geq \dots \geq \sigma_r > 0$ and $\sigma_{r+1} = \dots = \sigma_p = 0$.
- $\|A\|_2 = \sigma_1$.

The positive singular values of A are the positive square roots of the eigenvalues of AA^T or $A^T A$.

$$\bullet U^T U = I.$$

To compute SVD of A , we pick the smaller of $A^T A$ and AA^T and then:

1. Solve $\det(A^T A - \lambda I) = 0$ to get λ_1, \dots
2. Take sqrts to get σ_1, \dots . Skip $\lambda_i = 0$.
3. Find v_i by solving $(A^T A - \lambda_i I)v_i = 0$.
4. v_i are cols of V . **MUST BE orthonormal**.
5. Construct S from σ_i (same shape as A).
6. $\forall \sigma_i \neq 0, u_i = \frac{1}{\sigma_i} A v_i$.
7. u_i are cols of U . **If cols are missing use gram schmidt or cross product**.

If you pick $A^T A$ then calc U before V in the same way. Also:

- $\text{rk}(A) = \text{num of } \sigma_i \neq 0$.
- $\|A\|_2 = \sigma_{\max}(A)$.
- $\text{im}(A) = \text{span}\{u_1, \dots, u_r\}, r = \text{rk}(A)$.

5.1 Principal Component Analysis

Assume $A \in \mathbb{R}^{m \times n}$ represents m samples of n -dim data. The **principal axes** of A are the cols of V , and the cols of US (i.e. $\sigma_i u_i$) are the **principal components** of A .

$w^T A^T A w$ is maximised $V^T w = e_1, w = v_1$. Assuming $\|w\|_2 = 1$.

$A = \sum_{i=1}^r \sigma_i u_i v_i^T$. Then to max $\|Ax\|_2$, we need to align $x = v_1$ (since v_1 corresponds to largest σ_r). This maxes the contributions.

6 Generalized Eigenvectors

For $A \in \mathbb{R}^{n \times n}$, non-zero $v \in \mathbb{C}^n$ is a **generalized eigenvector of rank m (v^m)** with assoc eigenvalue $\lambda \in \mathbb{C}$ iff $(A - \lambda I)^m v = 0$ and $(A - \lambda I)^{m-1} v \neq 0$.

If an eigenvalue λ has algebraic multiplicity k , there are k linearly independent generalized eigenvectors associated with λ . $A \in \mathbb{R}^{n \times n}$ has n linearly independent generalized eigenvectors - \exists a basis of \mathbb{C}^n of generalized eigenvectors of A .

They always solve $(A - \lambda_i I)v_g = v_i$.

6.1 Jordan Normal Form (JNF)

A matrix is in JNF if it is in the form:

$$\begin{bmatrix} J_{k_1}(\lambda_1) & & & & \\ & \ddots & & & \\ & & \ddots & & \\ 0 & & \cdots & J_{k_n}(\lambda_n) & \end{bmatrix}$$

where $J_{k_i}(\lambda_i)$ is a **jordan block** of size k_i :

$$\begin{bmatrix} \lambda_i & & & \\ 1 & \ddots & & \\ & \ddots & \ddots & \\ 0 & \lambda_i & \ddots & \end{bmatrix}$$

The algebraic multiplicity of λ is the sum of the sizes of blocks with λ on the diagonal. The geometric multiplicity of λ is the number of blocks with λ on the diagonal. To find the JNF of $A \in \mathbb{R}^{n \times n}$:

1. Compute eigenvalues λ_i and their algebraic multiplicity a_i for A .

2. Find E_{λ_i} , the geometric multiplicities g_i , and eigenvectors $v_{ij}^1, \dots, v_{ij}^{g_i}$.

3. If $g_i < a_i$, find $a_i - g_i$ missing generalized eigenvectors: $\forall v_{ij}^k \in \mathbb{C}^n$ associated with λ_i , find all $v_{ij}^k \in \mathbb{C}^n$ s.t. $(A - \lambda_i I)v_{ij}^k = v_{ij}^{k-1}$. (Gaussian elim.)
4. Change of basis $B = [v_{11}^1, \dots, v_{11}^{k_1}, \dots, v_{1g_1}^1, \dots, v_{1g_1}^{k_{g_1}}, \dots, v_{m1}^{k_1}, \dots, v_{m1}^{k_{g_1}}, \dots, v_{mg_m}^1, \dots, v_{mg_m}^{k_{g_m}}]$.
5. Now $J = \begin{bmatrix} & & & \\ \vdots & & & \\ 0 & & \ddots & \\ & & & J_{k_{g_m}}(\lambda_m) \end{bmatrix}$.
6. $B^{-1}AB = J$.

