

So are we done? Not quite. Imagine the linear system of our robot car from lab, and consider the problem of bringing it to a particular state (i.e. assigning particular values to the wheel velocities) at a certain time. One way would be to apply a steady input to gradually ramp up the wheel velocities, so we reach the target state at the desired time. Another way, however, could be to apply large random inputs, accelerating and decelerating each wheel, until just before the target time t , at which point we would apply large controls to set the wheel velocities to their desired values. Though both approaches accomplish the same goal, the former should seem more “natural” than the latter.

More generally, in the case of arbitrary controllable linear systems, out of all the inputs \vec{u} that take us to \vec{x}^* , we will aim to choose the inputs that minimize some cost function, which expresses how “unnatural” or “bad” our inputs are. In our robotic car context, it would make sense to minimize the current and power drawn from the battery, which is directly correlated with the PWM input to our system. In this case, a plausible choice for this cost function would be the norm $\|\vec{u}\|$ of the inputs. With this cost function, the first example of applying steady inputs has a smaller norm than the large, randomly varying inputs of the second example, suggesting that this definition of the cost function is in accordance with our intuition. The problem of computing the control input that reaches a goal state while minimizing the norm of the inputs is known as *minimum energy control* or alternatively the *minimum norm solution* for the system of equations.

2 Minimum Norm Solution

The problem of minimum energy control is really just an example of solving *underdetermined* systems. We will consider the general case where we have an A matrix that is $n \times t$ where $n < t$ and we want to solve

$$A\vec{w} = \vec{y} \quad (6)$$

for some parameters \vec{w} . We assume that $\vec{y} \in \text{Col}(A)$ because otherwise, the equation is not satisfied. But since A is wide, it may have a null space, so it is possible for infinitely many solutions \vec{w} to exist. Thus, we wish to pick the solution with minimum norm as the best solution.

Let’s try to make sense of this problem by the geometric viewpoint. Let \vec{w}_0 be a particular solution to the above equation (that is not necessarily of minimum norm), and recall that the null space of A is a subspace of t -dimensional space. Furthermore, observe that for any vector $\vec{z} \in \text{Null}(A)$, $\vec{w}_0 + \vec{z}$ is a solution to our equation, since

$$A(\vec{w}_0 + \vec{z}) = A\vec{w}_0 + A\vec{z} = A\vec{w}_0 = \vec{y}. \quad (7)$$

Furthermore, from the same calculation, any possible solution \vec{w} where $A\vec{w} = \vec{y}$ can be written in the form $\vec{w}_0 + \vec{z}$ for some $\vec{z} \in \text{Null}(A)$. Intuitively, we don’t want to waste any of our parameter vector \vec{w} on parts which live in $\text{Null}(A)$ since they are not affecting the result. Thus the control input \vec{w} with minimum norm should be one that is entirely orthogonal to $\text{Null}(A)$. This would seem to be analogous to what we saw when studying least squares, where again we were interested in choosing a point on a subspace such that the error vector was orthogonal to said subspace.

Let’s try to prove this rigorously. Imagine that we had some orthonormal basis

$$\vec{v}_1, \vec{v}_2, \dots, \vec{v}_t \quad (8)$$

of our t -dimensional space of parameters, ordered such that the first r vectors $\vec{v}_1, \dots, \vec{v}_r$ form a basis for those vectors orthogonal to $\text{Null}(A)$, and the remaining $t - r$ vectors $\vec{v}_{r+1}, \dots, \vec{v}_t$ form a basis for $\text{Null}(A)$. We will worry about constructing this basis explicitly later – for now, let’s just assume that it exists. (It

must exist since we can get a basis for the nullspace using 16A techniques and then orthonormalize it using Gram-Schmidt. Then, we can extend it to a full basis for the space again using Gram-Schmidt.)

