8.7 Singular Values

Rectangular matrices do not have e-vals (why?).

Gram matrices $\mathbf{K} = \mathbf{A}^T\mathbf{A}$ are square and symmetric for *any* $\mathbf{A}$.

So how do the e-vals of $\mathbf{K}$ relate to $\mathbf{A}$?

Definition: The singular value (s-val) $\sigma_1, \ldots, \sigma_r$ of $\mathbf{A}^{m\times n}$ are the positive square roots, $\sigma_i = \sqrt{\lambda_i} > 0$, of the nonzero e-vals of the Gram matrix $\mathbf{K} = \mathbf{A}^T\mathbf{A}$. The corresponding e-vecs of $\mathbf{K}$ are known as the singular vectors (s-vecs) of $\mathbf{A}$.

But what if $\lambda_i < 0$? It can't happen, recall that Gram matrices $\mathbf{K}$ are positive semidefinite ($\lambda_i \geq 0$), which justifies positivity of s-vals of $\mathbf{A}$ (independently of whether $\mathbf{A}$ itself has positive, negative, or even complex e-vals; or is rectangular and has no e-vals at all!).

We will label s-vals in decreasing order: $\sigma_1 \geq \sigma_2 \geq \ldots \geq \sigma_r > 0$.

Concretely: Let $\mathbf{A} = \begin{bmatrix} 3 & 5 \\ 4 & 0 \end{bmatrix}$. Observe: $\mathbf{K} = \mathbf{A}^T\mathbf{A} = \begin{bmatrix} 3 & 4 \\ 5 & 0 \end{bmatrix}\begin{bmatrix} 3 & 5 \\ 4 & 0 \end{bmatrix} = \begin{bmatrix} 25 & 15 \\ 15 & 25 \end{bmatrix}$.

$\mathbf{K}$ has $\lambda_1 = 40$, $\lambda_2 = 10$, and e-vecs: $\vec{v}_1 = (1,1)$, and $\vec{v}_2 = (1,-1)$.

Thus, s-vals of $\mathbf{A}$ are $\sigma_1 = \sqrt{40} \approx 6.325$ and $\sigma_2 = \sqrt{10} \approx 3.162$ with s-vecs $\vec{v}_1, \vec{v}_2$.

In particular, $\mathbf{A}$'s s-vals are **not** $\mathbf{A}$'s e-vals, which are $\lambda_1 \approx 6.217$ and $\lambda_2 \approx -3.217$, **nor are $\mathbf{A}$'s s-vecs the e-vecs of $\mathbf{A}$**.

Indeed, the e-vecs of $\mathbf{A}$ are $(-0.8043, 1)$ and $(1.554, 1)$.

Proposition: If $\mathbf{A} = \mathbf{A}^T$, then $\mathbf{A}$'s s-vals are the absolute values of its nonzero e-vals: $\sigma_i = |\lambda_i| > 0$. Also, $\mathbf{A}$'s s-vecs coincide with its non–null e-vecs.

Proof: When $\mathbf{A}$ is symmetric, $\mathbf{K} = \mathbf{A}^T\mathbf{A} = \mathbf{A}^2$.

So, if $\mathbf{A}\vec{v} = \lambda\vec{v}$, then $\mathbf{K}\vec{v} = \mathbf{A}^2\vec{v} = \mathbf{A}(\lambda\vec{v}) = \lambda\mathbf{A}\vec{v} = \lambda^2\vec{v}$.

So, $\sigma = \sqrt{\lambda^2} = |\lambda| > 0$.

Thus, every e-vec $\vec{v}$ of $\mathbf{A}$ is also an e-vec of $\mathbf{K}$ with $\mathbf{K}$ e-val λ^2. ■

Also, observe that the e-vec basis of symmetric $\mathbf{A}$ (guaranteed by previous thm) is also an e-vec basis for $\mathbf{K}$,

and hence forms a complete system of s-vecs for $\mathbf{A}$.

Singular Value Decomposition (SVD)

Recall spectral (e-basis) factorization of symmetric matrices: $\mathbf{A} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{-1} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^T$.

We can generalize this to nonsymmetric matrices, this is known as **singular value decomposition**.

Theorem: A nonzero real $\mathbf{A}^{m\times n}$ of rank $r > 0$ can be factored,

$$\mathbf{A} = \mathbf{P}\mathbf{\Sigma}\mathbf{Q}^T \qquad (*)$$

where $\mathbf{P}^{m\times r}$ has orthonormal columns, so $\mathbf{P}^T\mathbf{P} = \mathbf{I}$. The diagonal $\mathbf{\Sigma}^{r\times r} = diag(\sigma_1, \dots, \sigma_r)$ has the s-vals of $\mathbf{A}$ as diagonal entries, and $\mathbf{Q}^T$ is $r \times n$ with orthonormal rows, so $\mathbf{Q}^T\mathbf{Q} = \mathbf{I}$, where $\mathbf{Q} = [\vec{q}_i]$ and the $\vec{q}_i$ are orthonormal e-vecs of the Gram matrix $\mathbf{K} = \mathbf{A}\mathbf{A}^T$.

