

7 Eigenvalues, eigenvectors, diagonalization

7.1 Overview of diagonalizations

We have seen that a transformation matrix looks completely different in different bases (the matrices (5.10) and (5.11) express the same transformation, but in different bases). It is natural to search for the “best” basis for a given matrix, i.e. for a basis in which the matrix is very simple. The “best” form of a matrix is the diagonal form, since it is extremely easy to manipulate with diagonal matrices (even we could take its exponential in Example II in the Introduction). Also look at the pictures of images of simple 2×2 matrices in Section 5.1 of [D], especially digest the picture on p.128.

There is a difference between square matrices and rectangular ones. Rectangular matrices act between different spaces, and one has the freedom to choose bases in these spaces independently. It turns out that with this freedom any matrix can be brought into a diagonal form by an orthogonal change of basis both in the domain space and in the image space. In other words, if one is allowed to change bases independently, then one gets the best possible form one could hope for. (It is another question how to compute it). This is called the **singular value decomposition (SVD)**, which we discuss shortly in Section 7.4. Roughly speaking it gives a factorization of A as $A = UDV^{-1} = UDV^t$, where D is diagonal and V, U are orthogonal on the appropriate subspaces.

The same construction works for square matrices as well. However, when discussing square matrices, it is desirable to have one single basis instead of two, since we are eventually in the same space. This means that we are looking for a matrix V such that $A = VDV^{-1}$ where D is diagonal (see Corollary 5.18). This is called the **diagonalization** of the matrix A . In other words we are looking for a diagonal matrix that is similar to A (see Definition 5.19). This reduces the freedom compared to SVD (you are allowed to play with only one V instead of U and V) and it turns out that not every square matrix can be written into a diagonal form.

EXERCISE: Show that the matrix $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ cannot be written as VDV^{-1} . (Hint: let $V = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ and $D = \begin{pmatrix} s & 0 \\ 0 & t \end{pmatrix}$ with any unknown entries and try to determine them; you will get contradiction).

There are two ways out, and both of them are used:

(I.) Insist on the diagonal form, and try to determine the most general class of matrices which can be diagonalized. It turns out that every symmetric matrix A (i.e. matrix such that $A = A^t$) with real entries is diagonalizable, moreover, it is diagonalizable by an orthogonal matrix V (remember, we like orthogonal matrices!!). This is one of the most important theorem in mathematics, and it usually runs under the name of **spectral theorem** (see Theorem 7.10). Its importance is comparable to the Fundamental Theorem of calculus. One could think that the restriction for symmetric matrices is very serious, but it is just a fact of life that in many many applications one automatically has to deal with a symmetric matrix from the nature of the problem.

Nevertheless, one can try to go beyond symmetric matrices, and it turns out that “most” matrices are actually diagonalizable, although the conjugation matrix V will not be orthogonal.

(II.) If a matrix is nondiagonalizable (like the one in the Exercise above), it is due to an unfortunate coincidence (namely that it happen to have multiple eigenvalues and something else also goes wrong). However, this phenomenon is very important; in the theory of differential equations this is the source of *resonances*, which could destroy the regular behaviour of a system (e.g. your old car sometimes starts terribly trembling at a certain speed, but only at that speed. This is a resonance effect).

For such matrices another factorization is known: **Jordan canonical form**. We will not discuss it here, but we remark, that it brings the matrix into an “almost” diagonal form. Namely it presents matrices V and D such that $A = VDV^{-1}$, but D is only almost diagonal,

i.e. apart from the diagonal it is also allowed to have nonzero elements at the entries just one step above the diagonal.

7.2 Full diagonalization of square matrices

We aim at $A = VDV^{-1}$ with some diagonal matrix D with numbers $\lambda_1, \lambda_2, \dots, \lambda_n$ in its diagonal (traditionally they are denoted by λ). Hence $AV = VD$, i.e.

$$\left[A\mathbf{v}_1 \mid A\mathbf{v}_2 \mid \dots \mid A\mathbf{v}_n \right] = \left[\lambda_1\mathbf{v}_1 \mid \lambda_2\mathbf{v}_2 \mid \dots \mid \lambda_n\mathbf{v}_n \right] \quad (7.1)$$

are the column vectors of the both sides of the equation $AV = VD$, hence $A\mathbf{v}_j = \lambda_j\mathbf{v}_j$ for the columns of V (the vertical lines just separate the columns). This gives rise to the

Definition 7.1 *A nonzero vector \mathbf{v} is an **eigenvector** of an $n \times n$ matrix A with **eigenvalue** λ if*

$$A\mathbf{v} = \lambda\mathbf{v}$$

WARNING: The condition that \mathbf{v} is not zero is crucial. The reason is that we want to characterize the eigenvalues, but if the zero vector were allowed then every number λ could be an eigenvalue. However $\lambda = 0$ is allowed.

REMARK: Eigenvalues and eigenvectors come together in pairs. Every eigenvector has a unique eigenvalue, just compute $A\mathbf{v}$ and see how it is related to \mathbf{v} . However, the eigenvector belonging to an eigenvalue is not unique. Clearly every nonzero multiple of an eigenvector is also an eigenvector (belonging to the same eigenvalue), so it would be better to talk about “eigendirections” or “eigenlines”, but traditionally we use eigenvectors with the understanding that there is always a freedom of a constant multiple. [In parenthesis: you have seen such thing, the indefinite integral was defined only up to an additive constant].

IMPORTANT REMARK: The true importance of the concept of eigenvalue is that it is independent of the basis; it is determined by the underlying linear transformation only. Recall that the same linear transformation is expressed by very differently looking matrices with respect to different bases. Of course all these matrices are *similar* (Definition 5.19). The key point is

Theorem 7.2 *Let A, B be similar matrices, i.e. there is an invertible matrix V with $B = VAV^{-1}$. Then the set of eigenvalues (with multiplicity) of A and B coincide.*

REMARK: The converse of this theorem is not always true. However, if A is diagonalizable, then B is so (CHECK(*)), and in this case the set of eigenvalues coincide.

Proof of Theorem 7.2: The simplest proof uses a fact about determinants, namely that

$$\det(AB) = \det(A)\det(B) \tag{7.2}$$

which we mention without proof. Now we can easily show that the characteristic polynomials of A and B coincide (hence so do the eigenvalues)

$$\begin{aligned} \det(B - \lambda I) &= \det[V(A - \lambda I)V^{-1}] = \det(V)\det(A - \lambda I)\det(V^{-1}) \\ &= \det(A - \lambda I)\det(V)\det(V^{-1}) = \det(A - \lambda I)\det(VV^{-1}) \\ &= \det(A - \lambda I) \end{aligned}$$

Notice that we used (7.2) twice.

An alternative proof for the case of simple eigenvalues easily follows from Theorem 7.5. (DO IT(*)) This proof avoids (7.2). \square

We can reformulate the diagonalization problem into the language of eigenvalues. The proof of the following Lemma is hopefully clear from (7.1) (THINK IT OVER(*))!

Lemma 7.3 *An $n \times n$ matrix is diagonalizable, i.e. it can be written as $A = VDV^{-1}$ for some invertible matrix V and diagonal matrix D if and only if there exists a basis $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ consisting of eigenvectors (this is also called **eigenbasis**). In this case the matrix V is formed from these vectors as columns and the diagonal matrix has the eigenvalues in its diagonal.*

A can be diagonalized by an orthogonal conjugating matrix V if and only if the eigenbasis can be chosen orthonormal.

