

CSE439 Fall 2025 Week 13: Matrix Algebra and the SVD

What happens if we try to take the "inner product" of two $m \times n$ matrices A and B by first "unrolling" them as vectors? Remembering to conjugate the entries of A , we get

$$\langle A, B \rangle = \sum_{\substack{i=1, m \\ j=1, n}} \overline{A[i, j]} B[i, j].$$

Now let $C = A^*B$. Since A^* is $n \times m$, this is an $n \times n$ square matrix. From

$$C[r, s] = \sum_{k=1}^m A^*[r, k]B[k, s] = \sum_{k=1}^m \overline{A[k, r]} B[k, s]$$

we get that the diagonal entries of C are $C[r, r] = \sum_{k=1}^m \overline{A[k, r]} B[k, r]$. Hence the diagonal sum gives

$$\sum_{r=1}^n C[r, r] = \sum_{\substack{k=1, m \\ r=1, n}} \overline{A[k, r]} B[k, r] = \langle A, B \rangle$$

as we defined it above. The diagonal sum at left is called the **trace**, with notation $Tr(C)$. Now for a vector v , the self inner-product $\langle v, v \rangle$ gives the squared Euclidean norm of v , written $\|v\|_2^2$, so

$\|v\|_2 = \sqrt{\langle v, v \rangle}$. The analogous concept for matrices is the **Frobenius norm**, named for Ferdinand Georg Frobenius:

$$\|A\|_F = \sqrt{Tr(A^*A)}.$$

Or you can simply say it's the Euclidean 2-norm of the vector obtained by "unrolling" the matrix. This norm, however, overstates the *action* of the matrix in Euclidean space, which involves its $m \times n$ dimensions. This is

$$\|A\|_2 = \sup\{\|Av\|_2 : v \text{ is a unit vector of length } n\}.$$

For some further remarks: Since our vectors are finite-dimensional, the "ball surface" of unit vectors is **compact**, which actually *means* that there is a definite vector v that maximizes $\|Av\|_2$ rather than just having a limit---so we can write "max" in place of "sup" for "supremum." The task of *finding* such a vector v is the main algorithmic need of computing the **singular value decomposition (SVD)** as treated below. It tumbles out of the SVD Theorem that $\|A\|_2 \leq \|A\|_F$ for every matrix A . But the intuition is that $\|A\|_2$ tells the most that A can "stretch" a vector along the fixed dimensions it operates on, whereas $\|A\|_F$ is the maximum amount of "stretch" that the entries of A could give under any configuration of dimensions.

The SVD

A matrix S is (**pseudo**-)diagonal if it is (**non**-)square and $S[i, j] = 0$ whenever $i \neq j$. It follows that both S^*S and SS^* are diagonal square matrices. Some of the diagonal entries may be 0.

SVD Theorem: For every $m \times n$ matrix A we can efficiently find:

- an $m \times m$ unitary matrix U ,
- an $m \times n$ pseudo-diagonal matrix Σ with non-negative entries $\Sigma[i, i] = \sigma_i$, and
- an $n \times n$ unitary matrix V ,

such that $A = U\Sigma V^*$. Furthermore, we can arrange that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{\min(m,n)}$, and in consequence:

- $\|A\|_F = \sqrt{\sum_i \sigma_i^2}$.
- $\|A\|_2 = \sigma_1$,
- $A^*A = V\Sigma^T U^* U\Sigma V^* = V \text{diag}(\sigma_i^2) V^*$, and
- $AA^* = U\Sigma V^* V\Sigma^T U^* = U \text{diag}(\sigma_i^2) U^*$,

so that the squares of the σ_i and associated vectors give the spectral decompositions of the Hermitian PSD matrices A^*A and AA^* , respectively.

The σ_i are the **singular values**. The number r of positive ones equals the **rank** of A . Whereas some of the λ_i can be negative in the Spectral Theorem---when the Hermitian matrix is not PSD---none of the σ_i is negative. The first r columns of U form an orthonormal basis for the subspace \mathbb{W} spanned by the columns of A (called the column space of A), while the first r columns of V form an orthonormal basis for the column space of A^* . The latter is identical to the row space of A when A is a real matrix---and in that case, U and V come out being real as well. The remaining $m - r$ columns of U form an orthonormal basis for the space \mathbb{W}^\perp , which is also the **nullspace** of A^* . As with the Spectral Theorem, the basis vectors are not unique when there is multiplicity or when we don't have $r = m = n$, but the values σ_i are unique (when sorted in nonascending order, so we can say the matrix Σ is unique too). Once U and V are specified, we get $\Sigma = U^*AV$ too.

Proof: The procedure works by recursion through subspaces and so resembles the proof of the Spectral Theorem. The first and top-level step is most emblematic. It begins by *finding* a unit vector v_1 that maximizes $\|Av_1\|_2$. Then $\sigma_1 = \|Av_1\|_2$ is the first and biggest singular value. It can't be zero (unless A is the all-zero matrix, in which case we've "hit triviality"), so

$$u_1 = \frac{Av_1}{\sigma_1}$$

is a unit vector. If there are more than one maximizing unit vectors v then we will get multiplicity, but let us first suppose that the v_1 and associated u_1 are unique. Before doing the recursion, we may

postulate that u_1 is arbitrarily extended to an orthonormal basis U_1 of \mathbb{C}^m (or of \mathbb{R}^m in the real case) and v_1 to an orthonormal basis V_1 of \mathbb{C}^n . In the resulting coordinates, we get

$$U_1^* A V_1 = \begin{bmatrix} \sigma_1 & \textcolor{red}{w}_1^* \\ \mathbf{0} & B \end{bmatrix} = S_1$$

for some vector w_1 of length $n - 1$ and $(m - 1) \times (n - 1)$ matrix B . The red $\mathbf{0}$ stands for $m - 1$ zeroes and is because $Av_1 = \sigma_1 u_1$ so there is no dependence on the other $m - 1$ coordinates. *The goal is to prove that w_1 must be all-zero too.* Then recursing on B hammers out the (pseudo-)diagonal matrix Σ .

Let $w = \begin{bmatrix} \sigma_1 \\ w_1 \end{bmatrix}$ as a column vector. Then $S_1 w = w' = \begin{bmatrix} \sigma_1^2 + w_1^* w_1 \\ Bw_1 \end{bmatrix}$. Ignoring the Bw part, we get $\|w'\| \geq \sigma_1^2 + w_1^* w_1$. The right-hand side equals $\|w\|_2^2$. Dividing by $\|w\|_2$ hence gives us

$$\frac{\|S_1 w\|_2}{\|w\|_2} \geq \|w\|_2 = \sqrt{\sigma_1^2 + w_1^* w_1}.$$

Now if w_1 is nonzero, then $w_1^* w_1$ is a positive real number, so $\frac{\|S_1 w\|_2}{\|w\|_2} > \sigma_1$. Under the definition of the 2-norm for matrices, this means $\|S_1\|_2 > \sigma_1$. But

$$\|S_1\|_2 = \|U_1^* A V_1\|_2 = \|A\|_2$$

because U_1 and V_1 are unitary. And $\|A\|_2 = \sigma_1$ by how we defined σ_1 . This is a contradiction saying " $\sigma_1 > \sigma_1$." The only way out is for w_1 to be a zero vector.