7 Cholesky Decomposition

$A \in \mathbb{R}^{n \times n}$ is **lower triangular** iff $\forall i < j. A_{ij} = 0$; **upper triangular** iff $\forall i > j. A_{ij} = 0$.

Let $A \in \mathbb{R}^{n \times n}$ be **symmetric**. If it is **positive definite**, all its diag elems are > 0 and $\forall i, j. \max(A_{ii}, A_{jj}) > |A_{ij}|$. If it is **positive semi-definite**, all its diag elems are ≥ 0 and $\forall i, j. \max(A_{ii}, A_{jj}) \geq |A_{ij}|$. As a consequence, the largest coefficient of A is **on its diagonal**. Also the $1 \times 1, \dots, n \times n$ matrices in the upper left corner of A are also positive (semi-)definite.

The **cholesky decomposition** of $A \in \mathbb{R}^{n \times n}$ is $A = LL^T$ where L is a **lower triangular matrix**. A is **positive semi-definite** $\Leftrightarrow \exists L \in \mathbb{R}^{n \times n}$ s.t. L is lower triangular and $A = LL^T$. A is **positive definite** $\Leftrightarrow \exists L \in \mathbb{R}^{n \times n}$ s.t. L is lower triangular and $A = LL^T$ and $\forall i. L_{ii} > 0$. To find a CD:

1. Consider $L = \begin{bmatrix} l_{11} & 0 & 0 & \cdots \\ l_{21} & l_{22} & 0 & \cdots \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & l_{n1} & l_{n2} & \cdots \end{bmatrix}$.

2. Solve $A = LL^T$. Example for 3×3 :

$$\begin{bmatrix} l_{11} & l_{11}l_{21} & l_{11}l_{31} \\ l_{11}l_{21} & l_{21}^2 + l_{22}^2 & l_{21}l_{31} + l_{22}l_{32} \\ l_{11}l_{31} & l_{21}l_{31} + l_{22}l_{32} & l_{31}^2 + l_{32}^2 + l_{33}^2 \end{bmatrix}$$

S is contained in the diagonal elems of L.

8 QR Decomposition

Let $A \in \mathbb{R}^{m \times n} = [a_1 \ \dots \ a_n]$ s.t. a_1, \dots, a_n are linearly independent. Apply gram schmidt to get (e_1, \dots, e_n) s.t. $\text{span}\{e_1, \dots, e_n\} = \text{span}\{a_1, \dots, a_n\}$. If we

let semi-orthogonal $Q = [e_1 \ \dots \ e_n] \in \mathbb{R}^{m \times n}$ then $A = QR$ with upper triangular:

$$R = \begin{bmatrix} (e_1 \cdot a_1) & (e_1 \cdot a_2) & \cdots & (e_1 \cdot a_n) \\ 0 & (e_2 \cdot a_2) & \cdots & (e_2 \cdot a_n) \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & (e_n \cdot a_n) \end{bmatrix}$$

8.1 Householder Maps

Suppose a hyperplane P goes through $\vec{0}$ with unit normal $u \in \mathbb{R}^m$ ($P = \{x \in \mathbb{R}^m \mid u \cdot x = 0\}$). The **householder matrix** $H_u = I - 2uu^T$ induces reflection wrt P .

9 Convergence in \mathbb{R}

Let $(a_n)_{n \in \mathbb{N}} \in \mathbb{R}^{\mathbb{N}}$ be a seq of reals and $l \in \mathbb{R}$. (a_n) **converges** to l ($\lim_{n \rightarrow \infty} a_n = l$) iff $\forall \epsilon > 0. \exists N \in \mathbb{N}$ s.t. $\forall n > N. |a_n - l| < \epsilon$. (a_n) is a **cauchy seq** iff $\forall \epsilon > 0. \exists N \in \mathbb{N}$ s.t. $\forall n, m > N. |a_n - a_m| < \epsilon$. (a_n) **converges iff it is cauchy**.

