

Lecture 15: Low Rank Approximation and the Singular Value Decomposition

Lecturer: *Sahil Singla*

Last updated: October 22, 2020

This unit is about finding compressed data representations – e.g. compressed representations of vectors $a_1, a_2, \dots, a_n \in \mathbb{R}^d$. The techniques we have discussed so far (Johnson-Lindenstrauss sketching and hashing for similarity search) are *oblivious* to any structure in the dataset. They compress each point a_i without looking at any other points. In some ways, this obliviousness is a strength. It makes the methods fast and, as we saw in the previous two lectures, we can obtain strong guarantees without making assumptions about our data.

At the same time, obliviousness can be a weakness. Oblivious dimensionality reduction methods don't take advantage of structure which might make better compression possible. Moreover, finding such structure might be interesting in its own right.

1 Low-rank structure

Today we will discuss a type of structure that can 1) allow for better dimensionality reduction, 2) lead to very interesting scientific discoveries and insights about data, and 3) is remarkably common across diverse application areas.

In particular, we are interested in datasets where most of our vectors a_1, \dots, a_n can be well approximated as a linear combination of a small ground set of vectors in \mathbb{R}^d , $\{b_1, \dots, b_k\}$. I.e. for some set of k coefficients $\{C_{1i}, C_{2i}, \dots, C_{ki}\}$ we approximate a_i by:

$$a_i \approx \sum_{j=1}^k C_{ji} b_j.$$

Let $A \in \mathbb{R}^{d \times n}$ be a *data matrix* which contains each a_i as a column. If A is a rank k matrix (i.e. a “low-rank” matrix), then it would be possible to find some ground set and coefficients so that, for all i , this is actually an equality: $a_i = \sum_{j=1}^k C_{ji} b_j$. Geometrically, this would mean that all points lie on a low dimensional hyperplane in \mathbb{R}^d (see Figure 1)




Figure 1: If a_1, \dots, a_n lie on a low dimensional hyperplane, then $A = [a_1, \dots, a_n]$ is exactly low rank, so it can be written as the product of two matrices, $B \in \mathbb{R}^{d \times k}$ and $C \in \mathbb{R}^{k \times n}$.

We can view the columns of $C = [c_1, \dots, c_n]$, which are each in \mathbb{R}^k , as dimensionality reductions of a_1, \dots, a_n . In the case when A is low-rank, consider choosing b_1, \dots, b_k to be an orthogonal span for the hyperplane containing a_1, \dots, a_n . Then each c_i is simply a_i represented in a coordinate basis for the hyperplane. So, for example, $\|c_i\|_2 = \|a_i\|_2$ for all i , $\|c_i - c_j\|_2 = \|a_i - a_j\|_2$ for all i, j , and in general, $[c_1, \dots, c_n]$ captures all the geometric information about our original dataset.

Low-rank approximation

Now of course, it's very rare to find matrices that are *actually* low rank. On the other hand, it's very common to find matrices that are *approximately* low rank. I.e. where there exists some set of ground set of vectors $\{b_1, \dots, b_k\}$ and some coefficients C_{ij} so that, for example, the following error measure is small:

$$\sum_{i=1}^n \left\| a_i - \sum_{j=1}^k c_{ij} b_j \right\|_2^2 = \|A - BC\|_F^2. \quad (1)$$

Here $\|\cdot\|_F$ denotes the Frobenius norm of a matrix (i.e. the square root of its sum of squared entries):

$$\|M\|_F = \sqrt{\sum_{i,j} M_{ij}^2}.$$

Using the Frobenius norm gives a convenient way of using matrices to express our error measure. There are other ways to measure if a matrix is “close to a low-rank matrix”, but Frobenius norm distance is a popular measure and the one we will focus on today.

Computation

If A were exactly low-rank, BC could be found using any standard orthogonalization procedure for finding a span for A 's columns. On the other, when A is only approximately low-rank, we ideally want to solve:

$$\min_{B \in \mathbb{R}^{d \times k}, C \in \mathbb{R}^{k \times n}} \|A - BC\|_F^2. \quad (2)$$

This is a non-linear, non-convex optimization problem, but surprisingly it can be solved very efficiently (in low polynomial time). We will discuss one particular method for doing so in this lecture.

2 Example Applications

Before getting into algorithms, let's see a few places where low-rank approximation is important in practice.

Latent Semantic Analysis (LSA)

Let's return to the bag-of-words model discussed in previous lectures. Here each a_i represents the i^{th} document in a database of n documents. The vector has an entry for all possible words in a given language. At that location, it contains a count for the number of times that word appeared in document i .

$A = [a_1, \dots, a_n]$ is called the “term-document” matrix – each row corresponds to a term and each column a document. Term document matrices tend to be close to low-rank, and algorithms for low-rank approximation can recover that structure. In particular, we can find an approximation BC for the term-document matrix where B has k columns $b_1, \dots, b_k \in \mathbb{R}^d$ (here d is the number of words in our language) and $\|A - BC\|_F^2$ is small.