The recursion then takes place on the perpendicular subspace of v_1 , or in general, the perpendicular subspace of the span of the orthogonal unit vectors v_j chosen thus far. The final point is that the corresponding vectors u_j also come out orthogonal. This is because, when $i \neq j$ (and at stages where σ_i and σ_j are both nonzero---else we are in the base case of completing orthonormal bases on the nullspaces):

$$\sigma_i \sigma_j \langle u_i | u_j \rangle = \langle \sigma_i u_i | \sigma_j u_j \rangle = \langle Av_i | Av_j \rangle = v_i^* A^* A v_j = v_i^* \sigma_j^2 v_j = \sigma_j^2 \langle v_i | v_j \rangle = 0,$$

finally using the orthogonality of the v_i vectors. The fourth equality happens because v_j is an eigenvector of $A^* A$ with eigenvalue σ_j^2 . The reason given by the (short!) proof in the MIT notes (https://math.mit.edu/classes/18.095/2016IAP/lec2/SVD_Notes.pdf) is that

$$A^* A = (U \Sigma V^*)^* (U \Sigma V^*) = V \Sigma^T U^* U \Sigma V^* = V \Sigma^T \Sigma V^*,$$

which in turn converts to the way we have been writing the spectral decomposition since V is unitary. However, substituting $U^*U = I$ strikes me as assuming what one is trying to prove about the u_i vectors.

To tie up the loose end, we *choose to restart the proof*. We apply the original Spectral Theorem to the Hermitian PSD matrix A^*A to get nonnegative eigenvalues $\lambda_1, \dots, \lambda_n$ ---listed in nonincreasing order---and orthonormal eigenvectors v_1, \dots, v_n such that

$$A^*A = \lambda_1|v_1\rangle\langle v_1| + \dots + \lambda_n|v_n\rangle\langle v_n| = V^*\text{diag}(\lambda_i)V,$$

taking V as the matrix with the eigenvectors as its columns. Now define σ_i to be the nonnegative square root of λ_i for each i . Since the rank r of A equals the rank of A^*A , we get $\sigma_i > 0$ for $i = 1$ to r . For these i , define

$$u_i = \frac{Av_i}{\sigma_i}.$$

Now the above demonstration that $\langle u_i | u_j \rangle = 0$ is logically valid, because we arranged that $\sigma_i^2 = \lambda_i$ is an eigenvalue of A^*A with eigenvector v_i in advance. What we've lost, however, is the original proof's definition of σ_i so that u_i is a unit vector. We recover it, however, this way:

$$\langle u_i | u_i \rangle = \left\langle \frac{Av_i}{\sigma_i} \middle| \frac{Av_i}{\sigma_i} \right\rangle = \frac{1}{\sigma_i^2} v_i^* A^* A v_i = \frac{1}{\sigma_i^2} v_i^* \lambda_i v_i = v_i^* v_i = 1.$$

And u_i is an eigenvector of AA^* because

$$AA^*u_i = AA^*\frac{Av_i}{\sigma_i} = \frac{1}{\sigma_i} A(A^*A)v_i = \frac{1}{\sigma_i} A\sigma_i^2 v_i = \sigma_i Av_i = \sigma_i^2 u_i.$$

For $i > r$, we can arbitrarily complete the basis by choosing orthonormal vectors that span the nullspace.

So now the only thing we've "lost" compared to the first proof strategy is the fact that at the first and each later step of the recursion, the choice of unit vector v_i maximizes $\|Av_i\|_2$. However, now we can appeal to the uniqueness of the λ_i and "quasi-uniqueness" of the eigenvectors up to the flex of multiplicity. The squares of the σ_i and the λ_i **must** coincide. What comes out is a deep fact that the largest eigenvalues of A^*A naturally pick out the directions in which A stretches the most. \square

Corollary: For a square matrix A already of the form E^*E (and that goes for any Hermitian PSD matrix), the SVD and spectrum of E coincide with $U = V$.

Proof. The diagonal form $E = U\Lambda U^*$ has the specified properties; because E is PSD, the λ_i are nonnegative, and we can arrange U so that the diagonal is in nonincreasing order. \square

In all other cases where A is diagonalizable, there are reasons for saying the SVD gives *more* information than the diagonalization. This is especially so with upper or lower triangular matrices---see example below. And of course, there are many square matrices that can't be diagonalized...to say nothing of non-square matrices...for which the SVD is the only game in town.

Our two-pronged proof suggests two different algorithms for *computing* the SVD of a matrix A :

- Diagonalize A^*A to get λ_i 's and V , then $\sigma_i = \sqrt{\lambda_i}$ and $u_i = \frac{Av_i}{\sigma_i}$.
- Find a unit vector v maximizing $\|Av\|_2$ and recurse.

Other methods come into play when A has certain particular features. Niloufer Mackey developed new methods in her 1993 UB CS PhD dissertation under Patricia Eberlein. Other remarks:

- The version giving $A = U\Sigma V^*$ with U and V both unitary, is called the **full SVD**.
- When the $m \times n$ matrix A has rank $r < \min(m, n)$, then we can also do $A = U\Sigma V^*$ with Σ being an $r \times r$ matrix with positive values on the main diagonal, U being $m \times r$, and V being $n \times r$. This is called the **reduced** or **compact SVD**.
- Some sources give a third version where U is $m \times r$ but Σ is $r \times n$ and V is $n \times n$ (and unitary). Let's call this the **semi-reduced** version.

Our proof and notes use the style of diagonalizing A^*A , getting V from the unit eigenvectors v_i of that, and then getting $u_i = Av_i$, dividing by σ_i to normalize u_i . There is also a symmetrical style of diagonalizing AA^* instead, forming its orthogonal unit eigenvectors as the columns of U , and getting V at the end. The nicely verbose applet

<https://www.emathhelp.net/calculators/linear-algebra/svd-calculator/>

does that. The most portable applets handle real numbers only, so they write A^T instead of A^* (or A^\dagger). There are some Java applets that allow complex numbers (but I haven't tried them). They all have limitations on m , n , and/or the magnitudes of matrix entries. The applet

<https://www.omnicalculator.com/math/svd#is-singular-value-decomposition-unique>

seems to do things the $A^T A$ way, with V first, but only does up to 3×3 and doesn't show intermediate steps. There are also differences in output caused by not sorting the singular values in nonascending order (so with the largest one at upper left) and the non-uniqueness of V and U .

Examples and Applications

In any upper or lower triangular matrix A , the elements of the diagonal *are* the eigenvalues. They are thus independent of all the off-diagonal entries at upper right. Those entries have information that does get picked up by the SVD. The two examples in the MIT notes are good for this.

Example 1:

$$A = \begin{bmatrix} 3 & 0 \\ 4 & 5 \end{bmatrix} \quad A^* = A^T = \begin{bmatrix} 3 & 4 \\ 0 & 5 \end{bmatrix}$$

$$A^*A = \begin{bmatrix} 9+16 & 20 \\ 20 & 25 \end{bmatrix} = \begin{bmatrix} 25 & 20 \\ 20 & 25 \end{bmatrix} \quad AA^* = \begin{bmatrix} 9 & 12 \\ 12 & 16+25 \end{bmatrix} = \begin{bmatrix} 9 & 12 \\ 12 & 41 \end{bmatrix}.$$

Abstracting this, consider $A = \begin{bmatrix} a & 0 \\ b & c \end{bmatrix}$. The eigenvalues are a and c with $[1, 0]^T$ as one of the eigenvectors. This has no dependence on the entry b . How much A can stretch a (unit) vector does depend on b . The SVD employs this information. We have

$$A^*A = \begin{bmatrix} a^* & b^* \\ 0 & c^* \end{bmatrix} \cdot \begin{bmatrix} a & 0 \\ b & c \end{bmatrix} = \begin{bmatrix} a^*a + b^*b & b^*c \\ cb^* & c^*c \end{bmatrix} = \begin{bmatrix} |a|^2 + |b|^2 & c\bar{b} \\ \bar{c}b & |c|^2 \end{bmatrix}.$$

