Graph Theory: CMSC 27530/37530 Lecture 13

Lecture by László Babai
Notes by Geoffrey West
Revised by instructor

May 14, 2019

Email the instructor a list of challenge problems submitted, except $\alpha(C_7 * C_7)$.

MULTISET, SPECTRUM

Definition 13.1. A **multiset** is a set where we allow an element to appear multiple times. To distinguish it from a set, we often use double braces, for example,

$$\{\{2, 2, 2, 3, 7, 7, 7, 7\}\}.$$

In this case we say that the element 2 has multiplicity 3, the element 3 has multiplicity 1, and the element 7 has multiplicity 4. Two multisets are equal if they have the same entries and each entry has the same multiplicity (order doesn't matter).

Example 13.2. A monic polynomial is characterized by the multiset of its complex roots.

Example 13.3. For $A \in M_n(\mathbb{C})$, the **spectrum** of A is the multiset of its eigenvalues. For example, the spectrum of the identity matrix is $\{\{\underbrace{1, 1, ..., 1}_{n \text{ times}}\}\}$.

Exercise 13.4. Consider a previous homework problem, to calculate the determinant of the matrix

$$A = \begin{pmatrix} a & & & & \\ & a & & b & \\ & & a & & \\ & b & & \ddots & \\ & & & & a \end{pmatrix}$$

First solution. We use basic determinat operations. For two columns $\mathbf{a}_i, \mathbf{a}_j$ $(j \neq i)$, we can replace column $\mathbf{a}_i$ with $\mathbf{a}_i + \beta \mathbf{a}_j$ without affecting the determinant. Let us add each of the first $n-1$ columns to the last column, obtaining the following.

$$\det A = \det \begin{pmatrix} a & b & \dots & a+(n-1)b \\ b & a & \dots & a+(n-1)b \\ \vdots & \vdots & \ddots & \vdots \\ b & b & \dots & a+(n-1)b \end{pmatrix} = (a+(n-1)b)\cdot \det \begin{pmatrix} a & b & \dots & 1 \\ b & a & \dots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ b & b & \dots & 1 \end{pmatrix}$$
$$= (a+(n-1)b)\cdot \det \begin{pmatrix} a-b & 0 & \dots & 1 \\ 0 & a-b & \dots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{pmatrix}$$

The last equality is by subtracting b times the last column from each of the first $n-1$ columns. Since now the matrix has triangular form, its determinant is the product of its diagonal elements. We obtain

$$\det A = (a+(n-1)b)\cdot (a-b)^{n-1}. \tag{1}$$

□

HW 13.5. (4+2 points) Determine the characteristic polynomial and the spectrum of the adjacency matrix of K_n and $\overline{K_n}$.

Next we determine the spectrum of the all-ones matrix and use it for a more elegant solution of Exercise 13.4.

Recall that an *eigenbasis* of a matrix $A \in M_n(\mathbb{R})$ (or $A \in M_n(\mathbb{C})$) is a basis of $\mathbb{R}^n$ (or alternatively $\mathbb{C}^n$) consisting of eigenvectors of A.

Definition 13.6. The **rank** of a matrix is the maximum number of linearly independent columns. We write $\mathrm{rk}(A)$ to denote this number.

Definition 13.7. The **kernel** of a matrix $A \in \mathbb{R}^{k\times\ell}$ is the set $\{\mathbf{x} \in \mathbb{R}^\ell \mid A\mathbf{x} = \mathbf{0}\}$, denoted $\ker(A)$.

Definition 13.8. The **nullity** of a matrix A is the dimension of $\ker(A)$, denoted $\mathrm{nullity}(A)$.

Theorem 13.9 (Rank–Nullity)**.** *For a matrix $A \in \mathbb{R}^{k\times\ell}$ we have* $\mathrm{rk}(A) + \mathrm{nullity}(A) = \ell$.

Consider the $n\times n$ all-ones matrix $J = (j_{ik})$ having each entry $j_{ik} = 1$. Clearly $\mathrm{rk}(J) = 1$, and by Rank–Nullity it follows that $\mathrm{nullity}(J) = n-1$. The eigenvalues of J are easy to find: any non-zero element of $\ker(J)$ is an eigenvector having eigenvalue equal to 0, so we already have $n-1$ of the eigenvalues (all zero), only one is missing. The missing eigenvalue can be calculated from the trace, which is n, so the spectruem of J is $\mathrm{spec}(J) = \{\{0^{n-1}, n\}\}$.

