18.303: Introduction to Green's functions and operator inverses

S. G. Johnson

October 9, 2011

Abstract

In analogy with the inverse A^{-1} of a matrix A, we try to construct an analogous inverse $\hat{A}^{-1}$ of differential operator $\hat{A}$, and are led to the concept of a Green's function $G(\mathbf{x}, \mathbf{x}')$ [the $(\hat{A}^{-1})_{\mathbf{x},\mathbf{x}'}$ "matrix element" of $\hat{A}^{-1}$]. Although G is a perfectly well-defined function, in trying to construct it and define its properties we are drawn inexorably to something that appears less well defined: $\hat{A}G$ seems to be some kind of function that is "infinitely concentrated at a point", a "delta function" $\delta(\mathbf{x} - \mathbf{x}')$ that does not have a proper definition as an ordinary function. In these notes, we show how to solve for G in an example 1d problem without resorting to delta functions, by a somewhat awkward process that involves taking a limit *after* solving the PDE. In principle, we could always apply such contortions, but the inconvenience of this approach will lead us in future lectures to a new definition of "function" in which delta functions are well defined.

0 Review: Matrix inverses

The inverse A^{-1} of a matrix A is already familiar to you from linear algebra, where $A^{-1}A = AA^{-1} = I$. Let's think about what this means.

Each *column* j of A^{-1} has a special meaning: $AA^{-1} = I$, so $A \cdot (\text{column } j \text{ of } A^{-1}) = (\text{column } j \text{ of } I) = \mathbf{e}_j$, where $\mathbf{e}_j$ is the unit vector $(0, 0, \ldots, 0, 1, 0, \ldots, 0)^T$ which has a 1 in the j-th row. That is, $(\text{column } j \text{ of } A^{-1}) = A^{-1}\mathbf{e}_j$, the solution to $A\mathbf{u} = \mathbf{f}$ with $\mathbf{e}_j$ on the right-hand-side. If we are solving $A\mathbf{u} = \mathbf{f}$, the statement $\mathbf{u} = A^{-1}\mathbf{f}$ can be written as a superposition of columns of A^{-1}:

$$\mathbf{u} = A^{-1}\mathbf{f} = \sum_j (\text{column } j \text{ of } A^{-1}) f_j = A^{-1} \sum_j f_j \mathbf{e}_j.$$

Since $\sum_j f_j \mathbf{e}_j = \mathbf{f}$ is just writing $\mathbf{f}$ in the basis of the $\mathbf{e}_j$'s, we can interpret $A^{-1}\mathbf{f}$ as *writing* $\mathbf{f}$ *in the basis of the* $\mathbf{e}_j$ *unit vectors and then summing (superimposing) the solutions for each* $\mathbf{e}_j$. Going one step further, and writing:

$$u_i = \left(A^{-1}\mathbf{f}\right)_i = \sum_j (A^{-1})_{i,j} f_j,$$

we can now have a simple "physical" interpretation of the matrix elements of A^{-1}:

$(A^{-1})_{i,j}$ is the **"effect" (solution) at i of a "source" at j** (a right-hand-side $\mathbf{e}_j$), [and column j of A^{-1} is the effect (solution) *everywhere* from the source at j].

Let us consider a couple of examples drawn from the finite-difference matrices for a discretized ∇^2 with Dirichlet (0) boundary conditions. Recall that one interpretation of the problem $\nabla^2 u = f$ is that u is a displacement of a stretched string (1d) or membrane(2d) and that $-f$ is a force density (pressure). Figure 1(top) shows several columns of A^{-1} for $\frac{\partial^2}{\partial x^2}$ in 1d, and it is very much like what you might intuitively expect if you pressed on a string at "one point." Figure 1(bottom) shows two column of A^{-1} for ∇^2 on a square 2d domain (e.g. a stretched square drum head), and again matches the intuition for what happens if you press at "one point." In both cases, the columns are, as expected, the solution from a source $\mathbf{e}_j$ at one point j (and the minimum points in the plot are indeed the points j).

