From *SIAM News*, Volume 44, Number 9, November 2011

ICIAM 2011

Understanding Nelder–Mead: The Game Is Afoot

By Barry A. Cipra

Anyone who owns a personal computer these days is accustomed to things that work in mysterious ways. Increasingly, we depend on slick programs that generally get the job done but sometimes fail spectacularly. We like to think that someone, somewhere actually understands the informational ballistics we trigger with our keyboard and touchpad interactions, but we can no longer be certain of that—if we ever could.

Not surprisingly, perhaps, these uncertainties apply even when what the computer is doing is utterly, algorithmically transparent. Take the Nelder–Mead "direct search" method, for instance. Developed in the mid-1960s, Nelder–Mead (named after the late John Nelder and Roger Mead) finds minima of functions that don't have easily computed derivatives. For a function of n variables, the algorithm takes a simplex of $n + 1$ points (e.g., the vertices of a triangle in the plane, when $n = 2$) and sequentially alters it by a systematic process of reflections, expansions, contractions, and reductions. The aim is to replace the "worst" point (that is, the point of the simplex where the function takes its largest value) with a better point. Nelder–Mead is widely used in practice, and very often (but not always) works well: Even if the starting simplex is far from any minimum, it typically lopes off in the direction of one and, once in the vicinity, will shrink down around it, converging to a desired solution. But researchers have been hard pressed to describe settings in which the algorithm will succeed—in part because of the many cases in which it's known *not* to work. Today, many experts on derivative-free methods look askance (or worse) at the Nelder–Mead method, preferring more modern methods with rigorous convergence proofs. Nonetheless, despite its well-known theoretical shortcomings, the method remains popular with practitioners because it is easy to implement and seems easy to understand.

In a session at ICIAM organized by the Association for Women in Mathematics, Margaret Wright of New York University described a new result that may offer some insight into Nelder–Mead's propensity to succeed. Wright and colleagues Jeff Lagarias of the University of Michigan and Bjorn Poonen of MIT have proved that a variant known as the restricted Nelder–Mead algorithm necessarily converges to the unique minimum for a specific and restricted class of functions of two variables. Their proof, Wright notes, is a kind of "Sherlock Holmes" argument: It eliminates the impossible, so that what remains must be the truth.

Researchers have been hard pressed to describe settings in which the Nelder–Mead algorithm will succeed—in part because of the many cases in which it's known not *to work.*

The restricted Nelder–Mead algorithm swears off expansions. In practice, expansion serves as an accelerant when the initial simplex starts far from a minimum. Each iteration begins by reflecting the "worst" point of the simplex through the centroid of the other points. If the reflected point turns out to be a new "best" point, the algorithm tries the point twice as far away, in hopes that it will be even better. In dimension 1, this amounts to doubling the size of successive steps until the algorithm knows it has a minimum trapped in an interval. This doesn't work very well, of course, on a function like e^x; even in one dimension, you need assumptions on the functions you're dealing with in order to say something about how well the algorithm works. By ruling out hiccups among the iterations, the restriction makes for amenable analysis.

A function to which the new result applies must have a positive definite Hessian (matrix of second derivatives) and bounded level sets. Such functions have unique minima. The positive definiteness of the Hessian is a strong assumption—it precludes functions like $x^4 + y^4$, for example—but the proof depends on it. That's because of a famous counterexample.

In 1996, Ken McKinnon of the University of Edinburgh showed why it's so hard to prove much about Nelder–Mead when he produced a class of functions for which the algorithm converges, but to the wrong point. One simple example is the function $f(x,y) = y + y^2 + cx^3$, where $c = -2400$ for $x < 0$, and $c = 6$ otherwise. If you started with a cunningly chosen simplex having vertices at (0,0), (1,1), and $((1 - \sqrt{33})/8)$, $((1 + \sqrt{33})/8)$, Nelder–Mead would produce a sequence of increasingly skinny triangles oriented along the x-axis and converging to the origin—which is not where the function takes its minimum. Basically, each trial reflection produces a point with a small negative value of y, but the coefficient –2400 produces such a large functional value there that it scares the algorithm back to an "interior contraction" in which the "worst" point is drawn halfway toward the centroid of the other two points. (You can of course try variants of Nelder–Mead that replace "halfway" by some other fraction, but all such variants are susceptible to cleverly concocted counterexamples.)

Because the reflected points in the McKinnon counterexample are always worse than what they would replace, expansion is never considered. The McKinnon function would thus fool even the restricted Nelder–Mead algorithm. It does have bounded level sets, and its Hessian is positive definite, but with one exception—at the origin. Sometimes all it takes is one exception.

In earlier work, also dating to 1996, Wright and Lagarias, with colleagues James Reeds and Paul Wright, all then at Bell Labs, had shown that the unrestricted Nelder–Mead algorithm acting on strictly convex functions with bounded level sets in dimensions 1 and 2 gets the diameter of the simplex to go to zero. In dimension 1, that happens only when the simplex (an interval) shrinks down to the minimum. But in dimension 2, as the McKinnon example shows, the simplex can shrink to a point—or, conceivably, a nonconvergent sequence of points—on some level set other than the minimum.

The new result relies on this earlier work. Positive definiteness of the Hessian is a stronger condition than strict convexity. In essence, positive definiteness means that, at any point other than the minimum, there is a linear affine transformation taking the given point to the origin and rescaling the function to have the form $f(x,y) = x^2 + y + ay^2$, with $a > 0$, plus an error term that is negligible compared to x^2 and y^2 (for values near 0). If there were a counterexample, Wright and colleagues show, the simplex would have to "flatten out" into the shape of a microscopic needle oriented in the direction of the level set corresponding to the non-minimizing value that's being attained. (The true minimum has a single point for a level set; all other values have convex curves, with well-defined tangents. A microscopic simplex "sees" the level set as essentially a straight line.)

For the diameter of the simplex to go to zero, the algorithm must invoke contraction from time to time: In two dimensions, reflection changes neither the size nor the shape of a triangle. (The only other way that Nelder–Mead makes things smaller—"reduction," in which all vertices are drawn halfway toward the current "best" vertex—is never invoked for functions that are strictly convex.) After introducing a technical measure of flatness, Lagarias, Poonen, and Wright prove two contradictory properties for a sequence of simplices that would constitute a counterexample. On the one hand, once the simplex is sufficiently close to its limiting, non-minimal level set, a contraction step is possible only if the flatness is less than 10. On the other hand, for every 14 iterations of the Nelder–Mead algorithm, flatness grows by a factor of at least 1.01. The technical delicacy of the proof, which is evident from the numbers 14 and 1.01, resides in the fact that the affine transformation that simplifies the form of the function in a neighborhood of the simplex must adapt to its wanderings along the level set.

A proper theory of Nelder–Mead is still a long way off, especially if you're working with functions of three or more variables. Even in two dimensions, it would be nice to have confidence in the unrestricted algorithm. As McKinnon pointed out in 1996, there was no proof that Nelder–Mead would necessarily locate the minimum of so simple a function as $f(x,y) = x^2 + y^2$. What was—and still is—lacking is proof that there can be no cunning choice of a starting triangle that makes the algorithm go astray. The new result shows that any such counterexample would necessarily invoke occasional expansion steps. Getting even that much is likely to require new ideas. As Sherlock Holmes might put it, "I never said it was elementary, my dear Watson."

Barry A. Cipra is a mathematician and writer based in Northfield, Minnesota.