Graph Theory

Linear algebra background

Instructor: *Nicolò Cesa-Bianchi* version of April 8, 2024

The material in this handout is taken from: Luca Trevisan, Lecture Notes on Graph Partitioning, Expanders and Spectral Methods, 2016.

Given a real $n \times n$ matrix M, if $M\boldsymbol{u} = \lambda \boldsymbol{u}$ for some $\lambda \in \mathbb{R}$ and $\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\}$, then $\boldsymbol{u}$ is an eigenvector of M with eigenvalue λ (we also say that $\boldsymbol{u}$ is an eigenvector of λ). Note that eigenvectors can be rescaled without changing the equation $M\boldsymbol{u} = \lambda \boldsymbol{u}$, hence we conventionally assume they have unit length.

Note that λ is an eigenvalue for M if and only if there exists $\boldsymbol{x} \neq \boldsymbol{0}$ such that $(M - \lambda I)\boldsymbol{x} = \boldsymbol{0}$, where I is the $n \times n$ identity matrix. The equation $(M - \lambda I)\boldsymbol{x} = \boldsymbol{0}$ holds for $\boldsymbol{x} \neq \boldsymbol{0}$ if and only if $M - \lambda I$ is singular, which is equivalent to $\det(M - \lambda I) = 0$. Since $\det(M - \lambda I)$ is a n-th degree univariate polynomial in λ, it has exactly n solutions by the fundamental theorem of algebra. This shows that every square matrix has n eigenvalues (not all necessarily distinct). Some of these eigenvalues, however, may correspond to solutions of $\det(M - \lambda I) = 0$ in the complex plane. The next result guarantees that at least one eigenvalue is real when M is symmetric.

Fact 1 (proof omitted) *If M is symmetric, then there exists $\lambda \in \mathbb{R}$ and $\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\}$ such that $M\boldsymbol{u} = \lambda \boldsymbol{u}$.*

Fact 2 *If M is symmetric then any two eigenvectors corresponding to distinct eigenvalues are orthogonal.*

Proof. Let $\boldsymbol{x}$ be an eigenvector of λ and $\boldsymbol{y}$ an eigenvector of λ' with $\lambda \neq \lambda'$. Since M is symmetric, $(M\boldsymbol{x})^\top \boldsymbol{y} = \boldsymbol{x}^\top M \boldsymbol{y}$. On the other hand, $(M\boldsymbol{x})^\top \boldsymbol{y} = \lambda \boldsymbol{x}^\top \boldsymbol{y}$ and $\boldsymbol{x}^\top M \boldsymbol{y} = \lambda' \boldsymbol{x}^\top \boldsymbol{y}$. Since $\lambda \neq \lambda'$, it must by $\boldsymbol{x}^\top \boldsymbol{y} = 0$, which means that $\boldsymbol{x}$ and $\boldsymbol{y}$ are orthogonal. □

Orthogonal projections. If $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_k \in \mathbb{R}^n$ is a set of $k \leq n$ orthonormal vectors in $\mathbb{R}^n$, then $A = \big[\boldsymbol{u}_1, \ldots, \boldsymbol{u}_k\big]$ maps any vector $\boldsymbol{x} \in \mathbb{R}^k$ onto the k-dimensional subspace $V \subseteq \mathbb{R}^n$ spanned by $\big\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_k\big\}$, so that $\|\boldsymbol{x}\| = \|A\boldsymbol{x}\|$. Instead, $AA^\top$ projects any vector in $\mathbb{R}^n$ onto $V \subseteq \mathbb{R}^n$. In particular, $AA^\top \boldsymbol{x} = \boldsymbol{x}$ for all $\boldsymbol{x} \in V$. Because the columns of A are orthonormal, we also have $A^\top A = I$ (the $k \times k$ identity matrix). In the special case $k = n$, A performs a change of basis from $\big\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_n\big\}$ to the canonical basis $\big\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n\big\}$, whereas $A^\top$ performs the inverse transformation. Moreover, $AA^\top = A^\top A = I$

We now use Fact 1 to prove that any symmetric matrix has n (not necessarily distinct) real eigenvalues.

