2026 Abacus Medal: Shayan Oveis Gharan

by Allyn Jackson

Shayan Oveis Gharan's problem-solving prowess and technical brilliance alone would suffice to mark him as a significant new talent in theoretical computer science. What has secured his position as a clear leader in the field are his ambition in taking on difficult open problems that have dramatic consequences, his ability to formulate a vision for solving those problems, and his tenacity in pursuing solutions, even if they take years to reach. He has an instinct for isolating the nub of a problem and a knack for finding, or creating, the needed tools. His work has spawned new ideas and techniques that others have gone on to develop and utilize to great effect.

Early in his career, Oveis Gharan became widely known for his work on the Traveling Salesperson Problem (TSP). This problem has stood at the center of theoretical computer science since the founding of the field in the 1970s. It's simple to state: Given an origin city and a collection of other cities connected to each other and to the origin by roads, what is the shortest route a salesperson can take that starts and ends at the origin city and that visits every other city at least once? The difficulty of the TSP lies in the way the number of possible routes explodes as the number of cities and roads increases.

For obvious reasons, the TSP arises in logistics and delivery planning, but it is also seen in many other contexts, such as designing efficient movement of robotic components for manufacturing. In some sense, the TSP is being solved all the time. But in another sense, it has *never* been solved: No one has ever found an algorithm that can efficiently find an optimal solution for *every* instance of the TSP. In fact, computer scientists believe that, as the notorious "P $\neq$ NP" conjecture suggests, no such algorithm exists. The algorithms used in logistics and manufacturing efficiently produce solutions that are good enough in those contexts but might not be optimal. At the same time, it's not difficult to concoct "worst case" instances on which those algorithms choke and become wildly inefficient.

Worst-case analysis of algorithms has been a major driver of developments in theoretical computer science. Early on researchers noticed that for some problems, including the TSP, efficient algorithms seemed out of reach. So they sought instead algorithms that would produce approximate solutions. One of the earliest approximation algorithms for the TSP was the Christofides algorithm, found in the mid-1970s (the same algorithm was found by independently by Serdyukov around the same time). The TSP can be represented abstractly as a graph—that is, a network of vertices connected by edges. Rather than finding the shortest route through a graph, the Christofides algorithm finds a route that, in the worst case, is 50 percent longer than the shortest route.

Researchers expected the Christofides algorithm would soon be surpassed. That didn't happen. Instead, its 50-percent barrier stood for four decades. By the time Oveis Gharan finished his PhD in 2013, that barrier was a long-standing conundrum that many people had struggled with and that seemed to require new ideas to overcome. It was just the kind of question that grips him.

In a series of papers, Oveis Gharan, together with Anna Karlin and Nathan Klein, pinpointed a fundamental inefficiency in the Christofides algorithm and showed how randomness leads to an improvement. The issue centers on finding a spanning tree, which is a minimal subset of the edges of a graph that connects all vertices. The Christofides algorithm generates a spanning tree of smallest cost and then adds a minimum-cost set of edges to create a good route. For some trees generated this way—for example, when the spanning tree has a palm-tree shape in which a single vertex is connected to every other vertex—then the algorithm must add many new edges. This is where Oveis Gharan saw an opening for greater efficiency.

He and his collaborators showed that fewer added edges are needed if the initial spanning tree is chosen at random; this is in part because a random tree is so unlikely to have a palm tree shape. Through a detailed analysis capitalizing on a new area called geometry of polynomials, they came up with a new algorithm that put a crack in the 50-percent barrier of the Christofides algorithm. The crack looks comically small: The new algorithm shaves only 10^{-36} off the 50-percent barrier for the worst cases. But any crack in the barrier is significant, and researchers have already widened the crack. In addition, in many practical applications the new algorithm actually performs far better than the Christofides algorithm. This demonstrates the power of the approach of Oveis Gharan and his co-authors and promises that yet more effective algorithms are on the horizon.

As spectacular as the TSP results were, they arguably have been outshined by Oveis Gharan's more-recent work on what are known as sampling problems. The techniques and analytical tools he created have revolutionized how researchers approach these problems and have set off a tidal wave of new developments.

Sampling problems arise across all of science, engineering, and the social sciences in situations in which random samples drawn from a population are used to estimate properties or behavior of the population as a whole. A canonical example comes from statistical physics. To study the behavior of a gas in a chamber, one can model the individual molecules' random motions using a probabilistic tool called a Markov chain. In any realistic scenario, the number of molecules is so large that it is not computationally feasible to enumerate all possible configurations of all molecules. Instead, one begins with an arbitrary sample of modeled molecules and repeatedly uses a Markov chain to make the sample more and more random. The motions of that randomized sample are then aggregated to obtain a picture of the overall behavior of the gas. This method is called a Monte Carlo Markov chain, or MCMC for short.

An obvious question is: How long do you need to let MCMC run to ensure it produces a uniform random sample? For decades, this question could be answered only for the simplest systems. When it came to many natural problems, like sampling gas molecules, the classical techniques for analyzing the relevant Markov chains were simply insufficient. What Oveis Gharan and his collaborators did was to develop new, more powerful techniques that give rigorous bounds on how long MCMC must run to ensure that it produces a random sample from the target distribution. They also provided a way to make intelligent choices about which variant of MCMC to use.

The starting point for these developments was Oveis Gharan's work on mathematical objects called matroids, which abstract the notion of linear independence of a set of vectors in a vector space. An example is the so-called graphic matroid, which is the collection of subsets of edges of a graph that do not include any cycles. A base for a matroid is analogous to the basis of a vector field. Bases of the graphic matroid, for example, are precisely the set of spanning trees. A matroid with thousands of elements can have an abundance of bases, more than the number of atoms in the universe. A longstanding open problem in the theory of computing asked: How long do you need to run MCMC to ensure it produces a uniform random base from a matroid?

Oveis Gharan, together with Nima Anari, Kuikui Liu, and Cynthia Vinzant, answered this question, thereby resolving the Mihail-Vazirani conjecture, which had been open for three decades. That conjecture centered on a natural algorithm used for sampling from a matroid, and the results of Oveis Gharan et al established a bound on how long the algorithm needs to run to sample sufficiently. That they found exactly the right bound is extraordinary; in new research, it is more common to find an approximate bound that is then improved over time.

Along the way they developed a new method called spectral independence, which can be used to bound the time it takes for a Markov chain to generate a uniformly random sample. This method has proved to be very fruitful and led to improved analysis of MCMC for many classic problems. By developing techniques in the idealized, theoretical setting of matroids, Oveis Gharan and his collaborators crystallized the essential difficulty in sampling problems and thereby pointed the way to adaptations to real-world problems like the gas molecule example.

The motion of a gas molecule is more strongly affected by nearby molecules than by molecules way off in another part of the chamber. This drop-off effect is called correlation decay. It turns out that spectral independence provides just the right way of accounting for correlation decay. Indeed, spectral independence is a very general notion that is applicable to any sampling problem that uses Markov chains. This work has set off an intellectual revolution that is having a transformative effect across many areas of research.

The powerful new approach Oveis Gharan has developed to address sampling problems was not found earlier, despite enormous efforts by researchers, partly because it built on the recently developed theory of high-dimensional expanders. In addition, it was not obvious that high-dimensional expanders would have a connection to sampling problems. Such connections are difficult to uncover. Oveis Gharan has an instinct for finding them.

That instinct, together with his bold and innovative vision, make Shayan Oveis Gharan a leading light in theoretical computer science.

Website of Shayan Oveis Gharan:
https://homes.cs.washington.edu/~shayan/