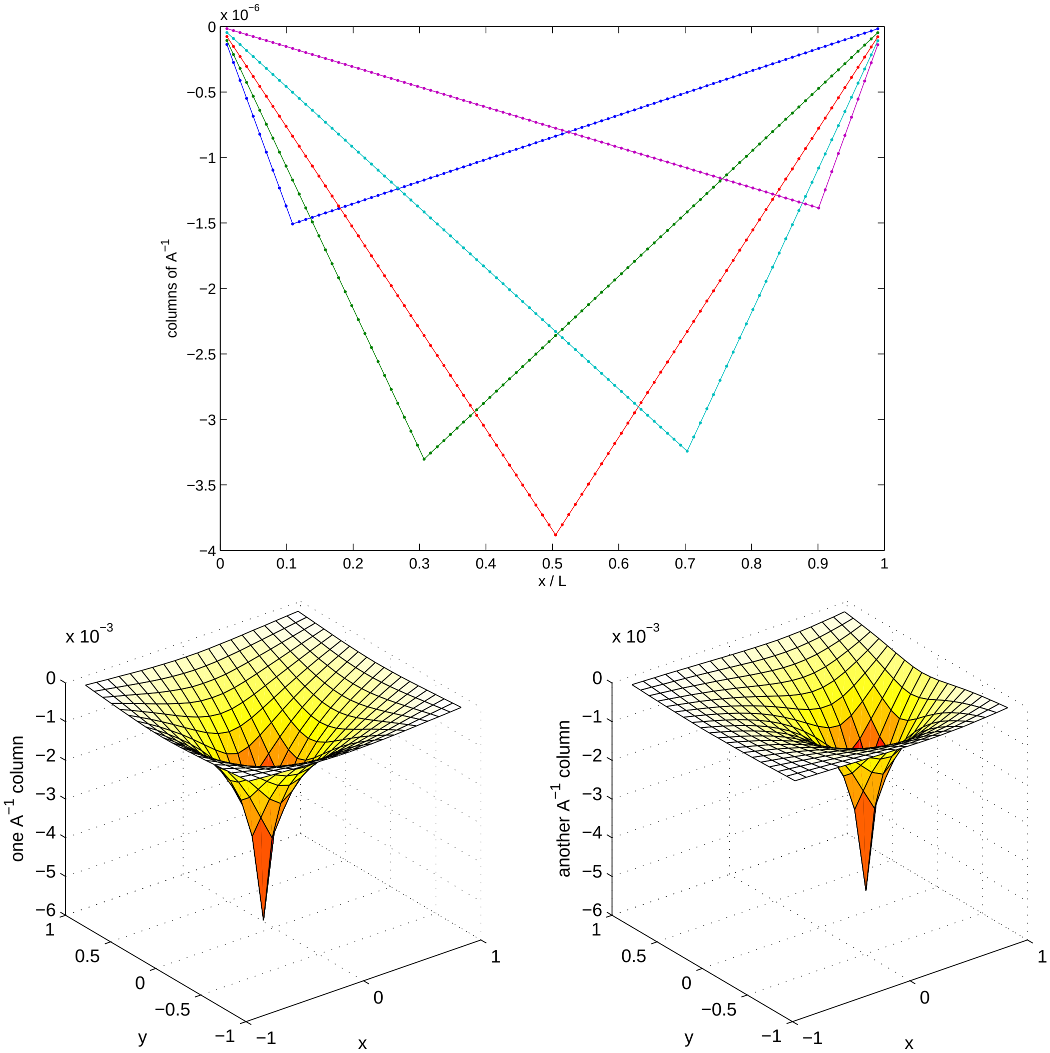

Figure 1: Columns of A^{-1} where A is a discretized ∇^2 with Dirichlet boundary conditions. *Top*: several columns in 1d ($\Omega = [0, L] = [0, 1]$). *Bottom*: two columns in 2d ($\Omega = [-1, 1] \times [-1, 1]$). In both 1d and 2d, the location of the minimum corresponds to the index of the column: this is the effect of a unit-vector "source" or "force" $= 1$ at that position (and $= 0$ elsewhere).