In the real case we can drop all the stars and bars. Then, solving $\det(A^*A - xI) = 0$ gives

$$0 = (a^2 + b^2 - x)(c^2 - x) - b^2c^2 = x^2 - (a^2 + b^2 + c^2)x + a^2c^2.$$

The two solutions given by

$$x = \frac{1}{2} \left(a^2 + b^2 + c^2 \pm \sqrt{(a^2 + b^2 + c^2)^2 - 4a^2c^2} \right)$$

do not simplify further in general. In the example $a = 3$, $b = 4$, and $c = 5$, the expression under the square root becomes $50^2 - 30^2 = 40^2$, so $x = \frac{1}{2}(50 \pm 40) = 45$ or just 5. Notice also that

$$Tr(A^*A) = |a|^2 + |b|^2 + |c|^2 = 9 + 16 + 25 = 50 = \lambda_1 + \lambda_2.$$

The singular values are the square roots, so $\sqrt{45} = 3\sqrt{5}$ and $\sqrt{5}$. The V matrix is formed from the eigenvectors of A^*A , so we solve:

$$\begin{bmatrix} 25 & 20 \\ 20 & 25 \end{bmatrix} \cdot \begin{bmatrix} y \\ z \end{bmatrix} = \begin{bmatrix} 25y + 20z \\ 20y + 25z \end{bmatrix} = \begin{bmatrix} 45y \\ 45z \end{bmatrix}, \quad \begin{bmatrix} 25 & 20 \\ 20 & 25 \end{bmatrix} \cdot \begin{bmatrix} y' \\ z' \end{bmatrix} = \begin{bmatrix} 25y' + 20z' \\ 20y' + 25z' \end{bmatrix} = \begin{bmatrix} 5y' \\ 5z' \end{bmatrix}.$$

This gives $\begin{bmatrix} y \\ z \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ for the vector v_1 and $\begin{bmatrix} y' \\ z' \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix}$ as one of a couple orthogonal choices for the vector v_2 . Then V becomes the Hadamard matrix. The U matrix is obtained by normalizing the columns of AV . We can normalize $\begin{bmatrix} 3 & 0 \\ 4 & 5 \end{bmatrix} \cdot \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} = \begin{bmatrix} 3 & 3 \\ 9 & -1 \end{bmatrix}$ columnwise as $\begin{bmatrix} 1/\sqrt{10} & 3/\sqrt{10} \\ 3/\sqrt{10} & -1/\sqrt{10} \end{bmatrix}$, so $U = \frac{1}{\sqrt{10}} \begin{bmatrix} 1 & 3 \\ 3 & -1 \end{bmatrix}$.

As a final check, $U\Sigma V^* =$

$$\frac{1}{\sqrt{20}} \begin{bmatrix} 1 & 3 \\ 3 & -1 \end{bmatrix} \begin{bmatrix} 3\sqrt{5} & 0 \\ 0 & \sqrt{5} \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} = \frac{1}{\sqrt{20}} \begin{bmatrix} 3\sqrt{5} & 3\sqrt{5} \\ 9\sqrt{5} & -\sqrt{5} \end{bmatrix} H = \frac{1}{2} \begin{bmatrix} 3 & 3 \\ 9 & -1 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 6 & 0 \\ 8 & 10 \end{bmatrix},$$

which equals A . We also get $\|A\|_2 = \sqrt{45}$ and $\|A\|_F = \sqrt{45+5} = \sqrt{50} = 5\sqrt{2}$.

To see that V is not unique, we could have chosen $\begin{bmatrix} y' \\ z' \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} -1 \\ 1 \end{bmatrix}$ as the second eigenvector instead. Then we'd get $V = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$, which Assignment 4 (in Fall 2024) called the "Damhard matrix H_4 ." The U matrix changes too: it comes by normalizing each column of $\begin{bmatrix} 3 & 0 \\ 4 & 5 \end{bmatrix} \cdot \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} 3 & -3 \\ 9 & 1 \end{bmatrix}$ to get $U = \frac{1}{\sqrt{10}} \begin{bmatrix} 1 & -3 \\ 3 & 1 \end{bmatrix}$. Note that this V is not Hermitian, so we have to remember to transpose it when we do the check that $U\Sigma V^* =$

$$\frac{1}{\sqrt{20}} \begin{bmatrix} 1 & -3 \\ 3 & 1 \end{bmatrix} \begin{bmatrix} 3\sqrt{5} & 0 \\ 0 & \sqrt{5} \end{bmatrix} \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 3 & -3 \\ 9 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 6 & 0 \\ 8 & 10 \end{bmatrix} = A$$

as before. (Nor does V square to the identity; $V^2 = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$, so this V is another square root of the matrix $B = -i\mathbf{Y}$.)

Low-Rank Approximation By SVD Truncation

Last, let's see what happens if we simply wipe out the smaller entry of Σ , which is $\sigma_2 = \sqrt{5}$:

$$\frac{1}{\sqrt{20}} \begin{bmatrix} 1 & 3 \\ 3 & -1 \end{bmatrix} \begin{bmatrix} 3\sqrt{5} & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} = \frac{1}{\sqrt{20}} \begin{bmatrix} 3\sqrt{5} & 0 \\ 9\sqrt{5} & 0 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 3 & 3 \\ 9 & 9 \end{bmatrix} = \begin{bmatrix} 1.5 & 1.5 \\ 4.5 & 4.5 \end{bmatrix}.$$

Is the resulting A' a reasonable approximation to A ? Note that A' stretches the first V vector v_1 by the same amount: $A' \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 3 \\ 9 \end{bmatrix}$, whose 2-norm is $\frac{1}{\sqrt{2}} \sqrt{3^2 + 9^2} = \sqrt{45} = \sigma_1$. But the second dimension v_2 gets zeroed out.

We can also preserve the trace by using $\Sigma' = \begin{bmatrix} 4\sqrt{5} & 0 \\ 0 & 0 \end{bmatrix}$ instead, which gives $A' = \begin{bmatrix} 2 & 2 \\ 6 & 6 \end{bmatrix}$. Then $A'v_1$ over-stretches, but in other contexts it may give better results. Or we might prefer to preserve the Frobenius norm by using $\Sigma'' = \begin{bmatrix} 5\sqrt{2} & 0 \\ 0 & 0 \end{bmatrix}$ instead, conserving $\sigma_1^2 + \sigma_2^2$. Well, the whole approximation idea looks better when the matrices are much larger to begin with.

[The Tue. 12/2 lecture ended here. I began Thursday with a challenge question:
What are the singular values of an $n \times n$ unitary matrix U ? Step 1 of our algorithm is:

1. Diagonalize U^*U and get the eigenvalues λ_i .
2. The σ_i are the positive square roots of the λ_i .

The answer is that since $U^*U = I$, they are all 1. No 0s. This also stands to reason from the rule that $\|Ux\|_2 = \|x\|_2$ for all vectors x , which means U does unit stretching in all directions.]

(Pseudo-)Inversion

In the invertible $n \times n$ Hermitian case where we get orthonormal diagonalization $A = U\Lambda U^*$ with all diagonal entries λ_i being nonzero, then using $\Lambda' = \text{diag}\left(\frac{1}{\lambda_i}\right)$ makes $U\Lambda'U^* = A^{-1}$. We can partly emulate this for any matrix by taking the reciprocals of the positive singular values.

Definition: The (Moore-(Bjerhammar)-Penrose) **pseudoinverse** of an arbitrary $m \times n$ matrix A with SVD $A = U\Sigma V^*$ is the $n \times m$ matrix given by $A^+ = V\Sigma^+U^*$, where Σ^+ transposes Σ and then replaces every nonzero σ_i by $1/\sigma_i$.

