On the Stability of Singular Values and Subspaces Under Perturbation

Arjun Veluri

Euler Circle

July 9, 2026

Today's agenda

Background and notation

Classical arc

SVD — Beltrami & Jordan (1870s)

Weyl (1912)

Eckart–Young (1936)

Davis–Kahan (1970)

Notation

$\tilde{A} = A + E$

$\|\cdot\|_2, \quad \|\cdot\|_F$

δ eigengap

$\sin\Theta(\tilde{V}, V)$ principal angles

The question

$$\tilde{A} = A + E \qquad (A = \text{signal}, \; E = \text{noise})$$

You run every method on $\tilde{A}$, not A. Two independent questions:

values: how far is $\tilde{\sigma}_i$ from σ_i? (*Weyl*)

subspaces: how far do they rotate? (*Davis–Kahan*)

The singular value decomposition

Theorem (Existence)

Every $A \in \mathbb{R}^{m\times n}$ *factors as* $A = U\Sigma V^{\mathsf{T}}$, *with* U, V *orthogonal and* Σ *diagonal,* $\sigma_1 \geq \sigma_2 \geq \cdots \geq 0$.

- exists for *every* matrix — no symmetry, no squareness
- σ_i: a ranked inventory of how strongly A acts in each direction

Constructing the SVD

Proof. $A^{\mathsf{T}}A$ is symmetric and positive semidefinite, so the spectral theorem gives $A^{\mathsf{T}}A = V\Lambda V^{\mathsf{T}}$ with $\lambda_i \geq 0$; set $\sigma_i = \sqrt{\lambda_i}$.

Constructing the SVD

Proof. $A^\mathsf{T} A$ is symmetric and positive semidefinite, so the spectral theorem gives $A^\mathsf{T} A = V \Lambda V^\mathsf{T}$ with $\lambda_i \geq 0$; set $\sigma_i = \sqrt{\lambda_i}$.

For $\sigma_i > 0$, put $u_i = \frac{1}{\sigma_i} A v_i$. These are orthonormal, since

$$u_i^\mathsf{T} u_j = \frac{\lambda_j}{\sigma_i \sigma_j} v_i^\mathsf{T} v_j = \delta_{ij}.$$

Constructing the SVD

Proof. $A^{\mathsf{T}}A$ is symmetric and positive semidefinite, so the spectral theorem gives $A^{\mathsf{T}}A = V\Lambda V^{\mathsf{T}}$ with $\lambda_i \geq 0$; set $\sigma_i = \sqrt{\lambda_i}$.

For $\sigma_i > 0$, put $u_i = \frac{1}{\sigma_i}Av_i$. These are orthonormal, since

$$u_i^{\mathsf{T}} u_j = \frac{\lambda_j}{\sigma_i \sigma_j} v_i^{\mathsf{T}} v_j = \delta_{ij}.$$

Extend $\{u_i\}$ to an orthonormal basis, giving an orthogonal U.

Constructing the SVD

Proof. $A^\top A$ is symmetric and positive semidefinite, so the spectral theorem gives $A^\top A = V \Lambda V^\top$ with $\lambda_i \geq 0$; set $\sigma_i = \sqrt{\lambda_i}$.

For $\sigma_i > 0$, put $u_i = \frac{1}{\sigma_i} A v_i$. These are orthonormal, since

$$u_i^\top u_j = \frac{\lambda_j}{\sigma_i \sigma_j} v_i^\top v_j = \delta_{ij}.$$

Extend $\{u_i\}$ to an orthonormal basis, giving an orthogonal U.

On the basis $\{v_i\}$ we have $AV = U\Sigma$, hence $A = U \Sigma V^\top$. □

Geometry and uniqueness

$$Av_i = \sigma_i u_i \qquad \text{(singular-vector relation)}$$

- image of the unit sphere = ellipsoid with semi-axes $\sigma_i u_i$
- the σ_i are *always* unique
- u_i, v_i unique up to sign *iff* σ_i is simple
- repeated σ_i: only the *subspace* is determined — this is what forces the subspace viewpoint later

Best low-rank approximation

Definition

Truncated SVD: $A_k = \sum_{i=1}^{k} \sigma_i u_i v_i^{\mathsf{T}}$.

Theorem (Eckart–Young)

Among all B with $\operatorname{rank}(B) \le k$, A_k is closest to A — in both $\|\cdot\|_2$ and $\|\cdot\|_F$ — with errors σ_{k+1} and $\left(\sum_{i>k} \sigma_i^2\right)^{1/2}$.

Why it is optimal

$$\max_{\operatorname{rank} P=k} \operatorname{tr}(P\,AA^{\mathsf{T}}) = \sigma_1^2 + \cdots + \sigma_k^2 \qquad \text{(Ky Fan)}$$

Frobenius case, in two moves:

- fix the column space C: the best fit is the projection $P_C A$
- maximize $\operatorname{tr}(P_C AA^{\mathsf{T}})$ over rank-k projectors $\Rightarrow$ the top-k eigenvectors

So the optimum is $\operatorname{span}\{u_1, \ldots, u_k\}$, i.e. A_k.