Theorem 3 (Spectral Theorem) *Let $M \in \mathbb{R}^{n \times n}$ be a symmetric matrix. Then there exists n (not necessarily distinct) real numbers $\lambda_1, \ldots, \lambda_n$ and n orthonormal real vectors $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_n$ such that $\boldsymbol{u}_i$ is an eigenvector of λ_i.*

PROOF. The proof is by induction on n. If $n = 1$, then M is a scalar. Hence, any nonzero $x \in \mathbb{R}$ is an eigenvector of M with eigenvalue M because $Mx = Mx$.

Assume now that the statement holds for $n-1$. By Fact 1, there exist an eigenvalue $\lambda_n \in \mathbb{R}$ with eigenvector $\boldsymbol{x}_n \in \mathbb{R}^n$.

Claim. $\boldsymbol{y}$ orthogonal to $\boldsymbol{x}_n$ implies $M\boldsymbol{y}$ is orthogonal to $\boldsymbol{x}_n$.

Indeed, $\boldsymbol{x}_n^\top M\boldsymbol{y} = (M\boldsymbol{x}_n)^\top \boldsymbol{y} = \lambda_n \boldsymbol{x}_n^\top \boldsymbol{y} = 0$.

Now let V be the $(n-1)$-dimensional subspace of $\mathbb{R}^n$ that contains all the vectors orthogonal to $\boldsymbol{x}_n$. Now choose an orthonormal basis $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{n-1}$ for V and let $B = \big[\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{n-1}\big]$. By construction, $BB^\top$ projects $\mathbb{R}^n$ onto V. In particular, $BB^\top \boldsymbol{z} = \boldsymbol{z}$ for all $\boldsymbol{z} \in V$. We now apply the inductive hypothesis to the $(n-1) \times (n-1)$ symmetric matrix $M' = B^\top MB$ and find real eigenvalues $\lambda_1, \ldots, \lambda_{n-1}$ and orthonormal eigenvectors $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_{n-1} \in \mathbb{R}^{n-1}$. For $i = 1, \ldots, n-1$ we have $M'\boldsymbol{y}_i = B^\top MB\boldsymbol{y}_i = \lambda_i \boldsymbol{y}_i$. Therefore, $BB^\top MB\boldsymbol{y}_i = \lambda_i B\boldsymbol{y}_i$. Since $\boldsymbol{y}_i \in \mathbb{R}^{n-1}$ and B maps $\mathbb{R}^{n-1}$ to V, $B\boldsymbol{y}_i$ is orthogonal to $\boldsymbol{x}_n$ and, by the above claim, $MB\boldsymbol{y}_i$ is orthogonal to $\boldsymbol{x}_n$ and so $MB\boldsymbol{y}_i \in V$. Therefore $\lambda_i B\boldsymbol{y}_i = BB^\top MB\boldsymbol{y}_i = MB\boldsymbol{y}_i$. If we now define $\boldsymbol{x}_i = B\boldsymbol{y}_i$ for $i = 1, \ldots, n-1$, then we have $M\boldsymbol{x}_i = \lambda_i \boldsymbol{x}_i$. To finish up, note that, by construction, $\boldsymbol{x}_n$ is orthogonal to $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{n-1}$. Moreover, for any $1 \le i < j \le n-1$, $\boldsymbol{x}_i^\top \boldsymbol{x}_j = (B\boldsymbol{y}_i)^\top (B\boldsymbol{y}_j) = \boldsymbol{y}_i^\top B^\top B\boldsymbol{y}_j = \boldsymbol{y}_i^\top \boldsymbol{y}_j = 0$. Hence we have found n eigenvalues with n eigenvectors. □