REMARK: The decomposition is not unique, there are several sources of ambiguity. For example every eigenvector can be multiplied by any number (but notice that this ambiguity is really irrelevant: if you multiply some eigenvector by 3, it multiplies a column of V by 3, but then on the other side it divides a row of V^{-1} by 3. CHECK(*!!)). Moreover, the order in which you labelled the eigenvectors is certainly not fixed. Finally, if two linearly independent vectors have the same eigenvalue, then any linear combination of them is also an eigenvector (CHECK (*!!)), hence you have a big freedom. However, the set of eigenvalues with their multiplicity is uniquely determined. This was proven in Theorem 7.2.

7.2.1 Finding eigenvalues, eigenvectors

Rewrite the eigenvalue equation as $(A - \lambda I)\mathbf{v} = 0$. Since \mathbf{v} is nonzero, it means that $A - \lambda I$ is a singular matrix (recall, a square matrix is singular if not invertible, see Theorem 4.16), hence

$$\det(A - \lambda I) = 0$$

(see (vi) of Theorem 4.16). If we compute this determinant, we get a polynomial of degree exactly n , which is called the **characteristic polynomial** of A :

$$p(\lambda) = \det(A - \lambda I)$$

and clearly the eigenvalues are its roots.

Recall the

Theorem 7.4 (Fundamental theorem of algebra) *Every polynomial $p(\lambda)$ of degree n has exactly n roots, including multiplicity and complex roots.*

REMARK: A root λ_0 has multiplicity k if $p(\lambda)$ is divisible by the factor $(\lambda - \lambda_0)^k$ but is not divisible by $(\lambda - \lambda_0)^{k+1}$. In other words, you can write $p(\lambda) = (\lambda - \lambda_0)^k q(\lambda)$, where $q(\lambda)$ is another polynomial – of degree $(n - k)$ –, where $q(\lambda_0) \neq 0$.

We also remark, that including complex roots is necessary even if the coefficients of the polynomial are real. Think of $p(\lambda) = \lambda^2 + 1$.

Applying this theorem to our case, it is clear that an $n \times n$ matrix has exactly n eigenvalues, counting multiplicity and complex eigenvalues. Hence finding eigenvalues is equivalent to find the roots of the characteristic polynomial. Recall that there is no general “root-formula” for polynomials of degree bigger than four, and even the formulas for the cubic and quartic equations are horrendous and not used (however, the quadratic formula is useful, and you should know it!). But Newton’s method can find at least the real roots of any polynomial very fast (for the complex root you have to be a bit smarter, but there are methods for that as well).

Once the eigenvalues are found, we search for eigenvectors. This has to be done separately for each eigenvalue and it consist of finding vectors in the nullspace of $A - \lambda I$. For real eigenvalues, the corresponding eigenvector is real as well. For complex eigenvalues you will have complex eigenvectors. They are in \mathbf{C}^n not in \mathbf{R}^n , which means that you have to use complex numbers everywhere. But they are just as good numbers as the real ones, everything we discussed for \mathbf{R}^n immediately goes through to \mathbf{C}^n .

For *multiple eigenvalues* you have to find as many linearly independent eigenvectors as the multiplicity of that eigenvalue. This step may not work in general, and this is the obstruction to general diagonalizability.

Example: The matrix $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ has characteristic polynomial $p(\lambda) = (\lambda - 1)^2$, hence the eigenvalue $\lambda = 1$ has multiplicity 2. But apart from $\mathbf{v}_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ there is no other eigenvector (lin. indep. of this one). This is why A is not diagonalizable (see Exercise in Section 7.1).

However, if the *eigenvalues are simple*, then this procedure always works and it leads to the following theorem:

Theorem 7.5 (Diagonalization for simple eigenvalues) *Let A be an $n \times n$ matrix and suppose that all eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$ are distinct (simple roots of $p(\lambda)$). Choose an eigenvector \mathbf{v}_i to the eigenvalue λ_i for each i . Then these eigenvectors form a basis. Hence the matrix is diagonalizable with the conjugating matrix $V = [\mathbf{v}_1, \mathbf{v}_2, \dots]$:*

$$A = VDV^{-1}.$$

with D being the diagonal matrix

$$D = \begin{pmatrix} \lambda_1 & 0 & 0 & \dots & 0 \\ 0 & \lambda_2 & 0 & \dots & 0 \\ 0 & 0 & \lambda_3 & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \dots & \lambda_n \end{pmatrix}$$

Proof: By Theorem 7.3 the only thing that remains to prove is that the eigenvectors are linearly independent (then they will form a basis by counting the dimension). Suppose, on contrary that some nontrivial linear combination of them is zero:

$$c_1\mathbf{v}_1 + c_2\mathbf{v}_2 + \dots + c_n\mathbf{v}_n = \mathbf{0} \tag{7.3}$$

and we will use a proof by contradiction. Choose that nontrivial linear combination which has the fewest nonzero coefficients (This is a good trick, watch out! You don't want to choose the

trivial case, when all of them are zero, but you still insist on choosing the closest possible). Suppose that $c_1 \neq 0$. Let A act on this equation:

$$A(c_1\mathbf{v}_1 + c_2\mathbf{v}_2 + \dots + c_n\mathbf{v}_n) = c_1\lambda_1\mathbf{v}_1 + c_2\lambda_2\mathbf{v}_2 + \dots + c_n\lambda_n\mathbf{v}_n = \mathbf{0}$$

recalling that $A\mathbf{v}_j = \lambda_j\mathbf{v}_j$. Subtract λ_1 times the equation (7.3) from this second equation, you get

$$c_2(\lambda_2 - \lambda_1)\mathbf{v}_2 + \dots + c_n(\lambda_n - \lambda_1)\mathbf{v}_n = \mathbf{0} \tag{7.4}$$

(notice that the first terms cancel). This is again a linear combination of the vectors. Could all coefficients be zero? Since the eigenvalues are distinct, this depends only on the c 's. But if all c_2, c_3, \dots were zero, then (7.3) would be $c_1\mathbf{v}_1 = \mathbf{0}$ hence c_1 were zero as well, i.e. (7.3) would have been a trivial linear combination which we excluded. Hence this new linear combination (7.4) is nontrivial as well. But it clearly has less nonzero coefficient than (7.3) had, since c_1 is not present any more. However, this contradicts to our choice of the original linear combination (7.3)! This contradiction shows that the *original* assumption, namely that the eigenvectors are linearly dependent was wrong. This proves what we wanted. (Try to digest this proof, it sounds quite sophisticated and perhaps messy, but if you put some effort in understanding it, it will make sense.) \square

Problem 7.6 *Diagonalize the matrix* $A = \begin{pmatrix} 2 & 5 \\ -1 & -4 \end{pmatrix}$

SOLUTION: We have to find eigenvalues and eigenvectors. The steps are as follows:

Step 1: Find the characteristic polynomial:

$$p(\lambda) = \det \begin{pmatrix} 2 - \lambda & 5 \\ -1 & -4 - \lambda \end{pmatrix} = (2 - \lambda)(-4 - \lambda) - 5(-1) = \lambda^2 + 2\lambda - 3$$

Step 2.: Find the roots of $p(\lambda)$.

$$\lambda_1 = \frac{-2 + \sqrt{2^2 - 4(-3)}}{2} = 1 \quad \lambda_2 = \frac{-2 - \sqrt{2^2 - 4(-3)}}{2} = -3$$

Step 3.: Find the eigenvectors. Must be done separately for each eigenvalue.