Proof: Let's begin by rewriting $(*)$ as $\mathbf{A}\mathbf{Q} = \mathbf{P}\mathbf{\Sigma}$.

(this is allowed since the $\vec{q}_i$ are the orthonormal e-vecs of $\mathbf{K}$ corresponding to the nonzero e-vals. So, $\mathbf{Q}$ is invertible)

The individual columns of this equation are: $\mathbf{A}\vec{q}_i = \sigma_i\vec{p}_i$, where $i = 1, \dots, r$. $(**)$

This eq. relates orthonormal columns of $\mathbf{Q} = \left[\vec{q}_1, \dots, \vec{q}_r\right]$ to orthonormal columns of $\mathbf{P} = \left[\vec{p}_1, \dots, \vec{p}_r\right]$.

Thus, our goal is to find orthonormal $\vec{p}_1, \dots, \vec{p}_r$.

Recall that the $\mathbf{K}$ e-vecs $\vec{q}_i$ (according to a previous proposition), form a basis for $img\,\mathbf{K} = coimg\,\mathbf{A}$ of dimension $r = rank\,\mathbf{A}$.

Thus, by the definition of the s-vals: $\mathbf{A}^T\mathbf{A}\vec{q}_i = \mathbf{K}\vec{q}_i = \sigma_i^2\vec{q}_i$, where $i = 1, \dots, r$. $(***)$

We claim that the image vecs $\vec{w}_i = \mathbf{A}\vec{q}_i$ are automatically orthogonal.

Indeed, in view of the orthonormality of the $\vec{q}_i$ combined with $(***)$, we have:

$$\vec{w}_i \cdot \vec{w}_j = \vec{w}_i^T\vec{w}_j = \left(\mathbf{A}\vec{q}_i\right)^T\mathbf{A}\vec{q}_j = \vec{q}_i^T\mathbf{A}^T\mathbf{A}\vec{q}_j$$

$$= \vec{q}_i^T \sigma_j^2 \vec{q}_j = \sigma_j^2 \vec{q}_i^T \vec{q}_j = \sigma_j^2 \vec{q}_i \cdot \vec{q}_j = \begin{cases} 0, & i \neq j, \\ \sigma_i^2, & i = j. \end{cases}$$

Consequently, $\vec{w}_1, \ldots, \vec{w}_r$ form an orthogonal system of vecs having: $\left|\vec{w}_i\right| = \sqrt{\vec{w}_i \cdot \vec{w}_i} = \sigma_i$.

So, the associated unit vecs: $\vec{p}_i = \frac{\vec{w}_i}{\sigma_i} = \frac{\mathbf{A}\vec{q}_i}{\sigma_i}$, $\qquad (****)$

where $i = 1, \ldots, r$, form an orthonormal set of vecs.

Rearranging this equation, we find: $\mathbf{A}\vec{q}_i = \sigma_i \vec{p}_i$, satisfying $(**)$. ■

Corollary: $\mathbf{A}$ and $\mathbf{A}^T$ have the same s-vals.

Proof: Observe that taking the transpose of $(*)$ (and noting $\mathbf{\Sigma}^T = \mathbf{\Sigma}$ is diagonal), we obtain: $\mathbf{A}^T = \mathbf{Q\Sigma P}^T$, which is a SVD of $\mathbf{A}^T$. ■

Observe that the s-vecs are not the same. Indeed, those of $\mathbf{A}$ are the orthogonal columns of $\mathbf{Q}$, where as those of $\mathbf{A}^T$ are the orthonormal columns of $\mathbf{P}$.

The SVD serves to diagonalize the Gram matrix $\mathbf{K}$. Indeed, since $\mathbf{P}^T\mathbf{P} = \mathbf{I}$, we have: $\mathbf{Q}^T\mathbf{K}\mathbf{Q} =$

$$= \mathbf{Q}^T(\mathbf{A}^T\mathbf{A})\mathbf{Q}$$

$$= \mathbf{Q}^T\mathbf{A}^T(\mathbf{P}\mathbf{P}^T)\mathbf{A}\mathbf{Q}$$

$$= (\mathbf{P}^T\mathbf{A}\mathbf{Q})^T(\mathbf{P}^T\mathbf{A}\mathbf{Q}) = \mathbf{\Sigma}^T\mathbf{\Sigma} = \mathbf{\Sigma}^2. \qquad (**) \qquad (\text{since } \mathbf{A} = \mathbf{P\Sigma Q}^T)$$

If $\mathbf{A}$ has rank n, then $\mathbf{Q}$ is an $n \times n$ orthogonal matrix and so $(**)$ implies that the linear transformation of $\mathbb{R}^n$ by $\mathbf{K}$ is diagonalized when expressed in terms of the orthonormal basis formed by the s-vecs.

