Geometry of the SVD II

Here we give a geometric interpretation of the outer product form that will be useful for the PCA.

$$A = \sigma_1 u_1 v_1^T + \cdots + \sigma_r u_r v_r^T \qquad A = \begin{pmatrix} | & & | \\ d_1 & \cdots & d_n \\ | & & | \end{pmatrix}$$

$$A v_i = \sigma_i u_i \qquad A^T u_i = \sigma_i v_i$$

Expand out:

$$\sigma_i v_i = A^T u_i = \begin{pmatrix} — d_1^T — \\ \vdots \\ — d_n^T — \end{pmatrix} u_i = \begin{pmatrix} d_1 \cdot u_i \\ \vdots \\ d_n \cdot u_i \end{pmatrix}$$

$$\Rightarrow \sigma_i u_i v_i^T = u_i (\sigma_i v_i)^T = u_i \begin{pmatrix} d_1 \cdot u_i & \cdots & d_n \cdot u_i \end{pmatrix}$$

$$= \begin{pmatrix} | & & | \\ (d_1 \cdot u_i) u_i & \cdots & (d_n \cdot u_i) u_i \\ | & & | \end{pmatrix}$$

Since $\|u_i\| = 1$,

$(d \cdot u_i) u_i =$ orthogonal projection of d onto $\text{Span}\{u_i\}$.

> The columns of $\sigma_i u_i v_i^T$ are the
> orthogonal projections
> of the columns of A onto $\text{Span}\{u_i\}$.

Now look at the sum:

$$A = \sigma_1 u_1 v_1^T + \cdots + \sigma_r u_r v_r^T$$

The i^{th} column of this sum is:

$$\underset{\text{of } A}{i^{th}\text{ col}} \rightarrow d_i = (d_i \cdot u_1)u_1 + \cdots + (d_i \cdot u_r)u_r$$

Since $\{u_1, \ldots, u_r\}$ is an orthonormal basis of $\mathrm{Col}(A)$, this is just the projection formula as applied to d_i: the projection of d_i onto $\mathrm{Col}(A)$ is just d_i since $d_i \in \mathrm{Col}(A)$ (it's the i^{th} column of A).

Eg: $A = \begin{pmatrix} 3 & -4 & 7 & 1 & -4 & -3 \\ 7 & -6 & 8 & -1 & -1 & -7 \end{pmatrix}$ (from last time)

$$A = \sigma_1 u_1 v_1^T + \sigma_2 u_2 v_2^T$$

$\sigma_1 \approx 16.9 \qquad \sigma_2 \approx 3.92$

$u_1 \approx \begin{pmatrix} 0.561 \\ 0.828 \end{pmatrix} \qquad u_2 \approx \begin{pmatrix} 0.828 \\ -0.561 \end{pmatrix}$

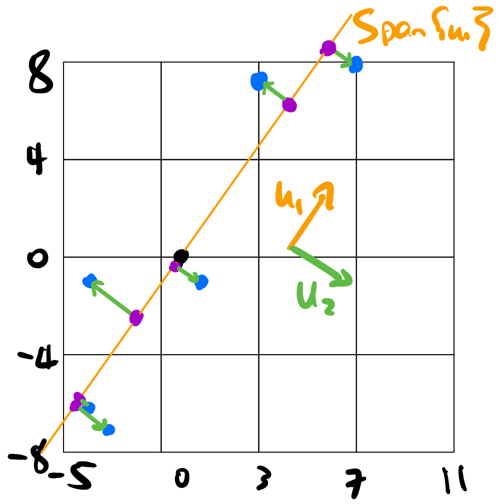

• $= d_i = \begin{pmatrix} 3 \\ 7 \end{pmatrix}, \begin{pmatrix} -4 \\ -6 \end{pmatrix}, \ldots$

• $=$ columns of $\sigma_1 u_1 v_1^T$

$=$ projections of • onto / $= \mathrm{Span}\{u_1\}$

↖ $=$ columns of $\sigma_2 u_2 v_2^T$

$=$ projections of • onto \ $= \mathrm{Span}\{u_2\}$

NB: • = • + ↖

So SVD "pulls apart" the columns of A in $u_1, \ldots, u_r$-components

Review: PCA so far

$A_0 = \begin{pmatrix} | & & | \\ \vec{d}_1 & \cdots & \vec{d}_n \\ | & & | \end{pmatrix}$:

$m \times n$ data matrix with n samples (data points) of m measurements each in the columns

$$A = \begin{pmatrix} | & & | \\ \bar{d}_1 & \cdots & \bar{d}_n \\ | & & | \end{pmatrix} = A - \begin{pmatrix} \mu_1 & \cdots & \mu_1 \\ \vdots & \ddots & \vdots \\ \mu_m & \cdots & \mu_m \end{pmatrix} \quad \mu_i = \text{mean of row } i$$

recentered data matrix obtained from A by subtracting the means of the measurements (rows)

$$S = \frac{1}{n-1} AA^T = \frac{1}{n-1} \begin{pmatrix} (\text{row }1)\cdot(\text{row }1) & (\text{row }1)\cdot(\text{row }2) & \cdots \\ (\text{row }2)\cdot(\text{row }1) & (\text{row }2)\cdot(\text{row }2) & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix}$$

$m \times m$ covariance matrix containing the variances of the measurements on the diagonal:

$$\frac{1}{n-1} (\text{row } i)\cdot(\text{row } i) = \frac{1}{n-1}\left(\bar{x}_{1i}^2 + \cdots + \bar{x}_{ni}^2\right) = s_i^2$$

$\longrightarrow$ total variance is $s^2 = s_1^2 + \cdots + s_m^2 = \mathrm{Tr}(S)$

NB: total variance is just

$$s^2 = s_1^2 + \cdots + s_m^2 = \frac{1}{n-1}\left(\bar{x}_{11}^2 + \cdots + \bar{x}_{n1}^2\right) + \cdots + \frac{1}{n-1}\left(\bar{x}_{1m}^2 + \cdots + \bar{x}_{nm}^2\right)$$

$$= \frac{1}{n-1}(\text{sum of squares of all entries of } A)$$

$$= \frac{1}{n-1}\left(\|\bar{d}_1\|^2 + \cdots + \|\bar{d}_n\|^2\right)$$

For $u \in \mathbb{R}^m$, $\|u\| = 1$, the variance in the u-direction is

$$s(u)^2 = u^T S u = \frac{1}{n-1}\left[(\bar{d}_1 \cdot u)^2 + \cdots + (\bar{d}_n \cdot u)^2\right]$$

If σ_1^2 is the largest eigenvalue of S then this is maximized at a unit q^2-eigenvector u_1 w/maximum value σ_1^2.

u_1 is the first principal component of A

Eg: $A_0 = \begin{pmatrix} x_1 & \cdots & x_n \\ y_1 & \cdots & y_n \end{pmatrix} = \begin{pmatrix} 8 & 1 & 12 & 6 & 1 & 2 \\ 15 & 2 & 16 & 7 & 7 & 1 \end{pmatrix}$ $\quad \mu_1 = 5, \ \mu_2 = 8$

$A = \begin{pmatrix} \bar{x}_1 & \cdots & \bar{x}_n \\ \bar{y}_1 & \cdots & \bar{y}_n \end{pmatrix} = \begin{pmatrix} 3 & -4 & 7 & 1 & -4 & -3 \\ 7 & -6 & 8 & -1 & -1 & -7 \end{pmatrix}$

$S = \frac{1}{5} AA^T = \begin{pmatrix} 20 & 25 \\ 25 & 40 \end{pmatrix}$ $\quad s_1^2 = 20, \ s_2^2 = 40$ $\quad s^2 = 20 + 40 = 60$

