CS 542G: Rayleigh-Ritz

Robert Bridson

November 3, 2009

1 Acoustics

To motivate the last of our eigenalgorithms, consider a significantly nastier variant on the symmetric eigenproblem we've been studying so far: the case where the space is infinite dimensional. Our example will be from acoustics: solving for the evolution of sound waves in a fluid like air.

To get the ball rolling we need to know a bit about the dynamics of fluid flow. Since we're focused on sound waves, which are relatively minute variations in density that travel extremely fast relative to the fluid's molecules themselves (ignoring supersonic planes and shock waves) and travel far, we can use a fairly simplified set of equations.

Our two primary variables are velocity $\vec{u}$ and density ρ (mass per unit volume) measured at any point in space. We ignore temperature and other thermodynamic quantities, again because we are looking at only the most important parts of the physics for sound waves.

Our first equation is the basics of dynamics: conservation of momentum, or in Newton's classic form $\vec{F} = m\vec{a}$. Consider a molecule of fluid with its own velocity $\vec{v}$: its acceleration is just

$$\vec{a} = \frac{d\vec{v}}{dt}$$

However, we're not tracking the velocities of individual molecules—we're instead tracking the velocity field as a function of space, i.e. the velocity at a fixed position as the fluid streams past. If our molecule is at position $\vec{x}$ at the time we measure its velocity $\vec{v}$, then

$$\vec{v} = \vec{u}(\vec{x}, t)$$

Using the chain rule, its acceleration is

$$\begin{aligned}\vec{a} &= \frac{d\vec{v}}{dt} \\ &= \frac{\partial \vec{u}}{\partial t} + \frac{\partial \vec{u}}{\partial \vec{x}} \cdot \frac{d\vec{x}}{dt} \\ &= \frac{\partial \vec{u}}{\partial t} + (\nabla \vec{u}) \cdot \vec{v} \\ &= \frac{\partial \vec{u}}{\partial t} + \vec{u} \cdot \nabla \vec{u}\end{aligned}$$

This is a bit nasty—there's a nonlinear second term accounting for the motion of the fluid itself (sometimes called the "advection", "convection" or "transport" term)—but we'll simplify it later.

So much for the acceleration: what about the mass? We'll be doing this on a continuum, per-unit-volume basis, so we use density ρ instead.

Finally, we have to think about the forces that act on a fluid. Here is where we make a lot of simplifying assumptions, considering only the dominant two for acoustics: pressure and gravity. Pressure exerts a force from high pressure regions to low pressure regions, as measured by the gradient of the pressure field p. Gravity is a constant acceleration $\vec{g}$, or per-unit-volume force $\rho\vec{g}$.

Putting all this together, $F = m\vec{a}$ becomes:

$$-\nabla p + \rho\vec{g} = \rho\left(\frac{\partial \vec{u}}{\partial t} + \vec{u} \cdot \nabla \vec{u}\right)$$

This is often rearranged, and divided through by ρ (to put it in "non-conservative" form) to get:

$$\frac{\partial \vec{u}}{\partial t} + \vec{u} \cdot \nabla \vec{u} + \frac{1}{\rho}\nabla p = \vec{g}$$

This, in either form, is called the **momentum equation** for an "ideal" fluid (or the momentum part of the "Euler equations").

We're not done, however, as this is only three equations (or one vector-valued equation) but we have four primitive variables and a new quantity pressure p: we need more to close the system.

For the pressure part we go to thermodynamics, for example the ideal gas law. In fact, since we've simplified away considerations of heat, and are looking at very small perturbations in density and pressure, we can assume an even simpler linear relationship between the perturbation of pressure from its rest value p_0 and the perturbation of density from its rest value ρ_0:

$$p = p_0 + k(\rho - \rho_0)$$

where k is the positive constant of proportionality. This is a particularly simple example of an **equation of state** (eos). Of course, we also need to know what the rest values are: we define them as whatever

will cancel out the force of gravity and result in no motion, i.e. solving $\nabla p_0 = \rho_0 \vec{g}$, and will ignore what exactly this gives. (This is called the "hydrostatic" case, where pressure gradient cancels off gravity to leave things static.)