Consider our \vec{w}_0 expressed in this orthonormal basis, as follows:

$$\vec{w}_0 = \alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \dots + \alpha_t \vec{v}_t = V \vec{\alpha}. \quad (9)$$

Observe that the norm of \vec{w}_0 is

$$\|\vec{w}_0\| = \sqrt{\vec{w}_0^\top \vec{w}_0} = \sqrt{\vec{\alpha}^\top V^\top V \vec{\alpha}} \quad (10)$$

$$= \sqrt{\vec{\alpha}^\top \vec{\alpha}} = \|\vec{\alpha}\|, \quad (11)$$

since the \vec{v}_i are orthonormal, so $V^\top V = I$.

Now, observe that in this basis, any $\vec{z} \in \text{Null}(A)$ can be represented as

$$\vec{z} = \beta_{r+1} \vec{v}_{r+1} + \beta_{r+2} \vec{v}_{r+2} + \dots + \beta_t \vec{v}_t, \quad (12)$$

since (by definition) it has components only in the the null space of A . Thus, we can write any solution $\vec{w} = \vec{w}_0 + \vec{z}$ as

$$\vec{w} = \vec{w}_0 + \vec{z} = \alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \dots + \alpha_r \vec{v}_r + (\alpha_{r+1} + \beta_{r+1}) \vec{v}_{r+1} + \dots + (\alpha_t + \beta_t) \vec{v}_t. \quad (13)$$

In a similar manner to what we did before, we see immediately that the norm of this expression is

$$\|\vec{w}\| = \sqrt{\alpha_1^2 + \alpha_2^2 + \dots + \alpha_r^2 + (\alpha_{r+1} + \beta_{r+1})^2 + \dots + (\alpha_t + \beta_t)^2}. \quad (14)$$

To minimize the norm of \vec{w} , therefore, we should set $\beta_i = -\alpha_i$ for all valid $i > r$. Therefore, our minimum energy control must be

$$\vec{w} = \vec{w}_0 + \vec{z} \quad (15)$$

$$= \alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \dots + \alpha_r \vec{v}_r + (\alpha_{r+1} + \beta_{r+1}) \vec{v}_{r+1} + \dots + (\alpha_t + \beta_t) \vec{v}_t \quad (16)$$

$$= \alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \dots + \alpha_r \vec{v}_r + (\alpha_{r+1} - \alpha_{r+1}) \vec{v}_{r+1} + \dots + (\alpha_t - \alpha_t) \vec{v}_t \quad (17)$$

$$= \alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \dots + \alpha_r \vec{v}_r. \quad (18)$$

Observe that this solution is entirely orthogonal to $\text{Null}(A)$, as we had expected. All that remains now is to demonstrate the existence of the \vec{v}_i .

3 Constructing an Orthonormal Basis

From Gaussian elimination, we know how to compute a basis of $\text{Null}(A)$, which we can orthonormalize using Gram-Schmidt to produce the $\vec{v}_{r+1}, \dots, \vec{v}_t$. We can then extend this basis again using Gram-Schmidt to span all of t -dimensional space, to produce the $\vec{v}_1, \dots, \vec{v}_r$, demonstrating the existence of our desired basis. Although this description is not fully precise, with enough effort we can make it rigorous and solve the problem of minimum energy control computationally. We don't get nice expressions, however, and so there is less insight to be had with this purely procedural approach.

Instead, we will choose to attack this problem from a different direction. By definition, the \vec{v}_i are an orthonormal basis of t -dimensional space. But recall from earlier that, by the real spectral theorem, the

eigenvectors of a real symmetric matrix can give such an orthonormal basis. Wouldn't it be interesting if some symmetric matrix Q had exactly these eigenvectors $\vec{v}_1, \dots, \vec{v}_t$? To be extra useful, it would be nice if $\vec{v}_{r+1}, \dots, \vec{v}_t$ are in fact a basis for $\text{Null}(Q)$. That is, we want Q to have the same nullspace as A .