Instead of employing the trace trick, we could also have noticed that the all-ones vector $\mathbf{x} = \begin{pmatrix}1 & 1 & \dots & 1\end{pmatrix}^T$ is an eigenvector with eigenvalue n.

Next we try to find an eigenbasis of J.

The elements $\mathbf{x} = (x_1, \ldots, x_n)$ of the kernel are the nonzero vectors satifying $J\mathbf{x} = \mathbf{0}$, or, in other words, $\sum x_i = 0$. Here are $n-1$ linearly independent vectors satisfying this condition.

$$\begin{aligned}
\mathbf{a}_1 &= \begin{pmatrix} 1 & -1 & 0 & \ldots & 0 \end{pmatrix}^T \\
\mathbf{a}_2 &= \begin{pmatrix} 1 & 0 & -1 & \ldots & 0 \end{pmatrix}^T \\
&\vdots \\
\mathbf{a}_{n-1} &= \begin{pmatrix} 1 & 0 & 0 & \ldots & -1 \end{pmatrix}^T .
\end{aligned}$$

These vectors form a basis of $\ker(J)$ (because they are linearly independent and their number is right.) The vectors $\mathbf{a}_1, ..., \mathbf{a}_{n-1}$ together with $\mathbf{x}$ constitute an eigenbasis to J.

DO 13.10. Let $A \in M_n(\mathbb{R})$ and let $B = (\mathbf{b}_1, ..., \mathbf{b}_n)$ be an eigenbasis of A with corresponding eigenvalues $\lambda_1, \ldots, \lambda_n$, so $A\mathbf{b}_i = \lambda_i \mathbf{b}_i$. Let $\alpha, \beta \in \mathbb{R}$.

1. B is an eigenbasis of αA with corresponding eigenvalues $\alpha\lambda_i$.
2. B is an eigenbasis of $A + \beta I$ with corresponding eigenvalues $\lambda_i + \beta$.
3. B is an eigenbasis of $\alpha A + \beta I$ with corresponding eigenvalues $\alpha\lambda_i + \beta$.

So we find that if A is diagonalizable and $\operatorname{spec}(A) = \{\{\lambda_1, \ldots, \lambda_n\}\}$ then

$$\operatorname{spec}(\alpha A + \beta I) = \{\{\alpha\lambda_1 + \beta, \ldots, \alpha\lambda_n + \beta\}\}. \tag{2}$$

The same holds over $\mathbb{C}$.

BONUS 13.11. (5 points) (Due Tuesday) Prove that over $\mathbb{C}$, equation (2) always holds, regardless of whether A is diagonalizable. You may use the result, to be explained later, that over $\mathbf{c}$, every square matrix is similar to a triangular matrix.

Second solution to Exercise 13.4. Consider again the matrix A discussed in Exercise 13.4. Notice that $A = bJ + (a-b)I$. Using the result that J is diagonalizable and $\operatorname{spec}(J) = \{\{0, ..., 0, n\}\}$, together with Eq. 2, we obtain that

$$\operatorname{spec}(A) = \{\{a-b, \ldots, a-b, a+(n-1)b\}\}. \tag{3}$$

Now the determinant is the product of the eigenvalues, so we again proved Eq. (1), this time without any tedious, ad hoc calculations. □

SIMILARITY OF MATRICES

Recall the multiplicativity of the determinants: if $A, B \in M_n(\mathbb{R})$ then

$$\det(AB) = \det(A) \cdot \det(B). \tag{4}$$

Recall that a square matrix A is invertible if and only if $\det(A) \neq 0$.

DO 13.12.

$$\det(A^{-1}) = \frac{1}{\det(A)}.$$

Definition 13.13. $A, B \in M_n(\mathbb{R})$ are **similar** if $(\exists S, S^{-1})(B = S^{-1}AS)$. In this case we write $A \sim B$.

Remark 13.14. The similarity relation gains intuitive understanding from the following idea: two matrices are similar if they represent the same linear transformation in different bases. The matrix S describes the change of basis.

DO 13.15. Similarity is an equivalence relation on $M_n(\mathbb{R})$.