We still need one more equation. Conservation of mass supplies this one for us. Consider a fixed but arbitrary region in space, Ω, and the mass of the fluid in Ω. Since mass can't be destroyed or created, any change in that mass must be equal to the amount of mass that entered or exited (or rather, the balance between entry and exit) through the boundary. More mathematically, the mass in the region is:

$$M = \int_\Omega \rho dV$$

and the rate at which it changing is (using the fact Ω is fixed, not changing in time):

$$\begin{aligned} \frac{dM}{dt} &= \frac{d}{dt}\int_\Omega \rho dV \\ &= \int_\Omega \frac{\partial \rho}{\partial t} dV \end{aligned}$$

Meanwhile, at any point on the boundary the fluid is moving out of the region at speed $\vec{u} \cdot \hat{n}$ where $\hat{n}$ is the unit-length outward-pointing normal. The rate at which mass is leaving, per unit area, is then $\rho \vec{u} \cdot \hat{n}$, and integrating this over the boundary gives the total rate of mass loss through the boundary:

$$\int_{\partial\Omega} \rho \vec{u} \cdot \hat{n} dA$$

Conservation of mass just says these two integrals are balanced:

$$\int_\Omega \frac{\partial \rho}{\partial t} dV + \int_{\partial\Omega} \rho \vec{u} \cdot \hat{n} dA = 0$$

This equation is a bit inconvenient, however, since it mixes a volume integral with a surface integral. Applying the Divergence Theorem to the surface integral simplifies it a bit:

$$\begin{aligned} \int_\Omega \frac{\partial \rho}{\partial t} dV + \int_\Omega \nabla \cdot (\rho \vec{u}) dV &= 0 \\ \int_\Omega \left[\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{u}) \right] dV &= 0 \end{aligned}$$

This is true for any arbitrary volume Ω, so assuming the integrand isn't crazy (discontinuous) it must be exactly zero everywhere:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{u}) = 0$$

This is the last equation we need, called the **continuity equation**.

To sum up, here are our equations:

$$\begin{aligned} \frac{\partial \vec{u}}{\partial t} + \vec{u} \cdot \nabla \vec{u} + \frac{1}{\rho} \nabla p &= \vec{g} \\ p &= p_0 + k(\rho - \rho_0) \\ \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{u}) &= 0 \end{aligned}$$

These are very nasty: several nonlinear terms, vector-valued. There's no hope of solving them analytically in general.

However, we can make further simplifying assumptions: that the velocity is small compared to the acoustic fluctuations, that the fluctuations in density are much smaller than the rest density, etc. This will remove all nonlinear terms, leaving us with:

$$\begin{aligned} \frac{\partial \vec{u}}{\partial t} + \frac{1}{\rho_0} \nabla (p_0 + \Delta p) &= \vec{g} \\ \Delta p &= k(\rho - \rho_0) \\ \frac{\partial \rho}{\partial t} + \rho_0 \nabla \cdot \vec{u} &= 0 \end{aligned}$$

Here I've introduced Δp as the difference between the current pressure and the rest pressure. The first equation simplifies just a little more, since the rest pressure cancels off the effect of gravity:

$$\begin{aligned} \frac{\partial \vec{u}}{\partial t} + \frac{1}{\rho_0} \nabla (\Delta p) &= 0 \\ \Delta p &= k(\rho - \rho_0) \\ \frac{\partial \rho}{\partial t} + \rho_0 \nabla \cdot \vec{u} &= 0 \end{aligned}$$

We don't really care that much about velocity, just density, so we can actually eliminate it from the equations by differentiating the last one in time:

$$\frac{\partial^2 \rho}{\partial t^2} + \rho_0 \nabla \cdot \frac{\partial \vec{u}}{\partial t} = 0$$

and then substituting in the first:

$$\frac{\partial^2 \rho}{\partial t^2} + \rho_0 \nabla \cdot \left(-\frac{1}{\rho_0} \nabla (\Delta p) \right) = 0$$

Plugging in the equation of state and throwing out a few more negligible terms, we get to a single simple equation for density:

$$\frac{\partial^2 \rho}{\partial t^2} - k \nabla \cdot \nabla \rho = 0$$

or more conventionally:

$$\frac{\partial^2 \rho}{\partial t^2} = k \nabla \cdot \nabla \rho$$

This is called the **wave equation**, derived in this case for acoustics, but also arising in many other physical phenomena (such as particular polarizations of light, certain waves in solid materials, small waves in shallow water (in 2D), etc.).