Eigenvector for $\lambda_1 = 1$:

This eigenvector $\mathbf{v}_1 = \begin{pmatrix} x \\ y \end{pmatrix}$ solves $A\mathbf{v}_1 = \lambda_1\mathbf{v}_1$, or $(A - \lambda_1 I)\mathbf{v}_1 = 0$. Hence

$$\begin{pmatrix} 1 & 5 \\ -1 & -5 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

i.e.

$$\begin{aligned} x + 5y &= 0 \\ -x - 5y &= 0 \end{aligned}$$

Of course these two equations are not independent, you can throw away one of them, and give any nontrivial solution to $x + 5y = 0$. For example $x = -5, y = 1$ would do, hence $\mathbf{v}_1 = \begin{pmatrix} -5 \\ 1 \end{pmatrix}$.

Eigenvector for $\lambda_2 = -3$:

Similarly, we need $\mathbf{v}_2 = \begin{pmatrix} x \\ y \end{pmatrix}$ solving $(A - \lambda_2 I)\mathbf{v}_2 = 0$ (these x, y are not the same ones as above), i.e.

$$\begin{pmatrix} 5 & 5 \\ -1 & -1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

i.e.

$$\begin{aligned} 5x + 5y &= 0 \\ -x - y &= 0 \end{aligned}$$

so for example $x = 1, y = -1$ would do. Hence $\mathbf{v}_2 = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$ is a possible eigenvector for $\lambda_2 = -3$.

Step 4.: Finally, we write up the diagonal form $A = VDV^{-1}$. Clearly

$$V = \begin{pmatrix} -5 & 1 \\ 1 & -1 \end{pmatrix}$$

$$D = \begin{pmatrix} 1 & 0 \\ 0 & -3 \end{pmatrix}$$

and we have to compute the inverse of V , which is

$$V^{-1} = \begin{pmatrix} -1/4 & -1/4 \\ -1/4 & -5/4 \end{pmatrix}$$

Hence

$$A = \begin{pmatrix} 2 & 5 \\ -1 & -4 \end{pmatrix} = \begin{pmatrix} -5 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & -3 \end{pmatrix} \begin{pmatrix} -1/4 & -1/4 \\ -1/4 & -5/4 \end{pmatrix}$$

is the diagonal decomposition of A .

Here is an example with complex eigenvalues-vectors:

Problem 7.7 *Diagonalize the matrix* $A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$

SOLUTION: This is the matrix of rotation with $+\frac{\pi}{2}$ in \mathbf{R}^2 . It is clear that a rotation does not take any vector into its constant multiple. So seemingly there is no eigenvector. It is true that there are no real eigenvectors (which can be seen in \mathbf{R}^2), but there are complex ones. These eigenvectors cannot be represented in \mathbf{R}^2 , but the algebraic question, find (possibly complex) eigenvectors-eigenvalues satisfying the eigenvalue equation $A\mathbf{v} = \lambda\mathbf{v}$ still makes sense. Recall that “vectors” are much more general objects than “visible” planar or spatial arrows. For example an n -tuple of *complex* numbers (z_1, z_2, \dots, z_n) can also be considered as a “vector”, and this is the element of the vectorspace \mathbf{C}^n with the “usual” operations (i.e. add and scalar-multiply vectors entrywise).

The characteristic polynomial of M is

$$p(\lambda) = \det \begin{pmatrix} -\lambda & -1 \\ 1 & -\lambda \end{pmatrix} = \lambda^2 + 1$$

hence the eigenvalues are $\lambda_1 = i$, $\lambda_2 = -i$.

The eigenvector $\mathbf{v}_1 = \begin{pmatrix} x \\ y \end{pmatrix}$ corresponding to $\lambda_1 = i$ solves

$$\begin{pmatrix} -i & -1 \\ 1 & -i \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

i.e.

$$\begin{aligned} -ix - y &= 0 \\ x - iy &= 0 \end{aligned}$$

These two equations are actually multiples of each other, namely the second is i -times the first (if you recall that $i^2 = -1$). Don't forget that the complex numbers are also "numbers" one can do calculations with them.

Hence you can forget about one of these equations (since they are the same), and just solve one of them. A nontrivial solution is $x = 1, y = -i$, hence

$$\mathbf{v}_1 = \begin{pmatrix} 1 \\ -i \end{pmatrix}$$

(CHECK that $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ -i \end{pmatrix} = \begin{pmatrix} i \\ 1 \end{pmatrix}$ which is really i -times \mathbf{v}_1 !!).

By a similar calculation you get that the eigenvector \mathbf{v}_2 for $\lambda_2 = -i$ is

$$\mathbf{v}_2 = \begin{pmatrix} 1 \\ i \end{pmatrix}$$

Now let

$$U = \begin{pmatrix} 1 & 1 \\ -i & i \end{pmatrix}$$

compute its inverse

$$U^{-1} = \frac{1}{2} \begin{pmatrix} 1 & i \\ 1 & -i \end{pmatrix}$$

and write up the diagonal matrix

$$D = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}$$

from the eigenvalues. Hence

$$A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ -i & i \end{pmatrix} \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix} \begin{pmatrix} \frac{1}{2} & \frac{i}{2} \\ \frac{1}{2} & -\frac{i}{2} \end{pmatrix}$$

is the diagonalization of A .

Finally we look at an example where diagonalization is possible despite a multiple eigenvalue. Multiple eigenvalue does not a priori exclude diagonalization. If you can find enough linearly independent eigenvectors (i.e. as many as the multiplicity) to the given eigenvalue, then you can still diagonalize. More precisely, if some λ is an eigenvalue of multiplicity k , then you need $N(A - \lambda I)$ be k -dimensional. Notice that the fact that λ is an eigenvalue only guarantees that $N(A - \lambda I)$ is not trivial. But it may contain only one eigenvector (more precisely only one eigendirection), which is not enough for higher multiplicity. This is exactly what happened for the nondiagonalizable matrix $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$, the eigenvalue $\lambda_1 = \lambda_2 = 1$ was double, but the nullspace of $A - 1 \cdot I = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ is only 1 dimensional (one free variable).

Problem 7.8 *Diagonalize the matrix*

$$A = \begin{pmatrix} 2 & -2 & -1 \\ 0 & 1 & 0 \\ 2 & -4 & -1 \end{pmatrix}$$

SOLUTION: Run the usual machinery:

$$p(\lambda) = \det \begin{pmatrix} 2 - \lambda & -2 & -1 \\ 0 & 1 - \lambda & 0 \\ 2 & -4 & -1 - \lambda \end{pmatrix} = -\lambda(1 - \lambda)^2$$

Hence $\lambda_1 = 0$, $\lambda_2 = \lambda_3 = 1$, i.e., the eigenvalue 1 has multiplicity 2.

First find the eigenvector \mathbf{v}_1 to $\lambda_1 = 0$. Since λ_1 is a simple root, this is easy, as before you solve

$$\begin{pmatrix} 2 & -2 & -1 \\ 0 & 1 & 0 \\ 2 & -4 & -1 \end{pmatrix} \mathbf{v}_1 = \begin{pmatrix} 2 & -2 & -1 \\ 0 & 1 & 0 \\ 2 & -4 & -1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

and easily find (one) solution $\mathbf{v}_1 = \begin{pmatrix} 1 \\ 0 \\ 2 \end{pmatrix}$. (Notice that the matrix here is A ; it really should be $A - \lambda I$, but now $\lambda = 0$).

Next, you try to find two linearly independent eigenvectors to $\lambda = 1$. I.e. you solve

$$\begin{pmatrix} 1 & -2 & -1 \\ 0 & 0 & 0 \\ 2 & -4 & -2 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

After Gauss

$$\left(\begin{array}{ccc|c} 1 & -2 & -1 & 0 \\ 0 & 0 & 0 & 0 \\ 2 & -4 & -2 & 0 \end{array} \right) \implies \left(\begin{array}{ccc|c} 1 & -2 & -1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{array} \right)$$

i.e. you have TWO free variables! Hence you will have two linearly independent solutions.

