Notes on the convergence of trapezoidal-rule quadrature

Steven G. Johnson

September 17, 2008

1 Introduction

Numerical quadrature is another name for numerical integration, which refers to the approximation of an integral $\int f(x)dx$ of some function $f(x)$ by a discrete summation $\sum w_i f(x_i)$ over points x_i with some weights w_i. There are many methods of numerical quadrature corresponding to different choices of points x_i and weights w_i, from Euler integration to sophisticated methods such as Gaussian quadrature, with varying degrees of accuracy for various types of functions $f(x)$. In this note, we examine the accuracy of one of the simplest methods: the trapezoidal rule with uniformly spaced points. In particular, we discuss how the convergence rate of this method is determined by the smoothness properties of $f(x)$—and, in practice, usually by the smoothness at the endpoints. (This behavior is the basis of a more sophisticated method, Clenshaw-Curtis quadrature, which is essentially trapezoidal integration plus a coordinate transformation to remove the endpoint problem.)

For simplicity, without loss of generality, we can take the integral to be for $x \in [0, 2\pi]$, i.e. the integral

$$I = \int_0^{2\pi} f(x)dx,$$

which is approximated in the trapezoidal rule[1] by the summation:

$$I_N = \frac{f(0)\Delta x}{2} + \sum_{n=1}^{N-1} f(n\Delta x)\Delta x + \frac{f(2\pi)\Delta x}{2},$$

where $\Delta x = \frac{2\pi}{N}$.

We now want to analyze how fast the error $E_N = |I - I_N|$ decreases with N. Many books estimate the error as being $O(\Delta x^2) = O(N^{-2})$, assuming $f(x)$ is twice differentiable on $(0, 2\pi)$—this estimate is correct, but only as an upper bound. For many interesting functions, the error can decrease much, much faster than that, as discussed below.

[1]Technically, this is the "composite" trapezoidal rule, where the "trapezoidal rule" by itself refers to the approximation $[f(x) + f(x + \Delta x)]\Delta x/2$ for a single Δx interval.

2 A Simple, Pessimistic Upper Bound

A simple, but perhaps too pessimistic, upper bound is as follows. The trapezoidal rule corresponds to approximating $f(x)$ by a straight line on each interval Δx, which means that the error is the integral of a quadratic remainder (to lowest order). The integral of a quadratic over a Δx is $O(\Delta x^3)$: this is the *local truncation error* over each interval. There are $N - 1 = O(\Delta x^{-1})$ intervals, so the total error is $O(N\Delta x^3) = O(\Delta x^2) = O(N^{-2})$. (The same bound can be derived in a number of ways, more formally via integration by parts.) However, it is important to emphasize that this is only an upper bound: we didn't take into account the possibility of cancellations in the errors between different intervals.

3 Quadrature error via Fourier analysis

One way to analyze the error more thoroughly is to consider the Fourier-series expansion of the function $f(x)$.

$$f(x) = \frac{1}{2\pi} \sum_{m=-\infty}^{\infty} c_m e^{imx}$$

with

$$c_m = \int_0^{2\pi} f(x) e^{-imx} dx.$$

Obviously, $I = c_0$. But now I_N is easy to evaluate.

$$I_N = \sum_{n=0}^{N-1} \frac{1}{2\pi} \sum_{m=-\infty}^{\infty} c_m e^{imn\Delta x} \Delta x = \sum_{m=-\infty}^{\infty} c_m \frac{1}{N} \sum_{n=0}^{N-1} e^{\frac{2\pi i}{N} mn}.$$

We have assumed that $f(x)$ has a convergent Fourier series at the points $x_n = 2\pi n/N$, which is true if $\int_0^{2\pi} |f(x)|^p dx < \infty$ for some $p > 1$ and if the periodic extension of $f(x)$ is continuous at those points x_n. At the endpoints $x = 0$ and $x = 2\pi$, the Fourier series will converge to $[f(0) + f(2\pi)]/2$ (i.e. we generally have an effective jump discontinuity at the endpoints), but this exactly matches how the endpoints are handled in the trapezoidal rule, which is why we were able to replace this term with the Fourier series at $x = 0$ in the above expression. The final summation simplifies enormously because $\sum_n e^{2\pi imn/N}$ is zero unless m is an integer multiple of N, in which case the sum is N. Therefore

$$I_N = \sum_{k=-\infty}^{\infty} c_{kN},$$

and the error in the trapezoidal rule is

$$E_N = |I - I_N| = \left| \sum_{k\neq 0} c_{kN} \right| = \left| \sum_{k=1}^{\infty} (c_{kN} + c_{-kN}) \right|,$$

which transforms the question of error analysis into a question of the convergence rate of the Fourier series expansion of $f(x)$. But the convergence rate of the Fourier series is determined by the smoothness of the function $f(x)$...or rather, of its periodic extension, so we have to include the periodicity of $f(x)$ and its derivatives at the endpoints of the integration interval.