0.1 Inverses in other bases

This is not the only way to write down A^{-1}, of course, because $\mathbf{e}_j$ are only one possible choice of basis. For example, suppose that $\mathbf{q}_1, \ldots, \mathbf{q}_m$ is *any* orthonormal basis for our vector space, and that $\mathbf{u}_j$ solves $A\mathbf{u}_j = \mathbf{q}_j$. Then we can write any $\mathbf{f} = \sum_j \mathbf{q}_j(\mathbf{q}_j^* \mathbf{f})$ and hence the solution $\mathbf{u}$ to $A\mathbf{u} = \mathbf{f}$ by $\mathbf{u} = A^{-1}\mathbf{f} = \sum_j \mathbf{u}_j(\mathbf{q}_j^* \mathbf{f})$. In other words, we are writing $A^{-1} = \sum_j \mathbf{u}_j \mathbf{q}_j^* = \sum_j (A^{-1}\mathbf{q}_j)\mathbf{q}_j^*$ in terms of this basis. In matrix form, if Q is the (unitary) matrix whose columns are the $\mathbf{q}_j$, then we are simply writing $A^{-1} = (A^{-1}Q)Q^*$, since $QQ^* = I$. [Similarly, if the columns of B are any basis, not necessarily orthonormal, we can write $A^{-1} = (A^{-1}B)B^{-1}$: we solve any right-hand-side $\mathbf{f}$ by a superposition of the solutions for the columns of B with coefficients $B^{-1}\mathbf{f}$.]

For example, if A is Hermitian ($A = A^*$), then we can choose the $\mathbf{q}_j$ to be orthonormal eigenvectors of A ($A\mathbf{q}_j = \lambda_j \mathbf{q}_j$), in which case $\mathbf{u}_j = \mathbf{q}_j/\lambda_j$ and $\mathbf{u} = A^{-1}\mathbf{f} = \sum_j \mathbf{u}_j(\mathbf{q}_j^*\mathbf{f}) = \sum_j \mathbf{q}_j(\mathbf{q}_j^*\mathbf{f})/\lambda_j$. But this is just writing $A^{-1} = Q\Lambda^{-1}Q^*$ in terms of the (hopefully) familiar diagonalization $A = Q\Lambda Q^*$.

1 Operator inverses

Now we would like to consider the inverse $\hat{A}^{-1}$ of an operator $\hat{A}$, if such a thing exists. Suppose $N(\hat{A}) = \{0\}$ so that $\hat{A}u = f$ has a unique solution u (for all f in some suitable space of functions). $\hat{A}^{-1}$ is then **whatever operator gives the solution** $u = \hat{A}^{-1}f$ from any f. This must clearly be a *linear* operator, since if u_1 solves $\hat{A}u_1 = f_1$ and u_2 solves $\hat{A}u_2 = f_2$ then $u = \alpha u_1 + \beta u_2$ solves $\hat{A} = \alpha f_1 + \beta f_2$.

In fact, we already know how to write down $\hat{A}^{-1}$, for self-adjoint $\hat{A}$, in terms of an orthonormal basis of eigenfunctions u_n: $\hat{A}^{-1}f = \sum_n u_n \langle u_n, f\rangle/\lambda_n$. But now we would like to write $\hat{A}^{-1}$ in a different way: analogous to the entries $(A^{-1})_{ij}$ above, we would like to write the solution $u(\mathbf{x})$ as a superposition of **effects at x from sources at x′** for all $\mathbf{x}'$ (where $'$ just means another point, not a derivative):