There are a number of justifications for why this might be the case. For example, one simple generative model for documents would be to simply look at all of the words in a language, look at the empirical frequency with which those words occur (e.g. “the” is used much more frequently than “president”) and then assume that a typical document is generated by drawing words at random according to their relative frequencies. Of course this isn't a very good model, but we can think about how to make it better.

For example, we might notice that documents about politics tend to use the word “president” much more frequently than those about sports. Moreover, political documents written in the US tend to use the word “president” more than political articles written in the UK. Similarly, sports articles written in the US tend to use “touchdown” more than sports articles written in the UK. We can construct a better model by assuming a document can be assigned to a mix of potentially overlapping categories – for example, a document about US environmental policy might have categories “politics”, “for US audience”, and “environment”. We could predict the words contained in that document by looking at the global distribution of words used in the “politics”, “for US audience”, and “environment” categories, and drawing words at random according to a mixture of those distributions.

This model turns out to be very powerful in predicting the frequency of words in a document. Since it does so in a linear way (drawing from a linear combination of word distributions), we expect that A will have a good k -rank approximation. Instead of hand choosing categories, columns in B can be seen as representing a set of “optimal” categories.

This is the idea behind what's known as “latent semantic analysis” in natural language processing. Each column $c_i \in C$ is viewed as an indicator vector for the presence of different categories in document a_i . As a “semantic representation” for a document, c_i 's can be used to compare or cluster documents. For a more in-depth discussion, check out Chapter 18 in [2] (freely available online) or [3].

Word embeddings

If columns in C give some sort of meaningful representation of words, what about rows in B ? We have one row for each word in our language. These rows are sometimes called “word embeddings” – they can give very powerful semantic representations of words. Words that tend to appear in similar categories of documents (e.g. in a similar context) tend to have similar word embedding vectors. These embeddings can be used for finding similar words or synonyms, but also for tasks like solving analogy problems.

There are many ways to generate better word embeddings. For example, instead of looking at term-document co-occurrence, it's more common to look at more local measures, like term-sentence co-occurrence. Current algorithms are also based on more complex models than the simple linear model discussed above (see e.g. <https://nlp.stanford.edu/projects/glove/> or <https://code.google.com/archive/p/word2vec/>), but it all started with basic LSA and low-rank approximation!

Visualizing and understanding genetic data

Genetic data tends to be low-rank. Consider a vector a_i for each individual in a population that holds a numerical representation of that individual's genome (e.g. a bit vector, with every two bits representing the expression of a single nucleotide, which can take values A, T, G, or C.). Why might this be the case?

At a course level genetics are controlled largely by ancestry – historically isolated populations (geographically, culturally, etc.) tend to have very similar genomes, with the exception of a relatively small number of genes that distinguish individuals. Accordingly, if we let our set of ground vectors b_1, \dots, b_k contain representative individuals from different ancestral populations, we can do a pretty good job reconstructing every a_i vector up to small error. If we take the best low-rank approximation, we can do even better.

One particularly dramatic exposition of the low-rank natural of genetic data is given in [1]. After a few basic data transformations (mean centering, removing outliers, etc.) they took a rank-2 approximation of a genetic data set from populations in Europe. This produced a set of two dimensional vectors, c_1, \dots, c_n for each individual. When plotting these points on a two-dimensional grid, the location of each point roughly reflects the ancestral origin of each individual!

In other words, each a_i is well represented by a two dimensional linear model where each dimension represents the East/West and North/South coordinates of where that individual is from. In fact, since it was derived from an optimal rank-2 approximation, this is the best two dimensional linear model for reconstructing A , indicating how important geography is in genetic variation. I would encourage you to check out the paper (e.g. at https://www.researchgate.net/publication/23469983_Genes_Mirror_Geography_within_Europe).

3 The Singular Value Decomposition

It turns out that optimal low-rank approximations can be computed using what's known as the “singular value decomposition”.

Theorem 1 (Singular Value Decomposition (SVD)). *Consider $A \in \mathbb{R}^{d \times n}$ and let $r = \min(d, n)$. A can always be written as the product of three matrices, $A = U\Sigma V^T$, where:*

- $U \in \mathbb{R}^{d \times r}$ is a matrix with orthonormal columns,
- $\Sigma = \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_r \end{bmatrix}$ is a non-negative diagonal matrix with entries $\sigma_1 \geq \dots \geq \sigma_r \geq 0$,

- $V \in \mathbb{R}^{n \times r}$ is a matrix with orthonormal columns.

U 's columns are called the “left singular vectors” of A , V 's columns are its “right singular vectors”, and $\sigma_1, \dots, \sigma_r$ are its “singular values”. When the SVD is computed after mean centering A 's columns or rows, the singular vectors are sometimes call “principal components”.

We will prove Theorem 1 in Section 4, and will first try to develop some intuition for it. The singular value decomposition is closely related to eigendecomposition: $U\Sigma^2U^T$ is the eigendecomposition of AA^T and $V\Sigma^2V^T$ is the eigendecomposition of A^TA . Using eigendecomposition algorithms (e.g., the QR algorithm), it can be computed in approximately $O(nd^2)$ time (assuming $d < n$).

The existence of the SVD may be surprising – it says that no matter what A looks like, it can be represented at the product of 3 simple matrices – two orthogonal spans and a diagonal scaling. The SVD has countless applications in linear algebra, but one of the most useful is that it can be used to read-off the optimal k -rank approximation for A , for any k .

