

Overview

In this note, we will examine a key application of the singular value decomposition (SVD), known as *principal component analysis* (PCA). We will explore its role in understanding large datasets with many dimensions, where PCA lets us determine the most important k -dimensional subspace for this data. To verify that the PCA approach is reasonable, we prove results involving *low-rank approximations* to matrices.

The key new idea here is that of approximation. When we collect data, there are potentially lots of real things that are influencing the exact values that we see. Our goal is often to figure out the most important of these. This allows us to compress or distill our data down to its important essence. How can we do this using the linear-algebraic tools that we have built?

1 Analyzing Matrix-type Data

Imagine we are conducting some sort of statistical study — say, we are interested in how a group of n students enjoy a set of m movies. Each student gives a numerical rating of each movie, indicating how much they enjoyed it. We place these ratings into a matrix R , where $R[i][j]$ is the rating the i th student gives to the j th movie. Our goal will be to draw conclusions from this dataset that enable us to make future predictions.

For instance, if we are told how much a new student enjoys a subset of our movies, we should be able to predict how much they will enjoy the other movies in our set. Alternatively, if we produce a new movie and are told how much a subset of our students enjoy it, we should be able to predict how much the other students in our study will enjoy it.

How on earth can we come up with these predictions? Aren't we all *unique and special*, so knowing how a student enjoys some movies provides no information on how they will enjoy the rest? In such a case, then yes, nothing could be done and our problem would be unsolvable.

However, in practice, it turns out that students are not as completely unique and special as all that! Typically, there is an underlying structure that lets us predict how much a student will enjoy each movie. For instance, imagine that every film has a certain intrinsic “goodness” $g[i]$, and every student has a “sensitivity”. Then the i th student's rating of the j th movie can be computed as the product of the student's sensitivity and the film's goodness, yielding

$$R[i][j] = s[i]g[j].$$

In such an ideal scenario, our ratings matrix would simply be

$$R = \begin{bmatrix} s[1] \\ s[2] \\ \vdots \\ s[n] \end{bmatrix} \begin{bmatrix} g[1] & g[2] & \cdots & g[m] \end{bmatrix} = \vec{s}\vec{g}^T.$$

To take a slightly more complex scenario, imagine for simplicity that movies can be described by a mix of three properties - their levels of “romance,” “comedy,” and “action”, represented by the vectors \vec{r} , \vec{c} , and \vec{a} . And imagine that each student computes how much they enjoy each movie by taking a linear combination of the “romance”, “comedy”, and “action” of each movie, where the coefficients of this linear combination vary with each student. Let these coefficients be stored in the vectors \vec{s}_r , \vec{s}_c , and \vec{s}_a respectively. We’ll call these coefficients the student’s “sensitivities” to each of the three genres.

In such a case, we have that

$$R[i][j] = s_r[i]r[j] + s_c[i]c[j] + s_a[i]a[j],$$

so the overall ratings matrix R can be expressed as

$$R = \vec{s}_r \vec{r}^T + \vec{s}_c \vec{c}^T + \vec{s}_a \vec{a}^T.$$

The point is that, despite being an $m \times n$ matrix, an idealized R according to this model can be represented as the sum of only a constant number of outer products of n - and m -dimensional vectors.

What could we do if we knew these vectors? Well, imagine that we were given a new student and their ratings on a subset of the movies, stored in the vector \vec{p} . What do we want to know about them? Well, we’ve said that their sensitivities to romance, comedy, and action fully describe their movie preferences. So if we can recover their sensitivities to these genres, then we can make predictions about their ratings of a future movie, so long as we know what genres this future movie is in.