How to find them? Exactly as you found a basis in the nullspace of a matrix: set a free variable 1, the rest zero, and repeat for all free variables. E.g. here y, z are free, set first $y = 1, z = 0$, get $x = 2$, i.e. $\mathbf{v}_2 = \begin{pmatrix} 2 \\ 1 \\ 0 \end{pmatrix}$; then set $y = 0, z = 1$, get $x = 1$, i.e. $\mathbf{v}_3 = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}$.

Finally you will have to invert the matrix

$$V = \begin{pmatrix} 1 & 2 & 1 \\ 0 & 1 & 0 \\ 2 & 0 & 1 \end{pmatrix}$$

the result is

$$V^{-1} = \begin{pmatrix} -1 & 2 & 1 \\ 0 & 1 & 0 \\ 2 & -4 & -1 \end{pmatrix}$$

and the diagonalization of A

$$A = \begin{pmatrix} 1 & 2 & 1 \\ 0 & 1 & 0 \\ 2 & 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} -1 & 2 & 1 \\ 0 & 1 & 0 \\ 2 & -4 & -1 \end{pmatrix}$$

Problem 7.9 Let again $A = \begin{pmatrix} 2 & 5 \\ -1 & -4 \end{pmatrix}$ and let $\mathbf{w} = \begin{pmatrix} 3 \\ 1 \end{pmatrix}$. Compute $A^{2000}\mathbf{w}$.

SOLUTION: Of course you could just run a machine to multiply this matrix 2000 times. But if the matrix is huge not just 2×2 , then it could take long time. Here is a better way. For that we have to recall the diagonalization above $A = VDV^{-1}$. Notice the remarkable telescopic cancellation property:

$$\begin{aligned} A^{2000} &= (VDV^{-1})(VDV^{-1})(VDV^{-1}) \dots (VDV^{-1}) \\ &= VD^{2000}V^{-1} \end{aligned}$$

since all the $V^{-1}V$ dropped out from the middle. Now computing D^{2000} is trivial:

$$D^{2000} = \begin{pmatrix} 1^{2000} & 0 \\ 0 & (-3)^{2000} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 3^{2000} \end{pmatrix}$$

since diagonal matrices can be multiplied and raised to power “entrywise” (unlike general matrices). Hence

$$A^{2000}\mathbf{w} = VD^{2000}V^{-1}\mathbf{w} = \begin{pmatrix} -5 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 3^{2000} \end{pmatrix} \begin{pmatrix} -1/4 & -1/4 \\ -1/4 & -5/4 \end{pmatrix} \begin{pmatrix} 3 \\ 1 \end{pmatrix} = \begin{pmatrix} 5 - 2 \cdot 3^{2000} \\ -1 + 2 \cdot 3^{2000} \end{pmatrix}$$

7.2.2 Spectral theorem: diagonalization of real symmetric matrices

Finally, we discuss the case of symmetric matrices, $A = A^t$. First notice that if V is orthogonal, then the conjugated matrix VAV^t is also symmetric (CHECK (*)). This means that the property whether a matrix is symmetric or not is actually independent of the orthonormal basis, hence this is really a property of the underlying linear map. It is always good to rely on properties which are independent of the chosen basis.

The following theorem is VERY IMPORTANT:

Theorem 7.10 [*Spectral Theorem for real symmetric matrices*] *Let A be an $n \times n$ symmetric matrix with real entries. Then it has n real eigenvalues (with possible multiplicity) and one*

can choose an orthonormal eigenbasis $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$. Using this basis to form the matrix $V = [\mathbf{v}_1 \ \mathbf{v}_2 \ \dots \ \mathbf{v}_n]$, the matrix A is diagonalizable as

$$A = VDV^t \tag{7.5}$$

where the diagonal matrix D contains the eigenvalues in the diagonal:

$$D = \begin{pmatrix} \lambda_1 & 0 & 0 & \dots & 0 \\ 0 & \lambda_2 & 0 & \dots & 0 \\ 0 & 0 & \lambda_3 & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \dots & \lambda_n \end{pmatrix} \tag{7.6}$$

We will prove this theorem after mentioning some VERY IMPORTANT consequences.

REMARK: Sometimes the following form of this theorem is useful. For any vector \mathbf{u}

$$A\mathbf{u} = \sum_{i=1}^n \lambda_i (\mathbf{v}_i^t \cdot \mathbf{u}) \mathbf{v}_i = \left(\sum_{i=1}^n \lambda_i \mathbf{v}_i \mathbf{v}_i^t \right) \mathbf{u} \tag{7.7}$$

or

$$A = \sum_{i=1}^n \lambda_i \mathbf{v}_i \mathbf{v}_i^t$$

Recall that the matrix $\mathbf{v}_i \mathbf{v}_i^t$ is the *orthogonal projection* onto the vector \mathbf{v}_i (See Section 6). In words: the action of a real symmetric matrix is that it projects the given vector into the eigendirections, scales the projected vectors by the corresponding eigenvalues, then adds up these scaled vectors.

MAKE SURE YOU TRULY UNDERSTAND IT. This is one of the most important tool whenever linear algebra is applied.

In particular it enables us to define *functions* of matrices. We can take integer powers of matrices (also negative powers, if A^{-1} exists). But what is e^A , or does $\sin(A)$ make sense? Recall that e^A was important (Section 1.2 Example II). The following is the right definition (sometimes this is called the spectral theorem).

Definition 7.11 Let A be a real symmetric $n \times n$ matrix and $f : \mathbf{R} \rightarrow \mathbf{R}$ a function. Then $f(A)$ is also an $n \times n$ matrix defined by its action on any vector \mathbf{u} as

$$f(A)\mathbf{u} = A\mathbf{u} = \sum_{i=1}^n f(\lambda_i)(\mathbf{v}_i^t \cdot \mathbf{u})\mathbf{v}_i$$

In other words

$$f(A) = \sum_{i=1}^n f(\lambda_i)\mathbf{v}_i\mathbf{v}_i^t$$

or, again in other words

$$f(A) = Vf(D)V^t$$

if $A = VD V^t$ is the diagonalization of A (see (7.5)), where $f(D)$ for a diagonal matrix D is defined as

$$f(D) = \begin{pmatrix} f(\lambda_1) & 0 & 0 & \dots & 0 \\ 0 & f(\lambda_2) & 0 & \dots & 0 \\ 0 & 0 & f(\lambda_3) & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \dots & f(\lambda_n) \end{pmatrix}$$

In short, you can take any function of a real symmetric matrix by taking the function of the eigenvalues and use these numbers $f(\lambda_i)$ as scaling factors on the eigenspaces. Moreover, the new matrix $f(A)$ is again symmetric (CHECK(*)), it has the same eigenvectors as A and eigenvalues $f(\lambda_1), f(\lambda_2), \dots, f(\lambda_n)$, where λ_i are the eigenvalues of A .

One of the most important case is when $f(A) = A^{-1}$ (inverse matrix). In particular, the eigenvalues of the inverse matrix are the inverses of the eigenvalues of A (for regular A).