Claim 2 (Truncated SVD). *For any $k \in 1, \dots, \min(n, d)$, let $U_k \in \mathbb{R}^{d \times k}$ contain the first k columns of U , let $V_k \in \mathbb{R}^{n \times k}$ contain the first k columns of V , and let Σ_k be a $k \times k$ diagonal matrix containing A 's first k singular values. Then:*

$$\|A - U_k \Sigma_k V_k^T\|_F^2 = \min_{B \in \mathbb{R}^{d \times k}, C \in \mathbb{R}^{k \times n}} \|A - BC\|_F^2.$$

In other words, there is no better rank k approximation for A than $U_k \Sigma_k V_k^T$.

Note that, a solution to our original low-rank approximation problem, (2), can be obtained either by setting $B = U_k \Sigma_k$ and $C = V_k^T$, or by setting $B = U_k$ and $C = \Sigma_k V_k^T$ – the product BC is the same.

One thing surprising about Claim 2 is that it implies that we can find a basis set for an optimal k rank approximation in a *greedy* way. The best ground set of a rank- k approximation, u_1, \dots, u_k , just adds one vector to the best basis set for a rank- $(k - 1)$ approximation, u_1, \dots, u_{k-1} .

I'm going to give a proof of Claim 2, but if you have already seen this before in a linear algebra class, feel free to skip it. Or try to reprove it on your own.

Proof. Case: $k = 1$.

This statement is easiest to prove for $k = 1$. For rank 1 approximation our goal is to choose $b \in \mathbb{R}^d$ and $c \in \mathbb{R}^n$ to minimize:

$$\|A - bc^T\|_F^2 = \sum_{i=1}^n \|a_i - c_i \cdot b\|_2^2,$$

where c_i is the i^{th} entry of c (and a_i is the i^{th} column of A). Without loss of generality, we may assume that b is a unit vector. For a given b , it's clear that we should choose c_i so that $c_i \cdot b$ is the projection of a_i onto b . I.e. we should set $c_i = \langle a_i, b \rangle$ (see Figure 2 for the geometric intuition). Equivalently, we should set:

$$c = A^T b. \tag{3}$$

So, our rank 1 optimization problem actually reduces to:

$$\min_{b \in \mathbb{R}^n, \|b\|_2=1} \|A - bb^T A\|_F^2. \quad (4)$$




Figure 2: For an optimal rank 1 approximation with a fixed b , we should always choose c_i so that $c_i b$ is the projection of a_i onto b .

In this case, by Pythagorean theorem, $\sum_{i=1}^n \|a_i - c_i \cdot b\|_2^2 = \sum_{i=1}^n (\|a_i\|_2^2 - \|c_i b\|_2^2) = \|A\|_F^2 - \sum_{i=1}^n \|c_i b\|_2^2$. So, in fact, solving (4) is actually equivalent to solving:

$$\max_{b \in \mathbb{R}^n, \|b\|_2=1} \|bb^T A\|_F^2. \quad (5)$$

From this point of view, it is clear that u_1 is the optimal choice for b . Writing A using the SVD, we have

$$\|bb^T A\|_F^2 = \|b^T A\|_2^2 = \|b^T U \Sigma V^T\|_2^2 = \|b^T U \Sigma\|_2^2 = \sum_{i=1}^r (b^T u_i)^2 \sigma_i^2.$$

To see the first equality, observe that $b^T A$ is a $1 \times n$ matrix, so $\|b(b^T A)\|_F^2$ is simply $\sum_{i=1}^d b_i^2 \|b^T A\|_2^2 = \|b^T A\|_2^2$, as b is a unit vector. The second equality is simply applying the SVD of A . The third equality follows as V is orthonormal (so multiplying on the right by V^T just changes bases in a norm-preserving manner). To see the final equality, recall that b^T is a $1 \times d$ vector, and U is a $d \times r$ matrix. So $b^T U$ is a $1 \times r$ matrix with entries $(b^T u_1), \dots, (b^T u_r)$, and recall that Σ is diagonal.

Now, because U is orthonormal and $\|b\|_2^2 = 1$, $\sum_{i=1}^r (b^T u_i)^2 = 1$. Accordingly, since $\sigma_1 \geq \dots \geq \sigma_r$, $\sum_{i=1}^k (b^T u_i)^2 \sigma_i^2$ is maximized when $(b^T u_1)^2 = 1$, which can be accomplished by setting $b = u_1$.

Case: $k > 1$.

The proof for $k > 1$ is similar. It will be helpful to use the following:

Claim 3 (Matrix Pythagorean Theorem). *If M and N are matrices with mutually orthogonal columns, i.e. $M^T N = 0$, then,*

$$\|M + N\|_F^2 = \|M\|_F^2 + \|N\|_F^2.$$

This is a direct consequence of writing $\|M + N\|_F^2 = \sum_i \|m_i + n_i\|_2^2$ and applying the Pythagorean theorem to each column $m_i + n_i$ separately.