Concretely: For $\mathbf{A} = \begin{bmatrix} 3 & 5 \\ 4 & 0 \end{bmatrix}$ seen above, find SVD.

We had calculated $\mathbf{K} = \mathbf{A}^T\mathbf{A} = \begin{bmatrix} 25 & 15 \\ 15 & 25 \end{bmatrix}$, with $\sigma_1 = \sqrt{40}$ and $\sigma_2 = \sqrt{10}$.

Normalizing the $\mathbf{K}$ e-vecs found above gives orthonormal basis of $\mathbf{A}$ s-vecs: $\vec{q}_1 = \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix}$, $\vec{q}_2 = \begin{bmatrix} -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix}$.

Thus, $\mathbf{Q} = \begin{bmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix}$.

Next, according to $(***)$, we have: $\vec{p}_1 = \frac{\mathbf{A}\vec{q}_1}{\sigma_1} = \frac{1}{\sqrt{40}}\begin{bmatrix} 4\sqrt{2} \\ 2\sqrt{2} \end{bmatrix} = \begin{bmatrix} \frac{2}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} \end{bmatrix}$,

$\vec{p}_2 = \frac{\mathbf{A}\vec{q}_2}{\sigma_2} = \frac{1}{\sqrt{10}}\begin{bmatrix} \sqrt{2} \\ -2\sqrt{2} \end{bmatrix} = \begin{bmatrix} \frac{1}{\sqrt{5}} \\ -\frac{2}{\sqrt{5}} \end{bmatrix}$, and thus $\mathbf{P} = \begin{bmatrix} \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} & -\frac{2}{\sqrt{5}} \end{bmatrix}$.

Checking my work: $\mathbf{A} = \begin{bmatrix} 3 & 5 \\ 4 & 0 \end{bmatrix} = \begin{bmatrix} \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} & -\frac{2}{\sqrt{5}} \end{bmatrix}\begin{bmatrix} \sqrt{40} & 0 \\ 0 & \sqrt{10} \end{bmatrix}\begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} = \mathbf{P\Sigma Q}^T$. ✓

Proposition: Given the SVD $\mathbf{A} = \mathbf{P\Sigma Q}^T$, the columns $\vec{q}_1, \ldots, \vec{q}_r$ of $\mathbf{Q}$ form an orthonormal basis for *coimg* $\mathbf{A}$, while the columns $\vec{p}_1, \ldots, \vec{p}_r$ of $\mathbf{P}$ form and orthonormal basis for *img* $\mathbf{A}$.

Proof: The first part of the proposition is automatic, since the $\vec{q}_1, \ldots, \vec{q}_r$ were defined to be the orthonormal e-vecs of $\mathbf{K} = \mathbf{A}^T\mathbf{A}$, and therefore in *coimg* $\mathbf{A}$, which has the same dimensions r as *img* $\mathbf{A}$.

Moreover, $\vec{p}_i = \mathbf{A}\left(\frac{\vec{q}_i}{\sigma_i}\right)$ for $i = 1, \ldots, r$ were shown in the above proof to be mutually orthogonal, of unit length, and belong to *img* $\mathbf{A}$. They therefore form an orthonormal basis for the image. ■

For SVD ($\mathbf{A} = \mathbf{P\Sigma Q}^T$), matrix $\mathbf{Q}^T$ represents an orthogonal projection from $\mathbb{R}^n$ to *coimg* $\mathbf{A}$, then $\mathbf{\Sigma}$ represents a stretching transformation within the r-dim subspace, while $\mathbf{P}$ maps the results to *img* $\mathbf{A} \subset \mathbb{R}^m$.

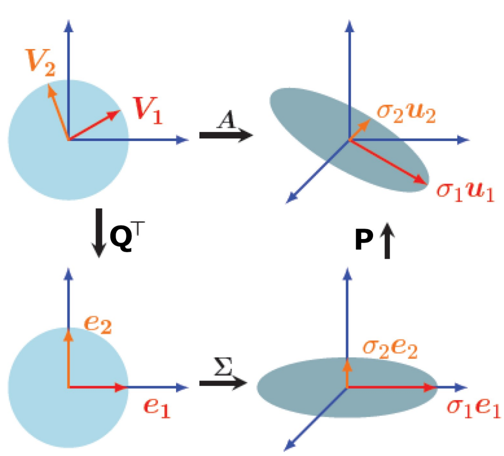

$\mathbf{A}^{3\times 2}$ SVD when $r = n$.