You can also check the *consistency* of this definition for powers of A . We know how to compute

$$A^2 = AA = \left(\sum_{i=1}^n \lambda_i \mathbf{v}_i \mathbf{v}_i^t \right) \left(\sum_{i=1}^n \lambda_i \mathbf{v}_i \mathbf{v}_i^t \right) \tag{7.8}$$

by simple matrix multiplication. Now Definition 7.11 gives another expression for $A^2 = f(A)$ with the function $f(x) = x^2$:

$$A^2 = \sum_{i=1}^n \lambda_i^2 \mathbf{v}_i \mathbf{v}_i^t \tag{7.9}$$

It would be quite confusing if (7.8) and (7.9) were not be the same. CHECK (*) that they are.

More generally you can check that the usual arithmetic operations with functions extend to matrices. For example $f(A)g(A) = (fg)(A)$. For this you have to check (DO IT(*)) that

$$\left(\sum_{i=1}^n f(\lambda_i) \mathbf{v}_i \mathbf{v}_i^t \right) \left(\sum_{i=1}^n g(\lambda_i) \mathbf{v}_i \mathbf{v}_i^t \right) = \sum_{i=1}^n f(\lambda_i) g(\lambda_i) \mathbf{v}_i \mathbf{v}_i^t$$

In fact Definition 7.11 also works for any diagonalizable matrix; if $A = VDV^{-1}$, then one can define $f(A) = Vf(D)V^{-1}$ (notice that here V is not necessarily orthogonal, hence we have to use V^{-1} instead of V^t). The trouble is that it is not so easy to decide about a matrix A whether it is diagonalizable without actually trying to diagonalize it. But the symmetry of a matrix is very easy to recognize.

Proof of Theorem 7.10. We know that the matrix has n eigenvalues. First we show that the eigenvalues are real. Suppose $A\mathbf{v} = \lambda\mathbf{v}$. Take the scalar product of this equation with $\overline{\mathbf{v}}$ (Recall: the overbar denotes complex conjugation for a complex number. The conjugation of a vector is just the entrywise conjugation of the elements). On the right hand side you get

$$\lambda \overline{\mathbf{v}}^t \cdot \mathbf{v} = \lambda (|v_1|^2 + |v_2|^2 + \dots) = \lambda \|\mathbf{v}\|^2$$

On the left hand side we have $\overline{\mathbf{v}}^t \cdot A\mathbf{v} = (A^t \overline{\mathbf{v}})^t \cdot \mathbf{v}$ (recall basic matrix multiplication). But $A^t = A$, hence we get

$$(A \overline{\mathbf{v}})^t \cdot \mathbf{v} = \lambda \|\mathbf{v}\|^2$$

Now take the conjugate of the eigenvalue equation $A\mathbf{v} = \lambda\mathbf{v}$. Since A has real entries, you don't have to worry about conjugating them, and you get

$$A\bar{\mathbf{v}} = \bar{\lambda}\bar{\mathbf{v}}.$$

Plug this into the previous equation, you get

$$\bar{\lambda}\|\mathbf{v}\|^2 = \lambda\|\mathbf{v}\|^2$$

i.e. $\bar{\lambda} = \lambda$, which exactly means that λ is real.

The proof that one can choose an orthonormal eigenbasis consisting of real vectors would go beyond the scope of this note for the general case, but in a remark after the proof we give the idea. In any case, we emphasize that the theorem is true without any assumption on the multiplicity. In particular the trouble with the Example $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ in Section 7.2.1 is excluded simply by the symmetry of the matrix A .

So we restrict our attention to the simpler case, when we assume, in addition, that the eigenvalues are simple, as in Theorem 7.5. Then clearly the eigenvectors are real, since they are in the nullspace of $A - \lambda I$, which is a real matrix.

So in this case all we have to show that the eigenvectors are pairwise orthogonal. Then one can easily normalize them to get an orthonormal set. So let

$$A\mathbf{v} = \lambda\mathbf{v}$$

$$A\mathbf{v}' = \lambda'\mathbf{v}'$$

be two eigenvalues-eigenvectors, and we assume that $\lambda \neq \lambda'$ (eigenvalues are distinct). Compute the scalar product $\mathbf{v}^t \cdot A\mathbf{v}'$ in two different ways. On one hand

$$\mathbf{v}^t \cdot A\mathbf{v}' = \lambda'(\mathbf{v}^t \cdot \mathbf{v}')$$

On the other hand

$$\mathbf{v}^t \cdot A\mathbf{v}' = (A^t\mathbf{v})^t \cdot \mathbf{v}' = (A\mathbf{v})^t \cdot \mathbf{v}' = \lambda(\mathbf{v}^t \cdot \mathbf{v}')$$

Since $\lambda \neq \lambda'$, we obtain that $\mathbf{v}^t \cdot \mathbf{v}' = 0$, which we wanted to show. This completes the proof of the spectral theorem in the case of simple eigenvalues. \square

REMARK: The idea behind the proof for the general case is similar. In nutshell: let \mathbf{v} an eigenvector corresponding to the eigenvalue λ . Consider the subspace S in \mathbf{R}^n of all vectors that are orthogonal to \mathbf{v} (sometimes we denote it by $S = \mathbf{v}^\perp$). (CHECK (*) that subspace!) Of course the eigendirection given by \mathbf{v} is invariant under the multiplication by A , i.e. $A\mathbf{v}$ is in the same direction as \mathbf{v} (just the definition of the eigenvector). The good news is that if A is symmetric, then the orthogonal complement S is also invariant under A . In other words, if $\mathbf{u} \in S$, i.e. $\mathbf{u}^t \cdot \mathbf{v} = 0$, then $A\mathbf{u} \in S$ as well (CHECK (*) that $(A\mathbf{u})^t \cdot \mathbf{v} = 0$). Choose an orthonormal basis $\{\mathbf{u}_1, \dots, \mathbf{u}_{n-1}\}$ in S , and write up the linear map given by the matrix of A in the basis $\{\mathbf{v}, \mathbf{u}_1, \dots, \mathbf{u}_{n-1}\}$ (in other words, conjugate A by the matrix U containing these basis vectors in the columns). Check that this is an orthonormal basis of \mathbf{R}^n , i.e. U is orthogonal. Check that map in this basis looks like

$$\begin{pmatrix} \lambda & 0 & 0 & \dots & 0 \\ 0 & * & * & \dots & * \\ 0 & * & * & \dots & * \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & * & * & \dots & * \end{pmatrix}$$

where $*$ could be anything. (i.e. check that U^tAU has this form). Consider the $(n-1) \times (n-1)$ matrix \tilde{A} in the lower right corner. Show that it is symmetric as well. Now you can run the same machinery as before for the matrix \tilde{A} , and get successively the diagonal form of A .

Problem 7.12 Find the spectral decomposition of the matrix $A = \begin{pmatrix} 3 & -1 \\ -1 & 3 \end{pmatrix}$.