As in the rank 1 case, without loss of generality we can view the low rank approximation problem as choosing an optimal orthonormal matrix B to minimize $\|A - BC\|_F^2$. Using an identical projection argument (fixing B , what is the optimal C ?), the optimal C for a given B is $B^T A$. So our goal is to solve:

$$\min_{B \in \mathbb{R}^{d \times k}, B^T B = I} \|A - BB^T A\|_F^2. \quad (6)$$

By matrix Pythagorean theorem applied to $\|(A - BB^T A) + BB^T A\|_F^2$, we have

$$\|A - BB^T A\|_F^2 = \|A\|_F^2 - \|BB^T A\|_F^2$$

and thus (6) is equivalent to

$$\max_{B \in \mathbb{R}^{d \times k}, B^T B = I} \|BB^T A\|_F^2. \quad (7)$$

Since A can be written as $U\Sigma V^T$, this is equivalent to solving:

$$\max_{B \in \mathbb{R}^{d \times k}, B^T B = I} \|BB^T U\Sigma V^T\|_F^2 = \max_{B \in \mathbb{R}^{d \times k}, B^T B = I} \|B^T U\Sigma\|_F^2.$$

Note that $Q = B^T U \in \mathbb{R}^{k \times r}$ has orthonormal rows, so its columns cannot have norm greater than 1. Also the sum of Q 's squared column norms is k (its Frobenius norm squared). It follows that $\|Q\Sigma\|_F^2 = \sum_{i=1}^r \|q\|_i^2 \sigma_i^2 \leq \sigma_1^2 + \dots + \sigma_k^2$. This maximum is obtained when Q 's first k columns are the standard basis vector – i.e. when $B = U_k$. \square

4 Greedily constructing low-rank approximations

As mentioned, one thing that's interesting about the SVD and Claim 2 is that it implies that we can construct an optimal low-rank approximation in a greedy way: if b_1 is the best basis vector for a rank 1 approximation, then there's an optimal rank 2 approximation that maintains b_1 as one of its basis vectors. In fact, this observation gives an approach to *proving* Theorem 1 that the SVD exists for all matrices. Again, if you've already seen this proven in another class, feel free to skip this section.

Consider the following iterative routine, which constructs a singular value decomposition for any matrix A :

- Let $A^{(1)} = A$.
- For $i = 1, \dots, r$:
 - Let $b_i, c_i = \arg \min_{b,c} \|A^{(i)} - bc^T\|_F^2$.
 - Let $A^{(i+1)} \leftarrow A^{(i)} - b_i c_i^T$ and set $u_i = b_i / \|b_i\|_2, v_i = c_i / \|c_i\|_2, \sigma_i = \|b_i\|_2 \cdot \|c_i\|_2$.
- Set $U = [u_1, \dots, u_r], V = [v_1, \dots, v_r]$, and $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_r)$.

Note that step one of the procedure requires an algorithm for computing an optimal *rank 1* approximation to a given matrix. At least to prove *existence* of the SVD, we do not need an actual implementation of this procedure. However, because we do care about computing the SVD and rank- k approximations, we will eventually see an algorithm for solving this rank 1 problem.

We first need to prove the following, which implies that our choice of U is orthonormal:

Claim 4. *If b_1, \dots, b_k are chosen as above, then $b_i^T b_j = 0$ for all i, j .*

Proof. Let's remove indices to keep notation simple, and consider an optimal rank 1 approximation bc^T for a matrix A . We claim:

1. b is always in the column span of A .
2. $b^T(A - bc^T) = 0$

The first point follows from a contradiction argument. If b is not in A 's column span, it can be written it as $Ax + y$ for some y orthogonal to all of A 's columns. Then:

$$\|A - bc^T\|_F^2 = \|A - Axc^T - yc^T\|_F^2 = \|A - Axc^T\|_F^2 + \|yc^T\|_F^2 \geq \|A - Axc^T\|_F^2.$$

The second equality follows from Claim 3 because every column in yc^T is orthogonal to every column in $A - Axc^T$. This is a contradiction because Axc^T is a rank 1 matrix that clearly achieves better error than bc^T , which we claimed was chosen to be optimal. So we conclude that b must in fact lie in A 's column span. The second claim follows from our earlier projection argument: $c(i)$ is chosen so that $c(i)b$ is the projection of a_i onto b , and thus $b^T(a_i - c(i)b) = 0$ for all i .

From these two claims, it follows that, for any i , $b_i^T b_{i+1} = 0$, because b_{i+1} is in the column span of $A^{(i+1)}$ but b_i is orthogonal to that span. Then, by induction $b_i^T b_j = 0$ for all $j > i + 1$ as well. We will just argue one step of the induction: $b_i^T A^{(i+2)} = b_i^T (A^{(i+1)} - b_{i+1}c_{i+1}^T) = 0 - 0$. So b_i is orthogonal to anything in the column span of $A^{(i+2)}$, and is thus orthogonal to b_{i+2} . \square

The same exact argument (applied to rows instead of columns) also lets us establish:

Claim 5. *If c_1, \dots, c_k are chosen using our greedy procedure above, then $c_i^T c_j = 0$ for all i, j . In other words, our ground basis is orthogonal.*

Finally, we note that, since b_1, \dots, b_r are all orthogonal to the column span of $A^{(r+1)}$, then it must be that $A^{(r+1)} = 0$. So $\sum_{i=1}^r b_i c_i^T = A$, and thus $\sum_{i=1}^r u_i \sigma_i v_i^T = A$ for the $U = [u_1, \dots, u_r]$, $V = [v_1, \dots, v_r]$, and $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_r)$ produced by the iterative algorithm. Combined with Claim 4 and Claim 5 this proves Theorem 1.