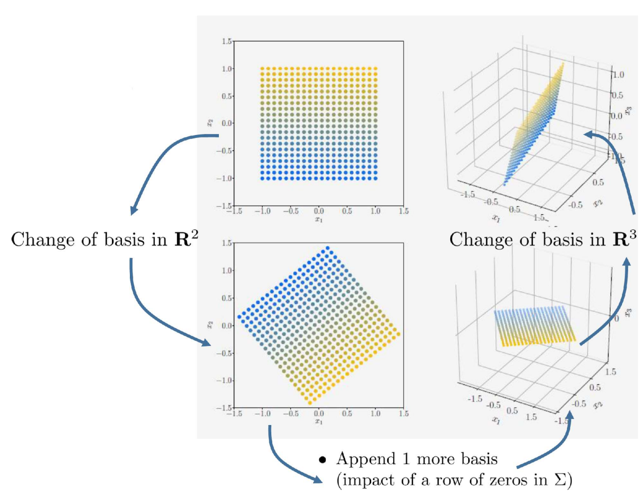

We have finally reached a complete understanding of the subtle geometry underlying the simple operation of multiplying a vector by a matrix!

Example: **True**/**False**? If **A** is symmetric, then its s-vals are the same as its e-vals.

False: $\sigma_i = |\lambda_i| > 0$; its s-vecs coincide with its **non-null** e-vecs.

Example: **True**/**False**? The s-vals of $\mathbf{A}^2$ are the squares of the s-vals of **A**.

False: $\mathbf{A}^2 = (\mathbf{P\Sigma Q}^T)^2 = \mathbf{P\Sigma Q}^T\mathbf{P\Sigma Q}^T$ $\ldots$??

Let: $\mathbf{A} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$. $\mathbf{K}_1 = \mathbf{A}^T\mathbf{A} = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$. $\sigma \in \{1\}$.

Observe: $\mathbf{A}^2 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$. $\mathbf{K}_2 = (\mathbf{A}^2)^T\mathbf{A}^2 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$. No S-vals!

The Pseudoinverse

Many matrices do not have an inverse, but we can generalize the idea of an inverse in a useful way.

Definition: The *pseudoinverse* of a nonzero $m \times n$ matrix with SVD $\mathbf{A} = \mathbf{P\Sigma Q}^T$ is
the $n \times m$ matrix $\mathbf{A}^+ := \mathbf{Q\Sigma}^{-1}\mathbf{P}^T$.

If $\mathbf{A}^{n\times n}$ is nonsingular, then $\mathbf{A}^+ = \mathbf{A}^{-1}$.

$$\mathbf{A}^{-1} = (\mathbf{P\Sigma Q}^T)^{-1} = (\mathbf{Q}^{-1})^T\mathbf{\Sigma}^{-1}\mathbf{P}^{-1} = \mathbf{Q\Sigma}^{-1}\mathbf{P}^T = \mathbf{A}^+.$$

But there is a quicker way:

Lemma: Let $\mathbf{A}^{m\times n}$ have rank n. Then $\mathbf{A}^+ = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T$.

Proof: Observe: $\mathbf{A}^T\mathbf{A} = (\mathbf{P\Sigma Q}^T)^T(\mathbf{P\Sigma Q}^T) = \mathbf{Q\Sigma P}^T\mathbf{P\Sigma Q}^T = \mathbf{Q\Sigma}^2\mathbf{Q}^T$,
since $\mathbf{\Sigma}=\mathbf{\Sigma}^T$ is a diagonal matrix, and $\mathbf{P}^T\mathbf{P} = \mathbf{I}$.

This is spectral factorization of Gram matrix $\mathbf{A}^T\mathbf{A}$ — which we already knew from original definition of s-vals and s-vecs.

If **A** has rank n, then **Q** is $n \times n$ orthogonal. So $\mathbf{Q}^{-1} = \mathbf{Q}^T$.

Therefore, $(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T = (\mathbf{Q\Sigma}^2\mathbf{Q}^T)^{-1}(\mathbf{P\Sigma Q}^T)^T$

$$= (\mathbf{Q\Sigma}^{-2}\mathbf{Q}^T)(\mathbf{Q\Sigma P}^T)$$

$= \mathbf{Q}\boldsymbol{\Sigma}^{-1}\mathbf{P}^T = \mathbf{A}^+.$ ■

Say we want to solve $\mathbf{A}\vec{x} = \vec{b}$. Rearranging, we want $\vec{x}$ such that $\mathbf{A}\vec{x} - \vec{b} = \vec{0}$.

In real life, this often is not possible. So, instead we look for $\vec{x}$ that minimizes $|\vec{r}| := |\mathbf{A}\vec{x} - \vec{b}|$.

