

CSE439 Week 14: Applications of the SVD and Classical Simulations

First, the version we proved last time, where $A = U\Sigma V^*$ 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 , which I show in the next example.

2 × 2 Example (based on https://math.mit.edu/classes/18.095/2016IAP/lec2/SVD_Notes.pdf)

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

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

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

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

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 $\|A\|_2 = \sqrt{5}$ and $\|A\|_F = \sqrt{45+5} = 5\sqrt{2}$.

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 called the "Hadamard 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 $B = -iY$ considered between the practice and actual Prelim II exams.)

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.

(Pseudo-)Inversion and Numerical Instability

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-(Bjerrhammer)-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---we will elaborate on this below. But such cavalier blipping of large values E betrays the fact of **numerical instability** lurking in concepts of inversion.

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 $B = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$? $B^TB = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$. 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

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

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 $\|Az - \mathbf{b}\|_2 \leq \|Ax - \mathbf{b}\|_2$ 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 $A^+AA^+ = A^+$ 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

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

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 (cf. <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.

The state in problem (3) of the Prelim II practice exam is $\frac{1}{2}(e_{000} + e_{001} + e_{110} - e_{111})$, which 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 $\begin{bmatrix} 1 & 1 \\ 0 & 0 \\ 0 & 0 \\ 1 & 1 \end{bmatrix}$ 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.

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 $\rho_A := Tr_B(|\phi\rangle\langle\phi|)$ and $\rho_B := Tr_A(|\phi\rangle\langle\phi|)$, 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 $Tr(\rho_A^2) < 1$ and to $Tr(\rho_B^2) < 1$.

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 partial trace ("traceout") was introduced in week 13, we get $Tr_B(|\phi\rangle\langle\phi|) = |\psi_A\rangle\langle\psi_A|$ and $Tr_A(|\phi\rangle\langle\phi|) = |\psi_B\rangle\langle\psi_B|$. 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 illustrated above 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. ☒

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. 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 Bob linked with Charlie. 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 so. In passing, let us note that the other reshaping,

$$A'_\phi = \begin{bmatrix} \sqrt{.17} & \sqrt{.17} \\ \sqrt{.125} & \sqrt{.125} \\ \sqrt{.125} & \sqrt{.125} \\ \sqrt{0.8} & \sqrt{0.8} \end{bmatrix},$$

just as obviously has rank only 1, so $|\phi\rangle$ is separable as (Alice+Charlie) \otimes Bob. The SVD of A'_ϕ is relatively boring. But let's go ahead and do the SVD of A_ϕ . The [emathhelp applet](#) actually allows entering square roots explicitly:

Size of the matrix: ×

Matrix: **A**

$\sqrt{.17}$	$\sqrt{.17}$	$\sqrt{1/8}$	$\sqrt{1/8}$
$\sqrt{1/8}$	$\sqrt{1/8}$	$\sqrt{0.08}$	$\sqrt{0.08}$

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^* :

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

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 $[\sqrt{.17}, \sqrt{.17}, \sqrt{.125}, \sqrt{.125}, \sqrt{.125}, \sqrt{.125}, \sqrt{0.08}, \sqrt{0.08}]$. 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 wiyth 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.

Problem Adding 10 Pts. to Prelim II

Use SVD truncation to find a completely separated three-qubit state that approximates this one:



0.25000000+0.00000000i	000>	6.2500%
0.25000000+0.00000000i	001>	6.2500%
0.75000000+0.00000000i	010>	56.2500%
-0.25000000+0.00000000i	011>	6.2500%
0.25000000+0.00000000i	100>	6.2500%
0.25000000+0.00000000i	101>	6.2500%
-0.25000000+0.00000000i	110>	6.2500%
-0.25000000+0.00000000i	111>	6.2500%

You are welcome to use an applet to compute the original SVD (again mindful of possible different ordering in operations and displays), but you are required to show the steps of truncating it and multiplying $U_1 \Sigma_1 V_1^*$ on-paper manually.

Can We Scale This Up?

Approximating entangled states by separable states---and telling properties of mixed states whether they are given as traceouts or not---goes into research that is plagued by **NP-hardness**. Bear in mind that the SVD representations have the same exponential " N " scaling as the underlying state vectors and matrices---as opposed to the order- n scaling of quantum circuits. Scott Aaronson makes these points pithily in his own [notes](#). [Added: I mentioned the analogy between *tensor contraction* and *database join* further down. The paper <https://arxiv.org/html/2209.12332v5>, from this past October, leverages this to show that even though certain problems of optimal contraction order are NP-hard, in the nice case of tree tensor networks and with a linearity condition, polynomial time algorithms are available. It also prominently references the dissertation work of Mahmoud Abo Khamis under Drs. Atri Rudra and Hung Ngo here at UB.]

Thus we cannot expect to be able to generate good approximations of arbitrary quantum states "given cold." This leaves two main possibilities as I see them:

1. Carry along succinct approximations to quantum states inductively as they are processed and built up in quantum circuits.
2. Focus on families of quantum states that have special structure that promotes classical approximations.

The main argument for 1 is evidently Nature computes efficiently, so has some way to avoid the exponential blowup that is ingrained in our explicit notation. Whether that applies to something as advanced as Shor's algorithm incurs other considerations---as the real-world quantum feasibility of Shor's algorithm is still not really established.

The rationale for 2 requires that the special structure does not impede the usefulness of quantum circuits/algorithms that abide by it. The major structural divide we have seen is between the Clifford family of gates: **H**, **X**, **Y**, **Z**, **S**, **CNOT**, **CZ**, versus the fact that adding any one of the gates **T**, **CS**, **CCZ**, or the Toffoli gate **CCX** gives the full power of quantum computation. The more fruitful structural limitations may apply to how gates are combined in circuits rather than which gates are allowed.

On the latter there is one major strand I know: Circuits that can be modeled as **tensor networks** that are close to being *trees* can be simulated classically with reasonable overhead. So we will say some words about tensor networks, as they are vital in classical machine learning as well.

A **tensor** T is a possibly higher-dimensional matrix. In the text's functional notation with tiered indexing, it is represented by a multi-ary function $T(i, j, k, \dots)$. The **order** is the number of tiers. In a **tensor**

network, each tensor is a node of degree equal to its order. Edges, commonly called "legs", do not have to go to another node; they can be "free". Those that do go to another node (and so become a shared leg of the other node) represent setting up a **contraction**. The allowed operations in a tensor network include:

1. Introduce a new tensor---this is implicitly a tensor product with the existing tensors.
2. **Reshape** a tensor in a way that changes its order.
3. **Contract** two nodes along one or more shared legs. *Matrix product* is the canonical simple example. The generalized concept was employed by Einstein via the [Einstein summation convention](#).

Segue to sources:

<https://arxiv.org/abs/1306.2164> (a representative research-level survey from 2014)

https://www.benasque.org/2020scs/talks_contr/106_tensor_networks_lecture1.pdf

(high-level but slides 15--20 are the best exposition of the SVD-based simulation idea in general contexts that I have found)

and especially

https://www.quantumcomputinglab.cineca.it/wp-content/uploads/2021/10/MPS_Lecture.pdf

(the most immediately accessible gateway to the main idea that I've found)

https://pennylane.ai/qml/demos/tutorial_tn_circuits

(shows how we could actually program this stuff)