5 Computing the best rank-1 approximation

The SVD gives one way of obtaining an optimal low-rank approximation for any rank parameter k . It can be computed in essentially time $O(ndr)$ time, where $r = \min(n, d)$.¹

¹We say ‘‘roughly’’ because technically there is no ‘‘exact’’ algorithm for the SVD, even in the Real RAM model of computation. This is consequence of the Abel-Ruffini theorem. Thus, all SVD algorithms are

In general, this is too slow from many large data matrices. Next we discuss the power method, which gives a much faster way of finding just the top singular vector of a matrix, which is typically the problem we want to solve in data applications. The power method runs in approximately $O(nd)$ time. It can thus find k singular vectors iteratively in $O(ndk)$ time, which is much faster than a full SVD when $k \ll \text{rank}(A)$.

The following material was outlined or not covered in class.

6 Connection to Other Matrix Decompositions

The singular value decomposition is closely related to other matrix decompositions:

Eigendecomposition The left singular vectors of A are eigenvectors of $AA^T = U\Sigma^2U^T$ and the right singular vectors are eigenvectors of A^TA . To see that this is the case, note that:

$$AA^T u_i = U\Sigma V^T V U^T u_i = U\Sigma e_i = \sigma_i u_i.$$

Here e_i is the i^{th} standard basis vector: $U^T u_i = e_i$ because u_i is orthogonal to all other columns in U .

The connection with eigendecomposition means that any algorithm for eigendecomposition can be used to compute an SVD. Suppose $d \leq n$. Then we can compute $A^T A$, from which we can compute V using an eigendecomposition algorithm. We then have $\Sigma U^T = AV^T$, so we can obtain Σ and U by normalizing the columns of this matrix and setting σ_i to be the normalization factor for column i . This procedure takes $O(nd^2)$ time to compute $A^T A$ and roughly $O(d^3)$ time to compute the eigendecomposition of this matrix²

On another note, you may recall that any real symmetric matrix M has eigendecomposition $U\Lambda U^T$ where U is orthonormal. Λ can have negative diagonal elements, so at least up to changing signs, M 's singular vectors are the same as its eigenvectors. Its singular values are the absolute values of its eigenvalues.

Principal Component Analysis (PCA) PCA is almost the same as the SVD, however, before computing singular vectors, we mean center A 's rows: $a_i \rightarrow a_i - \frac{1}{n} \sum_{j=1}^n a_j$. The right singular vectors of the resulting matrix are called the "principal components" of A .

7 The Power Method

For an $n \times d$ matrix with $n \leq d$, we cannot hope to do much better than $O(nd^2)$ time for computing an SVD. In theory, we can speed up the computation of $A^T A$ and the eigendecomposition of this $n \times n$ matrix with fast matrix multiplication. Doing so achieves

technically approximation algorithms. However, standard methods obtain *very* good ϵ dependence. E.g. the QR algorithm can compute a factorization $U\Sigma V^T$ with $\|U\Sigma V^T - A\| \leq \epsilon$ in $O(nd^2 + d^2 \log \log(1/\epsilon))$ time. The second term is ignored because it's always lower order in practice.

²We say roughly "roughly" because technically there is no "exact" algorithm for the SVD, even in the Real RAM model of computation. This is consequence of the Abel-Ruffini theorem. Thus, all SVD algorithms are technically approximation algorithms. However, standard methods obtain *very* good ϵ dependence. E.g. the QR algorithm can compute a factorization $V\Sigma^2 V^T$ with $\|V\Sigma^2 V^T - A^T A\| \leq \epsilon$ in $O(d^3 + d^2 \log \log(1/\epsilon))$ time. The second term is ignored because it is always lower order in practice.

a runtime of $O(nd^{\omega-1})$, where ω is the current best known exponent for $d \times d$ matrix multiplication ($\omega = 2.3728639\dots$ as of 2014 [4]). In practice, however, runtime still scales as $O(nd^2)$.

We want something faster. We are especially interested in algorithms that run more quickly when we only want to compute a few of A 's top singular vectors, not all n of them (as is often the case in applications). One such algorithm is the well known *power method*. We present a version below for approximately computing the top right singular vector of A , which can be used to find a best rank 1 approximation:

Power Method

- Initialize $z_0 \in \mathbb{R}^d$ to have every entry a random Gaussian variable. Set $z_0 = z_0/\|z_0\|_2$.
- Repeat: $z_{t+1} \leftarrow A^T(Az_t)$. $z_{t+1} \leftarrow z_{t+1}/\|z_{t+1}\|_2$.