Mirsky

The same A_k is optimal in *every* unitarily invariant norm.

Truncation, geometrically

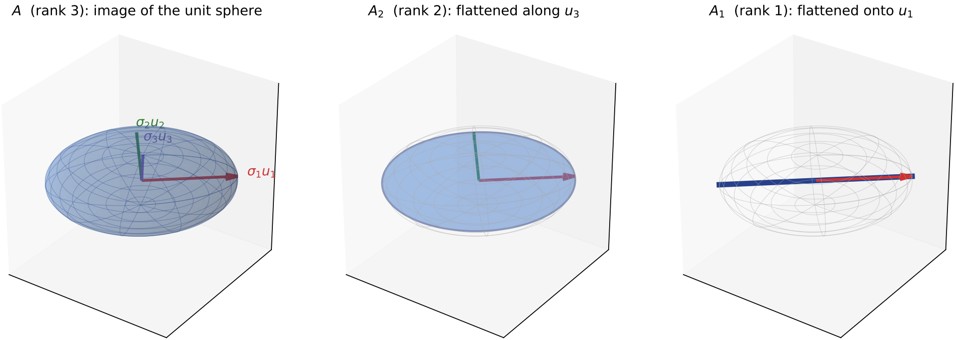

Discard the *shortest* axes; the error is the smallest σ_i. The rank-k matrices form a curved variety — yet the optimum is read straight off the SVD.

The engine: a variational formula

Theorem (Courant–Fischer / minimax)

$$\sigma_j = \max_{\dim S = j} \min_{\substack{x \in S \\ \|x\| = 1}} \|Ax\|.$$

This writes each singular value as a max–min of the single quantity $\|Ax\|$ — exactly what makes it stable under perturbation.

Weyl: the values barely move

Theorem (Weyl)

$$|\tilde{\sigma}_j - \sigma_j| \le \|\tilde{A} - A\|_2 \qquad \textit{for every } j.$$

Every singular value absorbs the size of the noise, and nothing more.

Weyl's Inequality

Proof. For a unit vector x, the triangle inequality gives

$$\|\tilde{A}x\| \leq \|Ax\| + \|Ex\| \leq \|Ax\| + \|E\|_2.$$

Weyl's Inequality

Proof. For a unit vector x, the triangle inequality gives

$$\|\tilde{A}x\| \leq \|Ax\| + \|Ex\| \leq \|Ax\| + \|E\|_2.$$

Fix a subspace S with $\dim S = j$ and minimize over unit $x \in S$:

$$\min_{x \in S} \|\tilde{A}x\| \leq \min_{x \in S} \|Ax\| + \|E\|_2.$$

Weyl's Inequality

Proof. For a unit vector x, the triangle inequality gives

$$\|\tilde{A}x\| \leq \|Ax\| + \|Ex\| \leq \|Ax\| + \|E\|_2.$$

Fix a subspace S with $\dim S = j$ and minimize over unit $x \in S$:

$$\min_{x \in S} \|\tilde{A}x\| \leq \min_{x \in S} \|Ax\| + \|E\|_2.$$

Maximizing over such S and applying the minimax formula, $\tilde{\sigma}_j \leq \sigma_j + \|E\|_2$.

Weyl's Inequality

Proof. For a unit vector x, the triangle inequality gives

$$\|\tilde{A}x\| \leq \|Ax\| + \|Ex\| \leq \|Ax\| + \|E\|_2.$$

Fix a subspace S with $\dim S = j$ and minimize over unit $x \in S$:

$$\min_{x \in S} \|\tilde{A}x\| \leq \min_{x \in S} \|Ax\| + \|E\|_2.$$

Maximizing over such S and applying the minimax formula, $\tilde{\sigma}_j \leq \sigma_j + \|E\|_2$.

Swapping $A \leftrightarrow \tilde{A}$ gives the reverse bound, so $|\tilde{\sigma}_j - \sigma_j| \leq \|E\|_2$. $\square$

Relative stability

- 1-Lipschitz, and the constant 1 is best possible (perturb along $u_1 v_1^{\mathsf{T}}$)

Relative stability

- 1-Lipschitz, and the constant 1 is best possible (perturb along $u_1 v_1^\mathsf{T}$)
- relative shift: $\dfrac{|\tilde{\sigma}_i - \sigma_i|}{\sigma_i} \leq \dfrac{\|E\|_2}{\sigma_i}$

Relative stability

- 1-Lipschitz, and the constant 1 is best possible (perturb along $u_1 v_1^\mathsf{T}$)
- relative shift: $\dfrac{|\tilde{\sigma}_i - \sigma_i|}{\sigma_i} \le \dfrac{\|E\|_2}{\sigma_i}$

The large singular values are the stable ones — which is exactly why keeping the top of the SVD (Eckart–Young) survives noise.

