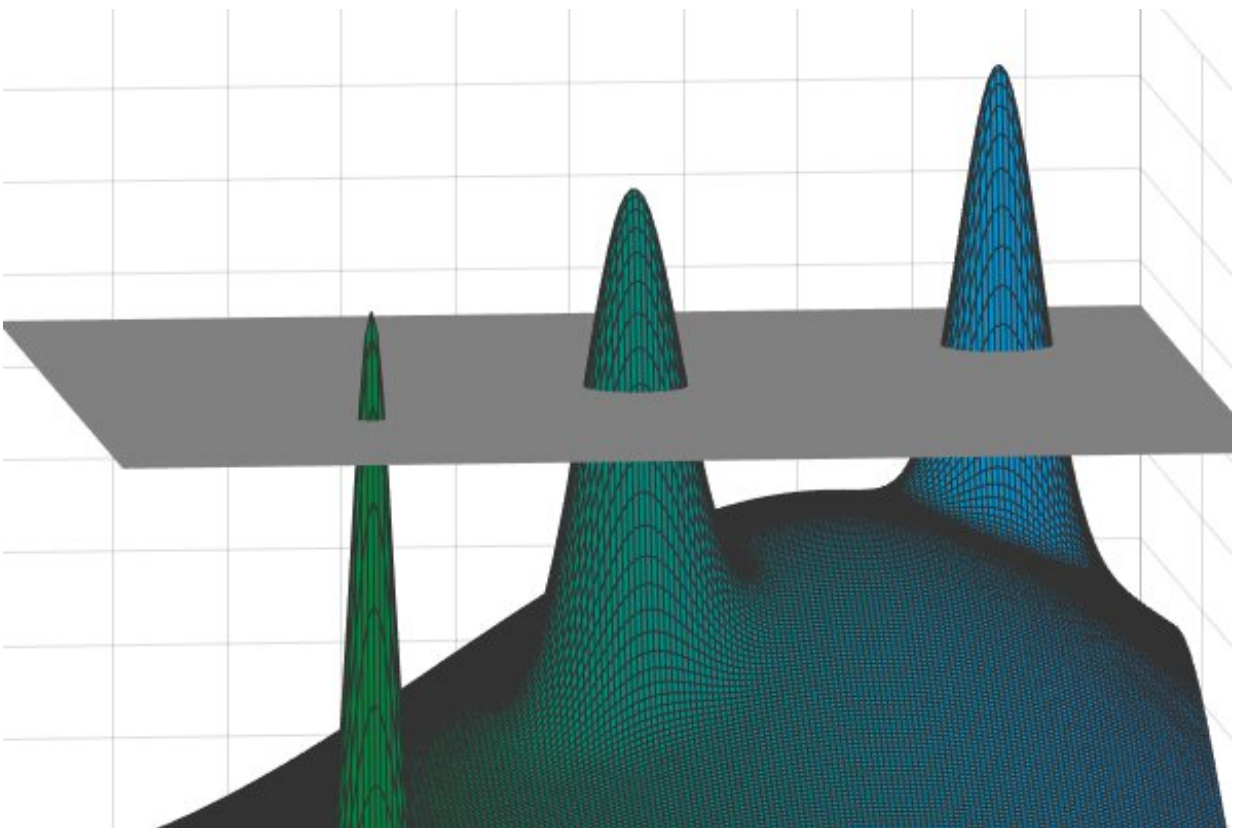# Researcher develops new tool for understanding hard computational problems that appear intractable

January 10 2022, by Steve Nadis

In some cases, the diameter of each peak will be much smaller than the distances between different peaks. Consequently, if one were to pick any two points on this sprawling landscape—any two possible "solutions"—they would either be very close (if they came from the same peak) or very far apart (if drawn from different peaks). In other words, there would be a telltale "gap" in these distances—either small or large, but nothing in between. Credit: David

Gamarnik et al.

The notion that some computational problems in math and computer science can be hard should come as no surprise. There is, in fact, an entire class of problems deemed impossible to solve algorithmically. Just below this class lie slightly "easier" problems that are less well-understood—and may be impossible, too.

David Gamarnik, professor of operations research at the MIT Sloan School of Management and the Institute for Data, Systems, and Society, is focusing his attention on the latter, less-studied category of problems, which are more relevant to the everyday world because they involve randomness—an integral feature of natural systems. He and his colleagues have developed a potent tool for analyzing these problems called the overlap gap property (or OGP). Gamarnik described the new methodology in a recent paper in the *Proceedings of the National Academy of Sciences*.

P ≠ NP

Fifty years ago, the most famous problem in theoretical computer science was formulated. Labeled "P ≠ NP," it asks if problems involving vast datasets exist for which an answer can be verified relatively quickly, but whose solution—even if worked out on the fastest available computers—would take an absurdly long time.

The P ≠ NP conjecture is still unproven, yet most computer scientists believe that many familiar problems—including, for instance, the traveling salesman problem—fall into this impossibly hard category. The challenge in the salesman example is to find the shortest route, in terms of distance or time, through N different cities. The task is easily

managed when N=4, because there are only six possible routes to consider. But for 30 cities, there are more than 10^{30} possible routes, and the numbers rise dramatically from there. The biggest difficulty comes in designing an [algorithm](#) that quickly solves the problem in all cases, for all integer values of N. Computer scientists are confident, based on algorithmic complexity theory, that no such algorithm exists, thus affirming that P ≠ NP.