$$\boxed{u(\mathbf{x}) = \hat{A}^{-1}f = \int_\Omega d^d\mathbf{x}' G(\mathbf{x},\mathbf{x}')f(\mathbf{x}')}, \tag{1}$$

where $G(\mathbf{x},\mathbf{x}')$ is the **Green's function** or the **fundamental solution** to the PDE. We need to figure out what G is (and whether it exists at all), and how to solve for it. This equation is the direct analogue of $u_i = \sum_j (A^{-1})_{i,j} f_j$ for the matrix case above, with sums turned into integrals, so this leads us to interpret $G(\mathbf{x},\mathbf{x}')$ as the "matrix element" $(\hat{A}^{-1})_{\mathbf{x},\mathbf{x}'}$. So, in analogy to the matrix case, we hope that $G(\mathbf{x},\mathbf{x}')$ **is the effect at x from a source "at" x′.** But what does this mean for a PDE?

1.1 $\hat{A}G$ and a "delta function"

If G exists, one way of getting at its properties is to consider what PDE $G(\mathbf{x},\mathbf{x}')$ might satisfy (as opposed to the integral equation above). In particular, we can try to apply the equation $\hat{A}u = f$, noting that $\hat{A}$ is an operator acting on $\mathbf{x}$ and *not* on $\mathbf{x}'$ so that it commutes with the $\int_\Omega d^d\mathbf{x}'$ integral:

$$\hat{A}u = \int_\Omega d^d\mathbf{x}'[\hat{A}G(\mathbf{x},\mathbf{x}')]f(\mathbf{x}') = f(\mathbf{x}).$$

This is telling us that $\hat{A}G(\mathbf{x},\mathbf{x}')$ must be some function that, when integrated against "any" $f(\mathbf{x}')$, gives us $f(\mathbf{x})$. What must such an $\hat{A}G(\mathbf{x},\mathbf{x}')$ look like? It must somehow be sharply peaked around $\mathbf{x}$ so that all of the integral's contributions come from $f(\mathbf{x})$, and in particular the integral must give some sort of **average of $f(\mathbf{x}')$ in an infinitesimal neighborhood of x**.

How could we get such an average? In one dimension for simplicity, consider the function

$$\delta_{\Delta x}(x) = \begin{cases} \frac{1}{\Delta x} & 0 \le x < \Delta x \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

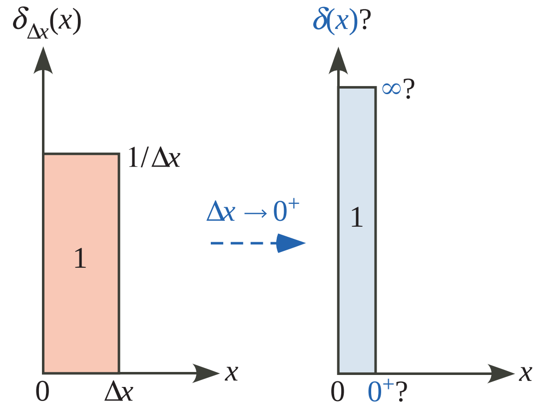

Figure 2: *Left:* A finite localized function $\delta_{\Delta x}(x)$ from eq. 2: a rectangle of area 1 on $[0, \Delta x]$. *Right:* We would like to take the $\Delta x \to 0^+$ limit to obtain a Dirac "delta function" $\delta(x)$, an infinitesimally wide and infinitely high "spike" of area 1, but this is not well defined according to our ordinary understanding of "function".

as depicted in figure 2(left): a box of width Δx and height $1/\Delta x$, with area 1. The integral

$$\int dx' \delta_{\Delta x}(x - x') f(x') = \frac{1}{\Delta x} \int_{x-\Delta x}^{x} f(x') dx$$

is just the average of $f(x')$ in $[x, x + \Delta x]$. As $\Delta x \to 0$, this goes to $f(x)$ (for continuous f), which is precisely the sort of integral we want. That is, we want

$$\hat{A}G(x, x') = \lim_{\Delta x \to 0} \delta_{\Delta x}(x - x') = \delta(x - x'), \tag{3}$$

where $\delta(x)$ is a "Dirac delta function" that is a box "infinitely high" and "infinitesimally narrow" with "area 1" around $x = 0$, as depicted in figure 2(right). Just one problem: **this is not a function**, at least as we currently understand "function," and the "definition" of $\delta(x)$ here doesn't actually make much rigorous sense.