How do we generate such a symmetric matrix Q ? Since it has the same null space as A , perhaps we could try writing it in the form

$$Q = DA, \tag{19}$$

where D is some unknown matrix of full column rank. With this choice of D , we get that Q and A have the same null space. Now, as Q is symmetric, we can take transposes to obtain

$$Q = Q^\top = A^\top D^\top \implies DA = A^\top D^\top. \tag{20}$$

Looking at the latter equality, a natural conjecture would be to try $D = A^\top$, so $Q = A^\top A$. It turns out that this is precisely the matrix that works. Remember that this was actually already proven to you in 16A when discussing the condition for when least squares works. We will prove it again here since it is an important result.

First, we show that $\text{Null}(A) \subseteq \text{Null}(A^\top A)$. Specifically, for any $\vec{v} \in \text{Null}(A)$, we'd like to show that $\vec{v} \in \text{Null}(A^\top A)$. Indeed,

$$A\vec{v} = \vec{0} \implies A^\top A\vec{v} = A^\top \vec{0} = \vec{0} \implies \vec{v} \in \text{Null}(A^\top A). \tag{21}$$

Thus $\text{Null}(A) \subseteq \text{Null}(A^\top A)$.

Let's try to prove the opposite relation, that is $\text{Null}(A) \supseteq \text{Null}(A^\top A)$, in order to show equality. Specifically, for any $\vec{v} \in \text{Null}(A^\top A)$, we'd like to show that $\vec{v} \in \text{Null}(A)$. Thus, we start off from a \vec{v} such that

$$Q\vec{v} = A^\top A\vec{v} = \vec{0}. \tag{22}$$

Now let us pre-multiply by \vec{v}^\top to get

$$\vec{v}^\top A^\top A\vec{v} = \vec{v}^\top \vec{0} = 0 \tag{23}$$

$$\|A\vec{v}\|^2 = 0 \tag{24}$$

$$\|A\vec{v}\| = 0 \tag{25}$$

$$\implies A\vec{v} = 0 \tag{26}$$

where by definition a norm is 0 if and only if the vector is the zero vector $\vec{0}$. Thus, we've shown $\text{Null}(A^\top A) \subseteq \text{Null}(A)$. Since we've proven this inequality in both directions, we have $\text{Null}(Q) = \text{Null}(A^\top A) = \text{Null}(A)$, as we desired. Thus, we can produce an orthonormal basis of the eigenspace corresponding to $\lambda = 0$ of Q that provides us with the $\vec{v}_{r+1}, \dots, \vec{v}_t$ that we wanted.

Considering the remaining eigenvectors of Q for $\lambda \neq 0$, by the real spectral theorem they are all mutually orthogonal and will all be orthogonal to $\text{Null}(Q)$. So we can choose the remaining r eigenvectors to form our $\vec{v}_1, \dots, \vec{v}_r$, completing our construction of the $\{\vec{v}_i\}$.

4 Singular Values

Now that we know how to construct this Q , and have demonstrated that its eigenvectors are exactly the $\vec{v}_1, \dots, \vec{v}_t$ that we had wanted, it is natural to speculate on the eigenvalues of each of these eigenvectors. We know that the eigenvalues of $\vec{v}_{r+1}, \dots, \vec{v}_t$ are all 0, since they lie in the null space of Q . But what about the eigenvalues for the first r eigenvectors? Let's try to work this out algebraically.

