CS 542G: Radial Basis Functions

Robert Bridson

September 15, 2008

1 Piecewise Linear Interpolation from Fundamental Solutions

Last time, we ended with a differential equation version of the interpolation problem. We began with the problem of finding $f(x)$ which interpolates the data, $f(x_i) = f_i$, that is as smooth as possible in the sense of minimizing

$$\int_{-\infty}^{\infty} |f'(x)|^2 \, dx$$

and through the calculus of variations reduced it to a differential equation:

$$\begin{aligned} f''(x) &= 0 \qquad \text{for} \quad x \neq x_i \quad (i = 1, \ldots n) \\ f(x_i) &= f_i \qquad \text{for} \quad i = 1, \ldots n \end{aligned}$$

We didn't highlight it, but we need to add to that a condition on f so that the integral is finite: $f'(x) \to 0$ at $x = \pm\infty$ in some appropriate sense.

This is fairly simple to solve, right off the bat: the only functions in 1D with zero derivatives beyond some order are polynomials, and so in this case $f(x)$ must be degree one (a straight line) on each interval $[x_i, x_{i+1}]$. In other words, we've rederived the piecewise linear interpolant we had before. However, this way of solving it doesn't generalize to higher dimensions anymore than our work from last time did, so we'll take a rather more convoluted approach that does. It's a lot of pointless work in 1D, but it makes the problem tractable in higher dimensions.

This is the method of **fundamental solutions**, a classic technique for analytically working with linear partial differential equations. The differential equation we have, $f''(x) = 0$, is formed from the second derivative operator d^2/dx^2, which is indeed linear:

$$\begin{aligned} \frac{d^2}{dx^2}(f+g) &= \frac{d^2}{dx^2} f + \frac{d^2}{dx^2} g \\ \frac{d^2}{dx^2}(\alpha f) &= \alpha \frac{d^2}{dx^2} f \end{aligned}$$

In some sense, it can be treated as an infinite-dimensional matrix: whereas matrices can represent any linear operator applied to vectors from finite dimensional spaces, differential operators apply to functions from possibly infinite dimensional spaces (like the space of all differentiable functions).

Take a look at a finite dimensional linear system, formed with a matrix A:

$$Ax = b$$

Here b is a fixed vector, and x is the unknown solution vector. One way to conceptually solve this is to introduce the matrix inverse: $x = A^{-1}b$. The matrix inverse can be defined as the matrix that satisfies $AA^{-1} = I$. In particular, if we look at the i'th column c_i of A^{-1}, from this equation it satisfies the linear system

$$Ac_i = e_i = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{pmatrix}$$

Here e_i is the i'th column of the identity matrix, zero everywhere except for its i'th entry. In other words, Ac_i is zero almost everywhere except in one spot, and from those columns c_i we can construct the solution to the original problem.

Bringing this intuition back to the differential equation, we can seek out an analogy of the matrix inverse for the linear operator d^2/dx^2, which boils down to finding functions $\phi(x)$ where $d^2\phi/dx^2 = 0$ almost everywhere except for one spot. It turns out the right analogy to the column of the identity e_i (which is zero everywhere except one entry, and sums to one) is the **Dirac delta** function[1] $\delta(x)$. Recall its defining properties:

$$\begin{aligned} \delta(x) &= 0 \qquad \text{for} \quad x \neq 0 \\ \int_{-\infty}^{\infty} \delta(x) &= 1 \end{aligned}$$

Our fundamental solution will thus solve:

$$\frac{d^2}{dx^2}\phi(x) = \delta(x)$$

It's easy to see this implies $\phi(x)$ is a straight line on either side of the origin, but has to have a kink at the origin; integrating once shows the difference in slopes has to be equal to 1. For simplicity, we take the

[1]Technically the Dirac delta is not a true function, but instead a **distribution**, a somewhat more general object.

fundamental solution to be the most symmetric possibility, which is:

$$\phi(x) = \frac{1}{2}|x|$$

up to an additive constant (more on this in a moment).