SOLUTION: The solution goes exactly in the same way as in Problem 7.6. The eigenvalues are $\lambda_1 = 2$, $\lambda_2 = 4$, eigenvectors $\mathbf{v}_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$, $\mathbf{v}_2 = \begin{pmatrix} -1 \\ 1 \end{pmatrix}$. These are orthogonal, but not yet normalized. So we replace them by

$$\mathbf{v}_1 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad \text{and} \quad \mathbf{v}_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ 1 \end{pmatrix}$$

Hence

$$V = \begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}$$

Notice that V is orthogonal, hence its inverse is just the transpose

$$V^{-1} = V^t = \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}$$

The diagonal matrix is of course

$$D = \begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix}$$

and the decomposition is

$$A = \begin{pmatrix} 3 & -1 \\ -1 & 3 \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}$$

Problem 7.13 Find the spectral decomposition of the matrix

$$A = \begin{pmatrix} 2 & 2 & 1 \\ 2 & -1 & -2 \\ 1 & -2 & 2 \end{pmatrix}$$

SOLUTION: The characteristic polynomial is

$$p(\lambda) = \det(\lambda - A) = \lambda^3 - 3\lambda^2 - 9\lambda + 27$$

It easily factorizes:

$$\lambda^3 - 3\lambda^2 - 9\lambda + 27 = \lambda^2(\lambda - 3) - 9(\lambda - 3) = (\lambda + 3)(\lambda - 3)^2$$

The eigenvalue $\lambda_1 = -3$ is single, the eigenvector \mathbf{v}_1 solves $A\mathbf{v}_1 = -3\mathbf{v}_1$, and one easily gets

$$\mathbf{v}_1 = \begin{pmatrix} -1 \\ 2 \\ 1 \end{pmatrix}$$

The other eigenvalue is double: $\lambda_2 = \lambda_3 = 3$. Hence we need to find an orthonormal basis in the solution space $A\mathbf{v} = 3\mathbf{v}$. Since A is symmetric, we know that this space is exactly two dimensional (as many as we need - recall that without symmetry it could be false). So we need a basis in the nullspace of the matrix $A - 3I$, then we can orthonormalize it. After row elimination on $A - 3I$ we get

$$A - 3I = \begin{pmatrix} -1 & 2 & 1 \\ 2 & -4 & -2 \\ 1 & -2 & -1 \end{pmatrix} \implies \begin{pmatrix} -1 & 2 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

and

$$N(A - 3I) = \left\{ \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 2 \\ 1 \\ 0 \end{pmatrix} \right\}$$

The vectors

$$\mathbf{v}_2 = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}, \quad \mathbf{v}_3 = \begin{pmatrix} 2 \\ 1 \\ 0 \end{pmatrix}$$

are clearly linearly independent, hence the matrix of eigenvectors is

$$V = \begin{pmatrix} -1 & 1 & 2 \\ 2 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}$$

and $A = VDV^{-1}$ with

$$D = \begin{pmatrix} -3 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 3 \end{pmatrix}$$

This is the diagonalization of A , but the eigenvectors are not orthonormal. Recall that A is symmetric, and eigenvectors belonging to different eigenvalues are orthogonal. Notice that

\mathbf{v}_1 is really orthogonal to $\mathbf{v}_2, \mathbf{v}_3$. Although \mathbf{v}_1 is not normalized, it is easy to normalize it and use

$$\mathbf{q}_1 = \frac{\mathbf{v}_1}{\|\mathbf{v}_1\|} = \frac{1}{\sqrt{6}} \begin{pmatrix} -1 \\ 2 \\ 1 \end{pmatrix}$$

instead of \mathbf{v}_1 as eigenvector for λ_1 .

But $\mathbf{v}_2, \mathbf{v}_3$ are not orthogonal yet. This is because we essentially chose an arbitrary (or most convenient) basis in $N(A - 3I)$. So we have to run Gram-Schmidt for $\mathbf{v}_2, \mathbf{v}_3$:

$$\mathbf{q}_2 := \frac{\mathbf{v}_2}{\|\mathbf{v}_2\|} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}$$

and

$$\mathbf{w}_3 = \mathbf{v}_3 - (\mathbf{v}_3 \cdot \mathbf{q}_2)\mathbf{q}_2 = \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}$$

hence

$$\mathbf{q}_3 = \frac{\mathbf{w}_3}{\|\mathbf{w}_3\|} = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}$$

Now $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3$ is an orthonormal eigenbasis, and we can form the matrix

$$Q = \begin{pmatrix} -1/\sqrt{6} & 1/\sqrt{2} & 1/\sqrt{3} \\ 2/\sqrt{6} & 0 & 1/\sqrt{3} \\ 1/\sqrt{6} & 1/\sqrt{2} & -1/\sqrt{3} \end{pmatrix}$$

One can check that $A = QDQ^t$ (notice that the D matrix is the same, the eigenvalues did not change, only the eigenvectors), and this is the spectral decomposition of A .

Notice that the eigensubspaces belonging to different eigenvalues are orthogonal (i.e. the one dimensional subspace spanned by \mathbf{v}_1 is orthogonal to the two dimensional subspace spanned by $\{\mathbf{v}_2, \mathbf{v}_3\}$), but the basis within each individual eigensubspace might not be orthogonal. When we learned how to find a basis in the nullspace of a matrix (Problem 4.7), we did not care about orthogonality. Hence first you have to find a basis in each nullspace

$N(A - \lambda_j I)$ (which is the same as the eigenspace belonging to the eigenvalue λ_j), then you have to orthonormalize this basis within each eigenspace by the Gram-Schmidt procedure.

7.3 Application of diagonalization to solve linear systems of differential equations

Recall Example II. from Section 1.2. The goal is to solve a linear system of ordinary differential equations written as

$$\mathbf{x}'(t) = A\mathbf{x}(t) \tag{7.10}$$

where A is a given $n \times n$ matrix and $\mathbf{x}(t)$ is an unknown vector-valued function, depending on a single variable t .

In the scalar case, $x'(t) = ax(t)$, where $a \in \mathbf{R}$ and $x(t)$ is just a scalar valued function, the solution is given as $x(t) = e^{at}x(0)$, where $x(0)$ is usually given (initial condition). Following this analogy, it is tempting to present the solution to (7.10) as

$$\mathbf{x}(t) = e^{At}\mathbf{x}(0)$$

and this is indeed correct if e^{At} is defined properly.

Suppose that A can be diagonalized, i.e. $A = VDV^{-1}$ (this is automatically true for symmetric A , but it is also true for many other matrices). Then we define

$$e^{At} := Ve^{Dt}V^{-1}$$

where e^{Dt} is a diagonal matrix with $e^{\lambda_j t}$ in the diagonal (here, as usual, λ_j 's are the eigenvalues, i.e. the diagonal entries of D). This definition is of course the same as given in the Spectral Theorem for symmetric matrices.

DELICATE REMARK: In fact, we could have defined $f(A)$ for arbitrary diagonalizable matrix as $f(A) = Vf(D)V^{-1}$ and not just for symmetric ones. The only trouble is that in the

general case the eigenvalues may be complex, so when forming the diagonal matrix $f(D)$ with $f(\lambda_1), f(\lambda_2), \dots$ in the diagonal, we should make sure that f is defined for complex numbers. Not every function given originally on the real line can be extended to the complex plane, but polynomials, rational functions, exponential functions etc. are OK. (The truth is that every function is OK which has convergent power series, these are called *analytic* functions.)

Now we can check that

$$\mathbf{x}(t) := e^{At}\mathbf{x}(0) = Ve^{Dt}V^{-1}\mathbf{x}(0)$$

indeed solves (7.10). Simply compute

$$\begin{aligned} \frac{d}{dt}\mathbf{x}(t) &= \frac{d}{dt}Ve^{Dt}V^{-1}\mathbf{x}(0) = \lim_{s \rightarrow t} \frac{Ve^{Ds}V^{-1}\mathbf{x}(0) - Ve^{Dt}V^{-1}\mathbf{x}(0)}{s - t} \\ &= V \left[\lim_{s \rightarrow t} \frac{e^{Ds} - e^{Dt}}{s - t} \right] V^{-1}\mathbf{x}(0) = VDe^{Dt}V^{-1}\mathbf{x}(0) \end{aligned}$$

Here we used that $e^{Ds} - e^{Dt}$ is a diagonal matrix with $e^{\lambda_j s} - e^{\lambda_j t}$ in the diagonal, and clearly

$$\lim_{s \rightarrow t} \frac{e^{\lambda_j s} - e^{\lambda_j t}}{s - t} = \lambda_j e^{\lambda_j t}$$

(recall the derivative of the $t \mapsto e^{\lambda t}$ function as a limit of the difference quotient). Hence

$$\left[\lim_{s \rightarrow t} \frac{e^{Ds} - e^{Dt}}{s - t} \right]$$

is a diagonal matrix with $\lambda_j e^{\lambda_j t}$ in the diagonal, which is exactly De^{Dt} .