Corollary 4 *Let $M \in \mathbb{R}^{n \times n}$ be a real symmetric matrix. Then*

$$M = U\Lambda U^\top = \sum_{i=1}^{n} \lambda_i \boldsymbol{u}_i \boldsymbol{u}_i^\top$$

where $U = \big[\boldsymbol{u}_1, \ldots, \boldsymbol{u}_n\big]$ and $\Lambda = \operatorname{diag}\big(\lambda_1, \ldots, \lambda_n\big)$. Here $\lambda_1 \le \lambda_2 \le \cdots \le \lambda_n$ are the real eigenvalues of M and $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_n \in \mathbb{R}^n$ are the corresponding eigenvectors.

PROOF. Note that $MU = \big[\lambda_1 \boldsymbol{u}_1, \ldots, \lambda_n \boldsymbol{u}_n\big]$ because $M\boldsymbol{u}_i = \lambda_i \boldsymbol{u}_i$ for each $i = 1, \ldots, n$. Hence $MU = U\Lambda$ where $\Lambda = \operatorname{diag}\big(\lambda_1, \ldots, \lambda_n\big)$. Since U is a square orthonormal matrix, $U^\top U = UU^\top = I$. Therefore $M = MUU^\top = U\Lambda U^\top$. □

Let $M \in \mathbb{R}^{n \times n}$ be a symmetric matrix. For any $\boldsymbol{x} \in \mathbb{R}^n$, the **Rayleigh quotient** is the ratio

$$R(M, \boldsymbol{x}) = \frac{\boldsymbol{x}^\top M\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}}$$

Theorem 5 (Variational characterization of eigenvalues — proof omitted) *Let $M \in \mathbb{R}^{n \times n}$ be a symmetric matrix, and $\lambda_1 \le \lambda_2 \le \cdots \le \lambda_n$ be its real eigenvalues. For $k < n$ let $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_k$ be orthonormal vectors such that $M\boldsymbol{u}_i = \lambda_i \boldsymbol{u}_i$ for $i = 1, \ldots, k$. Then*

$$\lambda_{k+1} = \min_{\substack{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\} \\ \boldsymbol{u} \perp \{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_k\}}} R(M, \boldsymbol{u})$$

and any minimizer $\boldsymbol{u}$ is an eigenvector of λ_{k+1}.

In particular,

$$\lambda_1 = \min_{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\mathbf{0}\}} R(M, \boldsymbol{u})$$

Also, because $-M$ has eigenvalues $-\lambda_n \leq -\lambda_{n-1} \leq \cdots \leq -\lambda_1$,

$$-\lambda_n = \min_{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\mathbf{0}\}} R(-M, \boldsymbol{u}) = -\max_{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\mathbf{0}\}} R(M, \boldsymbol{u})$$

and therefore

$$\lambda_n = \max_{\boldsymbol{u} \in \mathbb{R}^n \setminus \{\mathbf{0}\}} R(M, \boldsymbol{u})$$

A symmetric matrix M is **positive semidefinite** if $\boldsymbol{x}^\top M \boldsymbol{x} \geq 0$ for all $\boldsymbol{x} \in \mathbb{R}^n$.

Fact 6 *The eigenvalues of a positive semidefinite matrix are all nonnegative.*

PROOF. As the denominator of the Rayleigh quotient is clearly always positive, Theorem 5 implies that the sign of each eigenvalue is determined by the sign of $\boldsymbol{x}^\top M \boldsymbol{x}$. $\square$

We conclude with a different, but equally important characterization of eigenvalues.

Theorem 7 (Courant-Fischer — proof omitted) *Let M be a symmetric matrix with real eigenvalues $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$. Then*

$$\lambda_k = \min_{S \,:\, \dim(S)=k} \; \max_{\boldsymbol{u} \in S \setminus \{\mathbf{0}\}} R(M, \boldsymbol{u}) \qquad k = 1, \ldots, n$$

where the mininum is over all subspaces $S \subseteq \mathbb{R}^n$ of dimension k.