And how can we recover these sensitivities? We know that we can write each provided rating

$$p[j] = w_r r[j] + w_c c[j] + w_a a[j],$$

where w_r , w_c , and w_a are the sensitivities of our new student. Stacking, we obtain the vector equation

$$\begin{bmatrix} r[1] & c[1] & a[1] \\ r[2] & c[2] & a[2] \\ \vdots & \vdots & \vdots \\ r[m] & c[m] & a[m] \end{bmatrix} \begin{bmatrix} w_r \\ w_c \\ w_a \end{bmatrix} = \begin{bmatrix} p[1] \\ p[2] \\ \vdots \\ p[m] \end{bmatrix}.$$

Using least squares, we can therefore recover the sensitivities w_r , w_c , and w_a , and use them to make future predictions. In a similar manner, if we are presented with a new movie and know how a subset of our students rate it, we can set up a least squares problem to solve for its romance, comedy, and action components, and then use these inferred components to make predictions for the rest of the students.

2 Analyzing Data using the SVD

What did we rely on when making our predictions? Fundamentally, we relied on being able to write our large ratings matrix R as a low-rank sum of outer products. In reality, however, there will always be noise and variation in our data. We need some way to construct a “low-rank” approximation of a large matrix that preserves the fundamental structure of our data.

Let’s see if the SVD helps. After all, given an $m \times n$ wide matrix A with singular values

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_m \geq 0,$$

we saw earlier that the SVD lets us write it as the sum of outer products

$$A = \sum_{i=1}^m \sigma_i \vec{u}_i \vec{v}_i^T.$$

In general, all m of our singular values will be nonzero. However, we claim that when we have some linear, low-rank essential structure to our data as described above, most of our singular values will be very small. For instance, imagine that σ_1 and σ_2 are significantly larger than the others. Then we can write A as

$$(\sigma_1 \vec{u}_1 \vec{v}_1^T + \sigma_2 \vec{u}_2 \vec{v}_2^T) + \sum_{i=3}^m \sigma_i \vec{u}_i \vec{v}_i^T,$$

where the contribution of the summation is very small in comparison to the first two terms. This motivates approximating our dataset as just

$$\hat{A} = \sigma_1 \vec{u}_1 \vec{v}_1^T + \sigma_2 \vec{u}_2 \vec{v}_2^T,$$

and using this model to make predictions in the manner described.

As it turns out, empirically, this approach works quite well, and is the heart of what is known as *principal component analysis*. Still, it would be nice to somehow quantifiably justify its efficacy. In particular, why is it that taking just the first few terms of the SVD provides us with the “best” low-rank approximation to our dataset?

If you already believe the result, you might want to jump to the end and read Section 5.

3 Low-rank Approximations

Let’s make our intuition about the SVD more formal. Given a matrix A , we will set a goal of computing a “low-rank” approximation \hat{A} that is of rank $\leq k$, but that is still “approximately equal” to A .

What does “approximately equal” mean? Well, we wish to minimize the magnitude of the error $A - \hat{A}$, where the squared magnitude of a matrix is defined to be the sum of the squares of all its elements. Why? Because we have no a-priori reason to care about some entries more than other entries of the matrix. It turns out that this quantity is called the *Frobenius norm*, and we define it as

$$\|A\|_F^2 = \sum_{i=1}^m \sum_{j=1}^n A[i][j]^2.$$

Why do we care about minimizing this quantity $\|A - \hat{A}\|_F^2$? One way to think about it is that our low-rank approximation is essentially a model that we are trying to fit to our data. For instance, in the above “movies” example, we are essentially trying to fit a model of student movie preferences to observed data. So it makes sense that we should try to pick a model that minimizes the error between the predicted and observed data points.

Our conjecture is that the best such approximation of an $m \times n$ matrix A with SVD

$$A = U\Sigma V^T = \begin{bmatrix} | & & & | \\ \vec{u}_1 & \cdots & \vec{u}_m & \\ | & & & | \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & \sigma_2 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_m & 0 & \cdots & 0 \end{bmatrix} \begin{bmatrix} - & \vec{v}_1^T & - \\ - & \vec{v}_2^T & - \\ \vdots & \vdots & \\ - & \vec{v}_m^T & - \end{bmatrix} = \sum_{i=1}^m \sigma_i \vec{u}_i \vec{v}_i^T.$$

is

$$\hat{A} = \sum_{i=1}^k \sigma_i \vec{u}_i \vec{v}_i^T,$$

where $\sigma_1 \geq \sigma_2 \geq \dots$. In other words, we sum up only the largest k outer products, rather than all m of them. How can we prove this result? First, we should take the time to establish some properties of the Frobenius norm.