In fact, it is possible to define $\delta(x)$, and even $\lim_{\Delta x \to 0} \delta_{\Delta x}(x)$, in a perfectly rigorous fashion, but it requires a new level of abstraction: we will have to redefine our notion of what a "function" is, replacing it with something called a *distribution*. We aren't ready for that, yet, though—we should see how far we can get with ordinary functions first. In fact, we will see that it is perfectly possible to obtain G as an ordinary function, even though $\hat{A}G$ will not exist in the ordinary sense (G will have some discontinuity or singularity that prevents it from being classically differentiable). But the process is fairly indirect and painful: once we obtain G the hard way, we will appreciate how much nicer it will be to have $\delta(x)$, even at the price of an extra layer of abstraction.

1.2 G by a limit of superpositions

Let's consider the same problem from another angle. We saw for matrices that we could interpret the columns of A^{-1} as the response to unit vectors $\mathbf{e}_j$, and $A^{-1}\mathbf{f}$ as representing $\mathbf{f}$ in the basis of unit vectors and then superimposing (summing) the solutions for each $\mathbf{e}_j$. Let's try to do the same thing here, at least approximately. Consider one dimension for simplicity, with a domain $\Omega = [0, L]$. Let's approximate $f(x)$ by a piecewise-constant function as shown in figure 3: a "sum of $N + 1$ rectangles" of width $\Delta x = L/(N + 1)$. Algebraically, that could be written as a sum in terms of precisely the $\delta_{\Delta x}(x)$ functions from (2) above:

$$f(x) \approx f_N(x) = \sum_{n=0}^{N} f(n\Delta x) \delta_{\Delta x}(x - n\Delta x) \Delta x, \tag{4}$$

since $\delta_{\Delta x}(x - n\Delta x)\Delta x$ is just a box of height 1 on $[n\Delta x, (n + 1)\Delta x]$. In the limit $N \to \infty$ ($\Delta x \to 0$), $f_N(x)$ clearly goes to the exact $f(x)$ (for continuous f). But we don't want to

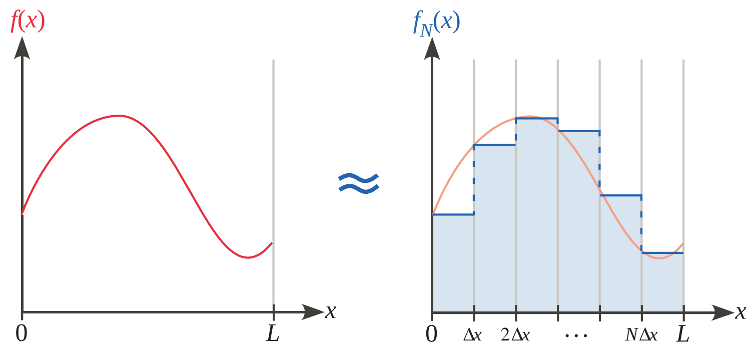

Figure 3: A function $f(x)$ (left) can be approximated by a piecewise-constant $f_N(x)$ (right) as in eq. (4), where each segment $[n\Delta x, (n+1)\Delta x]$ takes the value $f(n\Delta x)$.

take that limit yet. Instead, we will find the solution u_N to $\hat{A}u_N = f_N$, and *then* take the $N \to \infty$ limit.