9.1 Metric Spaces

A metric space is $\langle S, d \rangle$ where $S \neq \emptyset$ and $d : S \times S \rightarrow \mathbb{R}$ satisfies the following:

1. $\forall x, y \in S, d(x, y) \geq 0$.
2. $\forall x, y \in S, d(x, y) = 0 \Leftrightarrow x = y$.
3. $\forall x, y \in S, d(x, y) = d(y, x)$.
4. $\forall x, y, z \in S, d(x, y) \leq d(x, z) + d(z, y)$.

If V is a vector space equipped with the norm $\|\cdot\|$, then $\langle V, (x, y) \mapsto \|x - y\| \rangle$ is a metric space.

If $\langle S, d \rangle$ is a metric space and (a_n) is a seq in S , $\lim_{n \rightarrow \infty} a_n = l$ iff $\forall \epsilon > 0. \exists N \in \mathbb{N}$ s.t. $\forall n > N. d(a_n, l) < \epsilon$. If (a_n) is converging, the limit is **unique**. (a_n) is **cauchy** iff $\forall \epsilon > 0. \exists N \in \mathbb{N}$ s.t. $\forall n, m > N. d(a_n, a_m) < \epsilon$. (a_n) converges iff it is cauchy.

$\langle S, d \rangle$ is **complete** iff every cauchy seq in S converges in S . For any $k > 0$, \mathbb{R}^k with ℓ_1, ℓ_2 or ℓ_∞ is complete.

9.2 Fixed Point Equations

If $S \neq \emptyset$ and $f : S \rightarrow S$, then $p \in S$ is a **fixed point** of f if $f(p) = p$.

$f : S \rightarrow S$ is a **contraction** of S in $\langle S, d \rangle$ iff \exists **contraction constant** $0 \leq \alpha < 1$ s.t. $\forall x, y \in S. d(f(x), f(y)) \leq \alpha d(x, y)$. If $\langle S, d \rangle$ is **complete** and f is a **contraction** of S , then f has a **fixed point**.

10 Projectors

$\text{proj}_u(v) \triangleq \frac{u \cdot v}{u \cdot u} u$

The **hermitian conjugate** or **adjoint** of $A \in \mathbb{C}^{m \times n}$ is $A^* : A^T$ with all elems conjugated. **Hermatian Matrix**: $A = A^*$.

P is a **projector** iff $Pv = \text{proj}_u(v)$. Then:

- $P = P^2$ (Reprojection does nothing.)
- $(I - P)$ is also a projector.
- $\text{range}(P) = \text{null}(I - P) \Rightarrow S_1$.
- $\text{range}(P) = \text{null}(I - P) \Rightarrow S_2$.

• If S_1 orthogonal to S_2 , then P is an **orthogonal projector** (not orth matrix). P is an orth proj iff $P = P^*$. If \hat{Q} is an orthonormal basis for $\text{range}(P)$, then $P = \hat{Q}\hat{Q}^*$. Also, $P_u = uu^*$. To find orthonormal basis, gram schmidt.

11 Condition Numbers

The measure of the **sensibility** of a problem to small perturbations in its input. Let problem P take input d and some perturbation ϵ to give outputs $s(d)$, $s(d + \epsilon)$. Then the **condition number** $\kappa(P)$:

$$\kappa(P) = \max_{\epsilon} \frac{\|s(d) - s(d + \epsilon)\|}{\|\epsilon\|}$$

The **condition number measures the worst case scenario**. The relative condition number is defined in terms of relative diff:

$$\kappa(P) = \max_{\epsilon} \frac{\|s(d) - s(d + \epsilon)\|}{\|d\|} \frac{\|d\|}{\|s(d)\|}$$

When checking if a system is **stable**, we check the bound of the condition number $\forall d, \epsilon$. If it is bounded, it is stable. A large CN means a **ill-conditioned** problem. A small CN means a **well-conditioned** problem.

11.1 Matrix Condition Numbers

For a non-singular **square** matrix A , its condition number $\kappa(A) = \|A^{-1}\| \|A\|$. If the problem P is represented by the linear equation $Ax = b$, then $\kappa(P) = \kappa(A)$.

The condition number of a **non-square** matrix A is $\kappa(A) = \|A^\dagger\| \|A\|$, where $A^\dagger = (A^T A)^{-1} A^T$ is the **pseudo-inverse** of A .