3.1 The Frobenius Norm

Consider an $m \times n$ matrix M . As stated above, the Frobenius norm of this matrix is defined to be

$$\|M\|_F^2 = \sum_{i=1}^m \sum_{j=1}^n M[i][j]^2.$$

First, let's look at how the columns of M relate to this quantity. Let the columns of M be m -dimensional vectors \vec{v}_i such that

$$M = \begin{bmatrix} | & & | \\ \vec{v}_1 & \cdots & \vec{v}_n \\ | & & | \end{bmatrix}.$$

Observe that the squared norm of a particular vector is

$$\|\vec{v}_j\|^2 = \sum_{i=1}^m \vec{v}_j[i]^2.$$

In other words, it is the sum of squares of all the elements in the j th column. Summing this quantity over all the columns, we obtain the sum of squares of all the elements in M , which is the squared Frobenius norm. Expressed algebraically, it is the case that

$$\|M\|_F^2 = \sum_{i=1}^m \sum_{j=1}^n M[i][j]^2 = \sum_{j=1}^n \left(\sum_{i=1}^m M[i][j]^2 \right) = \sum_{j=1}^n \|\vec{v}_j\|^2.$$

A similar result can be derived in an exactly analogous manner for the rows of M .

Now, we will look at how the Frobenius norm of a matrix varies under an orthonormal change of basis. Let

Q be a matrix with m unit, orthogonal columns, such that $Q^T Q = I$. Observe that, using the above result,

$$\begin{aligned} \|QM\|_F^2 &= \left\| \begin{bmatrix} | & & | \\ Q\vec{v}_1 & \cdots & Q\vec{v}_n \\ | & & | \end{bmatrix} \right\|_F^2 \\ &= \sum_{j=1}^n \|Q\vec{v}_j\|^2 \\ &= \sum_{j=1}^n \vec{v}_j^T Q^T Q \vec{v}_j \\ &= \sum_{j=1}^n \vec{v}_j^T \vec{v}_j \\ &= \sum_{j=1}^n \|\vec{v}_j\|^2 \\ &= \|M\|_F^2, \end{aligned}$$

so the Frobenius norm of a matrix remains unchanged under pre-multiplication by Q . An intuitive way to interpret the above algebraic result is to notice that pre-multiplying by Q doesn't change the norm of any of the columns of M , so its Frobenius norm should similarly remain unchanged. It is interesting to note that here, we don't actually need Q to be square — that never came up in the above derivation. We just need orthonormal columns.

In a very similar manner, it can be shown that post-multiplication by a matrix Q^T with unit, orthogonal rows again does not change the Frobenius norm. An easy way of seeing this is to simply take transposes everywhere in the above calculation, recognizing that the Frobenius norm of a matrix does not change if we take the transpose.

Now, combining the above observations with our knowledge of the SVD, we find that, given a matrix A with SVD $A = U\Sigma V^T$,

$$\begin{aligned} \|A\|_F &= \|U\Sigma V^T\|_F \\ &= \|U^T U \Sigma V^T V\|_F \\ &= \|\Sigma\|_F, \end{aligned}$$

since U and V are both square¹ orthonormal matrices, so pre-multiplying by U^T and post-multiplying by V both do not change the Frobenius norm, as in both cases their rows and columns are all unit and orthogonal. Writing out the Frobenius norm of Σ in terms of the singular values explicitly, we find that

$$\|A\|_F = \sqrt{\sum_{i=1}^m \sigma_i^2}.$$

This seems like a useful result that we should hang on to.

¹Here, we are using the full form of the SVD, not the compact form.