DO 13.16. $A \sim B \implies \text{trace}(A) = \text{trace}(B)$.

Proof. We use the previous exercise that for matrices $K, L \in M_n(\mathbb{R})$ we have

$$\text{trace}(KL) = \text{trace}(LK).$$

Apply this with $K = S^{-1}$ and $L = AS$:

$$\text{trace}(B) = \text{trace}(S^{-1} \cdot AS) = \text{trace}(AS \cdot S^{-1}) = \text{trace}(A).$$

□

DO 13.17. $A \sim B \implies \det(A) = \det(B)$.

DO 13.18. $A \sim B \implies f_A = f_B$.
Hint. The matrix $tI - A$ is called the **characteristic matrix** of A. So the characteristic polynomial is the determinant of the characteristic matrix. Show that

$$\text{if } A \sim B \text{ then } (tI - A) \sim (tI - B).$$

Now apply DO 13.17.

Remark 13.19. We say that the characteristic polynomial is an **invariant** of the similarity relation, meaning that if $A \sim B$ then $f_A = f_B$. It is not a **complete invariant**, meaning that the converse does not hold: the characteristic polynomial does not characterize the similarity class of a matrix (see next exercise). The Jordan normal form (with appropriately ordered diagonal blocks) is a complete invariant over $\mathbb{C}$.

HW 13.20. (5 points) Find two matrices $A, B \in M_2(\mathbb{R})$ such that $f_A = f_B$ but $A \not\sim B$. Prove that they are not similar.

DO 13.21. $f_A(t) = t^n - \text{trace}(A)t^{n-1} + \ldots + (-1)^n \det(A)$.

Three of the statements we have made can be summarized as follows:

(1) $A \sim B \implies \text{trace}(A) = \text{trace}(B)$

(2) $A \sim B \implies \det A = \det B$

(3) $A \sim B \implies f_A = f_B$.

DO 13.22. Show that statement (3) implies (1) and (2).

Notation 13.23. We write $B = [\mathbf{b}_1, ..., \mathbf{b}_t]$ to express that $\mathbf{b}_1, ..., \mathbf{b}_t$ are the columns of the matrix B.

DO 13.24. Let $A \in \mathbb{R}^{r\times s}$ and $B \in \mathbb{R}^{s\times t}$. If $B = [\mathbf{b}_1, ..., \mathbf{b}_t]$, then $AB = [A\mathbf{b}_1, ..., A\mathbf{b}_t]$.

DO 13.25. If $A = [\mathbf{a}_1, ..., \mathbf{a}_s]$ and $\mathbf{x} = (x_1, \ldots, x_s)^T$ then then $A\mathbf{x} = \sum x_i \mathbf{a}_i$.

DO 13.26. Let $\mathbf{e}_i$ denote the i-th standard basis vector, i. e., the i-th column of the identity matrix. Then

$$A\mathbf{e}_i = \mathbf{a}_i, \quad \text{the } i\text{-th column of } A.$$

Proof. Immediate from DO 13.25. □

Definition 13.27. A matrix $D \in M_n(\mathbb{R})$ is **diagonal** if it has the form

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

We abbreviate using the notation $D = \operatorname{diag}(\lambda_1, \lambda_2, ..., \lambda_n)$.

DO 13.28. If $A = [\mathbf{a}_1, ..., \mathbf{a}_s]$ and $D = \operatorname{diag}(\lambda_1, \lambda_2, ..., \lambda_n)$, then $AD = [\lambda_1\mathbf{a}_1, ..., \lambda_s\mathbf{a}_s]$.

Recall our definition of diaginalizability.

Definition 13.29. We say that A is **diagonalizable** if it admits an eigenbasis.

The following result explains this terminology.

Theorem 13.30. *A is diagonalizable $\iff$ A is similar to a diagonal matrix.*

DO 13.31. If $A \sim \operatorname{diag}(\lambda_1, ..., \lambda_n)$ then $\{\{\lambda_1, ..., \lambda_n\}\} = \operatorname{spec}(A)$.

Proof. $A \sim D \implies f_A = f_D = \prod(t - \lambda_i)$. □

Definition 13.32. A matrix $A \in M_n(\mathbb{R})$ is called **nonsingular** if any of the following equivalent conditions hold.