11.2 Conditioning of a Problem

To decide if a condition number is big enough to say its problem is ill-conditioned, we say: for a condition number κ_0 , we lose $\approx \log_{10}(\kappa_0)$ significant figures in accuracy.

12 Stability

Fundamental Theory of Floating Point Arithmetic states all operations have rel error $\leq \epsilon_{\text{machine}}$. Define **problem** $f: \mathcal{X} \rightarrow \mathcal{Y}$ and **algorithm** $\tilde{f}: \mathcal{X} \rightarrow \mathcal{Y}$. $f(x)$ **well conditioned** if small changes in x mean small changes in $f(x)$. \tilde{f} is **backwards stable** iff $\tilde{f}(x) = f(\tilde{x}) \approx \frac{\|\tilde{x} - x\|}{\|x\|} = \mathcal{O}(\epsilon_{\text{machine}})$

A **backward stable algorithm** gives exactly the right answer to nearly the right question.

Flops & inner prod backwards stable. Outer prod not. To find if \tilde{f} is:

1. Show $\tilde{f}(x)$ is **exact** for perturbed \tilde{x} .
2. Verify rel err in \tilde{x} , $x \approx \epsilon_{\text{machine}}$.

e.g. **inner product** $x^T y = \sum_{i=1}^n x_i y_i$:

- $\text{fl}(x^T y) = \sum_{i=1}^n (x_i y_i)(1 + \epsilon_i)$, $|\epsilon_i| \leq \epsilon_m$.
- $\text{fl}(x^T y) = \tilde{x}^T y$, so **exact**.
- $\frac{|\tilde{x}_i - x_i|}{|x_i|} = \frac{|x_i \epsilon_i|}{|x_i|} = |\epsilon_i| \leq \epsilon_m = \mathcal{O}(\epsilon_m)$.

13 Systems of Equations

Basic gaussian elim splits $A = LU$, but this fails if diagonal elems ≈ 0 - not backwards stable. The orders of rows is arbitrary so we can **pivot** using a **orthonormal permutation matrix** P (1 non-zero entry per col = 1). PA permutes rows, AP permutes cols. They change $\kappa(A)$. This is stable iff $\rho = \frac{\max_{i,j} |u_{ij}|}{\max_{i,j} |d_{ij}|} = \mathcal{O}(1)$.

13.1 Gradient Optimisation

Func $f: \mathcal{X} \rightarrow \mathcal{Z}$ has **level curve** $\mathcal{X} \ni C = \{ \tilde{x} | f(\tilde{x}) = c \in \mathcal{Z} \}$, which can be projected onto \mathcal{X} to visualize f . $\frac{\partial f(\tilde{x})}{\partial u} =$

$$\begin{bmatrix} f_{x_1} & \cdots & f_{x_m} \end{bmatrix}^T u = \nabla f u. \quad \nabla^2 f = H = \begin{bmatrix} f_{x_1 x_1} & \cdots & f_{x_1 x_m} \\ \vdots & \ddots & \vdots \\ f_{x_m x_1} & \cdots & f_{x_m x_m} \end{bmatrix}.$$

The grad $\nabla f(\tilde{x})$ is the direction of the maximum rate of change $|\nabla f(\tilde{x})|$.

Conjugate Gradient Method solves $Ax = b$ by solving $\min_x f(x) = \frac{1}{2} x^T Ax - b^T x$ where $\nabla f(x) = Ax - b$. It chooses **dirs** $p^{(k)}$ s.t. they are conjugate: $\forall k, j, k \neq j \Rightarrow \langle p^{(k)}, p^{(j)} \rangle = p^{(k)T} A p^{(j)} = 0$ using **residuals** $r^{(k)} = -\nabla f = b - Ax^{(k)}$ and **step size** $\alpha^{(k)}$:

$$\begin{aligned} p^{(0)} &= -\nabla f = b - Ax^{(0)} \\ p^{(k)} &= r^{(k)} \sum_{i < k} \frac{p^{(i)T} A r^{(k)}}{p^{(i)T} A p^{(i)}} p^{(i)} \end{aligned}$$

$$\alpha^{(k)} = \arg \min_{\alpha} f(x^{(k)} + \alpha^{(k)} p^{(k)}) = \frac{p^{(k)T} r^{(k)}}{p^{(k)T} A p^{(k)}}$$

Converges in $\leq m$ iterations, as $r^{(k)}$ s are orthogonal, form a basis for \mathbb{R}^m , hence $r^{(m)} = \vec{0}$.

