*Proceedings of the 2025 Winter Simulation Conference*
*E. Azar, A. Djanatliev, A. Harper, C. Kogler, V. Ramamohan, A. Anagnostou, and S. J. E. Taylor, eds.*

# ASTROMORF: AN ADAPTIVE SAMPLING TRUST-REGION ALGORITHM WITH DIMENSIONALITY REDUCTION FOR LARGE-SCALE SIMULATION OPTIMIZATION

Benjamin Rees[1]

[1]School of Mathematical Sciences, University of Southampton, Southampton, UK

## ABSTRACT

High-dimensional simulation optimization problems have become prevalent in recent years. In practice, the objective function is typically influenced by a lower-dimensional combination of the original decision variables, and implementing dimensionality reduction can improve the efficiency of the optimization algorithm. In this extended abstract, we introduce a novel algorithm ASTROMoRF that combines adaptive sampling with dimensionality reduction, using an iterative trust-region approach. Within a trust-region algorithm, a series of surrogates are built to estimate the objective function. Using a lower-dimensional subspace reduces the number of design points needed for building a surrogate within each trust-region and consequently reduces the number of simulation replications. We introduce this algorithm and comment on its finite-time performance against other state-of-the-art solvers.

## 1 INTRODUCTION

Our contribution is the proposal of a novel trust-region optimization algorithm, called adaptive sampling trust-region optimization with moving ridge functions (ASTROMoRF), which combines the adaptive sampling and design point selection from ASTRO-DF (Shashaani et al. 2016) with techniques from OMoRF (Gross and Parks 2022) to construct local active subspace matrices to project the surrogate model into a lower-dimensional subspace $\mathbb{R}^d$.

The subspace dimension $d$ is selected a priori by the experimenter. In practice, the optimal choice for $d$ is unknown but knowledge of the problem may suggest a good choice. A smaller value for $d$ reduces the simulation budget expended per iteration but can lead to slower convergence if not appropriately selected. The application of a local active subspace allows the selection of coordinate directions that exhibit the most variability on the response surface, enabling more substantial steps in each iteration.

## 2 METHODOLOGY

ASTROMoRF improves upon ASTRO-DF in two significant ways: (i) it offers a novel sampling scheme that samples $\mathscr{O}(d)$ design points within a projected space to construct a stochastic polynomial surrogate model via interpolation; (ii) it introduces a modification of the ridge function recovery problem presented in Hokanson and Constantine (2018) to recover the surrogate model and active subspace at each iteration.

The $2d+1$ design points selected for interpolation are chosen using an underdetermined rotated coordinate stencil as outlined in Ha et al. (2024), which provides more precise estimates of the gradient at $\mathbf{x}_k$. Design points can be reused in this scheme to further conserve computational resources.

When constructing the surrogate model, we solve a ridge function recovery as presented in Hokanson and Constantine (2018). Our method for solving the ridge function recovery problem differs from Hokanson's approach in that we ensure the surrogate model constructed is an interpolation model. Finally, our method provides certifications on the constructed surrogate model through a criticality bound, which ensures that the model gradient is small enough relative to the trust-region radius and that the model satisfies first-order criticality conditions. If this bound is not satisfied, we apply a pivoting geometry-improvement algorithm from Conn et al. (2009) to produce models that are better suited to solving the trust-region subproblem.

Finally, we apply an adaptive sampling method presented in Ha et al. (2024) to estimate objective function values through sample average approximation (SAA) on a collection of responses of the simulation model at a particular design point. Applying this adaptive sampling rule allows for more accurate models to be constructed, improving the probability of candidate solutions being accepted.

## 3 NUMERICAL RESULTS

We compare ASTROMoRF against ASTRO-DF and OMoRF, on three high-dimensional test problems: (i) a 10-dimensional inventory problem with dynamic consumer substitution; (ii) a 15-dimensional nonlinear test function with additive noise; and (iii) a 15-dimensional noisy quadratic test function.

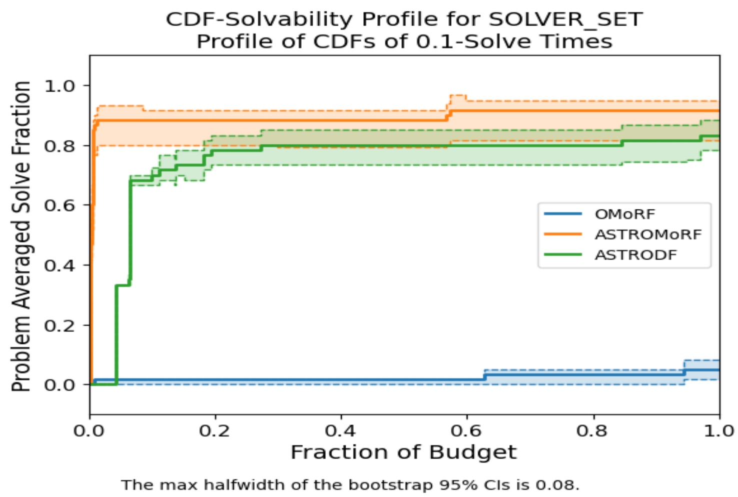


Figure 1: Fraction of problems within the SimOpt library solved to a 0.1-optimality.

## 4 CONCLUSION

ASTROMoRF offers a new method of efficiently computing good and stable solutions for high-dimensional simulation optimization problems. It also demonstrates the viability of active subspace dimensionality reduction in simulation optimization.

## ACKNOWLEDGMENTS

The author is grateful for the support of their supervisors Professor Christine S.M. Currie and Doctor Phan Tu Vuong at the University of Southampton School of Mathematical Sciences, and funding from the Engineering and Physical Sciences Research Council under grant EP/W524621/1 that provided support for this work.

## REFERENCES

Conn, A. R., K. Scheinberg, and L. N. Vicente. 2009. *Introduction to Derivative-Free Optimization*. SIAM.

Gross, J. C., and G. T. Parks. 2022. "Optimization by Moving Ridge Functions: Derivative-Free optimization for Computationally Intensive Functions". *Engineering Optimization* 54(4):553–575.

Ha, Y., S. Shashaani, and M. Menickelly. 2024. "Two-Stage Estimation and Variance Modeling for Latency-Constrained Variational Quantum Algorithms". *INFORMS Journal on Computing*.

Hokanson, J. M., and P. G. Constantine. 2018. "Data-Driven Polynomial Ridge Approximation using Variable Projection". *SIAM Journal on Scientific Computing* 40(3):A1566–A1589.

Shashaani, S., S. R. Hunter, and R. Pasupathy. 2016. "ASTRO-DF: Adaptive Sampling Trust-Region Optimization Algorithms, Heuristics, and Numerical Experience". In *2016 Winter Simulation Conference (WSC)*, 554–565.