For *finite* N, f_N is the superposition of a finite number of "box" functions that are "localized sources." Our strategy, similar to the matrix case, will be to solve [analogous to (3)]

$$\hat{A}g_n(x) = \delta_{\Delta x}(x - n\Delta x) \tag{5}$$

for each of these right-hand sides to obtain solutions $g_n(x)$, to write $u_N(x) = \sum_n f(n\Delta x)g_n(x)\Delta x$, and then to take the limit $\Delta x \to 0$ to obtain an integral

$$u(x) = \lim_{N\to\infty} u_N(x) = \lim_{\Delta x\to 0} \sum_{n=0}^{L/\Delta x} f(n\Delta x)g_n(x)\Delta x = \int_0^L G(x,x')f(x')dx',$$

where

$$G(x,x') = \lim_{\Delta x\to 0} g_{x'/\Delta x}(x). \tag{6}$$

We will find that, even though $\delta_{\Delta x}$ does *not* have a well-defined $\Delta x \to 0$ limit (at least, not according to the familiar definition of "function"), g_m *does* have a perfectly well-defined limit G as an ordinary function.

2 Example: Green's function of $-d^2/dx^2$

As an example, consider the familiar $\hat{A} = -\frac{d^2}{dx^2}$ on $[0, L]$ with Dirichlet boundaries $u(0) = u(L) = 0$. Following the approach above, we want to solve equation (5) and then take the $\Delta x \to 0$ limit to obtain $G(x, x')$. That is, we want to solve

$$-g_n''(x) = \delta_{\Delta x}(x - n\Delta x) \tag{7}$$

with the boundary conditions $g_n(0) = g_n(L) = 0$, and then take the limit (6) of g_n to get G.

As depicted in figure 4, we will solve this by breaking the solution up into three regions: (I) $x \in [0, n\Delta x)$; (II) $x \in [n\Delta x, (n+1)\Delta x)$; and (III) $x \in [(n+1)\Delta x, L]$. In region II, $-g_n'' = 1/\Delta x$, and the most general $g_m(x)$ with this property is the quadratic function $-\frac{x^2}{2\Delta x} + \gamma x + \kappa$ for some unknown constants γ and κ to be determined. In I and III, $g_n'' = 0$, so the solutions must be straight lines; combined with the boundary conditions at 0 and L, this gives a solution of the form:

$$g_n(x) = \begin{cases} \alpha x & x < n\Delta x \\ -\frac{x^2}{2\Delta x} + \gamma x + \kappa & x \in [n\Delta x, (n+1)\Delta x] \\ \beta(L - x) & x > (n+1)\Delta x \end{cases}, \tag{8}$$

for constants $\alpha, \beta, \gamma, \kappa$ to be determined. These unknowns are determined by imposing

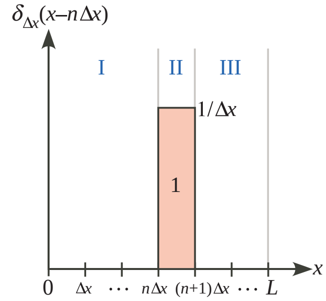

Figure 4: When solving $-g_n''(x) = \delta_{\Delta x}(x - n\Delta x)$, we will divide the solution into three regions: regions I and III, where the right-hand side is zero (and g_n is a line), and region II where the right-hand side is $1/\Delta x$ (and g_n is a quadratic).

continuity conditions at the interfaces between the regions, to make sure that the solutions "match up" with one another. Equation (7) tells us that both g_n and g_n' must be continuous in order to obtain a finite piecewise-continuous g_n''. Thus, we obtain four equations for our four unknowns:

$$\begin{aligned} \alpha n\Delta x &= -n^2\frac{\Delta x}{2} + \gamma n\Delta x + \kappa, \\ \alpha &= -n + \gamma, \\ \beta(L - [n+1]\Delta x) &= -(n+1)^2\frac{\Delta x}{2} + \gamma(n+1)\Delta x + \kappa, \\ -\beta &= -(n+1) + \gamma. \end{aligned}$$