Theorem 6. Let $\gamma = \frac{\sigma_1 - \sigma_2}{\sigma_1}$ be a parameter that measures the “gap” between A 's first and second singular values. After $t = O\left(\frac{\log(d/\epsilon)}{\gamma}\right)$ iterations, $\|v_1 - z_t\|_2^2 \leq \epsilon$. I.e. z_t is a very good approximate top right singular vector. The power method runs in $O(t \cdot nd)$ time.

Proof. Write $z_0 = \sum_{i=1}^d \alpha_i v_i$ where v_i is the i^{th} right singular vector of A . Each α_i represents “how much” of singular vector v_i is in z_0 . Let $\alpha \in \mathbb{R}^d$ be the vector containing these values. $\alpha = V^T g / \|g\|_2$ where g is a vector of independent Gaussians. By rotational invariance of the Gaussian distribution, $V^T g$ is also a random Gaussian vector. So at least to start, z_0 contains a random amount of every right singular vector in A .

It's not hard to check that $\alpha_1 > 1/\text{poly}(d)$ with high probability and, since z_0 has unit norm, $\max_i \alpha_i = 1$. So we at least have a non-negligible component of v_1 in z_0 .

The idea behind the power method is to boost this component so that, eventually, z_t is made up almost entirely of v_1 . This is accomplished by repeatedly multiplying by $A^T A$. After t steps, $z_t = c(A^T A)^t z_0$ for some scale factor c . Since $A^T A = V \Sigma^2 V^T$, after iteration t we have:

$$z_t = \sum_{i=1}^d w_i v_i$$

where each $w_i \sim \sigma_i^{2t} \alpha_i$. By our “gap” assumption, $\frac{\sigma_1}{\sigma_j} \geq 1 + \gamma$ for all $j \geq 2$. Accordingly, after t steps, for all $j \geq 2$,

$$\frac{w_j}{w_1} \leq (1 + \gamma)^{2t} \cdot \frac{\alpha_j}{\alpha_1} \leq (1 + \gamma)^{2t} \cdot \text{poly}(d).$$

If we set $t = O\left(\frac{\log(d/\epsilon)}{\gamma}\right)$ then we have $\frac{w_j}{w_1} \leq \sqrt{\epsilon/d}$, which means that $w_j \leq \sqrt{\epsilon/2d}$. Since $\|z\|_t = \sum_{j=1}^d w_j^2 = 1$, it follows that $w_1 \geq 1 - \epsilon/2$ and thus $z_t^T v_1 \geq (1 - \epsilon/2)$. So:

$$\|v_1 - z_t\|_2^2 = \|v_1\|_2^2 + \|z_t\|_2^2 - 2z_t^T v_1 \leq \epsilon.$$

□

So when γ is considered constant, power method converges in $\log(d/\epsilon)$ iterations. Accordingly, we can compute a good approximation to the top right singular vector in time $O(nd \log(d/\epsilon))$.

How about when γ is very small? In the most extreme case, when $\gamma = 0$, power method will never converge on v_1 and in fact the dependence on $1/\gamma$ is unavoidable. However, if γ is small, we don't typically care about finding v_1 ! Since $\sigma_1 = \sigma_2$, v_2 is just as "good" of an eigenvector as v_1 . It's a good exercise to prove the following:

Theorem 7. *After $t = O\left(\frac{\log(d/\epsilon)}{\epsilon}\right)$ iterations of power method, z_t satisfies:*

- $\|Az_t\|_2 \geq (1 - \epsilon)\sigma_1$
- $\|A - Az_t z_t^T\|_F \leq (1 + \epsilon)\|A - Av_1 v_1^T\|_F$
- $\|A - Az_t z_t^T\|_2 \leq (1 + \epsilon)\|A - Av_1 v_1^T\|_2$

In other words, after $O\left(\frac{\log(d/\epsilon)}{\epsilon}\right)$, by most common measures, projecting rows to z_t still gives a nearly optimal low-rank approximation for A . We've traded a $1/\gamma$ dependence for a $1/\epsilon$ dependence and a different, but arguably more natural approximation guarantee.

8 Beyond Power Method

Last class we discussed how low-rank approximations can be computed in a "greedy" way – i.e. we find a rank 1 approximation to A , subtract it off, then find a rank 1 approximation to the remainder, continuing for k steps. We sum up all of these rank-1 approximations to find a rank k approximation. This process is called "deflation" and it's possible to show that it works even when our rank-1 approximations are computed approximately (e.g. with power method).

Other ways of obtaining rank k approximations include "blocked" versions of the power method, where we derive k singular vectors from $(A^T A)^t Z$ where $Z \in \mathbb{R}^{d \times k}$ is a random Gaussian matrix (instead of just a vector).

In either case, these iterative methods take $O(t \cdot ndk)$ time to compute a nearly optimal rank- k approximation, where either $t = O\left(\frac{\log d}{\epsilon}\right)$ or depends on gaps between A 's singular vectors. In practice, this is typically much faster than computing a full SVD. As an added advantage, all of this runtime complexity comes from matrix-vector multiplications of the form $A^T A x$, which can be speed up beyond $O(nd)$ time when A is sparse or when parallel processing is available.