$\sigma_1^2 \approx 56.9$

$u_1 \approx \begin{pmatrix} 0.561 \\ 0.828 \end{pmatrix}$

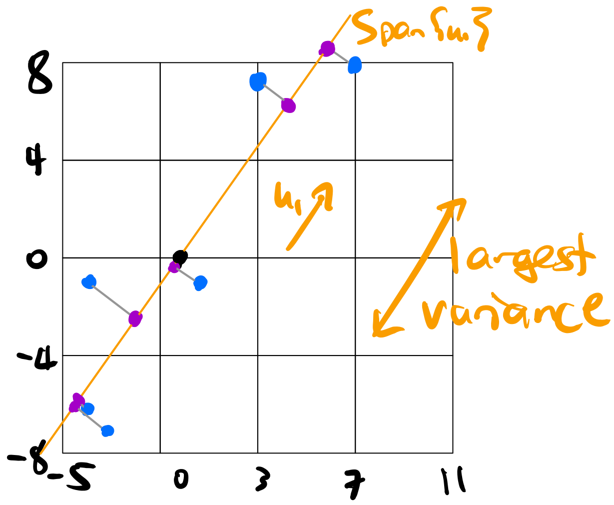

• $= \bar{d}_i$ $\quad$ • = projection of • onto Span$\{u_1\}$

So the first principal component is u_1, and the variance in that direction is $\approx 56.9 > 20, 40$.

Our data points are "stretched out" most in the u_1-direction.

NB: Here's how I should (but won't) grade the final exam:

- Put the scores of each problem in an $m \times n$ matrix A_0
  ($m = \#$problems, $n = \#$students)
- Subtract row averages $(\mu_1, \ldots, \mu_m)$ to recenter
  $\rightsquigarrow$ matrix $A = \begin{pmatrix} | & & | \\ d_1 & \cdots & d_n \\ | & & | \end{pmatrix}$
- Compute the 1st principal component u_1
- $D = \begin{pmatrix} D_1 \\ \vdots \\ D_m \end{pmatrix}$ $\quad D_j$ = max score on problem j
- The score for student i is

$$\frac{d_i \cdot u_1}{D \cdot u_1} \quad \text{(percent)}$$

This maximizes the standard deviation by reweighting the problems.

Key Point: Eigenvalues & eigenvectors of

$$S = \frac{1}{n-1} AA^T = \left(\frac{1}{\sqrt{n-1}} A\right)\left(\frac{1}{\sqrt{n-1}} A\right)^T$$

compute the SVD of $\frac{1}{\sqrt{n-1}} A$ and $\frac{1}{\sqrt{n-1}} A^T$!

$$\frac{1}{\sqrt{n-1}} A = \sigma_1 u_1 v_1^T + \cdots + \sigma_r u_r v_r^T \quad \& \quad \frac{1}{\sqrt{n-1}} A^T = \sigma_1 v_1 u_1^T + \cdots + \sigma_r v_r u_r^T$$

NB: the SVD of A is

$$A = \sqrt{n-1}\,\sigma_1 u_1 v_1^T + \cdots + \sqrt{n-1}\,\sigma_r u_r v_r^T$$

- $\sigma_1^2 \geq \cdots \geq \sigma_r^2 > 0$ = nonzero eigenvalues of S

  NB the singular values of A are $\sqrt{n-1}\,\sigma_1, \ldots, \sqrt{n-1}\,\sigma_r$

- The trace of a square matrix is the sum of its eigenvals (HW11)

$$\Longrightarrow \boxed{\text{total variance} = s^2 = \mathrm{Tr}(S) = \sigma_1^2 + \cdots + \sigma_r^2}$$

- $u_1, \ldots, u_r$ = orthonormal eigenvectors of S

  = left-singular vectors of $\frac{1}{\sqrt{n-1}} A$ (& of A)

- $v_i = \frac{1}{\sigma_i} \cdot \frac{1}{\sqrt{n-1}} A^T u_i$

  = right-singular vectors of $\frac{1}{\sqrt{n-1}} A$ (& of A)

Def: The i^{th} **principal component** of A is the direction of u_i.

We know the 1^{st} principal component is the direction of largest variance. What about the higher principal components?

QO with Extra Constraints:

- $s(u)^2 = u^T S u$ is maximized
  subject to $\|u\|=1$
  at u_1 with $s(u_1)^2 = \sigma_1^2$
  $\rightarrow$ u_1 is the direction with largest variance
- $s(u)^2$ is maximized
  subject to $\|u\|=1$ and $u \perp u_1$
  at u_2 with $s(u_2)^2 = \sigma_2^2$
  $\rightarrow$ u_2 is the direction with 2^{nd}-largest variance
- $s(u)^2$ is maximized
  subject to $\|u\|=1$ and $u \perp u_1, \ldots, u \perp u_{i-1}$
  at u_i with $s(u_i)^2 = \sigma_i^2$
  $\rightarrow$ u_i is the direction with i^{th}-largest variance

NB u_r is the direction with smallest variance if $r=n$.