This is known as the *least squares solution* to the linear system, because $|\vec{r}|^2 = r_1^2 + \ldots + r_n^2$ is the sum of the squares of the individual error components.

Theorem: Consider $\mathbf{A}\vec{x} = \vec{b}$. Let $\vec{x}^* = \mathbf{A}^+\vec{b}$. If $\ker \mathbf{A} = \{\vec{0}\}$, then $\vec{x}^*$ is the (Euclidean) least-squares solution to $\mathbf{A}\vec{x} = \vec{b}$. If, more generally, $\ker \mathbf{A} \neq \{\vec{0}\}$, then $\vec{x}^* = \mathbf{A}^+\vec{b} \in \textit{coimg}\,\mathbf{A}$ is the least-squares solution that has the minimal Euclidean norm ($|\vec{x}^*| \leq |\vec{x}|$) among all $\vec{x}$ that minimize the least-squares error $|\mathbf{A}\vec{x} - \vec{b}|^2$.

Proof: Relies on a section we did not cover this semester.

Concretely: Find the pseudoinverse of $\mathbf{A} = \begin{bmatrix} 2 & 0 \\ 0 & -1 \\ 0 & 0 \end{bmatrix}$.

Observe that this 3×2 matrix has $\textit{rank}\,\mathbf{A} = 2$. Therefore, the (quicker way) lemma above applies and:

$$\mathbf{A}^+ = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T = \left(\begin{bmatrix} 2 & 0 & 0 \\ 0 & -1 & 0 \end{bmatrix} \begin{bmatrix} 2 & 0 \\ 0 & -1 \\ 0 & 0 \end{bmatrix} \right)^{-1} \begin{bmatrix} 2 & 0 & 0 \\ 0 & -1 & 0 \end{bmatrix}$$

$$= \begin{bmatrix} 4 & 0 \\ 0 & 1 \end{bmatrix}^{-1} \begin{bmatrix} 2 & 0 & 0 \\ 0 & -1 & 0 \end{bmatrix}$$

$$= \begin{bmatrix} \frac{1}{4} & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 2 & 0 & 0 \\ 0 & -1 & 0 \end{bmatrix} = \begin{bmatrix} \frac{1}{2} & 0 & 0 \\ 0 & -1 & 0 \end{bmatrix}.$$

Now find the least-squares solution of $\mathbf{A}\vec{x} = \begin{bmatrix} -1 \\ 3 \\ -4 \end{bmatrix}$ that has the minimal Euclidean norm ($|\vec{x}^*| \leq |\vec{x}|$).

$$\mathbf{A}^+\vec{b} = \begin{bmatrix} \frac{1}{2} & 0 & 0 \\ 0 & -1 & 0 \end{bmatrix} \begin{bmatrix} -1 \\ 3 \\ -4 \end{bmatrix} = \begin{bmatrix} -\frac{1}{2} \\ -3 \end{bmatrix}.$$

The Euclidean Matrix Norm

Theorem: Let $|\bullet|_2$ denote the Euclidean norm on $\mathbb{R}^n$. Let $\mathbf{A}$ be nonzero with s-vals $\sigma_1 \geq \ldots \geq \sigma_r$. Then, the Euclidean norm of $\mathbf{A}$, defined as $|\mathbf{A}|_2 := \max\left\{|\mathbf{A}\vec{u}|_2 : |\vec{u}|_2 = 1\right\}$ equals its dominant (largest) s-val. So: $\max\left\{|\mathbf{A}\vec{u}|_2 : |\vec{u}|_2 = 1\right\} = \sigma_1$, while $|\mathbf{O}|_2 = 0$.

Proof: Note that we don't need to prove the definition, only that $\max\left\{|\mathbf{A}\vec{u}|_2 : |\vec{u}|_2 = 1\right\} = \sigma_1$.

Let $\vec{q}_1, \ldots, \vec{q}_n$ be an orthonormal basis of $\mathbb{R}^n$ consisting of the s-vecs $\vec{q}_1, \ldots, \vec{q}_r$ along with an orthonormal basis $\vec{q}_{r+1}, \ldots, \vec{q}_n$ of ker $\mathbf{A}$.

Thus by a thm in sect 8.6 (not covered by this class), $\mathbf{A}\vec{q}_i = \begin{cases} \sigma_i\vec{p}_i, & i = 1, \ldots, r, \\ 0, & i = r+1, \ldots, n \end{cases}$,

where $\vec{p}_1, \ldots, \vec{p}_r$ form and orthonormal basis for *img* $\mathbf{A}$.

