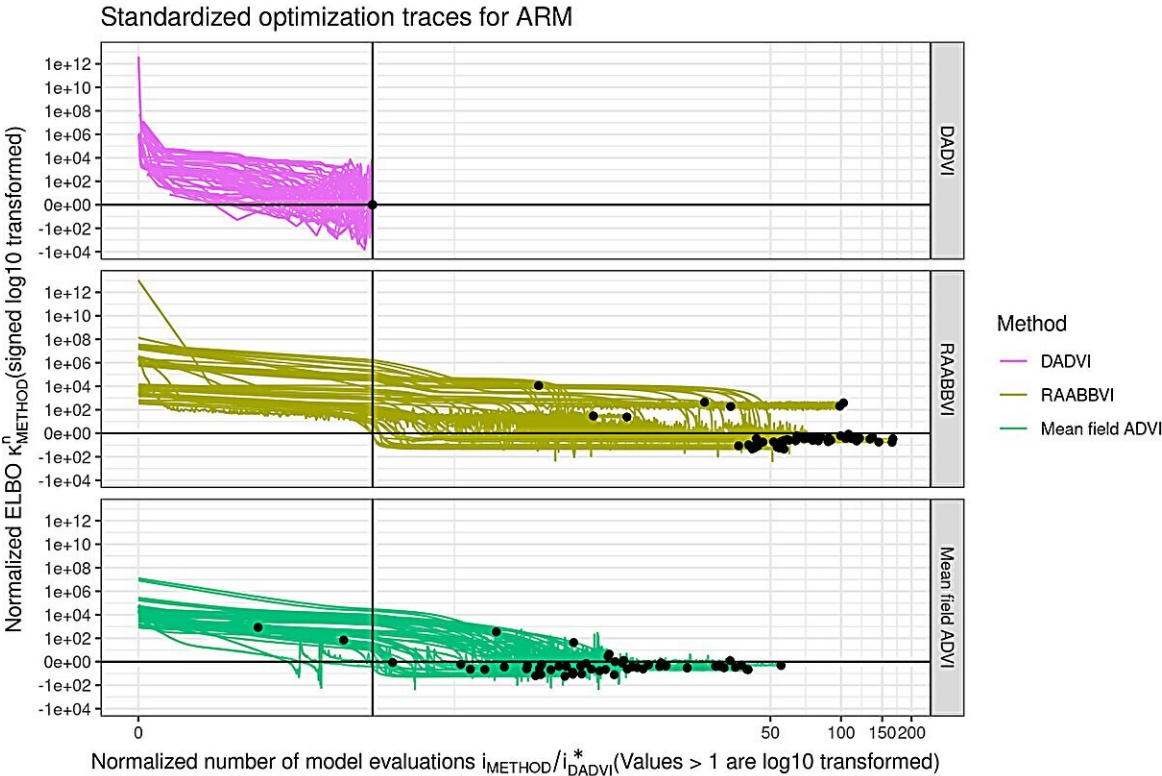


Automated method helps researchers quantify uncertainty in their predictions

February 21 2024, by Adam Zewe



Optimization traces for the ARM models. Black dots show the termination point of each method. Dots above the horizontal black line mean that DADVI found a better ELBO. Dots to the right of the vertical black line mean that DADVI terminated sooner in terms of model evaluations. Credit: <https://jmlr.org/papers/volume25/23-1015/23-1015.pdf>

Pollsters trying to predict presidential election results and physicists searching for distant exoplanets have at least one thing in common: They often use a tried-and-true scientific technique called Bayesian inference.

Bayesian inference allows these scientists to effectively estimate some unknown parameter—like the winner of an election—from data such as [poll results](#). But Bayesian inference can be slow, sometimes consuming weeks or even months of computation time or requiring a researcher to spend hours deriving tedious equations by hand.

Researchers from MIT and elsewhere have introduced an optimization technique that speeds things up without requiring a scientist to do a lot of additional work. Their method can achieve more accurate results faster than another popular approach for accelerating Bayesian inference.

Using this new automated technique, a scientist could simply input their model and then the optimization method does all the calculations under the hood to provide an approximation of some unknown parameter. The method also offers reliable uncertainty estimates that can help a researcher understand when to trust its predictions.