Finally, I'll mention that it is actually possible to improve the iteration complexity of the power method to $t = O\left(\frac{\log d}{\sqrt{\epsilon}}\right)$ using what is known as the Lanczos method. Variations on the Lanczos method are used almost everywhere in practice (e.g. if you run `svds` in MATLAB, Python, etc.). If you are interested, Chapter 10 in [5] gives a relatively simple analysis for the rank-1 case.

9 Matrix decomposition and graphs

In general, algorithms based on singular value decomposition or eigendecomposition are referred to as “spectral methods” – the singular values $\sigma_1, \dots, \sigma_r$ of a matrix or the eigenvalues $\lambda_1, \dots, \lambda_r$ are referred to as the “spectrum” of the matrix.

Beyond statistics, data analysis, and machine learning, spectral methods have been very important in developing faster algorithms for graphs, including for classic problems like minimum cut and max flow. Today we will see one particularly nice application.

A big reason for the connection between graphs and matrix decompositions is that the eigenvectors/singular vectors of certain matrix representations of a graph G contain a lot of information about *cuts* in the graph.

Let $G = (V, E)$ be an undirected graph on n nodes. Recall that G 's adjacency matrix A is defined by:

$$A_{u,v} = \begin{cases} 1 & \text{if } (u, v) \in E \\ 0 & \text{if } (u, v) \notin E \end{cases}$$

A is a symmetric matrix, so it has an eigendecomposition $U\Lambda U^T$ where U is orthonormal. For a given eigenvector u_i and corresponding eigenvalue λ_i , $u_i^T A u_i = \lambda_i$.

Consider a vector $z \in \{-1, 1\}^n$. It's not hard to check that $z^T A z$ equals:

$$z^T A z = \sum_{u,v \in V} \mathbb{1}[(u, v) \in E] z_u z_v.$$

Think of z as an indicator vector for a cut between two partitions of vertices, S and T . I.e. $z_u = 1$ for $u \in S$ and $z_u = -1$ for $u \in T$. Every edge within S or T adds a value of 1 to $z^T A z$, which every edge between the partitions adds -1 . So, in general, $z^T A z$ will be larger when z is an indicator for “good” partition of G that clusters the graph into two groups of well connected vertices.

In particular, this means that z correlates well with the top eigenvectors of A , which means that these eigenvectors are often useful in finding such cuts.

10 Planted Bisection/Stochastic Block Model/Community Detection

Unfortunately, most optimization problems involving balanced/sparse cuts are NP-hard, but there are many natural “average case” problems to study, which can justify why spectral methods work well in practice. Consider the follow:

Definition 1 (Stochastic Block Model). *Let $G(V, E)$ be a random graph with vertices $V = 1, \dots, n$. Let S, T form a bisection of V . I.e. $S, T \subset V$ with $S \cup T = V$, $S \cap T = \emptyset$ and $|S| = |T| = n/2$. For probabilities $p > q$, construct G by adding edge (i, j) independently with probability Y_{ij} , where:*

$$Y_{ij} = \begin{cases} p & \text{if both } i, j \in T \text{ or } i, j \in S \\ q & \text{if } i \in T, j \in S \text{ or } i \in S, j \in T. \end{cases}$$

We can think of S and T as disjoint “communities” in our graph. Nodes are connected randomly, but it is more likely that they are connected to members of their community than members outside their community.

Our goal is to design a spectral method to recover these underlying communities. Today we are just going to give a sketch of an algorithm/proof.

Let’s introduce another matrix $B \in \mathbb{R}^{n \times n}$ defined as follows:

$$B_{ij} = \begin{cases} p & \text{if } i, j \in T \text{ or } i, j \in S \\ q & \text{if } i \in T, j \in S \text{ or } i \in S, j \in T. \end{cases}$$

It is not hard to see that $B = \mathbb{E}[A] + pI$, where I is an $n \times n$ identity. Accordingly, at least in expectation, A has the eigenvectors as B . What are these eigenvectors?

B is rank two, so it only has two, u_1 and u_2 , where:

$$u_1(i) = \frac{1}{\sqrt{n}} \mathbf{1} \quad \forall i,$$

$$u_2(i) = \begin{cases} \frac{1}{\sqrt{n}} \mathbf{1} & \forall i \in S \\ \frac{1}{\sqrt{n}} - 1 & \forall i \in T. \end{cases}$$

$Bu_1 = \frac{n}{2}(p+q)u_1$ and $Bu_2 = \frac{n}{2}(p-q)u_2$. In this case, u_1 and u_2 are also B ’s singular vectors.

So, if we could compute B ’s eigenvectors, we could immediately recover our community by simply examining u_2 . Of course, we don’t have access to B , but we do have access to a perturbed version of the matrix via:

$$\hat{A} = A + pI.$$

Consider $R = B - \hat{A}$. Classic perturbation theory results in linear algebra tell us that if $\|R\|_2$ is small, then \hat{A} ’s eigenvalues and eigenvectors will be close to those of B .