If we specified that $A = U\Sigma V^*$ is the *reduced* SVD, then Σ would be an $r \times r$ diagonal matrix with positive diagonal entries, and we would simply get $A^+ = V\Sigma^{-1}U^*$. Saying it this way, however, would hide a highly important "pseudo" aspect. You might expect that for sake of *continuity*, a zero σ_i would be replaced by some large value, if not by (the IEEE representation of) **inf**. However, what happens more often instead is that when $\sigma_i < \epsilon$ for some threshold ϵ (e.g., $\epsilon = \epsilon_0 \max(m, n, \sigma_1)$ where ϵ_0 is the least positive hardware value), it is treated as zero and blipped---rather than put the large value $E = 1/\epsilon$ into the inverse. The rationale for this is that the dimensions and singular vectors associated to small σ_i can often be "cropped out" with minimal effect---as in the image compression application noted below. But blipping large E betrays the fact of **numerical instability** lurking in applications.

The pseudoinverse obeys the rule $(AB)^+ = B^+A^+$, and if A is invertible, then $A^+ = A^{-1}$. Thus

$$(A^+)^+ = (V\Sigma^+U^*)^+ = (U^*)^+(\Sigma^+)^+V^+ = (U^*)^{-1}\Sigma V^{-1} = U\Sigma V^* = A$$

back again, so this is a viable concept of inversion. However, $AA^+ = U\Sigma V^*V\Sigma^+U^*$ reduces to $U\Sigma\Sigma^+U^*$ but not necessarily to the identity matrix---because zeroes can occur in $\Sigma\Sigma^+$ from having $m < n$ even when all singular values are positive. It also obeys the rules:

- $AA^+A = U\Sigma V^*V\Sigma^+U^*U\Sigma V^* = U\Sigma\Sigma^+\Sigma V^* = U\Sigma V^* = A$;
- $A^+AA^+ = A^+$;
- AA^+ and A^+A are both Hermitian.

Indeed, A^+ is generally the unique matrix obeying these rules. Here are some more examples of SVDs and the resulting (pseudo-)inverses. Back to our 2×2 example:

$$A = \begin{bmatrix} 3 & 0 \\ 4 & 5 \end{bmatrix} = U\Sigma V^* = \frac{1}{\sqrt{10}} \begin{bmatrix} 1 & 3 \\ 3 & -1 \end{bmatrix} \begin{bmatrix} 3\sqrt{5} & 0 \\ 0 & \sqrt{5} \end{bmatrix} \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \text{ so}$$

$$A^+ = V\Sigma^{-1}U^* = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} \frac{1}{3\sqrt{5}} & 0 \\ 0 & \frac{1}{\sqrt{5}} \end{bmatrix} \frac{1}{\sqrt{10}} \begin{bmatrix} 1 & 3 \\ 3 & -1 \end{bmatrix} = \begin{bmatrix} \frac{1}{30} & \frac{1}{10} \\ \frac{1}{30} & \frac{1}{10} \end{bmatrix} \begin{bmatrix} 1 & 3 \\ 3 & -1 \end{bmatrix} = \begin{bmatrix} \frac{1}{3} & 0 \\ -\frac{4}{15} & \frac{1}{5} \end{bmatrix},$$

which is the same as A^{-1} . Of course, A is invertible by virtue of being square and having nonzero determinant, and we could have made life much easier using the **adjugate formula**

$$A^{-1} = \frac{1}{\det(A)} \begin{bmatrix} 5 & -4 \\ 0 & 3 \end{bmatrix}^T = \frac{1}{15} \begin{bmatrix} 5 & 0 \\ -4 & 3 \end{bmatrix}.$$

How about the pseudo-inverse of the matrix $B = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$? $B^T B = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$. We get

$v_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ with eigenvalue 1 and can choose $v_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ (orthonormal to v_1) for the eigenvalue 0. Then $u_1 = Bv_1 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$, while for u_2 we choose an orthonormal vector since $Bv_2 = 0$; $u_2 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ is the

natural choice. So we have $U = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$, $\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} = \Sigma^+$, and $V = I$. This makes

$$B^+ = V\Sigma^+U^* = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}. \text{ Then } B^+B = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \text{ while } BB^+ = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}.$$

A Second Example, With Numerical Instability

Now let's try the second MIT notes example: $A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{bmatrix}$. We get $A^T A = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 9 \end{bmatrix}$. Then

V is the identity matrix again while $U = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix}$ and $\Sigma = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{bmatrix}$ (ignoring the sorting order). So $A^+ = V\Sigma^+U^* = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1/2 & 0 \\ 0 & 0 & 0 & 1/3 \end{bmatrix} \cdot \begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1/2 & 0 & 0 \\ 0 & 0 & 1/3 & 0 \end{bmatrix}$. And

$$A^+A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \text{ while } AA^+ = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Regarding numerical instability, the MIT notes point out that if you make $A[4, 1]$ a small value δ so that A becomes invertible, the eigenvalues grow by more than expected. With $\delta = 1/60000$ the singular values stay 1, 2, 3, and $1/60000$ but the eigenvalues become $\left\{\frac{1}{10}, \frac{i}{10}, \frac{-1}{10}, \frac{-i}{10}\right\}$, as seen at

<https://www.emathhelp.net/calculators/linear-algebra/eigenvalue-and-eigenvector-calculator/>

The reason for using 60000 is that the determinant becomes (negative) $1/10000 = 1/10^4$, and that neatly spreads a factor of $1/10$ among four eigenvalues. The fact that the eigenvalues have equal magnitude is weird, given how the singular values match the sizes of the four positive matrix entries.

Applications to Solving Equations

Approximately Solving Linear Systems: When a matrix A is invertible, the solution to $Ax = \mathbf{b}$ is $\mathbf{x} = A^{-1}\mathbf{b}$. When A is not invertible, or not even square (thus denoting an overspecified or underspecified system), we can still use $\mathbf{z} = A^+\mathbf{b}$ as an "ersatz" solution.

How good a solution? It follows from the SVD theorem that $\|A\mathbf{z} - \mathbf{b}\|_2 \leq \|A\mathbf{x} - \mathbf{b}\|$ for all vectors \mathbf{x} . So this is the best approximation. When the system is underspecified, so that exact solutions exist, \mathbf{z} will be one of them---and moreover, *all* exact solutions have the form

$$\mathbf{z} + (I - A^+A)\mathbf{w}$$

for arbitrary vectors \mathbf{w} . This follows from the identity $A^+ A A^+ = A^+$ given in the "rules" above. **Least squares fitting** is essentially the same process, since we are using the $\|\cdot\|_2$ -norm.

In some cases we can combine A and \mathbf{b} into a matrix E such that $E\mathbf{x}$ measures the error in an attempted solution \mathbf{x} . Then we want to find the \mathbf{z} that *minimizes* $\|E\mathbf{z}\|_2$. This \mathbf{z} is given by the column of V that corresponds to the *least* singular value. (If 0 is a singular value of E , so that $E\mathbf{z} = 0$, this just means that \mathbf{z} is an exact solution.)

Succinct Approximation

This IMHO is the "signature" application of the SVD and will lead us back to quantum computing. Given a pseudodiagonal matrix Σ with $r > k$ positive entries (in sorted order), define Σ_k to be the result of zeroing out all but the k largest entries. If A has SVD $U\Sigma V^*$, then define $A_k = U\Sigma_k V^*$.

Eckart-Young-Mirsky Theorem: A_k minimizes both $\|A - B\|_F$ and $\|A - B\|_2$ over all matrices B of rank (at most) k .

The reason is that choosing the k largest singular values is both the way to maximize the sum of their squares (relevant to the Frobenius norm) and the way to minimize the size of any leftover singular value, i.e., of σ_{k+1} in sorted order (relevant to the 2-norm).

How good is the approximation? It depends on the size of $\sigma_{k+1}, \dots, \sigma_r$ (and their squares) in relation to the sizes of (the sum of squares of) the first k singular values. If the first k have the bulk of the magnitude, then the approximation can be quite good.