We can now use this to construct a solution to our problem by using versions of $\phi(x)$ shifted to be centred on the data points, analogous to the different columns of the matrix inverse. The general solution to our differential equation $f''(x) = 0$ except for $x = x_i, i = 1, \ldots n$, is:

$$f(x) = \sum_{i=1}^{n} \lambda_i \phi(x - x_i) + \lambda_{n+1}$$

Here the λ_i are constant coefficients, which we don't know yet. Note the inclusion of λ_{n+1}, a constant term: this reflects the fact that the roughness integral $\int |f'(x)|^2$ isn't changed by adding a constant to f, so we ought to include one in a general solution.

This general form of the solution has $n + 1$ unknowns. The interpolation condition $f(x_i) = f_i, i = 1, \ldots, n$ gives us n linear equations for them:

$$\begin{aligned}
f(x_1) &= \sum_{i=1}^{n} \lambda_i \phi(x_1 - x_i) + \lambda_{n+1} = f_1 \\
f(x_2) &= \sum_{i=1}^{n} \lambda_i \phi(x_2 - x_i) + \lambda_{n+1} = f_2 \\
&\vdots \\
f(x_n) &= \sum_{i=1}^{n} \lambda_i \phi(x_n - x_i) + \lambda_{n+1} = f_n
\end{aligned}$$

We need one extra equation to close the system. We get it by noting that when f minimizes the integral, the roughness integral had better be finite—which requires $f'(x) \to 0$ at $\pm\infty$. The derivative of f for x larger than all the data points is:

$$\begin{aligned}
f'(x) &= \sum_{i=1}^{n} \lambda_i \phi'(x - x_i) + 0 \\
&= \sum_{i=1}^{n} \lambda_i \frac{1}{2}
\end{aligned}$$

This has to be zero, i.e. $\sum_i \lambda_i = 0$, which is the extra equation we need. (You can verify the same condition pops up from looking at $f'(x)$ for x less than all data points too.)

Solving this linear system, of course, has to give us back the same piecewise linear interpolation we've been going over and over. Let's see if we can do something more interesting

2 Getting Smoother

Our next extension will be to get a smoother interpolant. Right now our roughness integral forces $f(x)$ to be as straight as possible between data points, i.e. it has to be formed of straight lines, but then at the data points themselves there are non-smooth kinks. To force $f(x)$ to be smooth even there, we'll use a higher order measure of roughness instead: we'll ask $f(x)$ to minimize

$$\int_{-\infty}^{\infty} |f''(x)|^2$$

while interpolating the data, $f(x_i) = f_i, i = 1, \ldots, n$. As before, we can do the calculus of variations by introducing some arbitrary but sensible function $g(x)$ that is zero at all the data points, $g(x_i) = 0, i = 1, \ldots, n$, and then looking at the roughness of $f + \epsilon g$ as a function of the scalar parameter ϵ:

$$h(\epsilon) = \int_{-\infty}^{\infty} |f''(x) + \epsilon g''(x)|^2$$

If $f(x)$ was indeed the minimum, $h'(\epsilon)$ had better be zero at $\epsilon = 0$. This works out to the condition

$$\int_{-\infty}^{\infty} f''(x)g''(x) = 0$$

We can then use integration-by-parts twice to put this in terms of $g(x)$ instead of $g''(x)$, noting that the boundary terms can be forces to vanish by requiring that $g(x)$ should drop to zero past some finite bound:

$$\begin{aligned} 0 &= \int_{-\infty}^{\infty} f''(x)g''(x) \\ &= -\int_{-\infty}^{\infty} f^{(3)}(x)g'(x) \\ &= \int_{-\infty}^{\infty} f^{(4)}(x)g(x) \end{aligned}$$

Since this is true for arbitrary $g(x)$ satisfying $g(x_i) = 0$, we conclude:

$$f^{(4)}(x) = 0 \qquad \text{for} \quad x \neq x_i, i = 1, \ldots, n$$

On top of this, $f(x)$ of course has to satisfy the interpolation condition, and for its roughness integral to be finite we see $f''(x) \to 0$ at $\pm\infty$.

