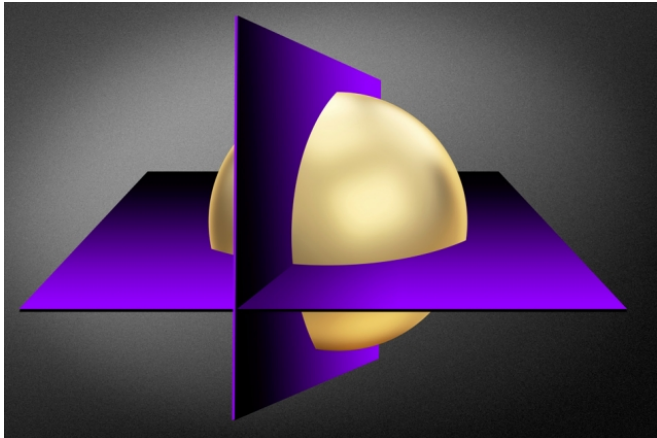


New general-purpose optimization algorithm promises order-of-magnitude speedups on some problems

23 October 2015, by Larry Hardesty



"Cutting plane" methods converge on the optimal values of a mathematical function by repeatedly cutting out regions of a much larger set of possibilities (gold sphere). Credit: Jose-Luis Olivares/MIT

Optimization problems are everywhere in engineering: Balancing design tradeoffs is an optimization problem, as are scheduling and logistical planning. The theory—and sometimes the implementation—of control systems relies heavily on optimization, and so does machine learning, which has been the basis of most recent advances in artificial intelligence.

This week, at the IEEE Symposium on Foundations of Computer Science, a trio of present and past MIT graduate students won a best-student-paper award for a new "cutting-plane" algorithm, a general-purpose algorithm for solving [optimization problems](#). The algorithm improves on the running time of its most efficient predecessor, and the researchers offer some reason to think that they may have reached the theoretical limit.

But they also present a new method for applying

their general algorithm to specific problems, which yields huge efficiency gains—several orders of magnitude.

"What we are trying to do is revive people's interest in the general problem the algorithm solves," says Yin-Tat Lee, an MIT graduate student in mathematics and one of the paper's co-authors. "Previously, people needed to devise different algorithms for each problem, and then they needed to optimize them for a long time. Now we are saying, if for many problems, you have one algorithm, then, in practice, we can try to optimize over one algorithm instead of many algorithms, and we may have a better chance to get faster algorithms for many problems."

Lee is joined on the paper by Aaron Sidford, who was an MIT graduate student in electrical engineering and computer science when the work was done but is now at Microsoft Research New England, and by Sam Wong, who earned bachelor's and master's degrees in math and electrical engineering and computer science at MIT before moving to the University of California at Berkeley for his PhD.

Inner circle

Optimization problems are generally framed as trying to find the minimum value of a mathematical function, called a "cost function." In car design, for example, the cost function might impose penalties for weight and drag but reward legroom and visibility; in an algorithm for object detection, the cost function would reward correct classification of various objects and penalize false positives.

At a very general level, finding the minimum of a cost function can be described as trying to find a small cluster of values amid a much larger set of

possibilities. Suppose that the total range of possible values for a cost function is represented by the interior of a circle. In a standard optimization problem, the values clustered around the minimum value would then be represented by a much smaller circle inside of the first one. But you don't know where it is.

Now pick a point at random inside the bigger circle. In standard optimization problems, it's generally possible to determine whether that point lies within the smaller circle. If it doesn't, it's also possible to draw a line that falls between it and the smaller circle.

Drawing that line cuts off a chunk of the circle, eliminating a range of possibilities. With each new random point you pick, you chop off another section of the circle, until you converge on the solution.

If you represent the range of possibilities as a sphere rather than a circle, then you use a plane, rather than a line, to cut some of them off. Hence the name for the technique: the cutting-plane method.

In most real optimization problems, you need a higher-dimensional object than either a circle or a sphere: You need a hypersphere, which you cut with a hyperplane. But the principle remains the same.

A matter of time

Theoretical computer scientists measure algorithms' running times not in seconds or hours, but in the number of operations required, relative to the number of elements being manipulated. With cutting-plane methods, the number of elements is the number of variables in the cost function—the weight of the car, the cost of its materials, drag, legroom, and so on. That's also the dimension of the hypersphere.

With the best general-purpose cutting-plane method, the time required to select each new point to test was proportional to the number of elements raised to the power 3.373. Sidford, Lee, and Wong get that down to 3.

But they also describe a new way to adapt cutting-plane methods to particular types of optimization problems, with names like submodular minimization, submodular flow, matroid intersection, and semidefinite programming. And in many of those cases, they report dramatic improvements in efficiency, from running times that scale with the fifth or sixth power of the number of variables (n^5 or n^6 , in [computer science](#) parlance) down to the second or third power (n^2 or n^3).

"This is indeed an astonishing paper," says Satoru Iwata, a professor of mathematical informatics at the University of Tokyo, who has published widely on the problem of submodular minimization. "For this problem," he says, "the running time bounds derived with the aid of discrete geometry and combinatorial techniques are by far better than what I could imagine."

More information: A Faster Cutting Plane Method and its Implications for Combinatorial and Convex Optimization. arxiv.org/abs/1508.04874

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

APA citation: New general-purpose optimization algorithm promises order-of-magnitude speedups on some problems (2015, October 23) retrieved 16 April 2021 from <https://phys.org/news/2015-10-general-purpose-optimization-algorithm-order-of-magnitude-speedups.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.