3.2 Reducing to an easier case via the SVD

Now, let's get back to the question of low-rank approximations. As stated before, we wish to choose a rank at most k matrix \hat{A} that minimizes $\|A - \hat{A}\|_F$. Looking at the SVD of A and applying simplifications very similar to those derived in the previous section, pre-multiplying by U^T and post-multiplying by V , we see that

$$\begin{aligned} \|A - \hat{A}\|_F &= \|U\Sigma V^T - \hat{A}\|_F \\ &= \|U^T U \Sigma V^T V - U^T \hat{A} V\|_F \\ &= \|\Sigma - U^T \hat{A} V\|_F. \end{aligned}$$

Let $X = U^T \hat{A} V$, for convenience. Observe that, since U and V are both invertible, given any X we can choose a corresponding \hat{A} , and vice-versa. Furthermore, it is clear that the rank of \hat{A} and X are the same, so X has rank no more than k . Thus, the problem reduces to choosing an X of rank no more than k that minimizes

$$\|\Sigma - X\|_F.$$

Essentially, the SVD tells us that all we really need to understand is the diagonal case. How can we best approximate a diagonal matrix with sorted entries down the diagonal using a rank at most k matrix?

Let's write out Σ and X explicitly, to try and get a better idea of what remains for us to prove. We wish to choose columns \vec{x}_i that minimize

$$\left\| \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & \sigma_2 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_m & 0 & \cdots & 0 \end{bmatrix} - \begin{bmatrix} | & & | & | & & | \\ \vec{x}_1 & \cdots & \vec{x}_m & \vec{x}_{m+1} & \cdots & \vec{x}_n \\ | & & | & | & & | \end{bmatrix} \right\|_F^2$$

Expressing the squared Frobenius norm as the sum of the squared norms of the columns, we see that we are tasked with minimizing

$$\|\sigma_1 \vec{e}_1 - \vec{x}_1\|^2 + \cdots + \|\sigma_m \vec{e}_m - \vec{x}_m\|^2 + \|\vec{0} - \vec{x}_{m+1}\|^2 + \cdots + \|\vec{0} - \vec{x}_n\|^2,$$

while keeping the rank of X less than or equal to some given constant k . Note that we use the notation \vec{e}_i here to represent the i th column of the identity matrix, working in \mathbb{R}_m throughout. It should be clear here that

$$\vec{x}_{m+1} = \vec{x}_{m+2} = \cdots = \vec{x}_n = \vec{0},$$

since making them nonzero would only hurt us by increasing the quantity that we are trying to minimize.

But what about the $\vec{x}_1, \dots, \vec{x}_m$? Because of the constraint on the rank of X , we can't make them all equal to the corresponding columns of Σ , as then the rank of X would be m , and not k .

There is a natural guess that seems somewhat obvious: just pick $\vec{x}_1 = \sigma_1 \vec{e}_1, \vec{x}_2 = \sigma_2 \vec{e}_2, \dots, \vec{x}_k = \sigma_k \vec{e}_k$ and set the rest of the $\vec{x}_i = \vec{0}$ for $i > k$. This definitely makes X have rank k and it does turn out to be the right answer, but is it really obvious? How can we be sure that there isn't a better choice out there?

This actually requires proof.

4 The Projection Perspective

In the previous section, we used the SVD to distill our problem down to the diagonal case. How can we choose an X of rank no more than k that minimizes

$$\|\Sigma - X\|_F.$$

We saw that what matters here is choosing the m vectors $\vec{x}_1, \vec{x}_2, \dots, \vec{x}_m$ that are matched up with the non-zero parts of Σ in the above subtraction, since the rest of X should clearly be all zeros. We have a guess as to what the optimal choice should be, but we need to prove that it is indeed optimal.

The projection perspective approaches this problem by stepping back a bit. Dealing with the constraint on the rank of X is annoying. What does this condition really mean from a geometric point of view? It means that all these m vectors $\vec{x}_1, \vec{x}_2, \dots, \vec{x}_m$ must lie in a k -dimensional subspace. This is what rank means.

