


From protein design to self-driving cars: UW team wins AI prize for optimization approach

August 13 2015



The UW algorithm reconstructed a three-dimensional scene from two-dimensional camera images, on average, between 100,000 and 10 billion times more accurately than current optimization techniques. Credit: University of Washington

The key to solving many of the most important problems in business, science and technology lies in optimization—finding the values for

variables that give you the highest benefit.

Whether those are which stocks to buy, which search results to return, what best predicts the outcome of the next presidential election or which amino acids to string together in a new drug to fight malaria or cancer, [optimization](#) is crucial to getting what we want. When a problem is simple, we can program a computer to solve it. When it's too complex for that, optimization is how the computer finds the solution by itself.

University of Washington machine learning researchers have developed a radically new approach to optimization, in part by borrowing a classic technique from artificial intelligence and computer science. The paper outlining their approach won the top prize in July at the 24th International Joint Conference on Artificial Intelligence, the world's largest AI conference.

In two applications that were tested experimentally, the new UW approach outperformed standard optimization techniques, in some instances by many orders of magnitude.

"In some ways optimization is the most important problem you've never heard of because it turns up in all areas of science, engineering and business. But a lot of [optimization problems](#) are extremely difficult to solve because they have a huge number of variables that interact in intricate ways," said senior author Pedro Domingos, UW professor of computer science and engineering.

For example, let's say you want to teach a computer to recognize the image of a cat by examining individual pixels. Pixels that denote orange, fur or whiskers increase the chances of yes. Pointy ears or claws help confirm. Blue feathers strongly suggest no. Optimization is the art and science of weighting those variables so the machine makes the correct choice as often as possible.

The UW [optimization algorithm](#), known by its acronym RDIS, progressively breaks an enormously complicated problem down into smaller, more manageable chunks—a simple idea commonly used when a problem consists of yes-or-no choices, but which had not previously been applied to numeric variables. RDIS can identify variables that, once set to specific values, break a larger problem into independent subproblems. Often, the problems are only nearly independent, but RDIS limits the error caused by treating them as fully independent.

"This approach is something that is very different than what people were doing before and it also does something magical, which is solve some problems exponentially faster. And anytime you can do that, that's when you get a big win," said Domingos.

The UW team tested RDIS against leading optimization methods in two real-world applications: determining the shape of folded proteins and accurately constructing three-dimensional objects and scenes from two-dimensional images.

For robotic arms to perform surgeries or to prevent self-driving cars from crashing, computers must accurately map the two-dimensional camera images that serve as the robot's "eyes" into realistic three-dimensional spaces. The UW team's optimization approach, on average, performed that task between 100,000 and 10 billion times more accurately than previous methods.

"You need to minimize the reconstruction error—the difference between what the algorithm predicts and what the image actually shows. Our algorithm did this very well," said lead author Abram Friesen, a UW computer science and engineering doctoral student. "This is crucial for a self-driving car—in order to safely navigate its environment, it first needs to be able to determine what objects are close to it, where the road is, and where the other cars are."

The researchers also tested RDIS against other optimization techniques for protein folding. A string of amino acids—protein building blocks that can occur in millions of different patterns—will generally fold and twist into a shape with the lowest energy. That configuration is important because it dictates how the protein interacts with viruses or invading cells or other parts of the immune system.

Enabling computers to accurately predict how proteins will fold can greatly speed the process of designing effective drugs. The UW team found its new optimization technique resulted in much lower-energy protein shapes than alternative methods, particularly for larger proteins.

That's because current optimization approaches fail as problems get more complex, Domingos said. Solving optimization problems is like being left blindfolded at the top of a hill and asked to walk to the ocean. One way to do that is to judge where to go by feeling around with your foot and taking one step a time in the steepest downward direction. That works if there's only one hill. But if you're at the top of the Himalayas, you'll quickly get stuck because there are thousands of peaks and foothills and flat parts. That's essentially what happens with current optimization algorithms.

"If you're lucky, maybe you'll wind up in the sea but more likely you'll wind up in a valley or a lake," said Domingos. "If you could see the whole landscape, you'd say 'oh, here's where I have to go,' but the problem is you can't see everywhere and neither can today's algorithms."

To address those deficiencies, the UW duo figured out how to apply decomposition techniques, commonly used by [artificial intelligence](#) researchers and puzzle solvers, to continuous optimization problems, which have historically been the bailiwick of engineers, applied mathematicians and physicists.

The next steps, researchers say, are to test the algorithm's performance on new and different applications. "This can be applied to pretty much any machine learning problem, but that's not to say it's going to be good for every machine learning problem," said Domingos. "That's what we have to work on and find out."

More information: Study paper:
homes.cs.washington.edu/~pedrod/papers/ijcai15.pdf

Provided by University of Washington

Citation: From protein design to self-driving cars: UW team wins AI prize for optimization approach (2015, August 13) retrieved 9 April 2024 from <https://phys.org/news/2015-08-protein-self-driving-cars-uw-team.html>

<p>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.</p>
--