Now we can use again the spectral theorem (for symmetric matrices), or the definition $f(A) = Vf(D)V^{-1}$ mentioned in the Delicate Remark above, for the function $f(x) = xe^{xt}$ to see that

$$VDe^{Dt}V^{-1}\mathbf{x}(0) = Ae^{At}\mathbf{x}(0) = A\mathbf{x}(t),$$

which completes the proof that $\mathbf{x}(t) = e^{At}\mathbf{x}(0)$ satisfies the equation (7.10).

Hence if you want to solve (7.10), first diagonalize the matrix A , i.e. find the $A = VDV^{-1}$ decomposition, then compute e^{At} as $Ve^{Dt}V^{-1}$, and this matrix acting on the initial condition vector $\mathbf{x}(0)$ will give the solution.

This method always works for symmetric matrices and it works for most general square matrices. If A is not diagonalizable, then one has to use the *Jordan canonical decomposition* instead, which was mentioned at the end of Section 7.1.

7.4 Singular value decomposition (SVD)

What is the “nice” form of a general $n \times k$ matrix A ? Since this matrix acts between two different spaces, you can choose a “good” basis in both. We know that good “basis” means orthonormal. The following theorem tells that this can always be achieved:

Theorem 7.14 (Singular value decomposition (SVD)) *Let A be an $n \times k$ matrix with $n \geq k$. Then there exist a matrix U of dimensions $n \times k$ and a matrix V dimensions $k \times k$ such that $U^tU = I$, $V^tV = I$ and*

$$A = UDV^t$$

where D is a diagonal $k \times k$ matrix with nonnegative decreasing diagonal entries $\sigma_1 \geq \sigma_2 \geq \dots \sigma_k \geq 0$. These numbers are called the **singular values** of A , while the columns $\{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k\}$ of U are called the **left singular vectors** and the columns of $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\}$ of V are the **right singular vectors**. The right singular vectors are in \mathbf{R}^n , the left ones are in \mathbf{R}^k . The decomposition in this form is unique.

Here is another form of this theorem, which is more similar to the form (7.7) of the spectral theorem:

$$A = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^t$$

or, as its action on a vector \mathbf{w} :

$$A\mathbf{w} = \sum_{i=1}^k \sigma_i (\mathbf{v}_i^t \cdot \mathbf{w}) \mathbf{u}_i$$

In plain words: For any matrix A one can choose orthonormal bases $\{\mathbf{u}\}$ and $\{\mathbf{v}\}$ and there are numbers σ_i such that the action of A on a given vector \mathbf{w} is the following: it finds the component of \mathbf{w} in the direction of the i -th right singular vector, multiplies this number by the singular value σ_i and takes this number as the i -th coordinate of the image vector $A\mathbf{w}$ in the $\{\mathbf{u}\}$ basis.

It is clear that SVD is a generalization of the Spectral Theorem 7.10.

We are not going to prove that SVD exists. But assuming it does, it is easy to find it from the Spectral theorem.

The key observation is to take the matrices AA^t and A^tA . It is clear that both of them are symmetric matrices (CHECK (*)). Moreover, using $A = UDV^t$ we see that

$$AA^t = UDV^tVDU^t = UD^2U^t$$

and

$$A^tA = VDU^tUDV^t = VD^2V^t$$

Hence the singular values of A can be obtained as the (positive) square roots of the eigenvalues of A^tA (or, which is the same, the eigenvalues of AA^t). The right singular vectors (columns of V) are the eigenvectors of A^tA , the left singular vectors are the eigenvectors of AA^t . Hence finding SVD reduces to diagonalizing the symmetric matrices AA^t and A^tA , which problem we have solved in Problem 7.12.

A minor care is needed, since the eigenvectors of a matrix are not unique. Even for simple eigenvalues, the corresponding normalized eigenvectors have a two-fold ambiguity: their sign. It is clear from the desired representation $A = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^t$ that one cannot flip say $\mathbf{u}_2 \rightarrow -\mathbf{u}_2$

(you can flip $\mathbf{u}_2 \rightarrow -\mathbf{u}_2$ and $\mathbf{v}_2 \rightarrow -\mathbf{v}_2$ simultaneously, then the above sum is unchanged). So one can choose one set of eigenvectors, say the right ones, $\{\mathbf{v}_i\}$, arbitrarily, but then the signs of the left vectors $\{\mathbf{u}_i\}$ are determined, and the proper choice may not be the one which you get just by arbitrarily choosing the eigenvectors of $A^t A$.

The way out is very simple, in fact it simplifies the whole algorithm. Suppose that we fixed the $\{\mathbf{v}_i\}$ vectors as arbitrary normalized eigenvectors of $A^t A$. Then, assuming the decomposition $A = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^t$, we compute

$$A\mathbf{v}_j = \sum_{i=1}^k \sigma_i \mathbf{u}_i (\mathbf{v}_i^t \cdot \mathbf{v}_j)$$

By orthonormality of the $\{\mathbf{v}_i\}$ vectors, only the $i = j$ term is nonzero in the sum, and we get

$$A\mathbf{v}_j = \sigma_j \mathbf{u}_j$$

i.e.

$$\mathbf{u}_j = \frac{1}{\sigma_j} A\mathbf{v}_j$$

or $U = AVD^{-1}$ (if σ_j 's are nonzero). This will be the right choice for the left singular vectors, given the right ones. If $\sigma_j = 0$, then the corresponding vectors can be chosen arbitrarily, they do not contribute to the SVD anyway.

If you happen to know the left vectors $\{\mathbf{u}_i\}$ and you want to find the right ones, then compute

$$\mathbf{u}_j^t A = \sum_{i=1}^k \sigma_i (\mathbf{u}_j^t \cdot \mathbf{u}_i) \mathbf{v}_i^t = \sigma_j \mathbf{v}_j^t$$

i.e.

$$\mathbf{v}_j = \frac{1}{\sigma_j} A^t \mathbf{u}_j \tag{7.11}$$

This latter formula is clearly the same as $V = A^t U D^{-1}$ which is clear from $A = U D V^t$.

Problem 7.15 Find the singular value decomposition of

$$A = \begin{pmatrix} 4/3 & 1/3 \\ -4/3 & 2/3 \\ 2/3 & 2/3 \end{pmatrix}$$

SOLUTION: By

$$AA^t = UDV^tVDU^t = UD^2U^t$$

we can find the the singular values of A which are the (positive) square roots of the eigenvalues of AA^t and the left singular vectors which are the normalized eigenvectors of AA^t with respect to the positive eigenvalues.

