CS 542G: FEM Optimality and Implementation

Robert Bridson

November 19, 2008

This lecture isn't properly written out, but I'm summarizing the important points for now as a substitute.

1 The Optimality of FEM

- We proved that the Galerkin FEM solution is the closest to the exact solution in the chosen subspace, measured in the "energy-norm" (i.e. with the integral of the gradient squared).
- This norm can be proven to be equivalent to the more usual L^2, i.e. different by only a constant
- Therefore, as long as we choose a subspace that can well approximate the solution, we will get a good answer back from Galerkin FEM
- Since we don't know what the solution will be before solving the problem, we instead look for spaces that can accurately approximate any reasonable function.

2 Choosing a Subspace

- Any of the spaces mentioned at the start of the course can work—polynomials, piecewise polynomials on grids and meshes, RBFs, even MLS-defined spaces—and many other "exotic" examples have been explored too (for example, if there will be a known type of singularity in the solution, such as the stress at the tip of a crack in a fracture mechanics problem, functions that match that type of singularity can be added to the space)
- We'll focus on the most common choice, piecewise polynomial functions built on meshes, in particular at the simplest piecewise linear functions in 2D on triangle meshes.

- Note that piecwise linear continuous functions aren't even differentiable, so it's a bit surprising they can be used to solve a PDE that needs second derivatives. However, the variational form only involved the integral of the first derivative squared—as long as the first derivative is sensible almost everywhere, and there aren't jump discontinuities in the functions that would result in Dirac delta spikes in the gradient, we're fine.
- The simplest basis is the "nodal basis", where $\phi_i(x)$ is 1 at interior mesh vertex i and 0 at the other mesh vertices. This results in "hat functions", which have their peak at vertex i and are exactly zero in every triangle not containing vertex i.
- Taylor's theorem basically guarantees a piecewise linear function can approximate a smooth function to within $O(\Delta x^2)$ error, where Δx is a measure of the distance between vertices; more on this next lecture.

3 Assembling the Linear System

- Once the mesh and the implied basis functions have been sorted out, we still need to compute the entries of A and f.
- From the method's history in civil engineering—understanding how strong structures are—the matrix is called the "stiffness matrix", and calculating its entries is called "assembling" it.
- Since our functions are defined piecewise over the mesh, it's natural to break the integrals defining A and f up into sums of integrals over the triangles. These triangle integrals are zero except where the nodal basis functions are nonzero—where both nodal basis functions are nonzero in the case of the entries of A.
- It's very natural to loop over mesh elements (triangles), calculating the nonzero contributions to every entry of A and f influenced by each triangle. These contributions can be thought of as forming a submatrix, called a "local stiffness matrix" which is added to the (global) stiffness matrix A.
- For our choice of nodal linear elements, the integrals in A are easy to work out exactly. The gradients are constant over each triangle and can be geometrically derived; the integral just needs to multiply the dot-product of gradient vectors by the area of the triangle.
- However, for f, and for more general FEM, the integrals can be extremely difficult to evaluate exactly. On the other hand, since we're making plenty of (hopefully small) errors elsewhere in solving the original problem, it's perfectly reasonable to just approximate these integrals with a similarly small error.

- **Quadrature** is the process of estimating an integral numerically. The simplest scheme is the midpoint method: evaluate the integrand at the "midpoint" of the element (for a triangle this is the centroid, the average of the vertex positions) and multiply by the area of the element. This is second-order accurate; higher accuracy schemes are available, such as the Gaussian quadrature approach, which just require more point samples, perhaps multiplied by special weights.