Example: $A = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$. Think of the rows as movies and the columns as users. Notice that

movie 1 is seen by everyone and user 1 is the most active. The emathhelp.net applet sorts the singular values in reverse order, giving (rounded to five places):

$$\Sigma \approx \begin{bmatrix} 0.29257 & & & & \\ & 0.72361 & & & \\ & & 1.16633 & & \\ & & & 1.33095 & \\ & & & & 3.04287 \end{bmatrix}$$

This has one distinctly low singular value and another one under 1. Its SVD comes with

$$U \approx \begin{bmatrix} -0.48209 & -0.23434 & 0.13187 & 0.44906 & 0.70258 \\ 0.55100 & -0.34647 & 0.30727 & -0.47362 & 0.50757 \\ -0.25405 & 0.76276 & -0.01555 & -0.45976 & 0.37687 \\ 0.34853 & 0.03548 & -0.86833 & 0.15420 & 0.31541 \\ 0.52722 & 0.49191 & 0.36599 & 0.58213 & 0.08507 \end{bmatrix}$$

and

$$V \approx \begin{bmatrix} 0.55847 & 0.30049 & -0.38132 & -0.24803 & 0.62521 \\ -0.63280 & 0.25146 & 0.36318 & -0.36389 & 0.52155 \\ 0.23554 & -0.80265 & 0.37652 & -0.01845 & 0.39770 \\ 0.15425 & 0.35595 & 0.42687 & 0.77478 & 0.25885 \\ -0.45650 & -0.27481 & -0.63143 & 0.45326 & 0.33455 \end{bmatrix}$$

Now suppose we delete the two smallest singular values at upper left. Then we also don't need the first two columns of U and V , the latter becoming the top two rows of V^* . We first compute

$$\begin{bmatrix} 0.13187 & 0.44906 & 0.70258 \\ 0.30727 & -0.47362 & 0.50757 \\ -0.01555 & -0.45976 & 0.37687 \\ -0.86833 & 0.15420 & 0.31541 \\ 0.36599 & 0.58213 & 0.08507 \end{bmatrix} \begin{bmatrix} 1.16633 & 0 & 0 \\ 0 & 1.33095 & 0 \\ 0 & 0 & 3.04287 \end{bmatrix} \approx \begin{bmatrix} 0.15381 & 0.59768 & 2.13787 \\ 0.35838 & -0.63036 & 1.54446 \\ -0.01814 & -0.61191 & 1.14676 \\ -1.01276 & 0.20523 & 0.95976 \\ 0.42687 & 0.77478 & 0.25885 \end{bmatrix}$$

Then multiplication with V^* gives

$$A_3 \approx \begin{bmatrix} 1.12973 & 0.95339 & 0.89712 & 1.08212 & 0.88901 \\ 0.67261 & 0.70630 & 0.73754 & 1.04116 & 0.57612 \\ 0.55827 & 0.38202 & 0.45161 & 0.77868 & 0.64955 \\ 0.16296 & 0.79370 & 0.75923 & 0.83976 & -0.22538 \\ -0.19311 & 0.00810 & 0.24937 & 0.84951 & 0.16823 \end{bmatrix}$$

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

Is this a reasonable approximation to $A = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$? The first and last rows are good.

The

entry in row 4, column 5 is way off, as are some others. But overall, not too shabby? Another reason this looks silly is that we not only need Σ_k but the relevant elements of U and V as well, which are all more complicated numbers than A has. However, the total number of entries is

$$km + k^2 + kn \text{ as compared with } mn \text{ entries in } A.$$

When $k \ll m, n$ this is a major savings. And when m, n are of order in the 1000s, $k = 100$ often gives a nice approximation.

Image Compression Examples. Companies that store user views of media content may have dimensions in the millions---and an even bigger motive to *calculate with reduced dimensions*. Then the approximations reflect the relative popularities of movies and other media content---while over in the column space of users, they indicate the patterns of frequent consumers.

We are most interested in compressing density-matrix representations of large quantum states.

Quantum Applications

(These notes draw on <https://www.math3ma.com/blog/understanding-entanglement-with-svd>)

First and simplest, SVD ideas give an easy way to tell whether a pure quantum state vector $|\phi\rangle$ is entangled. It finally leverages the relation between tensor product and outer product: Reshape $|\phi\rangle$ into the matrix A_ϕ that would occur if we really had $|\phi\rangle = |\phi_A\rangle \otimes |\phi_B\rangle$ from qubits held by Alice and Bob, respectively. Then we would have $A = |\phi_A\rangle \langle \phi_B|$ be of rank $r = 1$. So:

$|\phi\rangle$ is entangled between Alice and Bob if and only if A_ϕ has more than one nonzero singular value. The number of nonzero singular values quantifies the entanglement.

For the simplest example, $|\phi\rangle = \frac{1}{\sqrt{2}}[1, 0, 0, 1]^T$ gives $A_\phi = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. The matrix has rank $r = 2$.

So Alice and Bob are entangled. For vectors of length 4, this use of rank is equivalent to the "ratio test" that some of you have cited on homeworks---while for longer vectors and other possible divisions of qubits between "Alice" and "Bob", the use of rank is more general. In any event, when you give "Bob" the least significant (qu)bits---rightmost in big-endian form---the "rolling up" of the state vector is always in **row-major order**.

For example, the state $\frac{1}{2}(e_{000} + e_{001} + e_{110} - e_{111})$ gives the vector $[1, 1, 0, 0, 0, 0, 1, -1]^T$ (ignoring the

$\frac{1}{2}$). If Alice holds the first two qubits, it re-shapes as $\begin{bmatrix} 1 & 1 \\ 0 & 0 \\ 0 & 0 \\ 1 & -1 \end{bmatrix}$. This matrix has rank 2. But the state

$\frac{1}{2}(e_{000} + e_{001} + e_{110} + e_{111})$ becomes $\frac{1}{2} \begin{bmatrix} 1 & 1 \\ 0 & 0 \\ 0 & 0 \\ 1 & 1 \end{bmatrix}$ which has rank just 1 and so is *not* entangled. It is

$|\phi\rangle \langle +|$ with $|\phi\rangle$ as above. But if we gave Alice only the first qubit, then the shape would be $\begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}$. This does have rank $r = 2$, so qubit 1 is collectively entangled with Bob's "system" of qubits 2 and 3.

Believe-it-or-else, the following theorem is equivalent to one on the syllabus of MTH 309, but there it is not stated in our present quantum context. We may gloss over the statement and proof, since the applications can be understood by themselves.

Theorem: Let $|\phi\rangle$ be a pure state in the product $\mathbb{H}_A \otimes \mathbb{H}_B$ of two Hilbert spaces of dimensions d_A and d_B , respectively. Then we can find orthonormal bases $\{|i_A\rangle : 0 \leq i_A < d_A\}$ of \mathbb{H}_A and $\{|i_B\rangle : 0 \leq i_B < d_B\}$ of \mathbb{H}_B and positive numbers $\sigma_0, \dots, \sigma_{r-1}$ where $r \leq \min\{d_A, d_B\}$ such that

$$|\phi\rangle = \sum_{i=0}^{r-1} \sigma_i |i_A\rangle |i_B\rangle.$$

It follows that $\sum_i \sigma_i^2 = 1$ and that if we define $\rho_A := \text{Tr}_B(|\phi\rangle\langle\phi|)$ and $\rho_B := \text{Tr}_A(|\phi\rangle\langle\phi|)$, to be the density matrices resulting from tracing out \mathbb{H}_B , respectively tracing out \mathbb{H}_A , then

$$\rho_A = \sum_{i=0}^{r-1} \sigma_i^2 |i_A\rangle\langle i_A| \quad \text{and} \quad \rho_B = \sum_{i=0}^{r-1} \sigma_i^2 |i_B\rangle\langle i_B|.$$

The state $|\phi\rangle$ is separable over $\mathbb{H}_A \otimes \mathbb{H}_B$ if and only if this happens with $r = 1$. Otherwise, $|\phi\rangle$ is entangled with respect to $\mathbb{H}_A \otimes \mathbb{H}_B$, which is equivalent to $\text{Tr}(\rho_A^2) < 1$ and to $\text{Tr}(\rho_B^2) < 1$.