The columns of $\sqrt{n-1}\,\sigma_i u_i v_i^T$ are the orthogonal projections of the columns of A onto $\mathrm{Span}\{u_i\}$.

$\rightarrow$ $A = \sqrt{n-1}\,\sigma_1 u_1 v_1^T + \cdots + \sqrt{n-1}\,\sigma_r u_r v_r^T$

"breaks apart" your data points into components:

$$\boxed{\sqrt{n-1}\,\sigma_i u_i v_i^T = \text{piece in the direction of the } i^{th} \text{ largest variance}}$$

Eg: In our example, $\frac{1}{\sqrt{6-1}}A = \sigma_1 u_1 v_1^T + \sigma_2 u_2 v_2^T$

$\sigma_1^2 \approx 56.9$ $\quad\quad$ $\sigma_2^2 \approx 3.07$

$u_1 \approx \begin{pmatrix} 0.561 \\ 0.828 \end{pmatrix}$ $\quad\quad$ $u_2 \approx \begin{pmatrix} 0.828 \\ -0.561 \end{pmatrix}$

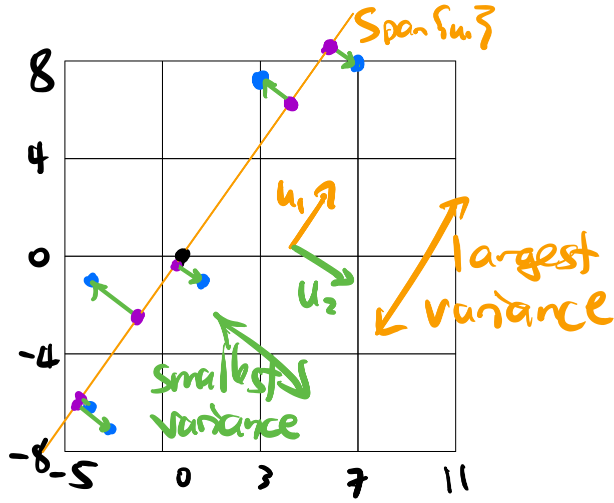

$S = \begin{pmatrix} 20 & 25 \\ 25 & 40 \end{pmatrix}$

$s_1^2 = 20 \quad s_2^2 = 40$

Total variance

$s^2 = 60 = \sigma_1^2 + \sigma_2^2$

• $= \bar{d}_i$

• = columns of $\sqrt{5}\,\sigma_1 u_1 v_1^T$

= projections of • onto $=\text{Span}\{u_1\}$

= columns of $\sqrt{5}\,\sigma_2 u_2 v_2^T$

= projections of • onto $=\text{Span}\{u_2\}$

NB: In this case, $s(u)^2$ is minimized at u_2 with minimum value σ_2^2 = smallest eigenvalue of S.

$$s(u_2)^2 = \frac{1}{n-1}\left[(d_1 \cdot u_2)^2 + \cdots + (d_n \cdot u_2)^2\right]$$

$$= \frac{1}{n-1}\left[\text{sum of squares of lengths of } \nwarrow \right]$$

Conclusion: The first principal component is the line of best fit in the sense of orthogonal least squares, and the

$$(\text{error})^2 = (n-1)s(u_2)^2 = (n-1)\sigma_2^2$$

Subspace(s) of Best Fit

What happens in general?

Def: Let V be a subspace of $\mathbb{R}^m$. The **variance along V** of our (recentered) data points $\bar{d}_1, \ldots, \bar{d}_n$ is

$$s(V)^2 = \frac{1}{n-1}\Big(\underbrace{\|(\bar{d}_1)_V\|^2 + \cdots + \|(\bar{d}_n)_V\|^2}_{\text{orthogonal projections}}\Big).$$

NB: If $V = \mathrm{Span}\{u\}$ for u a unit vector then $(\bar{d}_i)_V = (\bar{d}_i \cdot u)u$, so $\|(\bar{d}_i)_V\|^2 = (\bar{d}_i \cdot u)^2 \|u\|^2 = (\bar{d}_i \cdot u)^2$,

so
$$\boxed{s(V)^2 = \frac{1}{n-1}\left[(\bar{d}_1 \cdot u)^2 + \cdots + (\bar{d}_n \cdot u)^2\right] = s(u)^2}$$

Recall: if $u \perp v$ then $\|u+v\|^2 = \|u\|^2 + \|v\|^2$.