Let's reason about the problem in terms of a basis for this subspace. Imagine that we are already given this subspace as the span of some $m \times k$ matrix Q . What kind of matrix should we ask for? We know that orthonormal matrices are easy to work with, and any other matrix with linearly independent columns can be orthonormalized, so for convenience, let's assume that all the columns of Q have been orthonormalized, so $Q^T Q = I$. Then what should each of the \vec{x}_i be, assuming they are constrained to lie in the column space of Q ?

Well, once we're given Q , we simply need to choose each \vec{x}_i to minimize the term it is involved in — specifically, each \vec{x}_i should minimize

$$\|\sigma_i \vec{e}_i - \vec{x}_i\|^2.$$

And how can we choose such an \vec{x}_i ? We know this from 16A! It is simply the projection of $\sigma_i \vec{e}_i$ onto the column space of Q , which from least squares is just

$$\vec{x}_i = Q(Q^T Q)^{-1} Q^T (\sigma_i \vec{e}_i) = Q Q^T (\sigma_i \vec{e}_i) = \sigma_i Q Q^T \vec{e}_i,$$

simplifying by recalling that $Q^T Q = I$, as we chose the columns of Q to be an orthonormal basis.

This means that finding the right orthonormal Q is enough to solve our problem. In hindsight, this is not surprising since our initial motivation was to discover a subspace that best captures the data.

4.1 The finding Q problem

So our problem has been simplified further. Rather than trying to find an arbitrary $m \times n$ matrix X with rank- k , we can instead find an $m \times k$ orthonormal matrix Q , and then compute \vec{x}_i by projecting each of the columns of Σ onto the column space of Q .

Since we're now working with Q , rather than X , it makes sense to write the quantity we're interested in minimizing in terms of Q . Substituting in our least-squares solution for each of the \vec{x}_i into our earlier expression for the squared Frobenius norm, and setting all the $\vec{x}_i = \vec{0}$ for all $i > m$ (as we argued would be optimal), we find that the new quantity to minimize becomes

$$\begin{aligned} \|\sigma_1 \vec{e}_1 - \vec{x}_1\|^2 + \dots + \|\sigma_m \vec{e}_m - \vec{x}_m\|^2 + \|\vec{0} - \vec{x}_{m+1}\|^2 + \dots + \|\vec{0} - \vec{x}_n\|^2 &= \sum_{i=1}^m \|\sigma_i \vec{e}_i - \sigma_i Q Q^T \vec{e}_i\|^2 + \sum_{i=m+1}^n \|\vec{0} - \vec{0}\|^2 \\ &= \sum_{i=1}^m \sigma_i^2 \|\vec{e}_i - Q Q^T \vec{e}_i\|^2. \end{aligned}$$

Can we simplify this expression somewhat? Notice that $\vec{e}_i - QQ^T \vec{e}_i \perp QQ^T \vec{e}_i$, since we know from 16A that the least squares projection is orthogonal to the residual error vector. Thus, by the Pythagorean theorem, we can write

$$\|\vec{e}_i\|^2 = \|\vec{e}_i - QQ^T \vec{e}_i\|^2 + \|QQ^T \vec{e}_i\|^2 \implies \|\vec{e}_i - QQ^T \vec{e}_i\|^2 = \|\vec{e}_i\|^2 - \|QQ^T \vec{e}_i\|^2.$$

Making this substitution in our earlier expression, noticing that $\|\vec{e}_i\| = 1$ and rearranging, our problem reduces to finding an orthonormal Q that minimizes the quantity

$$\sum_{i=1}^m \sigma_i^2 (\|\vec{e}_i\|^2 - \|QQ^T \vec{e}_i\|^2) = \sum_{i=1}^m \sigma_i^2 - \sum_{i=1}^m \sigma_i^2 \|QQ^T \vec{e}_i\|^2.$$

Notice that we have broken the quantity we are minimizing into two summations. The first summation is simply the squared Frobenius norm of Σ , which clearly does not depend on Q . The second is the negation of the squared Frobenius norm of $QQ^T \Sigma$. Thus, in order to minimize the overall quantity, we should aim to choose our Q to *maximize* $\|QQ^T \Sigma\|_F^2$, since we are subtracting this quantity from another fixed value.