For example, suppose that the periodic extension of $f(x)$ is ℓ times differentiable and $f^{(\ell)}(x)$ is piecewise continuous with some jump discontinuities. In this case, it is straighforward to show via integration by parts that c_m goes asymptotically as $1/m^{\ell+1}$. In this case, E_N is $O(\frac{1}{N^{\ell+1}})$, since we can just (asymptotically, for large N) pull out the $1/N^{\ell+1}$ factor from each term in the sum (which then becomes some convergent series independent of N). However, it turns out that even this is an overestimate if the discontinuity occurs precisely at the *endpoints* $f^{(\ell)}(0) \neq f^{(\ell)}(2\pi)$. In this case, as we shall see below, when ℓ is *even* we get an additional cancellation and the convergence is $O(1/N^{\ell+2})$; when ℓ is odd the convergence is still $O(1/N^{\ell+1})$.

.As another example, suppose that the periodic extension of $f(x)$ is an analytic function (infinitely differentiable) with poles a nonzero distance from the real axis—in this case, the Fourier series converges exponentially fast, and hence E_N decreases *exponentially* with N. In general, it follows from above that any infinitely differentiable periodic $f(x)$ will have error that vanishes faster than any polynomial in $1/N$, but exactly how much faster will depend upon the nature of $f(x)$ and its singularities.

How does this error analysis compare with our $O(N^{-2})$ estimate from earlier? If $f(x)$ is an arbitrary differentiable function on $(0, 2\pi)$, then we must in general assume $f(0) \neq f(2\pi)$, and so the periodic extension of $f(x)$ is discontinuous. Hence we can apply our analysis from above, conclude that $\ell = 0$, and hence the error is $O(N^{-1})$ ...or rather, $O(N^{-2})$, as long as we apply the correction alluded to above, that we always round up the exponent $\ell + 1$ to the next *even* integer when the discontinuity occurs at the endpoints of the integration interval.

3.1 Convergence rate from the Fourier series

To obtain the convergence rate of the quadrature error, we need to find the asymptotic convergence rate of the Fourier series coefficients c_m. This is a rather standard analysis, but we repeate it here both as a review and because something interesting occurs when the first discontinuity occurs at the endpoints, as we alluded to above.

The most common case is where $f(x)$ [or rather, its periodic extension] has its first discontinuities in its ℓ-th derivative, and countably many jump discontinuities in general. That is, the periodic extension of $f^{(\ell)}(x)$ exists but is only piecewise continuous, with countably many jump discontinuities [most often a mismatch in the endpoints $f^{(\ell)}(0) \neq f^{(\ell)}(2\pi)$], while all lower derivatives are continuous and periodic. In this case, we simply integrate by parts $\ell + 1$ times, until we obtain delta functions from the jump discontinuities. Let

$$f^{(\ell+1)}(x) = \sum_j a_j \delta(x - x_j) + g(x)$$

for some bounded piecewise-continuous function $g(x)$ and delta functions corresponding to jump discontinuities at x_j in the periodic extension of $f^{(\ell)}(x)$. Then, integrating by parts in the Fourier integral for c_m (for $m \neq 0$), we obtain:

$$c_m = \int_0^{2\pi} f(x)e^{-imx}dx = \left.\frac{if(x)e^{-imx}}{m}\right|_0^{2\pi} - \frac{i}{m}\int_0^{2\pi} f'(x)e^{-imx}dx = \cdots$$
$$= \left(-\frac{i}{m}\right)^{\ell+1}\int_0^{2\pi} f^{(\ell+1)}(x)e^{-imx}dx = \left(-\frac{i}{m}\right)^{\ell+1}\left[\sum_{x_j\in[0,2\pi)} a_j e^{-imx_j} + \int_0^{2\pi} g(x)e^{-imx}dx\right]$$
$$= \left(-\frac{i}{m}\right)^{\ell+1}\left[\sum_{x_j\in[0,2\pi)} a_j e^{-imx_j} + \frac{i}{m}\left\{ g(x)e^{-imx}\Big|_0^{2\pi} - \int_0^{2\pi} g'(x)e^{-imx}dx\right\}\right],$$

where all of the boundary terms from integration by parts ℓ times are zero because of the assumed periodicity of the derivatives $< \ell$. If $f^{(\ell)}(0) \neq f^{(\ell)}(2\pi)$, then there will be a boundary term from the $(\ell+1)$-st integration by parts, but this is included above from the a_j term at $x_j = 0$.

We have also integrated by parts one last time on the $g(x)$ term; the $g'(x)$ integrand may include delta functions since $g(x)$ is only piecewise continuous, but the important point is that the $\{\cdots\}$ integral has a bounded magnitude. In particular, we can write $g'(x) = \sum_j b_j\delta(x-y_j)+h(x)$ for some bounded piecewise continuous function $h(x)$, and then

$$\left| g(x)e^{-imx}\Big|_0^{2\pi} - \int_0^{2\pi} g'(x)e^{-imx}dx\right| = \left|\sum_{y_j\in[0,2\pi)} b_j e^{-imy_j} + \int_0^{2\pi} h(x)e^{-imx}dx\right|$$
$$\leq \sum_{y_j\in[0,2\pi)} |b_j| + \int_0^{2\pi} |h(x)|dx.$$

Therefore, when we look at the asymptotic behavior as m grows large, we immediately find

$$|c_m| = \left|\left(-\frac{i}{m}\right)^{\ell+1}\left[\sum_{x_j\in[0,2\pi)} a_j e^{-imx_j} + O(1/m)\right]\right| \leq \frac{1}{m^{\ell+1}}\left[\sum_{x_j\in[0,2\pi)} |a_j| + O(1/m)\right]$$
$$= O(1/m^{\ell+1}).$$

We thus obtain an upper bound for the quadrature error $E_N = O(1/N^{\ell+1})$, as given in the previous section.