Let B have eigenvectors u_1, \dots, u_n and eigenvalues $\lambda_1, \dots, \lambda_n$. Let \hat{A} have eigenvectors $\hat{u}_1, \dots, \hat{u}_n$ and eigenvalues $\hat{\lambda}_1, \dots, \hat{\lambda}_n$. Using ideas from the past few lectures you could already prove the following result, which is a good exercise:

Claim 8. *If B and \hat{A} are real symmetric matrices with $\|B - \hat{A}\|_2 \leq \epsilon$, $\forall i$,*

$$|\lambda_i - \hat{\lambda}_i| \leq \epsilon.$$

In words, if \hat{A} and B are close in spectral norm, their eigenvalues are close. For our application, we further need that the matrices’ eigenvectors are close. Below is a classic result quantifying this – you can find a simple proof of a slightly weaker version in [6].

Claim 9 (Davis-Kahan, 1970 [7]). *Suppose B and \hat{A} are real symmetric matrices with $\|B - \hat{A}\|_2 \leq \epsilon$. Let θ_i denote the angle between u_i and \hat{u}_i . For all i ,*

$$\sin \theta_i \leq \frac{\epsilon}{\min_{j \neq i} |\lambda_i - \lambda_j|}.$$

Let's unpack this claim. It says that if B and \hat{A} are close in spectral norm, then their corresponding eigenvectors are close. However, the distance is effected by a factor of $1/|\lambda_i - \lambda_j|$. This makes sense – suppose $\lambda_i < \lambda_{i+1} + \epsilon$. Then a perturbation with spectral norm ϵ could cause the u_i and u_{i+1} to “swap” order – specifically just add $\epsilon u_{i+1} u_{i+1}^T$ to B to cause such a change. In the perturbed matrix, $\hat{u}_i = u_{i+1}$, which is orthogonal to u_i .

Fortunately, in our case, we have a gap between B 's eigenvalues – in particular, $|\lambda_2 - \lambda_1| \geq nq$ and $|\lambda_2 - 0| = \frac{n}{2}(p - q)$. Let's assume a challenging regime where q is close to p and thus $\frac{n}{2}(p - q) \leq nq$.

A simple corollary of Claim 9 is that $\|u_i - \hat{u}_i\|_2 \leq \frac{\sqrt{2}\epsilon}{\min_{j \neq i} |\lambda_i - \lambda_j|}$.

As an estimate for our community indicator vector u_2 , let's consider $\text{sign}(\hat{u}_2)$. Suppose this estimate differs from u_2 on k entries. Then it must be that:

$$\|\hat{u}_2 - \mu_2\|_2 \geq \sqrt{\frac{k}{n}}$$

So, by the eigenvector perturbation argument, we can bound

$$k \leq O\left(\frac{\epsilon^2}{n(p - q)^2}\right)$$

11 Eigenvalues of Random matrices

So we are left to bound $\|R\|_2$. $R = B - \hat{A}$ is a random matrix with half of its entries equal to p with probability $(1 - p)$ and $(p - 1)$ with probability p , and the other half equal to q with probability $(1 - q)$ and $(q - 1)$ with probability q .

It is possible to prove:

Theorem 10 (From [8]). *If $p \geq O(\log^4 n/n)$, then with high probability,*

$$\|R\|_2 \leq O(\sqrt{pn})$$

You will prove a very related (but slightly looser statement on the problem set).

With this bound in place, we immediately have that our spectral algorithm recovers the hidden partition with a number of mistakes bounded by:

$$k = O\left(\frac{p}{(p - q)^2}\right).$$

This is very good. Even when $q = p - O(1/\sqrt{n})$ (e.g. our probabilities are very close, so the communities should be hard to distinguish) we only make $O(n)$ mistakes – i.e. we can guess a large constant fraction of the community identities correctly.

References

- [1] John Novembre, Toby Johnson, Katarzyna Bryc, Zoltan Kutalik, Adam R Boyko, Adam Auton, Amit Indap, Karen King, Sven Bergmann, Matthew Nelson, Matthew Stephens, Carlos Bustamante, . (2008). Genes Mirror Geography within Europe. *Nature*. 456. 274.

- [2] Christopher D. Manning, Prabhakar Raghavan and Hinrich Schütze, Introduction to Information Retrieval, Cambridge University Press. 2008. <https://nlp.stanford.edu/IR-book/information-retrieval-book.html>.
- [3] T.K. Landauer, P.W. Foltz, and D. Laham. Introduction to Latent Semantic Analysis. Discourse Processes, 25, 259-284. (1998)
- [4] Le Gall, François. Powers of Tensors and Fast Matrix Multiplication. Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation, 296–303. 2014.
- [5] Sushant Sachdeva, Nisheeth K. Vishnoi. Faster Algorithms via Approximation Theory. Foundations and Trends in Theoretical Computer Science. 2013. https://theory.epfl.ch/vishnoi/Publications_files/approx-survey.pdf.
- [6] Daniel Spielman. Spectral Partitioning in a Stochastic Block Model. Lecture, Yale University. 2015. <http://www.cs.yale.edu/homes/spielman/561/lect21-15.pdf>.
- [7] Chandler Davis, William Morton Kahan. The rotation of eigenvectors by a perturbation. SIAM Journal on Numerical Analysis, 7(1):1–46, 1970.
- [8] Van Vu. Spectral norm of random matrices. Combinatorica, 27(6):721–736, 2007.