(a) $\det A \neq 0$

(b) $\operatorname{rk} A = n$

(c) $\exists A^{-1}$.

DO 13.33. Prove that these conditions are equivalent.

Proof of Theorem 13.30. We prove the "only if" direction. Suppose A has an eigenbasis, $(\mathbf{b}_1, \ldots, \mathbf{b}_n)$, such that $A\mathbf{b}_i = \lambda_i\mathbf{b}_i$. Let $B = [\mathbf{b}_1, ..., \mathbf{b}_n]$. Then $AB = [\lambda_1\mathbf{b}_1, ..., \lambda_n\mathbf{b}_n] = B \cdot \underbrace{\operatorname{diag}(\lambda_1, ..., \lambda_n)}_{D}$. (Here we used DO exercises 13.25 and 13.28.) So $AB = BD$ and therefore $B^{-1}AB = B^{-1}BD = D$. (Why is B invertible?) □

DO 13.34. Prove the converse: if A is similar to a diagonal matrix, then A has an eigenbasis.

HW 13.35. (5 points) If all eigenvalues of A are distinct over $\mathbb{C}$ (i. e., A has no multiple roots), then A is diagonalizable over $\mathbb{C}$.

HW 13.36. (4 points) Show that the converse is false: find a 2×2 matrix which is diagonalizable but has a multiple eigenvalue. Find all such matrices.

MIRACLES OF LINEAR ALGEBRA

I can't cease to view some of the basic facts of Linear Algebra with wonderment. I call the fact that $\dim(\mathbb{R}^n) = n$ the *First Miracle of linear algebra.* It expresses the impossibility of boosting linear independence. $\mathbb{R}^n$ is spanned by n vectors, and try as hard as you might, you will never be able to find more than n linearly independent vectors in $\mathbb{R}^n$. This feature distinguishes linear algebra from much of algebra. For instance, the free group with 2 generators contains a free group with infinitely many generators: two "independent" elements generate infinitely many.

I view the fact that $\text{rk}(A) = \text{rk}(A^T)$ the *Second Miracle.* This is true even if A is not a square matrix. Rows and columns live in different worlds. What do linear combinations of rows have to do with linear combinations of columns? Yet, the maximum number of linearly independent rows and columns is the same.

Next we draw some inferences from the first two miracles, in preparation for the third one.

Definition 13.37. Let $A \in \mathbb{R}^{k\times\ell}$ and $B, C \in \mathbb{R}^{\ell\times k}$. We say that B is a **right inverse** of A if $AB = I_k$ (the $k \times k$ identity matrix). We say that C is a **left inverse** of A if $CA = I_\ell$ (the $\ell \times \ell$ identity matrix).

DO 13.38. A matrix $A \in \mathbb{R}^{k\times\ell}$ has a right inverse if and only if $\text{rk}(A) = k$. This means its rows are linearly independent. In this case we say that A has *full row rank.*

DO 13.39. A has a left inverse if and only if $\text{rk}(A) = \ell$, i. e., its columns are linearly independent. In this case we say that A has *full column rank.*

A corollary of the last two results is that only square matrices can have inverses on both sides. A remarkable fact following from the first two miracles is that, for a square matrix, if it has either inverse, then it has *both*!

And these two inverses are necessarily equal.

DO 13.40. If B is a left inverse and C is a right inverse of the matrix A then $B = C$.

Proof. We have $BA = I$ and $AC = I$. Therefore

$$B = BI = B(AC) = (BA)C = IC = C.$$

Note that this proof works in any semigroup with identity. □

DO 13.41. Prove: if $A \in M_n(\mathbb{R})$ has a right inverse then this right inverse is unique.

The next exercise summarizes what we have found.

DO 13.42. For an $n \times n$ matrix the following are equivalent:

(1) $\mathrm{rk}(A) = n$ (A has full rank)

(2) A has a left inverse

(3) A has a right inverse

(4) A has an inverse (2-sided)

Moreover, in this case the inverse is unique, and there is no left or right inverse other than this unique inverse.

DO 13.43. Let $k \neq \ell$. Prove: if $A \in \mathbb{R}^{k\times\ell}$ has a right inverse then it has infinitely many right inverses.

Of course the same can be said about left inverses.

