Subspace distance and perturbation theory for symmetric matrices

Instructor: Anil Damle

In lecture we briefly discussed some perturbation theory for the eigenvalues and eigenvectors of symmetric matrices. As part of this discussion we also formalized how to talk about distances between subspaces.

Notation and assumptions

Throughout these notes we will assume that $A \in \mathbb{R}^{n\times n}$ and $A = A^T$. There are extensions of this theory to case where A is not symmetric, but we will not cover those here. We will denote the eigenvalues and vectors of a as $\lambda_1, \lambda_2, \ldots, \lambda_n$ and $v_1, v_2, \ldots, v_n$ respectively, and assume that the eigenvalues are ordered such that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$. $A = V\Lambda V^T$ denotes the spectral decomposition of A.

Subspace distance

When we talk about computing eigenvectors associated with simple eigenvalues, what we actually concerned with is computing the one-dimensional invariant subspace that eigenvector spans. Notably, this alleviates issues like the fact that even if we require $\|v\| = 1$ eigenvectors are not unique. Similarly, we are often interested with computing an ℓ-dimensional invariant subspace associated with ℓ eigenvalues of A. As such, we need a way to reason about how far apart two subspaces are.

Let $W \in \mathbb{R}^{n\times\ell}$ and $U \in \mathbb{R}^{n\times\ell}$ be matrices with orthonormal columns representing two subspaces of interest. We can define the distance between the range of W and the range of U as

$$\mathrm{dist}(W, U) = \|WW^T - UU^T\|_2, \tag{1}$$

where we have slightly abused notation to let W and V also represent their respective ranges. We remark that $0 \leq \mathrm{dist}(W, U) \leq 1$, with the distance being zero if the subspaces are the same and 1 if there exist vectors in W and U that are orthogonal (i.e., there exists some x and y such that $(Wx)^T(Uy) = 0$).

We should expect our notation of distance to be invariant to the specific orthonormal basis we choose to represent a subspace and this definition satisfies that condition. For any two orthogonal matrices Q_1 and Q_2 we see that

$$\|WQ_1Q_1^TW^T - UQ_2Q_2^TU^T\|_2 = \|WW^T - UU^T\|_2.$$

More generally, the basis independence follows from the fact that the orthogonal projector onto a subspace is unique.

In the case where $\ell = 1$ this reduces to

$$\mathrm{dist}(w, u) = \sqrt{1 - (w^Tu)^2},$$

where we have switched to a lower case w and u to highlight that they are just vectors. Since w and u are normalized we can use $w^Tu = \cos(\theta)$ to express the distance between the subspaces as $\mathrm{dist}(w, u) = \sin(\theta)$, where θ represents the angle between the subspaces.

In practice, if we want to compute the distance between two subspaces naively using () directly is unnecessarily expensive. Fortunately, we also have that

$$\mathrm{dist}(W, U) = \sqrt{1 - \sigma_{\min}(W^TV)^2}.$$

Eigenvalue perturbation bounds

It is useful to know how much we can change the eigenvalues of a matrix A by perturbing it by some "small" matrix E. For symmetric matrices, we can actually show that the most a single eigenvalue can change is bounded the size of perturbation. Specifically, given $A = A^T$ and $E = E^T$

$$|\lambda_i(A+E) - \lambda_i(A)| \leq \|E\|_2,$$

for $i = 1, 2, \ldots, n$. This is a simplification Weyl's inequality.

Notably, this is the best we could hope for—loosely speaking the conditioning of each eigenvalue is one. In contrast, for non-symmetric matrices the story can be much more complicated and for certain matrices there are eigenvalues that are way more sensitive to perturbations.

Eigenvalue perturbation bounds

Complementing the prior result about eigenvalues, we may also want to characterize how much the invariant supspaces of A can change when it is perturbed. Let i be the index of a simple eigenvalue of A and let $\gamma = \min(|\lambda_i - \lambda_{i+1}|, |\lambda_i - \lambda_{i-1}|)$ denote the gap between λ_i and the next closest eigenvalue.[1] In this setting, if $\|E\|_2 \leq \gamma/5$ then

$$\text{dist}(v_i, \hat{v}_i) \leq \frac{\|E\|_2}{\gamma}, \tag{2}$$

where $\hat{v}_i$ is an eigenvector associated with $\lambda_i(A+E)$. This is a simplified version of the Davis-Kahan Theorem.

The conditions under which (2) hold ensure that associating v_i and $\hat{v}_i$ is sensible. Since γ represents a gap between the eigenvalue of interest and others, we see that more well separated eigenvalues have more stable invariant subspaces for a fixed size perturbation (and we can characterize their behavior for larger perturbations).

[1] If $i = 1$ then $\gamma = \lambda_1 - \lambda_2$ and if $i = n$ then $\gamma = \lambda_{n-1} - \lambda_n$.