By the definition of eigenvectors, we can perform manipulations very similar to what we did earlier involving nullspaces, to see that

$$Q\vec{v}_i = \lambda_i\vec{v}_i \quad (27)$$

$$A^\top A\vec{v}_i = \lambda_i\vec{v}_i \quad (28)$$

$$\vec{v}_i^\top A^\top A\vec{v}_i = \lambda_i\vec{v}_i^\top \vec{v}_i \quad (29)$$

$$\|A\vec{v}_i\|^2 = \lambda_i\|\vec{v}_i\|^2 \quad (30)$$

$$\lambda_i = \frac{\|A\vec{v}_i\|^2}{\|\vec{v}_i\|^2} \quad (31)$$

$$= \|A\vec{v}_i\|^2. \quad (32)$$

Notice that $\lambda_i \geq 0$ due to the square. For $i > r$, the eigenvalues must be 0 since the eigenvectors are in the nullspace. But for $i \leq r$, since $\vec{v}_i \notin \text{Null}(A)$, $\lambda_i \neq 0$, so $\lambda_i > 0$.

This is interesting, as we can now order our \vec{v}_i such that

$$\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_r > \lambda_{r+1} = \lambda_{r+2} = \dots = \lambda_t = 0, \quad (33)$$

and place them as columns of the eigenvector matrix

$$V = \begin{bmatrix} | & & | \\ \vec{v}_1 & \dots & \vec{v}_t \\ | & & | \end{bmatrix}. \quad (34)$$

We can then write the eigendecomposition of Q as

$$Q = V\Lambda V^{-1} = V \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_t \end{bmatrix} V^\top. \quad (35)$$

Since all the $\lambda_i \geq 0$, we can define $\sigma_i = \sqrt{\lambda_i}$ for all i , which gives us

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > \sigma_{r+1} = \sigma_{r+2} = \dots = \sigma_t = 0 \quad (36)$$

Recall that $\lambda_i = \|A\vec{v}_i\|^2$, so $\sigma_i = \|A\vec{v}_i\|$. These σ_i are known as the *singular values* of A , and will prove important in the next section.

5 Constructing Another Orthonormal Basis For the Output

Now, let's try changing coordinates of our parameters \vec{w} to be in V -basis, so let $\vec{\tilde{w}} = V^{-1}\vec{w} = V^T\vec{w}$. Then to produce the same \vec{y} , we need to have $\tilde{A} = AV$ so that

$$\tilde{A}\vec{\tilde{w}} = (AV)(V^T\vec{w}) = A\vec{w} = \vec{y}. \quad (37)$$

How does $\tilde{A} = AV = [A\vec{v}_1 \ \dots \ A\vec{v}_t]$ look like? Since $\vec{v}_{r+1}, \dots, \vec{v}_t$ are all in the nullspace of A , we know that the last $t - r$ columns of \tilde{A} are all $\vec{0}$. So what about the first r columns of \tilde{A} , which are $A\vec{v}_i$ for $i = 1, \dots, r$?

Note that $\vec{v}_1, \dots, \vec{v}_r$ are exactly the vectors orthogonal to the null space of A , and thus a basis for the subspace of all inputs to A that create a non-zero output. Moreover, this means $A\vec{v}_1, \dots, A\vec{v}_r$ must exactly be a basis for all the outputs of applying A to any vector, meaning they are a basis for the column space of A .

Additionally, just as the \vec{v}_i were mutually orthogonal, we conjecture that this set of $A\vec{v}_i$ is mutually orthogonal too. Let's see why. For any valid choice of i, j , we see that

$$(A\vec{v}_j)^T(A\vec{v}_i) = \vec{v}_j^T A^T A \vec{v}_i \quad (38)$$

$$= \sigma_i^2 \vec{v}_j^T \vec{v}_i = \begin{cases} 0 & \text{if } i \neq j \\ \sigma_i^2 & \text{if } i = j \end{cases} \quad (39)$$

due to the eigenvector property of \vec{v}_i and since \vec{v}_i and \vec{v}_j are orthonormal.