There are many other examples of intractable problems like this. Suppose, for instance, you have a giant table of numbers with thousands of rows and thousands of columns. Can you find, among all possible combinations, the precise arrangement of 10 rows and 10 columns such that its 100 entries will have the highest sum attainable? "We call them optimization tasks," Gamarnik says, "because you're always trying to find the biggest or best—the biggest sum of numbers, the best route through cities, and so forth."

Computer scientists have long recognized that you can't create a fast algorithm that can, in all cases, efficiently solve problems like the saga of the traveling salesman. "Such a thing is likely impossible for reasons that are well-understood," Gamarnik notes. "But in real life, nature doesn't generate problems from an adversarial perspective. It's not trying to thwart you with the most challenging, hand-picked problem conceivable." In fact, people normally encounter problems under more random, less contrived circumstances, and those are the problems the OGP is intended to address.

Peaks and valleys

To understand what the OGP is all about, it might first be instructive to see how the idea arose. Since the 1970s, physicists have been studying spin glasses—materials with properties of both liquids and solids that have unusual magnetic behaviors. Research into spin glasses has given

rise to a general theory of complex systems that's relevant to problems in physics, math, computer science, materials science, and other fields. (This work earned Giorgio Parisi a 2021 Nobel Prize in Physics.)

One vexing issue physicists have wrestled with is trying to predict the energy states, and particularly the lowest energy configurations, of different spin glass structures. The situation is sometimes depicted by a "landscape" of countless mountain peaks separated by valleys, where the goal is to identify the highest peak. In this case, the highest peak actually represents the lowest energy state (though one could flip the picture around and instead look for the deepest hole). This turns out to be an optimization problem similar in form to the traveling salesman's dilemma, Gamarnik explains: "You've got this huge collection of mountains, and the only way to find the highest appears to be by climbing up each one"—a Sisyphean chore comparable to finding a needle in a haystack.

Physicists have shown that you can simplify this picture, and take a step toward a solution, by slicing the mountains at a certain, predetermined elevation and ignoring everything below that cutoff level. You'd then be left with a collection of peaks protruding above a uniform layer of clouds, with each point on those peaks representing a potential solution to the original problem.

In a 2014 paper, Gamarnik and his coauthors noticed something that had previously been overlooked. In some cases, they realized, the diameter of each peak will be much smaller than the distances between different peaks. Consequently, if one were to pick any two points on this sprawling landscape—any two possible "solutions"—they would either be very close (if they came from the same peak) or very far apart (if drawn from different peaks). In other words, there would be a telltale "gap" in these distances—either small or large, but nothing in between. A system in this state, Gamarnik and colleagues proposed, is

characterized by the OGP.

"We discovered that all known problems of a random nature that are algorithmically hard have a version of this property"—namely, that the mountain diameter in the schematic model is much smaller than the space between mountains, Gamarnik asserts. "This provides a more precise measure of algorithmic hardness."

Unlocking the secrets of algorithmic complexity

The emergence of the OGP can help researchers assess the difficulty of creating fast algorithms to tackle particular problems. And it has already enabled them "to mathematically [and] rigorously rule out a large class of algorithms as potential contenders," Gamarnik says. "We've learned, specifically, that stable algorithms—those whose output won't change much if the input only changes a little—will fail at solving this type of optimization problem." This negative result applies not only to conventional computers but also to quantum computers, and specifically, to so-called "quantum approximation optimization algorithms" (QAOAs), which some investigators had hoped could solve these same optimization problems. Now, owing to Gamarnik and his co-authors' findings, those hopes have been moderated by the recognition that many layers of operations would be required for QAOA-type algorithms to succeed, which could be technically challenging.

"Whether that's good news or bad news depends on your perspective," he says. "I think it's good news in the sense that it helps us unlock the secrets of algorithmic complexity and enhances our knowledge as to what is in the realm of possibility and what is not. It's bad news in the sense that it tells us that these problems are hard, even if nature produces them, and even if they're generated in a random way." The news is not really surprising, he adds. "Many of us expected it all along, but we now we have a more solid basis upon which to make this claim."

That still leaves researchers light-years away from being able to prove the nonexistence of fast algorithms that could solve these optimization problems in random settings. Having such a proof would provide a definitive answer to the P ≠ NP problem. "If we could show that we can't have an algorithm that works most of the time," he says, "that would tell us we certainly can't have an algorithm that works all the time."

Predicting how long it will take before the P ≠ NP problem is resolved appears to be an intractable problem in itself. It's likely there will be many more peaks to climb, and valleys to traverse, before researchers gain a clearer perspective on the situation.

More information: David Gamarnik, The overlap gap property: A topological barrier to optimizing over random structures, *Proceedings of the National Academy of Sciences* (2021). [DOI: 10.1073/pnas.2108492118](https://doi.org/10.1073/pnas.2108492118)

Provided by Massachusetts Institute of Technology