

18.303: The Min–Max/Variational Theorem and the Rayleigh Quotient

S. G. Johnson

October 5, 2011

Elementary linear-algebra classes always cover the fact that Hermitian matrices $A = A^*$ have real eigenvalues, eigenvectors that can be chosen orthonormal, and are diagonalizable. In 18.303, we learned that self-adjoint operators \hat{A} also have real values and orthonormal eigenfunctions, and are usually diagonalizable in the sense of a spectral theorem. However, these Hermitian eigenproblems (finite- or infinite-dimensional) have another important property that you may not have learned: they obey a **min–max theorem** or **variational principle**, which tells us that the eigensolutions *minimize or maximize* a certain quantity. This is useful in lots of ways, most immediately because it gives us an intuitive way to “guess” the qualitative features of the eigensolutions.

The Rayleigh Quotient

In particular, suppose we are looking at a (Sobolev) space V of functions (or vectors) u , with some inner product $\langle u, v \rangle$ under which $\hat{A} = \hat{A}^*$. Furthermore assume that \hat{A} is diagonalizable, i.e. we have a complete basis of eigenfunctions u_n satisfying $\hat{A}u_n = \lambda_n u_n$ and $\langle u_n, u_m \rangle = \begin{cases} 1 & n = m \\ 0 & n \neq m \end{cases}$. Consider the **Rayleigh quotient** for *any* $u \in V$ (not necessarily an eigenfunction):

$$R\{u\} = \frac{\langle u, \hat{A}u \rangle}{\langle u, u \rangle}.$$

We use $\{\dots\}$ brackets to denote that R is a *functional*: it maps functions u to numbers. Note that $R\{u\}$ is real because $\hat{A} = \hat{A}^*$: $\overline{R\{u\}} = \overline{\langle u, \hat{A}u \rangle / \langle u, u \rangle} = \langle \hat{A}u, u \rangle / \langle u, u \rangle = \langle u, \hat{A}u \rangle / \langle u, u \rangle = R\{u\}$. Note that if u is an eigenfunction u_n , then $R\{u_n\} = \lambda_n$.

The Min–Max Theorem

Now, the **min–max theorem** (also called the *variational principle*) tells us a remarkable fact: R is bounded above and below by the minimum and maximum λ (if any: λ may be unbounded in one direction). That is:

$$\lambda_{\min} \text{ (if any)} \leq R\{u\} \leq \lambda_{\max} \text{ (if any)}$$

for *any* $u \in V$ with $u \neq 0$. Equality is achieved if and only if $u = u_{\min/\max}$. That is, λ_{\min} and λ_{\max} and the corresponding eigenfunctions *minimize or maximize the Rayleigh quotient*, respectively:

$$\begin{aligned} \lambda_{\min} \text{ (if any)} &= \min_{u \neq 0 \in V} R\{u\}, \\ \lambda_{\max} \text{ (if any)} &= \max_{u \neq 0 \in V} R\{u\}, \end{aligned}$$

with the minima achieved at the eigenfunctions. (In fact, it is possible to go further and to show that *all* eigenvalues are *extrema* of R .)

The proof is very simple if we can assume a complete basis of eigenfunctions (although this is a nontrivial assumption for infinite-dimensional problems). In particular, write $u = \sum_n c_n u_n$ for coefficients $c_n = \langle u_n, u \rangle$. Then

$$\begin{aligned} R\{u\} &= \frac{\left\langle \sum_m c_m u_m, \hat{A} \sum_n c_n u_n \right\rangle}{\left\langle \sum_{m'} c_{m'} u_{m'}, \hat{A} \sum_{n'} c_{n'} u_{n'} \right\rangle} = \frac{\sum_{m,n} \overline{c_m} c_n \lambda_n \langle u_m, u_n \rangle}{\sum_{m',n'} \overline{c_{m'}} c_{n'} \langle u_{m'}, u_{n'} \rangle} \\ &= \frac{\sum_n |c_n|^2 \lambda_n}{\sum_{n'} |c_{n'}|^2}, \end{aligned}$$

which is just a weighted average of the λ 's and hence is bounded above and below by the maximum and minimum λ (if any):

$$\lambda_{\min} = \frac{\sum_n |c_n|^2 \lambda_{\min}}{\sum_{n'} |c_{n'}|^2} \leq \frac{\sum_n |c_n|^2 \lambda_n}{\sum_{n'} |c_{n'}|^2} \leq \frac{\sum_n |c_n|^2 \lambda_{\max}}{\sum_{n'} |c_{n'}|^2} = \lambda_{\max}.$$

Deflation

Suppose that we have an ascending sequence of eigenvalues $\lambda_1 < \lambda_2 < \lambda_3 < \dots$. We know that λ_1 minimizes $R\{u\}$, but what about λ_2 and so on? We can get at these higher eigenvalues by a process called **deflation**, where we “remove eigenvalues” from \hat{A} one by one.