To make each of these vectors of unit length, we should divide them by their length $\|A\vec{v}_i\|$, which we defined earlier as the singular value σ_i . Thus, we define the orthonormal vectors $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_r$ that form a basis for $\text{Col}(A)$ with

$$\vec{u}_i = \frac{A\vec{v}_i}{\sigma_i} \quad (40)$$

for all valid $i \leq r$. Rearranging, we get

$$A\vec{v}_i = \sigma_i \vec{u}_i. \quad (41)$$

Horizontally stacking this result for all $i \leq r$, we find that

$$A \begin{bmatrix} | & & | \\ \vec{v}_1 & \dots & \vec{v}_r \\ | & & | \end{bmatrix} = \begin{bmatrix} | & & | \\ \vec{u}_1 & \dots & \vec{u}_r \\ | & & | \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & \dots & 0 \\ 0 & \sigma_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_r \end{bmatrix} \quad (42)$$

$$AV_r = U_r \Sigma_r \quad (43)$$

where V_r is $t \times r$, U_r is $n \times r$, and Σ_r is $r \times r$. Is anything missing? Well, recall that each of the \vec{v}_i are t dimensional and each of the \vec{u}_i are n dimensional, but we only use $r \leq n < t$ of each here. To fix this, we can simply "pad" the V_r and U_r matrices with the remaining \vec{v}_i and \vec{u}_i for $i > r$.

We will first pad the U_r matrix. We want to extend the existing r orthonormal vectors to a full set of n orthonormal vectors in order to create an orthonormal basis. To do this, we can apply Gram-Schmidt

to get the remaining orthonormal vectors, for example by applying it to $\begin{bmatrix} \vec{u}_1 & \dots & \vec{u}_r & | & I_n \end{bmatrix}$ to get $\begin{bmatrix} \vec{u}_1 & \dots & \vec{u}_r & \vec{u}_{r+1} & \dots & \vec{u}_n \end{bmatrix}$ as our new U matrix. Then we need to add more zeros on the bottom of Σ_r to scale the new \vec{u}_i so that they don't change our product:

$$A \begin{bmatrix} | & & | \\ \vec{v}_1 & \dots & \vec{v}_r \\ | & & | \end{bmatrix} = \begin{bmatrix} | & & | \\ \vec{u}_1 & \dots & \vec{u}_n \\ | & & | \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & \dots & 0 \\ 0 & \sigma_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_r \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix} \quad (44)$$

$$= U \begin{bmatrix} \Sigma_r \\ \mathbf{0}_{(n-r) \times r} \end{bmatrix} \quad (45)$$

We then pad the V_r matrix with the other $\vec{v}_{r+1}, \dots, \vec{v}_t$. Then the new columns on the RHS should just be $\vec{0}$ since these \vec{v}_i are in the nullspace of A , so the diagonal matrix of the σ_i should have more zero columns on the right:

$$A \begin{bmatrix} | & & | \\ \vec{v}_1 & \dots & \vec{v}_t \\ | & & | \end{bmatrix} = \begin{bmatrix} | & & | \\ \vec{u}_1 & \dots & \vec{u}_n \\ | & & | \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & \dots & 0 \\ 0 & \sigma_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_r \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix} \mathbf{0}_{n \times (t-r)} \quad (46)$$

$$= U \begin{bmatrix} \Sigma_r & \mathbf{0}_{r \times (t-r)} \\ \mathbf{0}_{(n-r) \times r} & \mathbf{0}_{(n-r) \times (t-r)} \end{bmatrix} \quad (47)$$

Now, since V is a square matrix with orthonormal columns, $V^{-1} = V^\top$, so we can post-multiply both sides by V^\top to obtain

$$A = \begin{bmatrix} | & & | \\ \vec{u}_1 & \dots & \vec{u}_n \\ | & & | \end{bmatrix} \begin{bmatrix} \Sigma_r & \mathbf{0}_{r \times (t-r)} \\ \mathbf{0}_{(n-r) \times r} & \mathbf{0}_{(n-r) \times (t-r)} \end{bmatrix} \begin{bmatrix} | & & | \\ \vec{v}_1 & \dots & \vec{v}_t \\ | & & | \end{bmatrix}^\top = U \Sigma V^\top, \quad (48)$$

defining U to be the horizontally stacked \vec{u}_i , V to be the horizontally stacked \vec{v}_i , and Σ to be the rectangular diagonal matrix of the σ_i .