After straightforward but tedious algebra, letting $x' = n\Delta x$, one can obtain the solutions:

$$\begin{aligned} \beta &= \frac{x' + \frac{\Delta x}{2}}{L}, \\ \alpha &= 1 - \beta, \\ \gamma &= \frac{x'}{\Delta x} + \alpha, \\ \kappa &= \frac{-x'^2}{2\Delta x}. \end{aligned}$$

The resulting function $g_m(x)$ is plotted for three values of x' in figure 5(left). It looks much like what one might expect: linear functions in regions I and III that are smoothly "patched together" by a quadratic-shaped kink in region II. Physically, this can be interpreted as the shape of a stretched string when you press on it with a localized pressure $\delta_{\Delta x}(x - x')$.

We are now in a position to take the $\Delta x \to 0$ limit, keeping $x' = n\Delta x$ fixed. Region II disappears, and we are left with $\beta \to \frac{x'}{L}$ in region III and $\alpha \to 1 - \frac{x'}{L}$ in region I:

$$G(x, x') = \begin{cases} \left(1 - \frac{x'}{L}\right) x & x < x' \\ \frac{x'}{L}(L - x) & x \geq x' \end{cases} = \begin{cases} \left(1 - \frac{x'}{L}\right) x & x < x' \\ (1 - \frac{x}{L})x' & x \geq x' \end{cases},$$

a pleasingly symmetrical function [whose symmetry $G(x, x') = G(x', x)$ is called *reciprocity* and, we will eventually show, derives from the self-adjointness of $\hat{A}$]. This function is plotted in figure 5(right), and is continuous but with a discontinuity in its slope at $x = x'$. [Indeed, it looks just like what we got from the discrete Laplacian in figure 1(top), except with a sign flip since we are looking at $-d^2/dx^2$.] Physically, this intuitively corresponds to what happens when you press on a stretched string at "one point."

Thus, we have obtained the Green's function $G(x, x')$, which is indeed a perfectly well-defined, ordinary function. As promised in the previous section, $\hat{A}G$ does *not* exist according to our ordinary definition of derivative, not even piecewise, since $\frac{\partial G}{\partial x}$ is discontinuous: the second derivative is "infinity" at x'. Correspondingly, $\delta_{\Delta x}$ does not have a well-defined limit

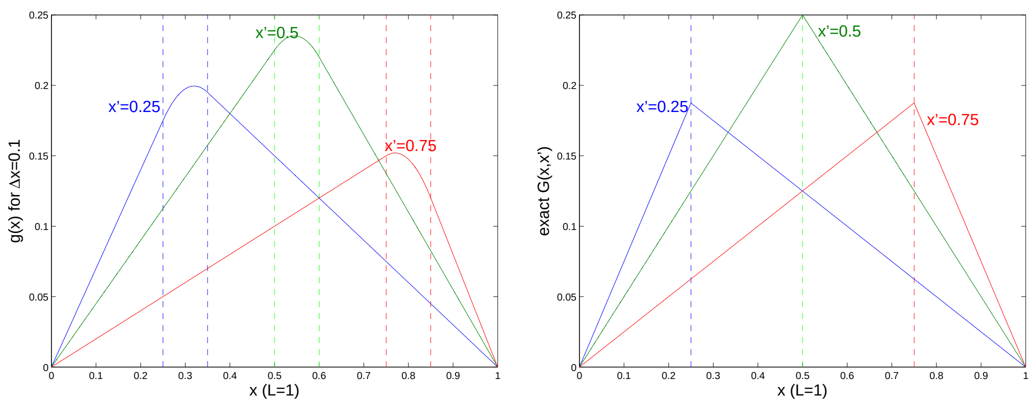

Figure 5: Green's function $G(x, x')$ for $\hat{A} = -\partial^2/\partial x^2$ on $[0, 1]$ with Dirichlet boundary conditions, for $x' = 0.25$, 0.5, and 0.75. *Left:* approximate Green's function, using our finite-width $\delta_{\Delta x}(x)$ as the right-hand-side. *Right*: exact Green's function, from $\Delta x \to 0$ limit.

as an ordinary function. So, even though the right-hand side of $\hat{A}g_n = \delta_{\Delta x}(x - n\Delta x)$ did not have a well-defined limit, the *solution* g_n did, and indeed we can then write the solution u for *any* f exactly by our integral (1) of G. If only there were an easier way to obtain G...

