Proceedings of the 2013 Winter Simulation Conference
R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, eds.

A TRUST REGION-BASED ALGORITHM FOR CONTINUOUS OPTIMIZATION VIA SIMULATION

Satyajith Amaran
Nikolaos V. Sahinidis

Department of Chemical Engineering
Carnegie Mellon University
Pittsburgh, PA 15213, USA

Bikram Sharda
Scott J. Bury

Engineering and Process Sciences, Core R&D
2301 N. Brazosport Blvd.
The Dow Chemical Company
Freeport, TX 77541, USA

ABSTRACT

Continuous Optimization via Simulation (COvS) involves the search for specific continuous input parameters to a stochastic simulation that yield optimal performance measures. Typically, these performance measures can only be evaluated through simulation. We introduce a new algorithm for solving COvS problems. The main idea is to use a nonparametric regression model that uses few samples, and embed it in an iterative trust-region framework. We name the proposed algorithm Simulation Optimization–Learning Via Trust Regions (SO-LViT). We discuss the algorithmic elements of this implementation, and hypothesize that this approach is especially suitable for situations where samples are expensive to obtain and the dimensionality of the problem is fairly large. We demonstrate promising results through computational experience, wherein we compare SO-LViT against several other approaches over a large test set under Gaussian noise conditions.

1 INTRODUCTION

There exist a large number of algorithms for continuous simulation optimization in the literature. Some of these algorithms have a focus that is more global in scope and try to explore many parts of the search space to look for solutions; other algorithms trade off between exploring the space while exploiting good solutions; yet another class of methods focuses on a more local scope, where the primary task is to start at a point and find search directions that are likely to result in descent.

Our current implementation of SO-LViT belongs to the third class of algorithms, where we restrict our attention to a local scope and emphasize on using few function evaluations to glean good search directions. Specifically, the class of algorithms it falls in attempts to embed some sort of metamodel or response surface within a trust-region framework (Conn, Gould, and Toint 2000). All of these algorithms differ in (1) the number of function evaluations used to build a metamodel; (2) the type of metamodel used; (3) the method of optimization of the metamodel; (4) the way subsequent sample points are chosen; and (5) the trust region update scheme.

2 APPROACH

The unique feature of the proposed algorithm is that it effectively combines nonparametric regression techniques (Rasmussen and Williams 2006) with trust regions and global optimization methods (Tawarmalani and Sahinidis 2005), while looking for search directions that ensure descent with high probability. Specifically, the algorithm uses samples at relatively few points (a number linear in the dimension of the problem) to construct a Gaussian process regression model. The reason for this choice is that GP regression is effective in modeling multi-modality and is capable of incorporating previously evaluated points in the model-building process.

Following this, global optimization routines are called to solve this nonconvex optimization problem to the global minimum, within a certain trust-region radius. This optimization is done using BARON (Branch-and-Reduce Optimization Navigator) (Tawarmalani and Sahinidis 2005), which is capable of performing deterministic global optimization efficiently. The simulation model is then sampled at this optimal point with replication to get a confidence interval for the expected output. Trust region tests are then used to determine whether (1) sufficient decrease has been achieved; and that (2) the decrease is ensured with high probability. Standard trust region criteria are then used to update the center and radius of the region. This process is repeated in the new trust region, and this procedure is continued until a computational budget is reached.

3 RESULTS

To benchmark the proposed algorithm, the test set that was used involved 502 deterministic black-box real-valued functions, whose output value was perturbed by a predetermined level of random Gaussian noise to represent stochasticity in the output. The test set involved diverse problems, with sizes ranging from 1–300 variables, and with functions belonging to the convex and nonconvex classes, as well as the smooth and nonsmooth classes. We benchmarked the proposed algorithm against six other algorithms—the algorithms we chose to compare include those that use response surfaces (as we relied on these as well), local search methods such as Nelder-Mead simplex procedures (as the proposed method has a local scope as well), and those that use trust regions.

The six algorithms and the proposed algorithm, SO-LViT, were run on the 502-problem test set. As we are interested in solving problems where simulation replications are expensive to obtain, a limit of 300 total function evaluations (simulations) was fixed for this study. The purpose is to assess the performance of the algorithms under this strict budget of simulations.

On this large test set, we ranked the six algorithms based on the fraction of problems for which they found the best solution, within a certain tolerance. We observe that SO-LViT finds the best solution on a larger fraction of problems ($> 50\%$) than all of the other algorithms (the next best solver found a better solution on about 40% of the problems, and the rest were better on less than 10% of the problems). We believe that the reason that SO-LViT performs better is that it is able to build surrogate regression models with fewer function evaluations, is able to optimize these models to global optimality, ensures descent with high probability for each subsequent trust region, and handles noise effectively through hypothesis tests.

4 DISCUSSION AND FUTURE WORK

The favorable results for the SO-LViT algorithm are encouraging as they show that it is still possible to push the envelope in the context of COvS algorithms. We conclude that this particular algorithmic framework performs very well and this necessitates further investigation. Towards this end, a more careful analysis of how close the algorithms were able to get to the true optimum given the simulation budget, comparisons with larger simulation budgets, comparisons with implementations of other families of solution approaches, and performance under heterogeneous noise of different levels would be revealing. We also intend to provide a proof of convergence for the algorithm, enhance some of the algorithmic elements to make it more generally applicable, and apply it to a real-world expensive simulation and analyze the results.

REFERENCES

Conn, A. R., N. I. M. Gould, and P. L. Toint. 2000. *Trust-region methods*. MOS-SIAM Series on Optimization.

Rasmussen, C. E., and C. K. I. Williams. 2006. *Gaussian Processes for Machine Learning*. The MIT Press.

Tawarmalani, M., and N. V. Sahinidis. 2005. “A polyhedral branch-and-cut approach to global optimization”. *Mathematical Programming* 103:225–249.