Now we compute AA^t and get

$$AA^t = \begin{pmatrix} 4/3 & 1/3 \\ -4/3 & 2/3 \\ 2/3 & 2/3 \end{pmatrix} \begin{pmatrix} 4/3 & 1/3 \\ -4/3 & 2/3 \\ 2/3 & 2/3 \end{pmatrix}^t = \begin{pmatrix} 17/9 & -14/9 & 10/9 \\ -14/9 & 20/9 & -4/9 \\ 10/9 & -4/9 & 8/9 \end{pmatrix}$$

The positive eigenvalues of AA^t are $\lambda_1 = 4$, $\lambda_2 = 1$ and the normalized eigenvectors are $\mathbf{u}_1 = \begin{pmatrix} -2/3 \\ 2/3 \\ -1/3 \end{pmatrix}$ and $\mathbf{u}_2 = \begin{pmatrix} -1/3 \\ -2/3 \\ -2/3 \end{pmatrix}$. Hence

$$U = \begin{pmatrix} -2/3 & -1/3 \\ 2/3 & -2/3 \\ -1/3 & -2/3 \end{pmatrix}$$

and

$$D = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}.$$

To find the right singular vectors, one can either run the same machinery for A^tA and then choose carefully the sign of the eigenvectors, or one can use the formula (7.11). This formula immediately gives

$$\mathbf{v}_1 = \frac{1}{2} \begin{pmatrix} 4/3 & -4/3 & 2/3 \\ 1/3 & 2/3 & 2/3 \end{pmatrix} \begin{pmatrix} -2/3 \\ 2/3 \\ -1/3 \end{pmatrix} = \begin{pmatrix} -1 \\ 0 \end{pmatrix}$$

$$\mathbf{v}_2 = \frac{1}{1} \begin{pmatrix} 4/3 & -4/3 & 2/3 \\ 1/3 & 2/3 & 2/3 \end{pmatrix} \begin{pmatrix} -1/3 \\ -2/3 \\ -2/3 \end{pmatrix} = \begin{pmatrix} 0 \\ -1 \end{pmatrix}$$

Alternatively, you could have find the eigenvectors of $A^t A$ directly since $A^t A = VDU^tUDV^t = VD^2V^t$. So we compute $A^t A$ and get

$$A^t A = \begin{pmatrix} \frac{4}{3} & \frac{1}{3} \\ -\frac{4}{3} & \frac{2}{3} \\ \frac{2}{3} & \frac{2}{3} \end{pmatrix}^t \begin{pmatrix} \frac{4}{3} & \frac{1}{3} \\ -\frac{4}{3} & \frac{2}{3} \\ \frac{2}{3} & \frac{2}{3} \end{pmatrix} = \begin{pmatrix} 4 & 0 \\ 0 & 1 \end{pmatrix}.$$

The positive eigenvalues of $A^t A$ are $\lambda_3 = 4$, $\lambda_4 = 1$ (check that we can also get the same D from the positive square roots of the eigenvalues of $A^t A$), and the normalized eigenvectors are $\mathbf{v}_1 = \begin{pmatrix} -1 \\ 0 \end{pmatrix}$ and $\mathbf{v}_2 = \begin{pmatrix} 0 \\ -1 \end{pmatrix}$. Hence

$$V = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Here we made a specific choice of the $\mathbf{v}_1, \mathbf{v}_2$ vectors, which turn out to be the right one, since

$$\begin{aligned} UDV^t &= \begin{pmatrix} -\frac{2}{3} & -\frac{1}{3} \\ \frac{2}{3} & -\frac{2}{3} \\ -\frac{1}{3} & -\frac{2}{3} \end{pmatrix} \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}^t = \begin{pmatrix} -\frac{2}{3} & -\frac{1}{3} \\ \frac{2}{3} & -\frac{2}{3} \\ -\frac{1}{3} & -\frac{2}{3} \end{pmatrix} \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix} \\ &= \begin{pmatrix} 4/3 & 1/3 \\ -4/3 & 2/3 \\ 2/3 & 2/3 \end{pmatrix} \end{aligned}$$

which is A .

But if you had chosen $\mathbf{v}_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$, $\mathbf{v}_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, then you would have gotten $-A$ instead of A . So using the formula (7.11) is always safer.

7.4.1 Application of SVD in image compression

[This section follows the presentation of the book Applied Numerical Linear Algebra by James Demmel, (SIAM, 1997), page 114-116.]

A picture on the computer screen can be encoded by a matrix A . If the screen has $n \times k$ pixels (say $n = 320$ rows and $k = 200$ columns), and the brightness of the pixel is expressed by a number from 0 to 1, then one can easily define the matrix A of a picture just by putting the brightness of the (i, j) -th pixel into the (i, j) -th entry of A (colors are also possible).

A big matrix A contains lots of information, it is expensive to store it and manipulate with it. How to reduce the size but still keep the relevant majority of the information? You could take every second pixel into account, for example. But this is not very good. The picture has “boring” parts (e.g. uniform background) where every tenth pixel would be enough to keep and still you could decode the picture quite easily just by approximation on the nearby pixels. On the other hand, on the “interesting” part of the picture (sharp contours etc.) you would like to keep all pixels. So we have to “teach” the computer to distinguish between “boring” and “interesting”.

There are very well developed methods for this, and we discuss a very “baby”-version of the main idea.

Take the SVD of A in the following form

$$A = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^t$$

(suppose $n \geq k$). To store the full U and V matrices are at least as expensive as storing the original A (A has nk entries, U and V together have $nk + k^2$ entries, plus you have to record the singular values, i.e. the SVD requires $(n + k + 1)k$ numbers to store.) But the idea is that the small singular values and their singular vectors probably do not matter too much. These belong to the “boring” part of the picture (more precisely, these express the small deviations of the true picture from the “boring” homogeneous approximation).

Fix a number $m \leq k$ and define the following matrix:

$$A_m = \sum_{i=1}^m \sigma_i \mathbf{u}_i \mathbf{v}_i^t$$

i.e. you truncate the SVD by retaining the m biggest singular values and throwing away the rest (with their singular vectors). You can store A_m in $m(n + k + 1)$ slots. If $m \ll k$, then it is much less memory space than storing A .

But how good is it? Of course A_m is not A , but it is close to it. In fact in some sense it is the closest possible matrix with that little memory requirement. The following theorem makes it more precise:

Theorem 7.16 *The matrix A_m is the best rank- m approximation to A in the sense that it minimizes the distance from A to the subspace of all matrices of rank at most m .*

Of course, we have not defined what it means that a matrix is “close” to another matrix. There are several ways to define it, and later on, when we discuss numerical methods we will be more precise. For the purpose of this theorem we just use the following definition:

Definition 7.17 *The distance between two matrices $A = (a_{ij})$ and $B = (b_{ij})$ of the same size is given by the number*

$$\sqrt{\sum_{i=1}^n \sum_{j=1}^k (a_{ij} - b_{ij})^2}$$

Back to the image compression, it turns out that if you store a usual picture of size 320×200 by a generally rank $k = 200$ matrix, and if you take a rank $m = 20$ approximation of it, i.e. you use only $A_{m=20}$ instead of A , then the picture is already fairly good, figures, faces etc. can be recognized. And the storage was 10 times smaller!

There is one more advantage of storing a picture in the SVD form. When you download a picture from the internet, it could take quite a long time. In many cases after having seen even a part of the picture, you can already decide that this is not what you wanted and click on the next page. Standard programs download the picture line by line, i.e. you see all the

details of the top of the picture (in many cases only the sky) before you see anything from the essential part. More sophisticated programs download a picture by some coarse-graining procedure, i.e. you immediately see the whole picture, but only a very coarsened version, and later, as the downloading goes on, the details are refined. The sequence of matrices A_1, A_2, \dots does exactly this. In fact, if you have already transmitted say A_{20} , then transmitting A_{21} is easy; you just transmit the vectors \mathbf{u}_{21} and \mathbf{v}_{21} and the number σ_{21} , and the receiver “adds” the matrix $\sigma_{21}\mathbf{u}_{21}\mathbf{v}_{21}^t$ to the existing picture.