This can be tackled directly as before, noting that if the fourth derivative of $f(x)$ is zero on the intervals, $f(x)$ must be a cubic polynomial in each interval, and on you can go. Instead, let's use fundamental solutions.

Now the fundamental solution satisfies

$$\frac{d^4}{dx^4}\phi(x) = \delta(x)$$

which means it is a cubic on either side of the origin, and by further requiring symmetry, it can be shown the fundamental solution is:

$$\phi(x) = \frac{1}{12}|x|^3$$

We can again write down the general form of a solution using $\phi(x)$, only this time we include an additional linear term to reflect the fact that adding a linear term won't change the second derivative and thus the roughness integral:

$$f(x) = \sum_{i=1}^{n} \lambda_i \phi(x - x_i) + \lambda_{n+1} + \lambda_{n+2} x$$

We have $n+2$ unknown coefficients, the λ_i. The interpolation conditions give n linear equations for them as before:

$$\begin{aligned} f(x_1) &= \sum_{i=1}^{n} \lambda_i \phi(x_1 - x_i) + \lambda_{n+1} = f_1 \\ &\vdots \\ f(x_n) &= \sum_{i=1}^{n} \lambda_i \phi(x_n - x_i) + \lambda_{n+1} = f_n \end{aligned}$$

To get two more equations, we'll look at making sure $f''(x) \to 0$ at $\pm\infty$, necessary for the roughness integral to be finite. The second derivative of f for x greater than all the data points is:

$$\begin{aligned} f''(x) &= \sum_{i=1}^{n} \lambda_i \phi''(x - x_i) + 0 + 0 \\ &= \sum_{i=1}^{n} \lambda_i \frac{1}{2}(x - x_i) \end{aligned}$$

This is a straight line. For it to have limit 0 at infinity it must in fact be zero, giving us two equations, that the linear term is zero and the constant term is zero:

$$\begin{aligned} \sum_{i=1}^{n} \lambda_i &= 0 \\ \sum_{i=1}^{n} \lambda_i x_i &= 0 \end{aligned}$$

You can double check the same equations arise by looking at x less than every data point. This completes our system, which we can solve to get a smoother cubic spline.

3 Extension to Higher Dimensions

Finally we are at the point of tackling interpolation in higher dimensions in an easy but optimal (in some sense) way. We'll again ask for $f(x)$ to interpolate the data but be as smooth as possible, defined as minimizing an integral measuring roughness.

Coming up with a single number to define how rough a multidimensional function is isn't immediately obvious: unlike in 1D, where there is only one fourth derivative, in k dimensions there are $O(k^4)$ different fourth order partial derivatives. There is a natural choice, however, based on the Laplacian (which we will return to later in the course). The Laplacian is the divergence of the gradient of a function, or in partial derivative notation:

$$\nabla \cdot \nabla f(x) = \sum_{i=1}^{k} \frac{\partial^2 f}{\partial x_i^2}$$

That is, the sum of the double second derivatives. One of the powerful things about the Laplacian is that it is rotationally invariant: if you rotate the coordinate system in which you measure it, its value doesn't change. It also appears in geometry as a natural way to measure mean curvature, and in a fundamental sense measures how different the value of a function is at a point from the average of the neighbouring function values.

We therefore will base our smoothness on minimizing the integral of the Laplacian squared:

$$\int_{\mathbb{R}^k} |\nabla \cdot \nabla f(x)|^2$$

With this specified, we can proceed with the same calculus of variations argument we have seen twice already in 1D, introducing an arbitrary $g(x)$ that goes through zero at the data points and examining the roughness of $f + \epsilon g$ as a function of ϵ. We get in exactly the same way as before that

$$\int_{\mathbb{R}^k} \nabla \cdot \nabla f(x) \nabla \cdot \nabla g(x) = 0$$

for any such $g(x)$. We can then apply integration by parts twice—the formula generalizes to multiple dimensions in the obvious ways—and again assume boundary terms will vanish by restricting $g(x)$ to drop to zero beyond a finite extent:

$$\begin{aligned} 0 &= \int_{\mathbb{R}^k} \nabla \cdot \nabla f(x) \nabla \cdot \nabla g(x) \\ &= -\int_{\mathbb{R}^k} \nabla \nabla \cdot \nabla f(x) \cdot \nabla g(x) \\ &= \int_{\mathbb{R}^k} \nabla \cdot \nabla \nabla \cdot \nabla f(x) g(x) \end{aligned}$$