We've numbered from 0 because $d_A = 2^m$ and $d_B = 2^n$ are powers of 2 when we talk about "Alice" holding m qubits and "Bob" holding n qubits, and while we've been numbering qubits from 1, we've been numbering the standard basis from 0 to leverage the correspondence between binary strings and binary numbers. It is less usual to number singular values from 0, but this serves to emphasize that we may have exponentially many of them when m and n get large. Also bear in mind that the dimension of $\mathbb{H}_A \otimes \mathbb{H}_B$ is $d_A \cdot d_B$ with *times*, not $d_A + d_B$ as it would be with an ordinary Cartesian product. The whole representation is called the **Schmidt decomposition** of $|\phi\rangle$.

To visualize the theorem statement, it helps to say what happens when $|\phi\rangle$ really is a tensor product $|\psi_A\rangle \otimes |\psi_B\rangle$ with $|\psi_A\rangle \in \mathbb{H}_A$ and $|\psi_B\rangle \in \mathbb{H}_B$. Then, as we observed when the parital trace ("traceout") was introduced in week 13, we get $\text{Tr}_B(|\phi\rangle\langle\phi|) = |\psi_A\rangle\langle\psi_A|$ and $\text{Tr}_A(|\phi\rangle\langle\phi|) = |\psi_B\rangle\langle\psi_B|$. Since we can trivially extend the pure state $|\psi_A\rangle$ to an orthonormal basis of all of \mathbb{H}_A and $|\psi_B\rangle$ likewise for \mathbb{H}_B , we get the theorem conclusion by taking $r = 1$ and $\sigma_0 = 1$. Moreover, if the theorem conclusion happens with $r = 1$, then we must have $\sigma_1 = 1$ to normalize, and so we get $\rho_A = |0_A\rangle\langle 0_A|$ and $\rho_B = |0_B\rangle\langle 0_B|$, from which it follows (these being pure states, so that $\rho_A^2 = \rho_A$ and $\rho_B^2 = \rho_B$) that $|\phi\rangle = |0_A\rangle \otimes |0_B\rangle$. This proves the conclusion about entanglement without having to invoke the SVD. But the general proof is really crisp doing so.

Proof: The state vector of $|\phi\rangle$ has length $d_A \cdot d_B$, so we can reshape it into a $d_A \times d_B$ matrix A_ϕ as done above---so that entry $A_\phi[i, j]$ equals entry $d_B i + j$ of $|\phi\rangle$ (again, numbering from 0). Take the full SVD $A_\phi = U\Sigma V^*$ with U and V unitary and Σ in nonincreasing order. Then the columns of U form the desired orthonormal basis for \mathbb{H}_A , the columns of V likewise for \mathbb{H}_B , and taking r to be the rank of A_ϕ gives the reduced SVD representation $A_\phi = U_r \Sigma_r V_r^*$ as well. Then Σ_r is a diagonal matrix, so the only nonzero terms $u_i \sigma_i v_j^T$ are those with $j = i$. So $|\phi\rangle = \sum_{i=0}^{r-1} \sigma_i |i_A\rangle |i_B\rangle$ follows.

For the rest, the mere fact that $|\phi\rangle$ is a unit vector forces $\sum_i \sigma_i^2 = 1$. Now when we trace out Bob from $|\phi\rangle\langle\phi|$ under this representation we get a 1 entry left over from each of his submatrices on the main diagonal only---but the σ_i becomes σ_i^2 in $|\phi\rangle\langle\phi|$ so we get $\rho_A = \sum_{i=0}^{r-1} \sigma_i^2 |i_A\rangle\langle i_A|$ - note that $\sum_i \sigma_i^2 = 1$ is exactly what's needed for this to have unit trace and so be a legal density matrix. Likewise for ρ_B . The final fact is that whenever a sum of squares is 1, the sum of the corresponding fourth powers is less than 1 unless the sum is just a single 1 and the rest zeroes. \square

A simple example that also resonates with our idea of *truncating* SVDs of quantum states is at <https://bpb-us-w2.wpmucdn.com/u.osu.edu/dist/7/36891/files/2023/04/SchmidtDecomposition.pdf>

Let $|\phi\rangle = [\sqrt{.17}, \sqrt{.17}, \sqrt{.125}, \sqrt{.125}, \sqrt{.125}, \sqrt{.125}, \sqrt{0.08}, \sqrt{0.08}]^T$. This is a pure state of a 3-qubit system we'll call Alice, Charlie, and Bob in that order. This state has the form $|\psi\rangle \otimes |+\rangle$ for some 2-qubit state $|\psi\rangle$ of Alice \otimes Charlie alone. We can see this by "rolling" it into a 4×2 matrix where Alice and Charlie have the row space and Bob has the column space:

$$A'_\phi = \begin{bmatrix} \sqrt{.17} & \sqrt{.17} \\ \sqrt{.125} & \sqrt{.125} \\ \sqrt{.125} & \sqrt{.125} \\ \sqrt{0.08} & \sqrt{0.08} \end{bmatrix}$$

This has rank 1; you could also say it passes the "ratio test" with a ratio of 1 over all four rows. However, we are going to group it the other way: $\mathbb{H}_A = \mathbb{C}^2$ representing Alice by herself and $\mathbb{H}_B = \mathbb{C}^4$ for Charlie linked with Bob. Is it separable that way? Well, "reshaping" with two rows for Alice and four columns for \mathbb{H}_B gives

$$A_\phi = \begin{bmatrix} \sqrt{.17} & \sqrt{.17} & \sqrt{.125} & \sqrt{.125} \\ \sqrt{.125} & \sqrt{.125} & \sqrt{0.08} & \sqrt{0.08} \end{bmatrix}.$$

It is easy to see that this has full rank---the second row is not a scalar multiple of the first row---so the Schmidt rank is 2, and so $|\phi\rangle$ is not separable as an Alice \otimes (Charlie+Bob) system. However, we will develop a sense in which it comes weirdly close to being separable that way after all, indeed three-way separable. We will do this via the SVD of A_ϕ (in passing, the SVD of A'_ϕ above is relatively boring). The [emathhelp applet](#) actually allows entering square roots explicitly---and makes a link from what you enter:

Size of the matrix: \times

Matrix: A

$\text{sqrt}(.17)$	$\text{sqrt}(.17)$	$\text{sqrt}(1/8)$	$\text{sqrt}(1/8)$
$\text{sqrt}(1/8)$	$\text{sqrt}(1/8)$	$\text{sqrt}(0.08)$	$\text{sqrt}(0.08)$

The exact calculations get quite freaky with nested radicals, but the numerics come out the same as in the [first source](#). With $r = 2$ for the reduced SVD, we get:

$$\Sigma_2 = \begin{bmatrix} 0.99985947 & 0 \\ 0 & 0.01676428 \end{bmatrix}$$

Wow: σ_0 has almost all the bulk. (These rounded numbers' squares sum to **1.000000008325993** on my Windows calculator.) This asymmetry isn't obvious if you just look at the U and V matrices:

$$U = \begin{bmatrix} 0.7681475 & -0.6402729 \\ 0.6402729 & 0.7681475 \end{bmatrix},$$

$$V^* = \begin{bmatrix} 0.5431623 & 0.5431623 & 0.4527413 & 0.4527413 \\ 0.4527413 & 0.4527413 & -0.5431623 & -0.5431623 \end{bmatrix}$$