2.1 Deriving the slope discontinuity of G

Why is $\partial G/\partial x$ discontinuous at $x = x'$? It is easy to see this just by integrating both sides of equation (7) for $g_n(x)$:

$$-\int_{n\Delta x}^{(n+1)\Delta x} g_n''(x)\,dx = g_n'(n\Delta x) - g_n'(n\Delta x + \Delta x) = \int_{n\Delta x}^{(n+1)\Delta x} \delta_{\Delta x}(x)\,dx = 1.$$

Thus, the slope g' drops by 1 when going from region I to region II. The limit $\Delta x \to 0^+$ of these expressions (the *integrals*) is perfectly well defined (even though the limit of the *integrands* is not), and gives

$$\left.\frac{\partial G}{\partial x}\right|_{x=x'^-} - \left.\frac{\partial G}{\partial x}\right|_{x=x'^+} = 1,$$

so the slope $\partial G/\partial x$ must drop discontinuously by 1 at $x = x'$.

In fact, we can use this to easily solve for $G(x, x')$. For $x < x'$ (region I) and for $x > x'$ (region III) we must have $\hat{A}G = 0$ and hence G is a line: $G(x, x') = \begin{cases} \alpha x & x < x' \\ \beta(L - x) & x > x' \end{cases}$. We have two unknown constants α and β, but we have two equations: continuity of G, $\alpha x' = \beta(L - x')$; and our slope condition $\alpha - (-\beta) = 1$. Combining these, one can quickly find that $\alpha = 1 - \frac{x'}{L}$ and $\beta = \frac{x'}{L}$ as above!

3 To infinity... and beyond!

It is perfectly possible to work with Green's functions and $\hat{A}^{-1}$ in this way, putting finite $\delta_{\Delta x}$-like sources on the right-hand sides and only taking the $\Delta x \to 0^+$ limit or its equivalent *after* solving the PDE. Physically, this corresponds to putting in a source in a small but non-infinitesimal region (e.g. a force in a small region for a stretched string, or a small ball of charge for electrostatics), and then taking the limit of the *solution* to obtain the limiting behavior as the source becomes "infinitely concentrated at a point" (e.g. to obtain the displacement for a string pressed at one "point", or the potential for a "point" charge). However, this is all terribly cumbersome and complicated.

It would be so much simpler if we could have an "infinitely concentrated" source at a "point" to start with. For example, when you study freshman mechanics or electromagnetism, you start with "point masses" and "point charges" and only *later* do you consider extended distributions of mass or charge. But to do that in a rigorous way seems problematic because "$\delta(x) = \lim_{\Delta x \to 0} \delta_{\Delta x}(x)$" apparently does not exist as a "function" in any sense that you have learned—it does not have a value for every x.

It turns out that the problem is not with $\delta(x)$, the problem is with the definition of a function—surprisingly, classical functions don't really correspond to what we want in many practical circumstances, and they are often not quite what we want to deal with for PDEs. A clue is provided in section 2.1, where we saw that even though our *integrands* might not always have well-defined limits, our *integrals* did. If we can only define a "function" in terms of its *integrals* rather than by its values at *points*, that might fix our problems. In fact, we can do something very much like this, which leads to the concept of a **distribution** (or *generalized function*) or a **weak solution**. At the cost of an increase in the level of abstraction, we can completely eliminate the tortured contortions that were required above (and only get worse in higher dimensions) to avoid delta functions.