14 Iterative Sol. to Linear Equations

Suppose we wish to solve $Ax = b$ for $A \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^n$. Gaussian elim is $\mathcal{O}(n^3)$, which is impractical. Instead, split $A = G + R$ s.t. $Ax = b \Leftrightarrow x = G^{-1}b - G^{-1}Rx$. Then if $M = -G^{-1}R$ and $c = G^{-1}b$ then $x = Mx + c$. To solve this **iteratively**, we must find a **fixed point** of $f(x) = Mx + c$.

Using **fixed point theorem**, we define seq (x_k) for some starting x_0 s.t. $x_{k+1} = Mx_k + c$.

Then for a **consistent** norm $\|\cdot\|$, (x_k) converges from any starting point x_0 if $\|M\| < 1$.

The **rate of conv** $r \propto -\log_{10} \|M\|$, so we pick G, R s.t. $-G^{-1}R$ and G^{-1} are easy to compute, and $\|M\|$ is small.

14.1 Common Splitting

Assume wlog that A has no 0s on the diag (if not so, perform a basis change). Then $A = D + L + U$ where D is the diagonal, and L, U are the strict lower and upper triangular parts of A .

14.2 Jacobi Method

If $A = D + L + U$, we say $A = D + R$ where $R = L + U$, then $Ax = b \Leftrightarrow x = Mx + c$ where $M = -D^{-1}R$ and $c = D^{-1}b$. Once again, consider seq $(x^{(k)})$ defined as $x^{(k+1)} = Mx^{(k)} + c$ and $x^{(0)} \in \mathbb{R}^n$. Then:

$$Mx = -D^{-1}Rx = \begin{bmatrix} -\frac{1}{a_{1,1}} \sum_{j \neq 1} a_{1,j} x_j \\ \vdots \\ -\frac{1}{a_{n,n}} \sum_{j \neq n} a_{n,j} x_j \end{bmatrix}$$

As $x^{(k+1)} = Mx^{(k)} + c$, it comes that:

$$x_i^{(k+1)} = \frac{1}{a_{i,i}} (b_i - \sum_{j \neq i} a_{i,j} x_j^{(k)})$$

So to compute i^{th} element of $x^{(k+1)}$, we only need b , $x^{(k)}$, and the i^{th} row of A , which is great for **parallelization**.

The conv rate is $\max \lambda$ of $D^{-1}R$.

14.3 Gauss-Seidel

Now, we split $A = (D + L) + U$ s.t. $Ax = b \Leftrightarrow x = Mx + c$ where $M = -(D + L)^{-1}U$ and $c = (D + L)^{-1}b$. Once again, consider seq $(x^{(k)})$ defined as $x^{(k+1)} = Mx^{(k)} + c$ and $x^{(0)} \in \mathbb{R}^n$. We can rewrite that as:

$$\begin{aligned} (D + L)x^{(k+1)} &= -Ux^{(k)} + b \\ \Downarrow \\ \sum_{j < i} a_{i,j} x_j^{(k+1)} &= -\sum_{j > i} a_{i,j} x_j^{(k)} + b_i \end{aligned}$$

So to compute the i^{th} elem of $x^{(k+1)}$, we only need A , $x^{(k)}$, b and the k^{th} elems of $x^{(k+1)}$ for $k < i$. Since the update is computed with more recent quantities, is faster than jacobi. The conv rate is $\max \lambda$ of $L^{-1}U$.

14.4 Convergence

If $A \in \mathbb{R}^{n \times n}$ is **strictly row diagonally dominant**: $\forall i, |a_{ii}| > \sum_{j \neq i} |a_{ij}|$, then Jacobi and Gauss-Seidel methods **will converge**.

When $\kappa(A)$ is small, conv is faster. If $\kappa(A)$ is too big, it may even diverge.

A matrix A is **irreducible** if by **symmetric permutations of rows and columns** (i.e. if you swap row a with row b , you must swap col a with col b), it can not take the form $\begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix}$ where A_{11} and 0 are block matrices.

If A is **weakly row diagonally dominant** and **irreducible**, Jacobi and G-S still converge.

15 Iteratively Finding Eigenvalues

Suppose we wish to find the eigenvalues and eigenspaces of $A \in \mathbb{R}^{n \times n}$: we must find roots of $\det(A - \lambda I)$ and do gaussian elimination for each eigenvector - an $\mathcal{O}(n^4)$ problem.