The wave equation is by itself not a well-posed problem to solve: we also need boundary conditions, and initial conditions. We'll ignore the delicate issue of boundary conditions for now, but argue that often a specific set of initial conditions isn't really the way to go anyhow: in many cases, what's important is not a single solution to the wave equation but more generally how it responds in general scenarios. This boils down to looking at *resonant* waves, special solutions to the equation (plus the boundary conditions we ignore) that are periodic in time. For example, if designing a car interior it would be important to avoid having a resonant wave that matches the vibration caused by typical driving speeds on typical highway surfaces—it would be very noisy for the people inside.

Assume the periodic-in-time function we seek is a "time harmonic" solution, of the form:

$$\rho(\vec{x}, t) = \bar{\rho}(\vec{x}) \sin(\omega t)$$

for example, where ω is the time frequency, and $\bar{\rho}$ gives the spatial shape of the wave. Plugging this into the wave equation gives:

$$\bar{\rho}(\vec{x}) \left(-\omega^2 \sin(\omega t)\right) = k \nabla \cdot \nabla \bar{\rho}(\vec{x}) \sin(\omega t)$$

Dividing out by the common factor of $\sin(\omega t)$ and rearranging gives:

$$\nabla \cdot \nabla \bar{\rho} = \left(\frac{-\omega^2}{k}\right) \bar{\rho}$$

If we label the Laplacian operator $\nabla \cdot \nabla$ with the letter A, the unknown scalar value $-\omega^2/k$ as λ, and the unknown spatial wave shape function $\bar{\rho}$ as x, we get the following familiar equation:

$$Ax = \lambda x$$

We're back to a good old eigenproblem! The only wrinkle, of course, is that it's in infinite dimensions.

2 Solving Infinite or Big Problems

We're in bad shape here, with our existing eigentools. Power Method and relatives like Orthogonal Iteration won't be much help: the Laplacian's spectrum of eigenvalues stretches all the way to $-\infty$, so there is no maximum magnitude eigenvalue we could possibly converge to. Shift-and-invert is more promising, except each iteration requires solving a full-fledged partial differential equation which can't be done analytically in general (it's almost as hard as the original eigenproblem). We need a new tool.

This is not the first case of a potentially infinite problem we've come across in this course. For interpolating scattered data we also were dealing with choosing one preferred function from an infinite dimensional space, and our key tool there was to pick a good finite dimensional subspace, then solve the problem restricted to that space. For example, with RBF's we derived some powerful radially symmetric basis functions, and the premise behind least squares was that you have prior knowledge of the phenomenon to pick a small but smart subspace. Let's do the same here.

For the acoustic wave equation example, the subspace is typically built from piecewise polynomial functions defined on an adequately high resolution mesh that fills the domain. This sort of subspace can accurately approximate any smooth function that doesn't vary faster than the mesh spacing, making it a great general-purpose choice. In fact, for a slightly different and easier problem of just solving regular partial differential equations, we'll develop this idea later in the course into the Finite Element Method. It's important to point out we'll only get an approximation to the correct answer, but we do have control (up to the limits of our computational resources) over how accurate that approximation will be by using higher and higher resolution meshes, but we won't go further with it here.

In fact, this idea is also useful for merely large but finite eigenproblems. If A is an $n \times n$ matrix with n very large, and we are only looking for a small set of k eigenvectors, if we can get a basis for a likely subspace that might approximate that set, we can use the same tool. We'll focus on this situation from now on, so as to avoid the technical nastiness involved in working with infinite dimensions.

3 Rayleigh-Ritz

Now we have the setting for the problem, let's get to specifics. Again let $A \in \mathbb{R}^{n \times n}$ be a symmetric matrix. Suppose we have a k-dimensional subspace with orthonormal basis in the columns of $U \in \mathbb{R}^{n \times k}$.[1] We are seeking an optimal set of k vectors from this space approximating eigenvectors of A. That is, we seek k vectors, each of which a is linear combination of the columns of U: putting those approximate eigenvectors as columns of another matrix we see it must be expressed as

$$UV$$

where $V \in \mathbb{R}^{k \times k}$ is a square matrix containing the coefficients of those linear combinations.