4.2 Working with this even simpler problem

Cool! We've managed to simplify our problem further, by taking advantage of properties of projections and orthonormal bases. Recalling earlier properties of the Frobenius norm, since $Q^T Q = I$, we see immediately that the quantity we are trying to maximize can be slightly simplified as

$$\|QQ^T \Sigma\|_F^2 = \|Q^T \Sigma\|_F^2. \quad (1)$$

Be aware that, since Q is not square, $QQ^T \neq I$, as the rows are not necessarily mutually orthogonal. So we can't use this property to make any further simplifications. However, it is nice to see that (1) feels very reminiscent of the maximizing correlation goal we've seen in 16A.

How can we show that our guessed solution (just choose the first k columns of the identity for Q) is indeed the best choice for maximizing (1)? That choice would give $\|Q^T \Sigma\|_F^2 = \sum_{j=1}^k \sigma_j^2$. One way to show that this is optimal is to show that we cannot possibly do any better than that.

To do this, let's just expand this product out directly. For notational convenience, let the columns of Q be \vec{q}_i , and so

$$Q^T = \begin{bmatrix} - & \vec{q}_1^T & - \\ - & \vec{q}_2^T & - \\ & \vdots & \\ - & \vec{q}_k^T & - \end{bmatrix}.$$

Remember that each of the \vec{q}_i are m -dimensional vectors, and that they are of unit norm and mutually orthogonal.

Now, writing out our product, we obtain

$$\begin{aligned} \|Q^T \Sigma\|_F^2 &= \left\| \begin{bmatrix} - & \vec{q}_1 & - \\ - & \vec{q}_2 & - \\ & \vdots & \\ - & \vec{q}_k & - \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & \sigma_2 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_m & 0 & \cdots & 0 \end{bmatrix} \right\|_F^2 \\ &= \sum_{i=1}^k \sum_{j=1}^m \vec{q}_i[j]^2 \sigma_j^2 \\ &= \sum_{j=1}^m \sigma_j^2 \sum_{i=1}^k \vec{q}_i[j]^2. \end{aligned}$$

Thus, we have expressed our desired Frobenius norm as a linear combination of the squared singular values σ_j^2 . For convenience and to focus our attention on what is going to matter, let the coefficients of our linear combination be denoted as

$$d_j = \sum_{i=1}^k \vec{q}_i[j]^2, \tag{2}$$

so we can express the term we want to maximize as

$$\|Q^T \Sigma\|_F^2 = \sum_{j=1}^m \sigma_j^2 d_j. \tag{3}$$

Now, observe that the Frobenius norm of Q^T itself is clearly $\|Q\|_F = k$, since Q has k unit-norm columns. Expressing this quantity in terms of the d_j , we see that

$$\sum_{j=1}^m d_j = \|Q^T\|_F^2 = \|Q\|_F^2 = k. \tag{4}$$

Furthermore, notice that, since they are defined as a sum of squares in (2), each of the $d_j \geq 0$. Clearly, as $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_m$, we could try to upper-bound the quantity we are trying to maximize by choosing $d_1 = k$ and setting the other $d_j = 0$.

That would give $k\sigma_1^2$ which is bigger than our guessed optimal solution. But is this extreme choice of d_j actually achievable by some Q ? The challenge is that we know a lot about the columns of Q since they are orthonormal. But d_j is related to the j -th row of Q . How can we understand anything about the rows of Q ?

If Q were square, then the fact that $Q^T Q = I$ would tell us that Q^T is the inverse of Q and so Q^T would also be orthonormal. But Q is decidedly not square, so what can we do?