15.1 Power Iteration

Let $A \in \mathbb{R}^{n \times n}$ be a diagonalizable matrix with λ s of distinct mod ($\forall a, b, |\lambda_a| \neq |\lambda_b|$). Let $\lambda \in \mathbb{R}$ be **dominant** ($\forall i, |\lambda| \geq |\lambda_i|$). Consider the seq (x_k) defined by $x_{k+1} = \frac{Ax_k}{\|Ax_k\|}$

and $x_0 \in \mathbb{R}^n \setminus \{0\}$. Then:

$$x_k \xrightarrow{k \rightarrow \infty} v \quad \|Ax_k\| \xrightarrow{k \rightarrow \infty} |\lambda|$$

Where $v \in \mathbb{R}^n$ is the normalized dominant eigenvector. Here, the notion of conv is not rigorous, so we say (x_k) converges to a corresponding eigenspace. However:

- With random x_0 is possible for the it. to give 2nd dominant EVs. **Make sure** there is at least one non-zero component in the corresponding eigenspace.
- For matrix with multiple EVs of max modulus, power it. will converge to a linear combination of the corresponding eigenvectors.
- conv may be slow if the dominant eigenvector is not "very dominant".
- conv Rate $\approx \frac{\lambda}{\lambda_1}$.

15.2 Inverse Power Iteration

Let $A \in \mathbb{R}^{n \times n}$ be a diagonalizable non-singular matrix with eigenvalues of distinct modulus ($\forall a, b, |\lambda_a| \neq |\lambda_b|$). This means that $\forall i, \lambda_i \neq 0$ and μ is an eigenvalue of A iff μ^{-1} is an eigenvalue of A^{-1} with the same eigenvectors. So if λ is the smallest eigenvalue of A , λ^{-1} is the dominant eigenvalue of A^{-1} .

Let $\lambda \in \mathbb{R}$ be the eigenvalue of the **smallest modulus**. Consider the seq (x_k) defined by

$$x_{k+1} = \frac{A^{-1}x_k}{\|A^{-1}x_k\|} \quad \text{and} \quad x_0 \in \mathbb{R}^n \setminus \{0\}.$$

$$x_k \xrightarrow{k \rightarrow \infty} v \quad \|A^{-1}x_k\| \xrightarrow{k \rightarrow \infty} \frac{1}{|\lambda|}$$

Where $v \in \mathbb{R}^n$ is a normalized eigenvector corresponding to λ . Conv rate $\approx \frac{\lambda_{n-1}}{\lambda_n}$.

15.3 Shifts

Let $A \in \mathbb{R}^{n \times n}$ and $s \in \mathbb{R}$. $\lambda \in \mathbb{R}$ is an eigenvalue of A iff $\lambda - s$ is an eigenvalue of **shifted** $A - sI$ with the same eigenvectors. If $A \in \mathbb{R}^{n \times n}$ is a diagonalizable matrix with eigenvalues of distinct modulus, and $A - sI$ is non-singular (s is not an eigenvalue of A), we can find the eigenvalue of A that is closest to s with inverse power iterations on $A - sI$.

15.4 Rayleigh Quotient

For $A \in \mathbb{R}^{n \times n}$ and $x \in \mathbb{R}^n \setminus \{0\}$, the **rayleigh quotient** $R(A, x) = \frac{x^T Ax}{x^T x}$ - an approx of an eigenvalue directly and not its modulus.

15.5 Deflation

To find eigenvalues apart form the dominant, we need to deflate the matrix $A \in \mathbb{R}^{n \times n}$ to $B \in \mathbb{R}^{(n-1) \times (n-1)}$ that has the same eigenvalues as A but the dominant one.