The above decomposition is called the full **Singular Value Decomposition (SVD)** of A as it decomposes A into two orthonormal matrices and a diagonal matrix of singular values.

6 Applications to Planning

The singular value decomposition of A can be interpreted as follows - the first r columns of U form a basis for the column space of A and the last $t - r$ columns of V form a basis for the null space of A .

Moreover, they map between each other in a very clean manner, with $A\vec{v}_i = \sigma_i\vec{u}_i$ for $i \leq r$, and $A\vec{v}_i = \vec{0}$ for $i > r$. Let's see how we can use this property of the SVD can help us find the minimum cost control input.

We will again want to solve $\mathcal{C}_t\vec{w} = \vec{x}^*$ such that $\|\vec{w}\|$ is minimized. We have replaced the conventional \vec{u} inputs with \vec{w} to not confuse the notation. Let the SVD of $\mathcal{C}_t = U\Sigma V^T$ Recall that we showed that our control vector \vec{w} must have components only along $\vec{v}_1, \dots, \vec{v}_r$, in order to minimize its norm. Additionally since $\vec{u}_1, \dots, \vec{u}_r$ are a basis for $\text{Col}(A)$, then \vec{x}^* must be written as a linear combination of them. We can use these 2 properties to change coordinates of \vec{w} and \vec{x}^* to in terms of these basis vectors. As a reminder, due to orthonormality, the coefficient of the basis vector is just the inner product so

$$\mathcal{C}_t\vec{w} = \vec{x}^* \tag{49}$$

$$A \left(\sum_{i=1}^r \langle \vec{w}, \vec{v}_i \rangle \vec{v}_i \right) = \sum_{i=1}^r \langle \vec{x}^*, \vec{u}_i \rangle \vec{u}_i \tag{50}$$

$$\sum_{i=1}^r \langle \vec{w}, \vec{v}_i \rangle (A\vec{v}_i) = \sum_{i=1}^r \langle \vec{x}^*, \vec{u}_i \rangle \vec{u}_i \tag{51}$$

$$\sum_{i=1}^r \sigma_i \langle \vec{w}, \vec{v}_i \rangle \vec{u}_i = \sum_{i=1}^r \langle \vec{x}^*, \vec{u}_i \rangle \vec{u}_i, \tag{52}$$

so

$$\sigma_i \langle \vec{w}, \vec{v}_i \rangle = \langle \vec{x}^*, \vec{u}_i \rangle \implies \langle \vec{w}, \vec{v}_i \rangle = \frac{\langle \vec{x}^*, \vec{u}_i \rangle}{\sigma_i} \quad \forall i \leq r. \tag{53}$$

Substituting this back into our expression for \vec{w} , we obtain

$$\vec{w} = \sum_{i=1}^r \frac{\langle \vec{x}^*, \vec{u}_i \rangle}{\sigma_i} \vec{v}_i, \tag{54}$$

which is a clean expression for the minimum energy control to reach our desired state in terms of the SVD.

7 Outer Products

Now, we will look at a new interpretation of matrix multiplication, in order to construct an alternative way of writing the SVD in terms of what are known as *outer products*.

Recall that, for real vectors \vec{x} and \vec{y} expressed as columns with n components, their inner product is defined as $\vec{y}^T \vec{x}$, which yields a 1×1 matrix typically treated as a scalar.