Yes, the squares of a column of U sum to **0.99999996823066** and squares in columns of V sum to **0.99999993773396** on my calculator. Now let us truncate by zeroing out the 0.01676428 entry. Since we want to preserve the property that the sum of σ_i^2 is 1, we also replace 0.99985947 simply by 1. This also allows us to discard the second column of U and the second row of V^* :

$$\begin{aligned} |\phi_1\rangle &= U_1 \Sigma_1 V_1^* = \begin{bmatrix} 0.7681475 \\ 0.6402729 \end{bmatrix} \begin{bmatrix} 0.5431623, & 0.5431623, & 0.4527413, & 0.4527413 \end{bmatrix} \\ &= [0.417229, 0.417229, 0.347772, 0.347772, 0.347772, 0.347772, 0.289878, 0.289878]. \end{aligned}$$

Rounded to six decimal places, these entries' squares sum to 1.000000042586, so this is legal like $|\phi\rangle = [0.412310, 0.412310, 0.353553, 0.353553, 0.353553, 0.353553, 0.282843, 0.282843]$ (also to

six decimal places). The differences in the second or third decimal place between entries of $|\phi_1\rangle$ and those of the original $|\phi\rangle$ are similar to how we truncated-and-rounded the singular values. But to compare probabilities, we need the entries' squares, which are under the square-root signs in $|\phi_1\rangle = [\sqrt{.174080}, \sqrt{.174080}, \sqrt{.120945}, \sqrt{.120945}, \sqrt{.120945}, \sqrt{.120945}, \sqrt{.084029}, \sqrt{.084029}]$ versus the original $[\sqrt{.17}, \sqrt{.17}, \sqrt{.125}, \sqrt{.125}, \sqrt{.125}, \sqrt{.125}, \sqrt{.08}, \sqrt{.08}]$. This is also not bad. The property of "Alice+Charlie" not being entangled with "Bob" is clear when we reshape $|\phi_1\rangle$ as

$$\begin{bmatrix} 0.417229 & 0.417229 \\ 0.347772 & 0.347772 \\ 0.347772 & 0.347772 \\ 0.289878 & 0.289878 \end{bmatrix}$$

since the columns are identical. For our drumroll conclusion---that Alice is not entangled with Bob+Charlie either---we also get separability under the reshaping

$$\begin{bmatrix} 0.417229 & 0.417229 & 0.347772 & 0.347772 \\ 0.347772 & 0.347772 & 0.289878 & 0.289878 \end{bmatrix}$$

because $\frac{0.417229}{0.347772} = \frac{0.347772}{0.289878} = 1.199720\dots$ So we have approximated the entangled state by the completely separable state

$$|\phi\rangle = \begin{bmatrix} 0.543162 \\ 0.839628 \end{bmatrix} \otimes \begin{bmatrix} 0.768148 \\ 0.640272 \end{bmatrix} \otimes |+\rangle.$$

The relationship to U and to one of the entries of V (equality up to the six-place rounding) is striking. Note also that the approximation did not affect Bob's qubit at all---it was separate and stayed separate.

We can conclude that although there is entanglement between Alice and (Charlie+Bob), as shown *qualitatively* by the rank of A_ϕ being 2, it is *quantitatively* weak. The appropriate metric for this is the **von Neumann entropy** of the *squares* of the singular values:

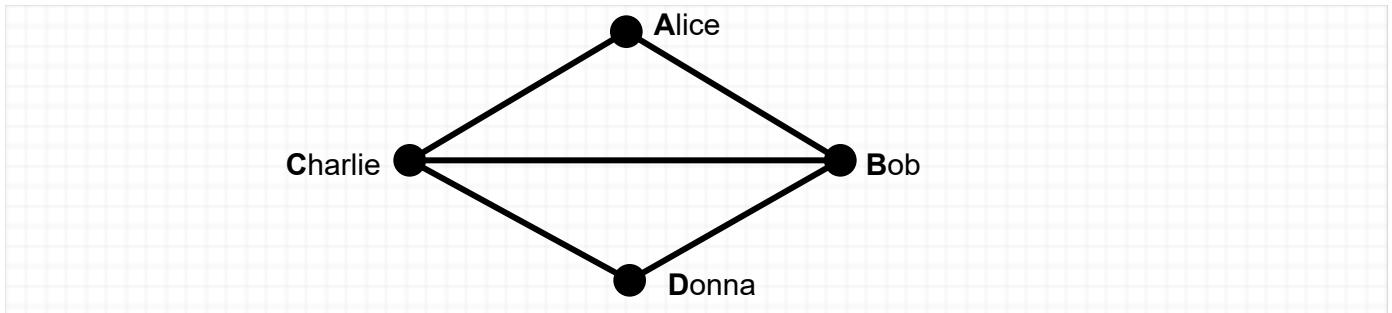
$$-\sum_{i=1}^r \sigma_i^2 \ln(\sigma_i^2).$$

(I am cutting corners here---the metric is really defined on the eigenvalues λ_i of density matrices that are Hermitian and PSD.) A term of this sum is near zero when σ_i is near zero **or** when σ_i is near 1. That is the case with both our singular values in Σ_2 above. So the entanglement in this case is near zero. For two-party breakdowns like "Alice" versus "Bob+Charlie", this metric, called **entanglement of formation (EF)**, when properly formalized) is recognized as the unique best measure. But for multi-partite entanglement---when you don't specify a binary division---there is as yet no such consensus. (I've tried to develop the idea that an "algebraic-geometric invariant" of n -qubit *circuits* C_n that produce

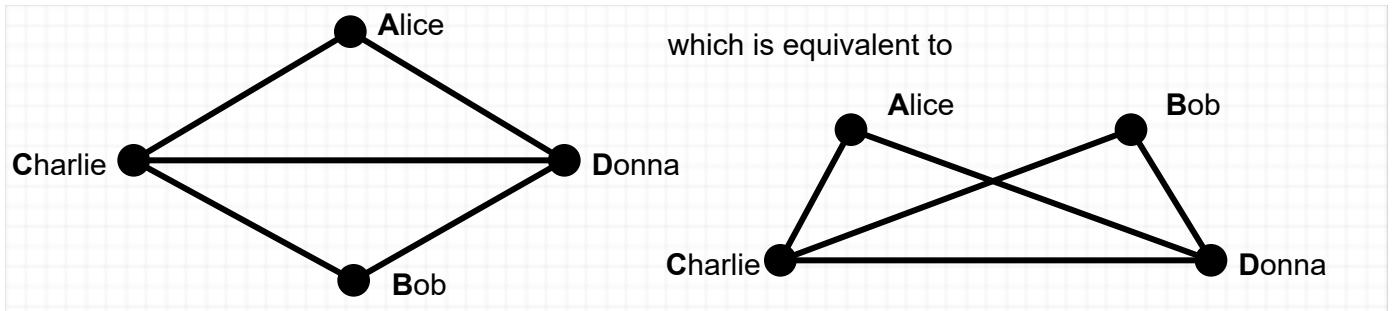
these states from input $|0^n\rangle$ can furnish such a measure---and also embody a complexity measure of the effort required to operate C_n coherently that is bigger than $O(s + n)$, where s is the count of gates in C_n ---but have not gotten it to work.)