Consider the orthogonal complement u_1^\perp : the subspace of all $u \in V$ with $\langle u_1, u \rangle = 0$. This modifies the argument in the previous section because, in this space, $c_1 = 0$ and the λ_1 term is missing from $R\{u\}$. That means that the minimum of $R\{u\}$ for $u \in u_1^\perp$ is λ_2 ! Similarly, to get λ_3 we should look at the space $\{u_1, u_2\}^\perp$: all u orthogonal to u_1 and u_2 (which removes λ_1 and λ_2 from R), and so on. Thus:

$$\begin{aligned} \lambda_1 &= \min_{u \neq 0 \in V} R\{u\}, \\ \lambda_2 &= \min_{u \neq 0 \in u_1^\perp} R\{u\}, \\ \lambda_3 &= \min_{u \neq 0 \in \{u_1, u_2\}^\perp} R\{u\}, \\ &\vdots \end{aligned}$$

where the minima are achieved at $u = u_1, u_2, u_3, \dots$

Not too much changes if we have a repeated eigenvalue, say $\lambda_2 = \lambda_3$. In that case, the second minimum will be achieved for two linearly independent u , giving u_2 and u_3 (or we can simply pick one as u_2 , and then minimizing $\perp u_1, u_2$ will give u_3 with the same λ).

Example: Guessing eigenfunctions of $-\nabla^2$

Consider the operator $\hat{A} = -\nabla^2$ on some Ω , with $u|_{\partial\Omega} = 0$. Integrating by parts, we find

$$R\{u\} = \frac{\int_\Omega |\nabla u|^2}{\int_\Omega |u|^2}.$$

Therefore, u_1 is the function that *oscillates as little as possible* (in the sense of minimizing its mean $|\nabla u|^2$). It has to oscillate a little, because it has to go back to zero at the boundaries, so it will be some function that starts at zero at the boundaries, goes up slowly to some maximum in the middle of Ω , and then goes back down to zero at the opposite boundaries. For example, just like $\sin(\pi x/L)$ in 1d!

Similarly, u_2 also “wants” to oscillate as little as possible, but must be $\perp u_1$. Therefore, it is *forced* to go through zero roughly where u_1 peaks, in order that $\langle u_1, u_2 \rangle = 0$. For example, like $\sin(2\pi x/L)$ in 1d. u_3 also wants to oscillate as little as possible, but must have additional oscillations to be \perp to u_1 and u_2 , and so on.

In a multidimensional domain like a rectangle, the u_{2+} functions have a “choice” about which direction they want to oscillate along in order to be \perp to the previous eigenfunctions.




Figure 1: Contour plots of the eigenfunctions for the smallest three eigenvalues λ of the $-\nabla^2$ operator on a triangular half of a square (cut diagonally), with Dirichlet boundary conditions.

Since they want to oscillate as little as possible (minimal R), they will *oscillate along the “long axes” of Ω first*.

This allows us to easily “guess” the qualitative features of the smallest- $|\lambda|$ eigenfunctions even for complicated domains Ω . For example, consider the triangular domain Ω in figure 1. The lowest λ solution should therefore just have a single peak in the center—it has to be zero at the boundaries, and must be nonzero somewhere, so the slowest it can oscillate is to go up to a peak in the center and then back down, as in figure 1(left). The second u must be orthogonal to the first, so it must flip sign, i.e. have two peaks of opposite sign, but in which direction? This triangle is *not equilateral*¹—it is longer in the $(-1, 1)$ direction (parallel to the diagonal of the square), so that is the direction in which a sign oscillation can occur most slowly. Hence, we should expect a single $+/-$ oscillation along this direction, as in figure 1(middle). The third u must be orthogonal to the first two, and the slowest oscillation that will do this turns out to be in the $(1, 1)$ direction, as in figure 1(right). However, this one is a little tricky, because it is not completely obvious that the third u does not instead oscillate three times in the $(-1, 1)$ direction (i.e. $-/+/-$ peaks along that direction). As $|\lambda|$ gets bigger, it gets harder to guess the exact ordering of the eigenfunctions, but you can still guess what they look like modulo some uncertainty which of certain pairs come first. [Exact calculations turn out to show that a 3-peak oscillation in the $(-1, 1)$ direction gives the fourth λ , which is about 30% bigger than the third λ , corresponding to an average “wavelength” that is about 12% smaller.]

Example: Guessing eigenfunctions of $-c\nabla^2$

It gets even more interesting for the case of nonconstant coefficients, e.g. $\hat{A} = -c(\vec{x})\nabla^2$ for some $c > 0$. In this case, \hat{A} is self-adjoint under the weighted inner product $\langle u, v \rangle = \int_{\Omega} \bar{u}v/c$, so the Rayleigh quotient is

$$R\{u\} = \frac{\int_{\Omega} |\nabla u|^2}{\int_{\Omega} |u|^2/c}.$$

Now, u_1 is trying to satisfy two *competing* concerns:

- u_1 “wants” to *oscillate as little as possible* to minimize the numerator of R .
- u_1 “wants” to *concentrate in low- c regions* to maximize the denominator of R .

Which of these concerns dominates depends on how much c varies. As c in one region gets smaller and smaller relative to other regions, u_1 will concentrate more and more in this region, even though by doing so it will have larger $|\nabla u|$.

As before, u_2 and higher minimize the same R but must be orthogonal to u_1 . So, u_2 is “forced out” of the low- c regions to some extent by orthogonality, since it must have a node in those regions to be $\perp u_1$.

¹In an equilateral triangle, the second two λ ’s turn out to be equal, i.e. oscillating in either of the two directions gives the same Rayleigh quotient. This equality turns out to be a deeper consequence of symmetry, but that is outside the scope of 18.303.