Similarly, we will define their *outer product* to be $\vec{x}\vec{y}^T$. Let's see what this means. Let

$$\vec{x} = \begin{bmatrix} x_1 & x_2 & \cdots & x_m \end{bmatrix}^T$$

$$\vec{y} = \begin{bmatrix} y_1 & y_2 & \cdots & y_n \end{bmatrix}^T,$$

where it is possible that $m \neq n$. Then, by the definition of matrix multiplication,

$$\vec{x}\vec{y}^\top = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{bmatrix} \begin{bmatrix} y_1 & y_2 & \cdots & y_n \end{bmatrix} = \begin{bmatrix} x_1y_1 & x_1y_2 & \cdots & x_1y_n \\ x_2y_1 & x_2y_2 & \cdots & x_2y_n \\ \vdots & \vdots & \ddots & \vdots \\ x_my_1 & x_my_2 & \cdots & x_my_n \end{bmatrix}. \quad (55)$$

So while the inner product took two vectors of the same dimension and produced a scalar, the outer product takes two vectors of possibly different dimensions and yields a matrix!

Furthermore, notice that this matrix cannot be any arbitrary matrix - since each of its columns are a scalar multiple of \vec{x} , it cannot be of a rank greater than 1. It is straightforward to show that any matrix of rank 0 or 1 can be produced by an outer product of two vectors, but we will not discuss the details here.

Now, why are we interested in the outer product? Well, recall that we can express real matrix multiplication in terms of inner products. Specifically, we know that

$$\begin{bmatrix} - & \vec{x}_1^\top & - \\ & \cdots & \\ - & \vec{x}_m^\top & - \end{bmatrix} \begin{bmatrix} | & & | \\ \vec{y}_1 & \cdots & \vec{y}_n \\ | & & | \end{bmatrix} = \begin{bmatrix} \vec{x}_1^\top \vec{y}_1 & \vec{x}_1^\top \vec{y}_2 & \cdots & \vec{x}_1^\top \vec{y}_n \\ \vec{x}_2^\top \vec{y}_1 & \vec{x}_2^\top \vec{y}_2 & \cdots & \vec{x}_2^\top \vec{y}_n \\ \vdots & \vdots & \ddots & \vdots \\ \vec{x}_m^\top \vec{y}_1 & \vec{x}_m^\top \vec{y}_2 & \cdots & \vec{x}_m^\top \vec{y}_n \end{bmatrix}, \quad (56)$$

where the element at the i th row and j th column of the product of two matrices X and Y is the dot product of the i th row of X and the j th column of Y .

However, what if we were interested in the columns of X and the rows of Y instead? As it turns out, it is the case that

$$\begin{bmatrix} | & & | \\ \vec{x}_1 & \cdots & \vec{x}_n \\ | & & | \end{bmatrix} \begin{bmatrix} - & \vec{y}_1^\top & - \\ & \cdots & \\ - & \vec{y}_n^\top & - \end{bmatrix} = \vec{x}_1\vec{y}_1^\top + \vec{x}_2\vec{y}_2^\top + \cdots + \vec{x}_n\vec{y}_n^\top. \quad (57)$$

This result can be proved by applying the definition of matrix multiplication, but it is tedious and so will be omitted.

Instead, we will look at an example that demonstrates the main ideas behind the proof. Consider the product

$$\begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \begin{bmatrix} 5 & 6 \\ 7 & 8 \end{bmatrix}. \quad (58)$$

From our knowledge of the matrix product as the composition of dot products representing each element in the result, we obtain

$$\begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \begin{bmatrix} 5 & 6 \\ 7 & 8 \end{bmatrix} = \begin{bmatrix} 1 \cdot 5 + 2 \cdot 7 & 1 \cdot 6 + 2 \cdot 8 \\ 3 \cdot 5 + 4 \cdot 7 & 3 \cdot 6 + 4 \cdot 8 \end{bmatrix}. \quad (59)$$