This versatile technique could be applied to a wide array of scientific quandaries that incorporate Bayesian inference. For instance, it could be used by economists studying the impact of microcredit loans in developing nations or sports analysts using a model to rank top tennis players.

"When you actually dig into what people are doing in the social sciences, physics, chemistry, or biology, they are often using a lot of the same tools under the hood. There are so many Bayesian analyses out there.

"If we can build a really great tool that makes these researchers lives easier, then we can really make a difference to a lot of people in many

different research areas," says senior author Tamara Broderick, an associate professor in MIT's Department of Electrical Engineering and Computer Science (EECS) and a member of the Laboratory for Information and Decision Systems and the Institute for Data, Systems, and Society.

Broderick is joined on the paper by co-lead authors Ryan Giordano, an assistant professor of statistics at the University of California at Berkeley; and Martin Ingram, a data scientist at the AI company KONUX. The paper was recently [published](#) in the *Journal of Machine Learning Research*.

Faster results

When researchers seek a faster form of Bayesian inference, they often turn to a technique called automatic differentiation variational inference (ADVI), which is often both fast to run and easy to use.

But Broderick and her collaborators have found a number of practical issues with ADVI. It has to solve an optimization problem and can do so only approximately. So, ADVI can still require a lot of computation time and user effort to determine whether the approximate solution is good enough. And once it arrives at a solution, it tends to provide poor uncertainty estimates.

Rather than reinventing the wheel, the team took many ideas from ADVI but turned them around to create a technique called deterministic ADVI (DADVI) that doesn't have these downsides.

With DADVI, it is very clear when the optimization is finished, so a user won't need to spend extra computation time to ensure that the best solution has been found. DADVI also permits the incorporation of more powerful optimization methods that give it an additional speed and

performance boost.

Once it reaches a result, DADVI is set up to allow the use of uncertainty corrections. These corrections make its uncertainty estimates much more accurate than those of ADVI.

DADVI also enables the user to clearly see how much error they have incurred in the approximation to the [optimization problem](#). This prevents a user from needlessly running the optimization again and again with more and more resources to try and reduce the error.

"We wanted to see if we could live up to the promise of black-box inference in the sense of, once the user makes their model, they can just run Bayesian inference and don't have to derive everything by hand, they don't need to figure out when to stop their algorithm, and they have a sense of how accurate their approximate solution is," Broderick says.

Defying conventional wisdom

DADVI can be more effective than ADVI because it uses an efficient approximation method, called sample average approximation, which estimates an unknown quantity by taking a series of exact steps.

Because the steps along the way are exact, it is clear when the objective has been reached. Plus, getting to that objective typically requires fewer steps.

Often, researchers expect sample average approximation to be more computationally intensive than a more popular method, known as stochastic gradient, which is used by ADVI. But Broderick and her collaborators showed that, in many applications, this is not the case.

"A lot of problems really do have special structure, and you can be so

much more efficient and get better performance by taking advantage of that special structure. That is something we have really seen in this paper," she adds.

They tested DADVI on a number of real-world models and datasets, including a model used by economists to evaluate the effectiveness of microcredit loans and one used in ecology to determine whether a species is present at a particular site.

Across the board, they found that DADVI can estimate unknown parameters faster and more reliably than other methods, and achieves as good or better accuracy than ADVI. Because it is easier to use than other techniques, DADVI could offer a boost to scientists in a wide variety of fields.

In the future, the researchers want to dig deeper into correction methods for uncertainty estimates so they can better understand why these corrections can produce such accurate uncertainties, and when they could fall short.

"In applied statistics, we often have to use approximate algorithms for problems that are too complex or high-dimensional to allow exact solutions to be computed in reasonable time. This new paper offers an interesting set of theory and empirical results that point to an improvement in a popular existing approximate algorithm for Bayesian inference," says Andrew Gelman, a professor of statistics and political science at Columbia University, who was not involved with the study. "As one of the team involved in the creation of that earlier work, I'm happy to see our algorithm superseded by something more stable."

More information: Black Box Variational Inference with a Deterministic Objective: Faster, More Accurate, and Even More Black Box, *Journal of Machine Learning Research*(2024).

jmlr.org/papers/volume25/23-1015/23-1015.pdf

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

Citation: Automated method helps researchers quantify uncertainty in their predictions (2024, February 21) retrieved 28 April 2024 from <https://techxplore.com/news/2024-02-automated-method-quantify-uncertainty.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.