Suppose $\vec{u}$ is any unit vector, so $\vec{u} = c_1\vec{q}_1 + \ldots + c_n\vec{q}_n$, where $|\vec{u}| = \sqrt{c_1^2 + \ldots + c_n^2} = 1$,

thanks to the orthonormality of the basis vecs and the Pythagorean formula. Then, $\mathbf{A}\vec{u} =$??

$\mathbf{A}\vec{u} = c_1\sigma_1\vec{p}_1 + \ldots + c_r\sigma_r\vec{p}_r$, and hence $|\mathbf{A}\vec{u}|_2 = \sqrt{c_1^2\sigma_1^2 + \ldots + c_r^2\sigma_r^2}$, since $\vec{p}_1, \ldots, \vec{p}_n$ are also orthonormal.

Now, since $\sigma_1 \geq \sigma_2 \geq \ldots \geq \sigma_r$, we have $|\mathbf{A}\vec{u}|_2 = \sqrt{c_1^2\sigma_1^2 + \ldots + c_r^2\sigma_r^2} \leq \sqrt{c_1^2\sigma_1^2 + \ldots + c_r^2\sigma_1^2} = \sigma_1\sqrt{c_1^2 + \ldots + c_n^2} = \sigma_1$.

Moreover, if $c_1 = 1,\ c_2 = \ldots = c_n = 0$, then $\vec{u} = \vec{q}_1$, and hence $|\mathbf{A}\vec{u}|_2 = |\mathbf{A}\vec{q}_1|_2 = |\sigma_1\vec{p}_1|_2 = \sigma_1$.

This implies the desired formula. ■

Concretely: Consider $\mathbf{A} = \begin{bmatrix} 0 & -\frac{1}{3} & \frac{1}{3} \\ \frac{1}{4} & 0 & \frac{1}{2} \\ \frac{2}{5} & \frac{1}{5} & 0 \end{bmatrix}$. What is its Euclidean norm?

Gram matrix $\mathbf{A}^T\mathbf{A} = \begin{bmatrix} \frac{89}{400} & \frac{2}{25} & \frac{1}{8} \\ \frac{2}{25} & \frac{34}{225} & -\frac{1}{9} \\ \frac{1}{8} & -\frac{1}{9} & \frac{13}{36} \end{bmatrix}$ has e-vals $\lambda_1 \approx 0.447,\ \lambda_2 \approx 0.267,\ \lambda_3 \approx 0.021$, and hence the s-vals of $\mathbf{A}$ are their square roots: $\sigma_1 \approx 0.669, \ \sigma_2 \approx 0.516, \ \sigma_3 \approx 0.145$.

The Euclidean norm of $\mathbf{A}$ is the largest s-val, and so $|\mathbf{A}|_2 \approx 0.669$.

Condition Number and Rank

Not only do s-vals provide a compelling geometric interpretation of the action of a matrix on a vector, s-vals also play a role in computer algorithms.

The magnitudes of s-vals can be used to distinguish a well behaved linear system from ill-conditioned ones, which are more challenging to solve accurately.

This information is quantified by the **condition number** κ of a matrix.

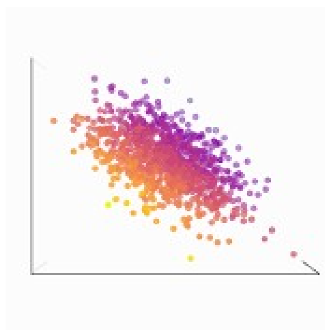

[see animation in class]

Definition: The *condition number* κ of a nonsingular $n \times n$ matrix is the ratio between its largest and smallest s-vals: $\kappa(\mathbf{A}) = \frac{\sigma_1}{\sigma_n}$.

Since the number of s-vals equals the matrix's rank, an $n \times n$ matrix with fewer than n s-vals is singular, and is said to have condition number ∞.

A matrix with a very large condition number is close to singular, and is designated as ill-conditioned.

In practical terms, this occurs when the condition number is larger than the reciprocal of the machine's precision, e.g., 10^7.

It is much harder to solve $\mathbf{A}\vec{x} = \vec{b}$ when its coefficient matrix is ill-conditioned, and hence close to singular.

Theorem: Let $\mathbf{A}^{m \times n}$ have rank r and SVD $\mathbf{A} = \mathbf{P}\mathbf{\Sigma}\mathbf{Q}^T$.
Given $1 \leq k \leq r$, let $\mathbf{\Sigma}_k$ denote the upper left $k \times k$ diagonal submatrix of $\mathbf{\Sigma}$ containing the largest k s-vals on $\mathbf{\Sigma}'$s diagonal.
Let $\mathbf{Q}_k$ be $n \times k$ formed from the first k columns of $\mathbf{Q}$, which are the first k orthonormal s-vecs of $\mathbf{A}$, and let $\mathbf{P}_k$ be $m \times k$ formed from the first k columns of $\mathbf{P}$.
Then $m \times n$ matrix $\mathbf{A}_k = \mathbf{P}_k\mathbf{\Sigma}_k\mathbf{Q}_k^T$ has rank k. Moreover, $\mathbf{A}_k$ is the closest rank k matrix to $\mathbf{A}$ in the sense that, among all $m \times n$ matrices $\mathbf{B}$ of rank k, the Euclidean matrix norm $|\mathbf{A} - \mathbf{B}|$ is minimized when $\mathbf{B} = \mathbf{A}_k$.