We will not simplify this result, for reasons that will become more clear in a moment. Now, calculating the

product of these matrices using outer products, we obtain

$$\begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \begin{bmatrix} 5 & 6 \\ 7 & 8 \end{bmatrix} = \begin{bmatrix} 1 \\ 3 \end{bmatrix} \begin{bmatrix} 5 & 6 \end{bmatrix} + \begin{bmatrix} 2 \\ 4 \end{bmatrix} \begin{bmatrix} 7 & 8 \end{bmatrix} = \begin{bmatrix} 1 \cdot 5 & 1 \cdot 6 \\ 3 \cdot 5 & 3 \cdot 6 \end{bmatrix} + \begin{bmatrix} 2 \cdot 7 & 2 \cdot 8 \\ 4 \cdot 7 & 4 \cdot 8 \end{bmatrix}. \quad (60)$$

Notice that the terms in our matrix multiplication evaluated using dot products correspond exactly to the terms in our sums of outer products, so our outer product definition of matrix multiplication is consistent with our previous definitions, at least in this example. A slightly more rigorous form of this argument can be used to prove this result in general, but it is best to understand this result at an intuitive level.

8 Outer Product Form of the SVD

Now, we will use this new interpretation of matrix multiplication as a sum of outer products in order to obtain an alternative way of expressing the SVD. Recall that the SVD of a matrix A represented it as the product

$$A = U\Sigma V^T, \quad (61)$$

where Σ was a matrix with nonzero entries only along its main diagonal. Let A be an $m \times n$ matrix, so U is a square $m \times m$ matrix and V^T is a square $n \times n$ matrix. Additionally, without loss of generality, assume that $m \geq n$, so A is a “tall” matrix (the same results can be obtained if A is “wide” by considering A^T , which would become “tall”).

Let the columns of U be \vec{u}_1 through \vec{u}_m , the nonzero diagonal entries of Σ be σ_1 through σ_n (moving from the top-left to the bottom-right entry), and the rows of V^T be \vec{v}_1^T through \vec{v}_n^T . From our understanding of outer products, we know how to express UV^T in terms of the \vec{u}_i and \vec{v}_i . But how does Σ affect this interpretation?

Let’s consider just the first two terms of the SVD - the product $U\Sigma$. By the linear combination of columns interpretation of matrix multiplication, we have that

$$U\Sigma = \begin{bmatrix} | & & | \\ \vec{u}_1 & \cdots & \vec{u}_m \\ | & & | \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ 0 & \sigma_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \cdots & \sigma_n \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix} \quad (62)$$

$$= \begin{bmatrix} | & | & & | \\ \sigma_1 \vec{u}_1 & \sigma_2 \vec{u}_2 & \cdots & \sigma_n \vec{u}_n \\ | & | & & | \end{bmatrix}, \quad (63)$$

since the σ_i are coefficients for the first n columns of \vec{u} , with the subsequent columns vanishing entirely. Now, we can multiply by V and apply the outer-product interpretation of matrix multiplication to obtain

$$A = U\Sigma V^T \quad (64)$$

$$= \begin{bmatrix} | & | & & | \\ \sigma_1 \vec{u}_1 & \sigma_2 \vec{u}_2 & \cdots & \sigma_n \vec{u}_n \\ | & | & & | \end{bmatrix} \begin{bmatrix} - & \vec{v}_1^\top & - \\ - & \vec{v}_2^\top & - \\ & \cdots & \\ - & \vec{v}_n^\top & - \end{bmatrix} \quad (65)$$

$$= \sigma_1 \vec{u}_1 \vec{v}_1^\top + \sigma_2 \vec{u}_2 \vec{v}_2^\top + \cdots + \sigma_n \vec{u}_n \vec{v}_n^\top \quad (66)$$

$$= \sum_{i=1}^n \sigma_i \vec{u}_i \vec{v}_i^\top \quad (67)$$

so any $m \times n$ matrix A (where $m \geq n$) can be expressed as the weighted sum of n rank-1 matrices of the form $\sigma_i \vec{u}_i \vec{v}_i^\top$. This interpretation of the SVD is known as the *outer product form*.

Contributors:

- Ashwin Vangipuram.
- Druv Pai.
- Rahul Arya.
- Anant Sahai.