Observe that the k columns of Q can be extended using Gram-Schmidt to form a full orthonormal basis for \mathbb{R}_m . Let this extended basis be

$$\widehat{Q} = \begin{bmatrix} | & & | & | & & | \\ \vec{q}_1 & \cdots & \vec{q}_k & \vec{q}_{k+1} & \cdots & \vec{q}_m \\ | & & | & | & & | \end{bmatrix}.$$

By the properties of orthonormal square matrices, we know that in addition to the columns, the rows of \widehat{Q}

also form an orthonormal basis, and so are all of unit norm. Thus, looking at the j th row, we can write

$$\sum_{i=1}^m \tilde{q}_i[j]^2 = 1$$

for all j .

This tells us that the rows of \widehat{Q} can't be too large, and thus, truncating the sum by removing the terms corresponding to the $k+1$ th column onwards, we see that

$$d_j = \sum_{i=1}^k \tilde{q}_i[j]^2 \leq \sum_{i=1}^m \tilde{q}_i[j]^2 = 1. \quad (5)$$

So in addition to requiring that they sum to k , we have established that each individual $0 \leq d_j \leq 1$ for all j .

Now, we have found a simpler problem inside our original problem.

4.3 Working with the inner problem

At this point, we have turned even the simplified problem into something² only involving numbers. From (3), we know we want to maximize $\sum_{j=1}^m \sigma_j^2 d_j$ over the choice of d_j that we know must satisfy certain constraints. From (4), we know their overall sum $\sum_{j=1}^m d_j = k$. We also from (5), that $0 \leq d_j \leq 1$ for all j .

A consequence of this is that our previous overly greedy assignment of values to d_j is not possible if $k > 1$, since we must have $d_1 \leq 1 < k$. Instead, if we wish to maximize our linear combination of the σ_j^2 , it is clear that we should set

$$d_1 = d_2 = \dots = d_k = 1$$

and

$$d_{k+1} = d_{k+2} = \dots = d_m = 0,$$

in order to place as much weight as possible on the largest squared singular values without violating the constraints on the d_j . This is very intuitive and gives us $\sum_{j=1}^k \sigma_j^2$ as the best we can do.

Why is this intuitive? We can think of each d_j as telling us how much money we want to spend to buy product j . The store only has up to 1 liter of each product in stock. Each item costs one dollar per liter. The product j contains σ_j^2 grams of gold per liter. If we want to get as much gold as possible, what should we do if we have k dollars to spend? The answer is clearly buy out all the stock of the first k products since they all cost the same, but the first k products have more gold per liter than any of the other products.

The rigorous proof of this claim is omitted here in this note, but only because you are walked through this exact proof on the homework. (It follows from a straightforward argument that doing anything else can't be optimal since anything else could be improved by moving it closer to this solution.)

But is *this* assignment of d_j above actually achievable in our original problem of interest? It certainly satisfies all of our inequalities, but that is not sufficient to show it is actually possible — perhaps we could have derived another inequality that makes this assignment impossible. After all, the d_j are defined in (2) from the underlying Q matrix. Maybe no Q matrix exists that satisfies our desired assignments of d_j ?

Consequently, the best way to show that it is in fact feasible is to present such a Q . We want the first k columns of Q^T to all be of unit norm, and the remaining columns to all be 0. Thus, an obvious choice would

²You will see in later courses like 127 and 170 that such problems are actually called “linear programs” and are extremely useful. They arise naturally in lots of settings. In 127 and 170, you will learn ways of thinking about them that are generally useful. In our particularly simple case here, we can reason about the problem directly.

be to write

$$Q^T = \begin{bmatrix} 1 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 & 0 & \dots & 0 \end{bmatrix}.$$

It is straightforward to verify that the rows of Q^T form an orthonormal basis of a k -dimensional subspace, so this choice of Q is valid, and so it is optimal.

So what low-rank approximation does this optimal choice of Q correspond to? Substituting back, we see that

$$\begin{aligned} \hat{A} &= UXV^T \\ &= UQQ^T\Sigma V^T. \end{aligned}$$

Now, observe that

$$QQ^T = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 & 0 & \dots & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 & 0 & \dots & 0 \end{bmatrix}.$$

In words, QQ^T is a diagonal matrix where the first k entries along the diagonal are 1, and the remainder are 0. Thus, we see that

$$\hat{A} = U \begin{bmatrix} \sigma_1 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & \sigma_2 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_k & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 & 0 & \dots & 0 \end{bmatrix} V^T = \sum_{i=1}^k \sigma_i \vec{u}_i \vec{v}_i^T,$$

as expected. This completes the proof.