For ease of proof, we'll restrict our search for approximations to orthonormal vectors: using the fact that U has orthonormal columns, this will imply V is an orthogonal $k \times k$ matrix.

[1] Technically, we can do this with any basis, not just orthonormal ones, but it will complicate this first introduction to the algorithm—the reduced dimension eigenproblem ends up being a "generalized" eigenproblem of the form $Ax = \lambda Bx$ which we won't touch in this course.

To say we are finding approximate eigenvectors means there will also be a $k \times k$ diagonal matrix D of the accompanying eigenvalues, and that

$$A(UV) \approx (UV)D.$$

Note that this will be impossible to satisfy exactly if the span of U doesn't have a basis of exact eigenvectors of A, so we will instead want to find a V and D that come as close as possible to satisfying it.

Yet again, the simplest route to a well-posed problem is to set this up as a least-squares minimization of some sort, say using the Frobenius norm of the difference $A(UV) - (UV)D$:

$$\min_{V,D} \|AUV - UVD\|_F^2$$

taken over all $k \times k$ orthogonal matrices V and diagonal matrices D.

We can figure out the answer to this by completing U with $\bar{U}$, an $n \times (n-k)$ matrix whose columns are orthonormal and orthogonal to U: this means the complete matrix $(U|\bar{U})$ is an orthogonal matrix. Note that applying Householder QR factorization to U and taking the last $n-k$ columns of the Q would be an adequate way of computing $\bar{U}$, but we only will use it to derive the algorithm.

Multiplying by an orthogonal matrix such as $(U|\bar{U})^T$ doesn't change the Frobenius norm—multiply it on the left:

$$\begin{aligned} \|AUV - UVD\|_F^2 &= \|(U|\bar{U})^T AUV - (U|\bar{U})^T UVD\|_F^2 \\ &= \left\| \begin{pmatrix} U^T AUV - U^T UVD \\ \bar{U}^T AUV - \bar{U}^T UVD \end{pmatrix} \right\|_F^2 \end{aligned}$$

Note that $U^T U = I$, but $\bar{U}^T U = 0$. We can also separate the Frobenius norm squared of the big matrix into the sum of the squared norms of the top and bottom parts, to get:

$$\|AUV - UVD\|_F^2 = \|U^T AUV - VD\|_F^2 + \|\bar{U}^T AUV\|_F^2$$

For the second term, note that multiplying on the right by the orthogonal matrix V^T doesn't change the Frobenius norm either, thus this term is the same as $\|\bar{U}^T AU\|_F^2$ which only depends on the (fixed) choice of U, not the matrices V and D over which we are minimizing. This term in fact represents how much error there will be in approximating eigenvectors from the subspace U—how close that subspace is to being spanned by k exact eigenvectors, in a sense.

We are then left with the equivalent mimimization problem

$$\min_{V,D} \|(U^T AU)V - VD\|_F^2$$

where V is orthogonal and D is diagonal. Obviously the minimum occurs when this is zero, achieved by taking V to be the set of eigenvectors of the $k \times k$ symmetric matrix $U^T AU$, and the diagonal entries of D the matching eigenvalues. Solving just a $k \times k$ eigenproblem gives us the optimal estimate UV for k of the eigenvectors of A, and D for the associated eigenvalues.

This procedure is called **Rayleigh-Ritz**, and the associated eigenvalue estimates are called **Ritz values**. It can be seen as a generalization of the Rayleigh quotient we saw earlier, which is the $k = 1$ case. Rayleigh-Ritz is an extremely powerful method for both making infinite eigenproblems finitely approximable, and getting accurate estimates of subsets of eigenpairs for finite yet large problems. It also can be profitably used to increase the power of the other methods we have seen. For example, one problem with the plain QR method (as we saw it) is its difficulty in not being able to separate equal magnitude but opposite sign eigenvalues—the algorithm instead converges to a non-eigenspace containing both eigenvectors. Rayleigh-Ritz applied to that small space leads to a 2×2 eigenproblem which can be solved "exactly" with the quadratic equation, precisely seperating the positive and negative eigenvalues.