We conclude as before with a PDE that $f(x)$ satisfies:

$$\nabla \cdot \nabla\nabla \cdot \nabla f(x) = 0 \qquad \text{except when} \quad x = x_i, i = 1, \dots n$$

This is called the "biharmonic equation", similar to how we term functions "harmonic" when their Laplacian is zero.[2]

Solving a fourth order partial differential equation in k dimensions might seem pretty daunting. However, this is where fundamental solutions actually make life easier! The fundamental solution is, as before, a function $\phi(x)$ satisfying:

$$\nabla \cdot \nabla\nabla \cdot \nabla \phi(x) = \delta(x)$$

We'll further require that $\phi(x)$ is radially symmetric, i.e. that $\phi(x) = \phi(y)$ whenever $||x|| = ||y||$, or in other words that $\phi(x)$ really only depends on the distance r that x is from the origin. Thus $\phi(r)$ can be thought of as a 1D function, and the partial differential equation in x can be reduced to a simple 1D differential equation in r. This is an example of a **Radial Basis Function** or RBF.

The full details of finding the fundamental solution belong in an applied math course on partial differential equations: we'll instead just quote the results. Interestingly enough, the fundamental solution depends critically on the dimension we're in—this is a reasonably common occurence and is important to have in mind when thinking about setting up model problems in lower dimensions. Here are the first three dimensions, ignoring scale factors and lower order terms:

$$\phi(r) = \begin{cases} r^3 & : \quad k = 1 \\ r^2 \log r & : \quad k = 2 \\ r & : \quad k = 3 \end{cases}$$

As before, we use this to write down the general form for $f(x)$, with the addition of a linear polynomial (an extra $k + 1$ terms):

$$f(x) = \sum_{i=1}^{n} \lambda_i \phi(x - x_i) + \lambda_{n+1} + \lambda_{n+2} x^{(1)} + \cdots \lambda_{n+k+1} x^{(k)}$$

The interpolation conditions give the usual n equations, and keeping the integral finite supplies the re-

[2]There's also a triharmonic equation, involving three Laplacians.

maining $k+1$ equations which are a clear generalization of the 1D versions:

$$\begin{aligned}
\sum_{i=1}^{n} \lambda_i &= 0 \\
\sum_{i=1}^{n} \lambda_i x_i^{(1)} &= 0 \\
&\vdots \\
\sum_{i=1}^{n} \lambda_i x_i^{(k)} &= 0
\end{aligned}$$

Once these $n+k+1$ linear equations have been solved for the coefficients λ_i, we have our Radial Basis Function interpolant.

4 Radial Basis Functions

The $\phi(r)$ which we gave in the previous section goes under the name **thin-plate spline**, due to its connection to simulating the bending of a thin metal plate (whose potential energy can be approximated with the integral of Laplacian squared), and is one of the most popular RBFs to use. However, there are plenty of others. They all have in common the idea of using a $\phi(x)$ which is radially symmetric, and typically add some low order polynomial terms (and matching extra equations).

Some are derived from optimality principles and partial differential equations, like the thin-plate spline: e.g. the triharmonic basis function $\phi(r) = r^3$ in 3D. Some arise from very different application areas, such as the Gaussian $\phi(r) = \exp(-r^2/c^2)$ which needs some length scale c to be specified. Others have no particular justification, other than that they seem to work really well in practice: chief among these is the multiquadric function, $\phi(r) = \sqrt{r^2 + c^2}$, where c is again a user-specified length scale.