Proof in book.

Can’t do better than this with matrix of lower rank: $|\mathbf{A} - \mathbf{B}|$ is minimized when $\mathbf{B} = \mathbf{A}_k$ among all matrices with $rank\,\mathbf{B} \leq k$.

So, when solving ill-conditioned $\mathbf{A}\vec{x} = \vec{b}$, a strategy is to eliminate "insignificant" s-vals below a cut off, replacing $\mathbf{A}$ by $\mathbf{A}_k$.

Applying the corresponding approximating pseudoinverse $\mathbf{A}_k^+ = \mathbf{Q}_k \mathbf{\Sigma}_k^{-1} \mathbf{P}_k^T$ to solve for $\vec{x}^* = \mathbf{A}_k^+ \vec{b}$ will usually circumvent the effects of ill-conditioning.

Image Compression

How do computers store images?

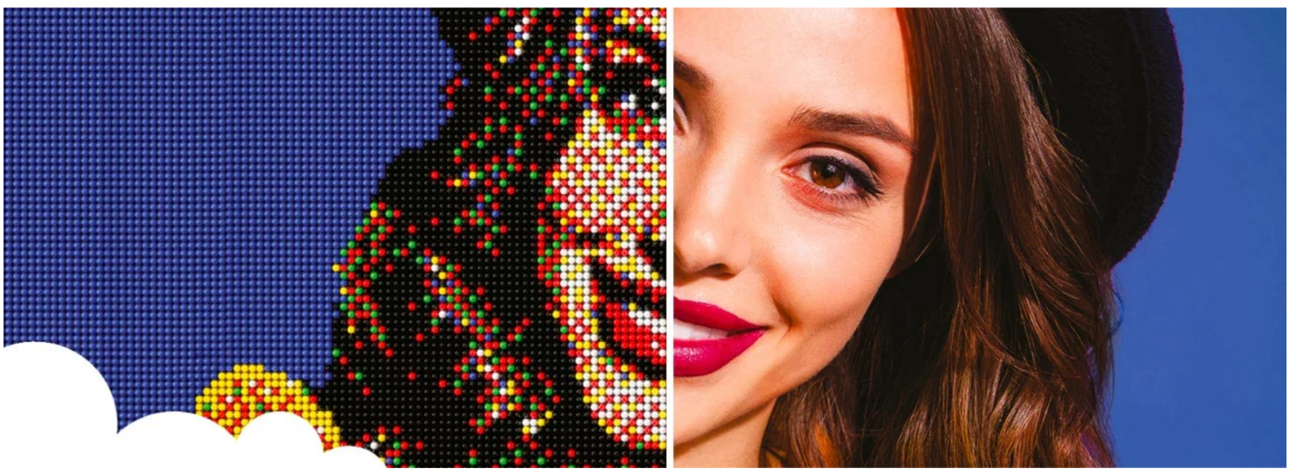

Divide an image into pixels, n pixels wide, and m pixels high. This gives us an $\mathbf{A}^{m\times n}$ matrix.

But how do we define each matrix component?

Pixel: Each color can be defined by some mixture of red/blue/green.

Mathematically, we will use hexadecimals $\{0, 1, \ldots, 9, A, B, \ldots, F\}$.

So, let each color be two digits, this gives us $16 \times 16 = 256$ shades of each color (where $00 =$ black and $FF =$ white).

And with three primary colors, this gives us $256^3 = 16,777,216$ numbers (colors) per pixel.

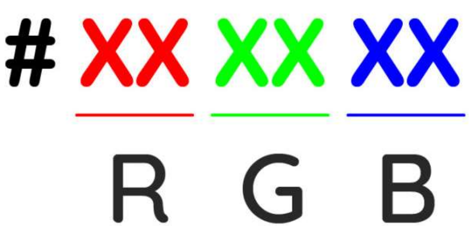

#000000 = Black
#FFFFFF = White
#A0A0A0 = Gray

#FF0000 = Red
#00FF00 = Green
#0000FF = Blue

For a 100 × 100 pixel image, you must store $10,000 \cdot 256^3 \approx 168$ billion choices of color.

For a 1280 × 720 pixel image (HD standard), that is $1280 \cdot 720 \cdot 256^3 \approx 15.5$ trillion choices of color, takes up a lot of resources!