The entire proof here is not something that we would expect you to be able to come up with on your own at the level of 16B. However, it is elementary enough that you should be able to follow the steps, and the goal is to help you begin to feel how each part follows naturally.

5 Using the PCA idea

The above theorem tells us that if we want to find the best k -dimensional subspace to use to approximate a set of (real) data, we can take the data, arrange it into a matrix M by stacking the columns of data, and then taking the SVD. Looking at the SVD in outer product form $M = \sum_{i=1} \sigma_i \vec{u}_i \vec{v}_i^T$, where the σ_i are in decreasing order, we can just choose to truncate the sum to the top k terms. We get an approximation of all the data

$\hat{M} = \sum_{i=1}^k \sigma_i \vec{u}_i \vec{v}_i^T$, but more importantly, we also get a basis for our subspace of interest. Namely, the \vec{u}_i give us such an orthonormal basis. The first such \vec{u}_1 is called the first principal component (for columns), the second such \vec{u}_2 is called the second principal component (for columns), and so on.

If we had arranged the data into rows instead, then we could do the same thing and we would use the first \vec{w}_i instead as the principal components. (Or if we want to keep them in row form, \vec{w}_i^T .) The choice depends on the problem context. What is the nature of the data and what are we trying to do.

Once we have learned such a subspace from training data, we can use it for dimensionality reduction when dealing with new raw data points. We just project the new raw data points onto the subspace. Although our goal is for the projections to be close to the original raw data points, what we work with are the coefficients of the projections in the subspace basis that we have found. For example, if the raw data points have $m = 100$ entries in them, and we decided to use $k = 3$ to choose a 3 dimensional subspace, then we have managed to distill the essential information into just 3 coefficients. These are also easy to calculate by the orthonormality

of the \vec{u}_i . For the case of $k = 3$, we just keep $\begin{bmatrix} \vec{u}_1^T \vec{x} \\ \vec{u}_2^T \vec{x} \\ \vec{u}_3^T \vec{x} \end{bmatrix} = \begin{bmatrix} \vec{u}_1^T \\ \vec{u}_2^T \\ \vec{u}_3^T \end{bmatrix} \vec{x}$ instead of the whole m -dimensional raw data point \vec{x} .

Such dimensionality reduction is often very useful in reducing the computational burden³ for doing classification, etc.

In some problem contexts, the data that we are given has DC offsets⁴ that we don't want to consider as directions of variation. When that happens, those DC offsets are often estimated and removed before one does PCA analysis. Such offsets can be found by taking appropriate means/averages of the training data, and then subtracting them off. However, it is useful to consider that "de-meaning" step as being something separate from the core task of subspace discovery. Instead, it should be considered an application-specific data cleaning or preprocessing step. It should only be done if you actually believe that such irrelevant DC offsets might exist.

Contributors:

- Rahul Arya.
- Anant Sahai.

³For reasons that you will learn in future courses like 189, the advantage is not simply computational. Distilling the data down by dimensionality reduction can also be helpful in eliminating irrelevant noise that could confuse subsequent stages of learning. Even more subtly, it can help make problems tractable that otherwise might seem impossible. For example, if the data that is being recorded is a microphone trace that is hundreds of samples long, then many approaches to learning how to classify words might require thousands of examples of each word. By reducing the dimensionality of the incoming microphone trace, it can become possible to learn from fewer examples using those approaches. All of this is not in scope for 16B.

⁴What do we mean? Consider two-dimensional data. Our approach to PCA will find the best line through the origin that goes through the data. But what if the data was actually along a line that didn't go through the origin? How would we find this? In that case, we would want to "center" the data first. The average of all the data points that actually lie along a straight line will also lie on that same straight line. Consequently, the same will also be approximately true if the points are approximately on a line.