Let $\lambda_1, \dots, \lambda_n$ by the eigenvalues of A s.t. $|\lambda_1| \geq \dots \geq |\lambda_n|$ with corresponding eigenvectors v_1, \dots, v_n . Define $H \in \mathbb{R}^{n \times n}$ as a non-singular matrix s.t. $Hx_1 = \alpha e_1$, $\text{proj}_{Hv_1} v_1 = v_1$, $\text{proj}_{Hv_1} v_2 = 0$, $\text{proj}_{Hv_1} v_3 = 0$, \dots , $\text{proj}_{Hv_1} v_n = 0$. Then $H = \begin{bmatrix} \alpha & & & \\ & \ddots & & \\ & & 1 & \\ & & & 0 & \end{bmatrix}$.

where $\alpha \in \mathbb{R} \setminus \{0\}$ and $e_1 = [1, 0, \dots, 0]^T$. Then, we have:

$$HAH^{-1}e_1 = HA \frac{x_1}{\alpha} = H \frac{\lambda_1}{\alpha} x_1 = \lambda_1 e_1$$

Since the 1st col of HAH^{-1} is $[\lambda_1, 0, \dots, 0]^T$, we can write:

$$HAH^{-1} = \begin{bmatrix} \lambda_1 & & & \\ 0 & & & B \end{bmatrix}$$

Where $B \in \mathbb{R}^{(n-1) \times (n-1)}$. For any eigenvalue λ and corresponding eigenvector v : $Ax = \lambda x \Leftrightarrow HAx = \lambda Hx \Leftrightarrow HAH^{-1}(Hx) = \lambda Hx$. So λ is an eigenvalue of A iff it is an eigenvalue of HAH^{-1} .

Let λ_2 be the second dominant eigenvector of A s.t. $\lambda_2 \neq \lambda_1$. The corresponding eigenvector $x_2 = H^{-1} \begin{bmatrix} \beta \\ z_2 \end{bmatrix}$ with $\beta = \frac{b^T z_2}{\lambda_2 - \lambda_1}$ and z_2 is the dominant eigenvector of B .

We assumed that $\lambda_1 \neq \lambda_2$, or that the dominant eigenvalue λ_1 had geometric multiplicity 1. Suppose that instead it had geometric multiplicity p and $\lambda_1 = \dots = \lambda_p$. Then, deflation works similarly but with blocks instead of vectors: $B \in \mathbb{R}^{(n-p) \times (n-p)}$.

H is the **householder transformation** $H = I - \frac{2u u^T}{u^T u}$, s.t. $H = H^T = H^{-1}$ is **symmetric** and **orthogonal**. To make sure $Hx_1 = \alpha e_1$, we set: $u = x_1 + \|x_1\|_2 e_1$ and $\alpha = -\|x_1\|_2$.

16 Linear Algebra Basics

- $(ABC)^T = C^T B^T A^T$
- $\vec{u}_1 \times \vec{u}_2 = \det \begin{vmatrix} i & j & k \\ u_{11} & u_{12} & u_{13} \\ u_{21} & u_{22} & u_{23} \end{vmatrix} = \dots$
- **Diagonalisable** means all algebraic multiplicities equal to geometric multiplicity.
- **Algmul** - number of times λ appears.
- **Geomul** - number of eigenvectors of λ .
- Eigenvector and eigenval: $Av = \lambda v$
- $\forall u, v \in \mathbb{R}^n, u, v \neq \vec{0} \Rightarrow \text{rk}(uv^T) = 1$.
- **Rank** is the number of eigenvalues $\neq 0$.
- $\det(A^{-1}) = \det(A)^{-1}$.
- $\forall A \in \mathbb{R}^{n \times n}$, complex λ come in pairs with their conjugates $\bar{\lambda}$.
- **Cauchy-Schwarz**: $|\vec{x} \cdot \vec{y}| \leq \|\vec{x}\|_2 \|\vec{y}\|_2$
- If A has σ, λ then A^{-1} has $\frac{1}{\sigma}, \frac{1}{\lambda}$.

17 Gram Schmidt

The **gram schmidt process** iteratively builds an **orthonormal basis** (e_1, \dots, e_n) for the n -dimensional subspace generated by $v_j \in \mathbb{R}^m$ for $1 \leq j \leq n$ as follows:

1. To find e_1 , let $u_1 = v_1$, let $e_1 = \frac{u_1}{\|u_1\|}$.
2. To find e_2 , let $u_2 = v_2 - \text{proj}_{u_1} v_2 = v_2 - (e_1 \cdot v_2) e_1$. Then $e_2 = \frac{u_2}{\|u_2\|}$.
3. To find e_3 , let $u_3 = v_3 - \text{proj}_{u_1} v_3 - \text{proj}_{u_2} v_3 = v_3 - (e_1 \cdot v_3) e_1 - (e_2 \cdot v_3) e_2$. Then $e_3 = \frac{u_3}{\|u_3\|}$, and so on ...