There must be a different way, right?

The Different Way

Throw away the less relevant data.

Recall (?) from dynamical systems that there are e-val/e-vec pairs that dominate the behavior of the system. That is, the ones associated with the e-vals of greatest magnitude.

Example: if $(\lambda_i, v_i) \in \left\{(10, \vec{v}_1), (5, \vec{v}_2), \left(\frac{1}{2} \vec{v}_3\right)\right\}$ for a discrete dynamical system, then for most initial states, what's the long-term behavior of the system? ...

Grows without bound, parallel to $\vec{v}_1$.

But the image matrix above won't be square (needed for e-vals), so we need SVD: $\mathbf{A} = \mathbf{P\Sigma Q}^T$.

This allows us to compress data by only using the first $\boldsymbol{k}$ columns of $\mathbf{P}$, the upper left $\boldsymbol{k} \times \boldsymbol{k}$ submatrix of $\mathbf{\Sigma}$, and the first $\boldsymbol{k}$ rows of $\mathbf{Q}^T$.

$$\mathbf{A} = \begin{bmatrix} \boldsymbol{u}_{11} & \boldsymbol{u}_{12} & u_{13} & \cdots & u_{1m} \\ \boldsymbol{u}_{21} & \boldsymbol{u}_{22} & u_{23} & \cdots & u_{2m} \\ \boldsymbol{u}_{31} & \boldsymbol{u}_{32} & u_{33} & \cdots & u_{3m} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{u}_{m1} & \boldsymbol{u}_{m2} & u_{m3} & \cdots & u_{mm} \end{bmatrix} \begin{bmatrix} \sqrt{\lambda_1} & & & & \\ & \sqrt{\lambda_2} & & & \\ & & \sqrt{\lambda_3} & & \\ & & & \ddots & \\ & & & & \sqrt{\lambda_m} \end{bmatrix} \begin{bmatrix} \boldsymbol{v}_{11} & \boldsymbol{v}_{12} & \boldsymbol{v}_{13} & \cdots & \boldsymbol{v}_{1n} \\ \boldsymbol{v}_{21} & \boldsymbol{v}_{22} & \boldsymbol{v}_{23} & \cdots & \boldsymbol{v}_{2n} \\ v_{31} & v_{32} & v_{33} & \cdots & v_{3n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ v_{n1} & v_{n2} & v_{n3} & \cdots & v_{nn} \end{bmatrix}$$

$m \times \mathbf{2}$ $\qquad$ $\mathbf{2} \times \mathbf{2}$ $\qquad$ $\mathbf{2} \times n$

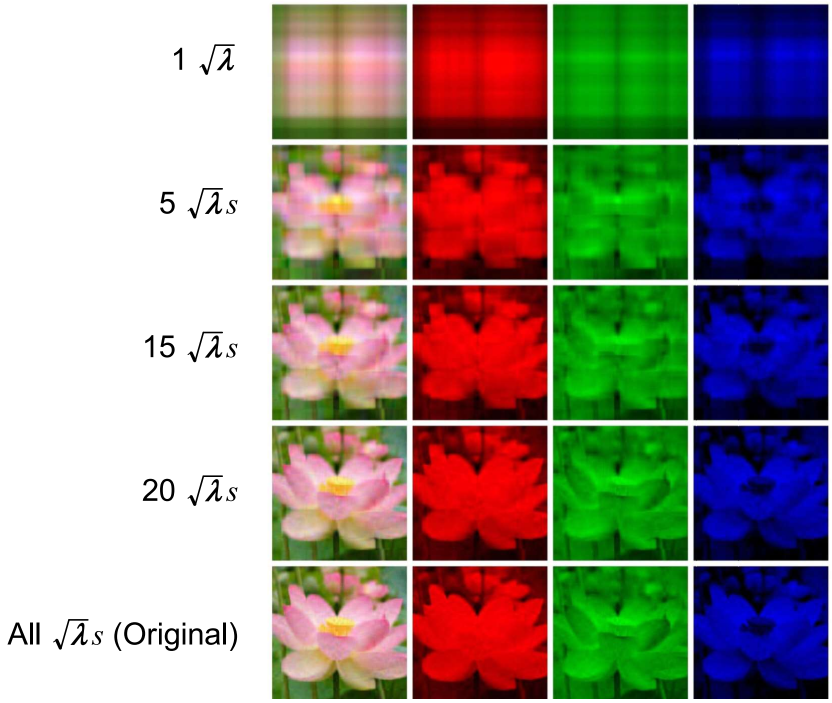

Learn more: overbye.engr.tamu.edu/wp-content/uploads/sites/146/2020/10/ECEN615_Fall2020_Lect17.pdf