However, as mentioned in the previous section, this bound is too pessimistic in one especially common case: suppose all of the discontinuities (or at least, the lowest-order discontinuities) are at the endpoints, i.e. the only $x_j \in [0, 2\pi)$ is $x_0 = 0$. In that case, we can exploit the fact that we are *not* interested in c_m alone, but rather our error is a summation of terms of the form $c_m + c_{-m}$. We therefore obtain, in this special case of endpoint discontinuities:

$$|c_m + c_{-m}| = \left|\left\{\left(-\frac{i}{m}\right)^{\ell+1} + \left(+\frac{i}{m}\right)^{\ell+1}\right\} a_0 + O\left(\frac{1}{m^{\ell+2}}\right)\right| = \begin{cases} O\left(\frac{1}{m^{\ell+2}}\right) & \ell \text{ even} \\ O\left(\frac{1}{m^{\ell+1}}\right) & \ell \text{ odd} \end{cases}.$$

That is, in this special case the $1/m^{\ell+1}$ terms exactly cancel when ℓ is even, and we increase the order of convergence by one!

An equivalent result states that, if $f(x)$ is infinitely differentiable on the interval $(0, 2\pi)$, then the error E_N can be expanded as an explicit power series in Δx in which only *even* powers of Δx are present. A proof (not using Fourier series) is given in, for example *An Introduction to Numerical Analysis* by Suli and Mayers.[2] However, I find the Fourier series approach much nicer—the convergence analysis of the Fourier series is very standard, straightforward (at least if you avoid the really pathological functions), and well known, including many more situations than the ones I have discussed above. (The analysis of Suli and Mayers involves a relatively unfamiliar set of polynomial basis functions.)

4 Clenshaw-Curtis quadrature

Even if $f(x)$ is a nice, smooth function inside the integration interval, it is often not periodic at its endpoints, which is what commonly reduces the error of the trapezoidal rule to the pessimistic $O(N^{-2})$ bound. However, this can be fixed by performing a change of variables, which is the basic idea behind Clenshaw-Curtis quadrature.

For simplicity let us assume we are integrating for $x \in [-1, 1]$ rather than $[0, 2\pi]$. In this case, we make the substitution $x = \cos\theta$, and obtain the integral:

$$I = \int_{-1}^{1} f(x)dx = \int_{0}^{\pi} f(\cos\theta)\sin\theta d\theta.$$

Now, $f(\cos\theta)$ is by construction nice and periodic, so we would like to extend this integral to $[-\pi, \pi]$ and use trapezoidal integration as above. However, this is spoiled by the $\sin\theta$ term, which would make the integral on $[-\pi, \pi]$ zero, so we use one additional trick: we replace $f(\cos\theta)$ by its cosine series, integrate each cosine term against sine analytically, and obtain the coefficients in the cosine series by trapezoidal quadrature. That is:

$$f(\cos\theta) = \frac{a_0}{2} + \sum_{m=1}^{\infty} a_{,m}\cos(m\theta),$$

$$a_m = \frac{2}{\pi}\int_{0}^{\pi} f(\cos\theta)\cos(m\theta)d\theta = \frac{1}{\pi}\int_{-\pi}^{\pi} f(\cos\theta)\cos(m\theta)d\theta,$$

$$I = a_0 + \sum_{k=1}^{\infty}\frac{2a_{2k}}{1-(2k)^2}.$$

Here, the integral to obtain a_m is over $[-\pi, \pi]$ of a periodic function $f(\cos\theta)\cos(m\theta)$ whose smoothness is determined by that of $f(x)$ on $x \in (-1, 1)$, regardless of whether $f(x)$ itself is periodic. Therefore, if $f(x)$ is sufficiently smooth, the coefficients a_m converge quite rapidly to zero as discussed above (even exponentially fast if $f(x)$ is analytic), and can be computed by trapezoidal-rule integration with error that converges

[2] See also the online course notes by Patch Kessler at `http://www.me.berkeley.edu/~watchwrk/math128a/extrap.pdf`.

to zero at the same rate. Moreover, it turns out that the trapezoidal-rule integration with $2N-1$ points is equivalent to a type-I discrete cosine transform, and can be evaluated very rapidly for $m = 0, \ldots, N$ simultaneously via fast Fourier transform methods. [Clenshaw-Curtis quadrature can also be viewed as expansion of $f(x)$ in Chebyshev polynomials $T_m(x)$, since by definition $T_m(x) = \cos(m \cos^{-1} x)$.]