Taking $u = (\bar{d}_i)_V$ & $v = (\bar{d}_i)_{V^\perp}$ gives $\underbrace{\bar{d}_i = (\bar{d}_i)_V + (\bar{d}_i)_{V^\perp}}_{\text{orthogonal decomposition}}$

$$\Longrightarrow \|\bar{d}_i\|^2 = \|(\bar{d}_i)_V\|^2 + \|(\bar{d}_i)_{V^\perp}\|^2$$

Sum over all i:

$$\boxed{\begin{aligned} &\text{For any subspace } V, \\ &s(V)^2 + s(V^\perp)^2 = \frac{1}{n-1}\left[\|\bar{d}_1\|^2 + \cdots + \|\bar{d}_n\|^2\right] \\ &\qquad \underset{(p.3)}{=} (\text{total variance}) = \sigma_1^2 + \cdots + \sigma_r^2 \end{aligned}}$$

NB: $s(V^\perp)^2 = \frac{1}{n-1}\left(\|(\bar{d}_1)_{V^\perp}\|^2 + \cdots + \|(d_n)_{V^\perp}\|^2\right)$ is $\frac{1}{n-1} \times$ the sum of the squares of the (orthogonal) distances of the $\bar{d}_i$ to V.

Def: The d-space of best fit in the sense of orthogonal least squares is the d-dimensional subspace V minimizing $s(V^\perp)^2$. The error^2 is $s(V^\perp)^2$.

NB: Minimizing $s(V^\perp)^2$ means maximizing $s(V)^2$ since $s(V)^2 + s(V^\perp)^2 =$ total variance.

Thm: Let A be a centered data matrix with SVD

$$\frac{1}{\sqrt{n-1}}A = \sigma_1 u_1 v_1^T + \cdots + \sigma_r u_r v_r^T.$$

The d-space of best fit to its columns is

$$V = \mathrm{Span}\{u_1, \ldots, u_d\}.$$

The variance along V is $s(V) = \sigma_1^2 + \cdots + \sigma_d^2$ and the error^2 is $s(V^\perp) = \sigma_{d+1}^2 + \cdots + \sigma_r^2$.

So you "split" the total variance $\sigma_1^2 + \cdots + \sigma_r^2 = s^2 = s(V)^2 + s(V^\perp)^2$ into the large part $s(V) = \sigma_1^2 + \cdots + \sigma_d^2$ and the small part $s(V^\perp) = \sigma_{d+1}^2 + \cdots + \sigma_r^2$.

Eg: The line of best fit is the first principal component $V = \mathrm{Span}\{u_1\}$. The error$^2 = \sigma_2^2 + \cdots + \sigma_r^2$.

Eg: The plane of best fit is the span of the first 2 principal components: $V = \mathrm{Span}\{u_1, u_2\}$ $\text{error}^2 = \sigma_3^2 + \cdots + \sigma_r^2$

Eg: Suppose

$$A = 10u_1v_1^T + 8u_2v_2^T + .2u_3v_3^T + .1u_4v_4^T$$

Then A fits the plane $V = \mathrm{Span}\{u_1, u_2\}$ to a small $\text{error}^2 = .2^2 + .1^2$.

But A does not fit the line $L = \mathrm{Span}\{u_1\}$ well: the $\text{error}^2 = 8^2 + .2^2 + .1^2$.

Upshot: If $\sigma_1, \ldots, \sigma_d$ are much larger than $\sigma_{d+1}, \ldots, \sigma_r$ then your data closely fit the d-space

$$V = \mathrm{Span}\{u_1, \ldots, u_d\}$$

(but not a smaller subspace like $\mathrm{Span}\{u_1, \ldots, u_{d-1}\}$).

NB: This is all applied to the recentered data points. Your original data points $d_1, \ldots, d_n$ = columns of A fit the translated subspace

$$V + \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_m \end{pmatrix} \quad \text{(add back the means).}$$

See the Netflix problem on HW15.