Subspaces: the vector is the wrong object

When σ_i is repeated (or nearly so), v_i is *not* determined by A: any basis of the eigenspace of $A^{\mathsf{T}}A$ serves.

Definition (distance between subspaces)

$\|\sin\Theta(\tilde{V}, V)\|$, where $\sin\theta_i$ are the singular values of $(I - VV^{\mathsf{T}})\tilde{V}$.

Track the subspace, not the vector.

Davis–Kahan: separation is stability

Theorem (Davis–Kahan, $\sin\Theta$)

For symmetric $H, \tilde{H}$ with eigengap δ around the block:

$$\|\sin\Theta(\tilde{V}, V)\|_F \le \frac{\|\tilde{H} - H\|_F}{\delta}.$$

A direction is stable exactly when its eigenvalue is *isolated*.

Davis–Kahan Theorem

Proof. Let $P = VV^{\mathsf{T}}$; it commutes with H, and $\|\sin\Theta(\tilde{V}, V)\|_F = \|(I - P)\tilde{V}\|_F$.

Davis–Kahan Theorem

Proof. Let $P = VV^{\mathsf{T}}$; it commutes with H, and $\|\sin\Theta(\tilde{V}, V)\|_F = \|(I - P)\tilde{V}\|_F$.

For each $\tilde{v}_j$, the eigen-relation gives $(H - \tilde{\lambda}_j I)\tilde{v}_j = (H - \tilde{H})\tilde{v}_j$.

Davis–Kahan Theorem

Proof. Let $P = VV^{\mathsf{T}}$; it commutes with H, and $\|\sin\Theta(\tilde{V}, V)\|_F = \|(I - P)\tilde{V}\|_F$.

For each $\tilde{v}_j$, the eigen-relation gives $(H - \tilde{\lambda}_j I)\tilde{v}_j = (H - \tilde{H})\tilde{v}_j$.

On range$(I - P)$ the matrix $H - \tilde{\lambda}_j I$ expands by at least δ, so

$$\delta \, \|(I - P)\tilde{v}_j\| \leq \|(\tilde{H} - H)\tilde{v}_j\|.$$

Davis–Kahan Theorem

Proof. Let $P = VV^{\mathsf{T}}$; it commutes with H, and $\|\sin\Theta(\tilde{V}, V)\|_F = \|(I-P)\tilde{V}\|_F$.

For each $\tilde{v}_j$, the eigen-relation gives $(H - \tilde{\lambda}_j I)\tilde{v}_j = (H - \tilde{H})\tilde{v}_j$.

On range$(I-P)$ the matrix $H - \tilde{\lambda}_j I$ expands by at least δ, so

$$\delta\,\|(I-P)\tilde{v}_j\| \le \|(\tilde{H} - H)\tilde{v}_j\|.$$

Summing over the block, $\delta\,\|\sin\Theta(\tilde{V}, V)\|_F \le \|\tilde{H} - H\|_F$. □

From eigenvectors to singular subspaces

Right singular vectors of A = eigenvectors of $A^\mathsf{T}A$. Apply Davis–Kahan to $H = A^\mathsf{T}A$:

$$\|\sin\Theta(\tilde{V}, V)\|_F \le \frac{(2\sigma_1 + \|E\|_2)\,\|E\|_F}{\delta}, \qquad \delta \text{ in the } \sigma_i^2.$$

Left subspaces: same with AA^T. A sharper σ-gap version comes from Wedin's dilation $\left[\begin{smallmatrix} 0 & A \\ A^\mathsf{T} & 0 \end{smallmatrix}\right]$ — omitted here to stay self-contained.

Gershgorin's Theorem

Theorem (Gershgorin)

Every eigenvalue lies in $\bigcup_i \{z : |z - a_{ii}| \le \sum_{j \ne i} |a_{ij}|\}$.

- entrywise — no norm, no SVD
- applied to $A^\top A$: *certifies* the gap δ that Davis–Kahan needs

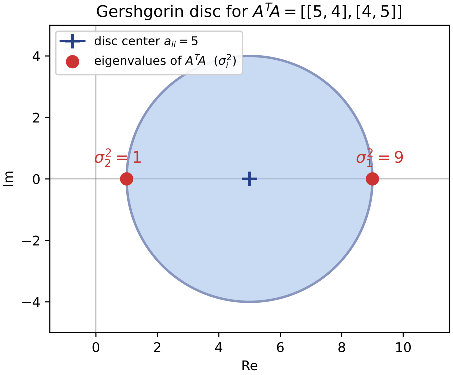

Why is this even useful

- **Machine learning:** principal component analysis / dimensionality reduction
- **Artificial intelligence:** low-rank adaptation of large language models (LoRA)
- **Data science:** covariance estimation
- **Computer science:** spectral clustering & community detection

Questions?

Eckart–Young (1936), Mirsky (1960), Davis–Kahan (1970),
Yu–Wang–Samworth (2015); Gershgorin (1931); Horn & Johnson, *Matrix Analysis*.