ORTHOGONAL MATRICES, THE THIRD MIRACLE

Definition 13.44. A matrix $A \in M_n(\mathbb{R})$ is an **orthogonal matrix** if the columns of A form an ONB of $\mathbb{R}^n$.

Contrary to most definitions before, this one does not work over $\mathbb{C}$, it is essential that our scalars are real numbers.

Example 13.45. The rotation matrix $R_\theta = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}$ is an orthogonal matrix.

Example 13.46. The reflection matrix $F_\theta = \begin{pmatrix} \cos\theta & \sin\theta \\ \sin\theta & -\cos\theta \end{pmatrix}$ is an orthogonal matrix.

BONUS 13.47. (6 points) (Due Tuesday) These two classes of examples comprise all orthogonal 2×2 matrices.

HW 13.48. (5 points) $F_\theta \sim \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ for all $\theta \in \mathbb{R}$.

The rotation matrix R_θ has eigenvalues $e^{\pm i\theta}$. What happens to the complex plane if we multiply every number by $\mathrm{e}^{i\theta}$, i. e., what is the transformation $z \mapsto \mathrm{e}^{i\theta}z$ ($z \in \mathbb{C}$)? The plane rotates by angle θ. We started from geometry (rotation), turned it into algebra (matrix, characteristic polynomial, eigenvalues), and the output gives us back the geoimetric transformation we started from.

Theorem 13.49. *A is orthogonal (its columns form an ONB) $\iff$ its rows form an ONB.*

I like to call this the *Third Miracle* of Linear Algebra.

Ortonormality of the columns is a set of equations on the elements of the matrix. Ortonormality of the rows is another set of equations. The equations for the columns combine entirely different sets of variables than those for the rows. What do the two sets of equations have to do with each other? Well, they are equivalent, and we shall derive this fact from the Second Miracle.

DO 13.50. A is orthogonal if and only if $A^TA = I$.

Proof. If $A = [\mathbf{a}_1, ..., \mathbf{a}_n]$, what is the (ij)th entry of A^TA? The answer is $\mathbf{a}_i^T\mathbf{a}_j$. Therefore, $A^TA = I \iff (\forall i, j)(\mathbf{a}_i^T\mathbf{a}_j = \delta_{ij}) \iff$ the $\mathbf{a}_k$ form an ONB. □

So $A \in M_n(\mathbb{R})$ is orthogonal if and only if A^T is the left inverse of A. But then, by exercise DO 13.42, we have the following corollaries.

DO 13.51. A is orthogonal if and only if A^T is orthogonal.

Proof. A is orthogonal if A^T is the left inverse of A, and A^T is orthogonal if A^T is the right inverse of A. But these two conditions are equivalent. □

DO 13.52. Notice that exercise DO 13.51 is just a restatement of the Third Miracle (Theorem 13.49), and completes its proof.

DO 13.53. A is orthogonal if and only if $A^T = A^{-1}$.

HW 13.54. (4 points) Prove: if A is an orthogonal matrix then $\det(A) = \pm 1$.

HW 13.55. (4 points) If A, B are orthogonal matrices, then AB is an orthogonal matrix.

HW 13.56. (4 points) If $A \in M_n(\mathbb{R})$ is an orthogonal matrix then $(\forall \mathbf{x} \in \mathbb{R}^n)(\|A\mathbf{x}\| = \|x\|)$.

BONUS 13.57. (5 points) Prove the converse: if $(\forall \mathbf{x} \in \mathbb{R}^n)(\|A\mathbf{x}\| = \|x\|)$ then A is orthogonal.

Remark 13.58. The last two exercises say that orthogonal matrices represent precisely those linear transformations that preserve norm and therefore, distance: they are those congruences of $\mathbb{R}^n$ that fix the origin.

ORTHOGONAL SIMILARITY. THE SPECTRAL THEOREM RESTATED

Definition 13.59. Let $A, B \in M_n(\mathbb{R})$. We say that A and B are **orthogonally similar** if there exists and orthogonal matrix S such that $B = S^{-1}AS$. We say that A is **orthogonally diagonalizable** if A is orthogonally similar to a diagonal matrix.

DO 13.60. Prove the following statement is equivalent to the Spectral Theorem.
Every symmetric real matrix is orthogonally diagonalizable.