Entanglement and Traceout (in Graph States)

For the graph on Assignment 7, let us suppose that qubit 1 is held by "Alice" (in big-endian order), qubit 2 by "Bob", qubit 3 by "Charlie", and qubit 4 by "Donna":



This is set up to do traceouts neatly for the divisions A versus B+C+D, A+B versus C+D, and A+B+C versus D. To handle the other three divisions, one must either take traces of non-contiguous square sub-matrices or permute the rows and columns so that the parties being traced out occupy the least significant (qu)bits. Likewise if you want to trace out A while leaving B+C+D intact, and so on. And actually, we will do that---so the labeling changes to



The graph state Γ is $\frac{1}{4}[1, 1, 1, -1, 1, -1, -1, -1, 1, -1, -1, -1, -1, 1, 1, 1, -1]^T$ (before the second Hadamard transform) and the whole 16×16 self-outerproduct density matrix $|\Gamma\rangle\langle\Gamma|$ is

$$\rho = \frac{1}{16} \begin{bmatrix} 1 & 1 & 1 & -1 & 1 & -1 & -1 & -1 & 1 & -1 & -1 & -1 & 1 & 1 & 1 & 1 & -1 \\ 1 & 1 & 1 & -1 & 1 & -1 & -1 & -1 & 1 & -1 & -1 & -1 & 1 & 1 & 1 & 1 & -1 \\ 1 & 1 & 1 & -1 & 1 & -1 & -1 & -1 & 1 & -1 & -1 & -1 & 1 & 1 & 1 & 1 & -1 \\ -1 & -1 & -1 & 1 & -1 & 1 & 1 & 1 & -1 & 1 & 1 & 1 & -1 & -1 & -1 & 1 & 1 \\ 1 & 1 & 1 & -1 & 1 & -1 & -1 & -1 & 1 & -1 & -1 & -1 & 1 & 1 & 1 & 1 & -1 \\ -1 & -1 & -1 & 1 & -1 & 1 & 1 & 1 & -1 & 1 & 1 & 1 & -1 & -1 & -1 & 1 & 1 \\ -1 & -1 & -1 & 1 & -1 & 1 & 1 & 1 & -1 & 1 & 1 & 1 & -1 & -1 & -1 & 1 & 1 \\ -1 & -1 & -1 & 1 & -1 & 1 & 1 & 1 & -1 & 1 & 1 & 1 & -1 & -1 & -1 & 1 & 1 \\ -1 & -1 & -1 & 1 & -1 & 1 & 1 & 1 & -1 & 1 & 1 & 1 & -1 & -1 & -1 & 1 & 1 \\ 1 & 1 & 1 & -1 & 1 & -1 & -1 & -1 & 1 & -1 & -1 & -1 & 1 & 1 & 1 & 1 & -1 \\ -1 & -1 & -1 & 1 & -1 & 1 & 1 & 1 & -1 & 1 & 1 & 1 & -1 & -1 & -1 & 1 & 1 \\ -1 & -1 & -1 & 1 & -1 & 1 & 1 & 1 & -1 & 1 & 1 & 1 & -1 & -1 & -1 & 1 & 1 \\ -1 & -1 & -1 & 1 & -1 & 1 & 1 & 1 & -1 & 1 & 1 & 1 & -1 & -1 & -1 & 1 & 1 \\ 1 & 1 & 1 & -1 & 1 & -1 & -1 & -1 & 1 & -1 & -1 & -1 & 1 & 1 & 1 & 1 & -1 \\ 1 & 1 & 1 & -1 & 1 & -1 & -1 & -1 & 1 & -1 & -1 & -1 & 1 & 1 & 1 & 1 & -1 \\ 1 & 1 & 1 & -1 & 1 & -1 & -1 & -1 & 1 & -1 & -1 & -1 & 1 & 1 & 1 & 1 & -1 \\ -1 & -1 & -1 & 1 & -1 & 1 & 1 & 1 & -1 & 1 & 1 & 1 & -1 & -1 & -1 & 1 & 1 \end{bmatrix}$$

For the first case where we trace out all of Bob, Charlie, and Donna, we make four 8×8 quadrants and take the trace of each. This yields $\frac{1}{16} \begin{bmatrix} 8 & 0 \\ 0 & 8 \end{bmatrix} = \begin{bmatrix} 0.5 & 0 \\ 0 & 0.5 \end{bmatrix}$. This is the matrix of the completely mixed state on one qubit. This does imply that Alice is entangled with the others in this state. We can also see this insofar as ρ is not a tensor product of a 2×2 matrix with an 8×8 matrix.

For Alice+Bob versus Charlie+Donna, tracing the latter two out, we divide into sixteen 4×4 blocks. Taking the trace of each gives

$$\rho_{A,B} = Tr_{C,D}(\rho) = \frac{1}{16} \begin{bmatrix} 4 & 0 & 0 & 4 \\ 0 & 4 & 4 & 0 \\ 0 & 4 & 4 & 0 \\ 4 & 0 & 0 & 4 \end{bmatrix} = \frac{1}{4} \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix}.$$

This is not a pure state: If you square it, you get $1/8$ not $1/4$ in the upper-left corner, so $\rho_{A,B}^2 \neq \rho_{A,B}$. It is not a directly separable matrix. It is, however, a sum (that is, a mixture) of separable pure states:

$$\rho_{A,B} = \frac{1}{8} \left(\begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} + \begin{bmatrix} 1 & -1 & -1 & 1 \\ -1 & 1 & 1 & -1 \\ -1 & 1 & 1 & -1 \\ 1 & -1 & -1 & 1 \end{bmatrix} \right) = \frac{1}{2} |++\rangle\langle++| + \frac{1}{2} |--\rangle\langle--|.$$

A *classical mixture* of separable states counts as separable. The context needs parsing out:

- The fact that $\rho_{A,B}$ is a mixed state traced out from a pure state---technically that $Tr(\rho_{A,B}^2) < Tr(\rho_{A,B})$ ---means that Alice+Bob are entangled *with Charlie and Donna*.
- But Alice and Bob are not entangled between themselves.

The second point leads to an ultimate subtlety. It is best viewed if we transform back to the standard basis and consider the state

$$\tau = \frac{1}{2}|00\rangle\langle 00| + \frac{1}{2}|11\rangle\langle 11| = \begin{bmatrix} 0.5 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.5 \end{bmatrix}.$$

This is not the same as the self-outerproduct of $\Phi^+ = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, which has 0.5 also in the upper-right and lower-left corners. It is not a separable matrix---and its *Schmidt rank* is 2 not 1. Its singular values are 0.5 twice then 0 twice. If you think of Alice and Bob each having a classical coin, it sure seems the coins are entangled by dint of saying that either they both get heads or both get tails. But as a mixture, it is a *single* coin that either comes up "double heads" or "double tails."

Well, there are other ways to get the matrix τ as a mixture involving entangled states, namely as

$$\frac{1}{2}\Phi^+ + \frac{1}{2}\Phi^- = \frac{1}{2}\begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 1 \end{pmatrix}.$$

But the fact that there is a way to get it as a mixture of states that are separable (with the same division of qubits) makes it separable.

Definition. A **mixed state** is **separable** if it can be written as a distribution over pure states that are separable the same way.

The fact that a pure state cannot be written as a proper mixture (proof for 1-qubit states and general intuition: all proper mixtures are within the Bloch sphere, not on its surface) makes this definition self-consistent. It would be nice to have a simple criterion by which we could decide whether an arbitrary legal *density matrix* of a **mixed** state is separable or entangled---the way we can do with the "rolled-up" matrices from a vector of a **pure** state. However, there likely isn't one: the general problem of determining whether a density matrix ρ is separable is **NP-hard**. There are cases where you can prove separability, such as ρ being a tensor product already or finding a separating mixture. There are also cases where one can prove entanglement. The [Peres-Horodecki criterion](#) says to divide ρ into blocks as if you were going to ``trace out Bob," but instead transpose each of the sub-matrices (*without* conjugating them) to get a matrix ρ' . If ρ' has a negative eigenvalue, then ρ is entangled. For two-qubit systems (and ones of a qubit and a three-way **qutrit**), this is an "only if" as well. So the fact that τ above is unchanged upon transposing its 2×2 "Bob blocks" and has all eigenvalues ≥ 0 means it is separable. But for more qubits, no such simple criterion is known and the whole question devolves into the mysteries of computational complexity.