A Trust Region Method for Bi-Objective Stochastic Optimization

Sujin Kim and Jong-hyun Ryu

Department of Industrial and Systems Engineering, National University of Singapore

Abstract

We consider black-box problems where the analytic forms of the objective functions are not available, and the values can only be estimated by output responses from computationally expensive simulations. We apply the sample average approximation method to multi-objective stochastic optimization problems and prove the convergence properties of the method under a set of fairly general regularity conditions. We develop a new algorithm, based on the trust-region method, for approximating the Pareto front of a bi-objective stochastic optimization problem. At each iteration of the proposed algorithm, a trust region is identified and quadratic approximate functions for the expected objective functions are built using sample average values. To determine non-dominated solutions in the trust region, a single-objective optimization problem is constructed based on the approximate objective functions. After updating the set of non-dominated solutions, a new trust region around the most isolated point is determined to explore areas that have not been visited. The numerical results show that our proposed method is feasible, and the performance can be significantly improved with an appropriate sample size.

Introduction

Multi-objective optimization arises in a wide variety of applications, whenever it is necessary to make a tradeoff between different important, but conflicting goals. The usual concept of optimality from single-objective optimization is not directly applicable in these settings, because it is impossible to optimize multiple conflicting objectives at the same time. Rather, we seek a good tradeoff among the multiple objectives, which can be formalized using the notion of Pareto optimality or Pareto efficiency. The general form of multi-objective optimization problem is as follows:

- Decision vector: $x \in X \subset \mathbb{R}^p$
- Objective function vector: $F: \mathbb{R}^p \to \mathbb{R}^q$
- Optimization problem:

$$\min_{x \in X} F(x) = [f_1(x), f_2(x), \dots, f_q(x)]$$

Pareto optimality for the problem with the above form is defined by the following dominance relationship:

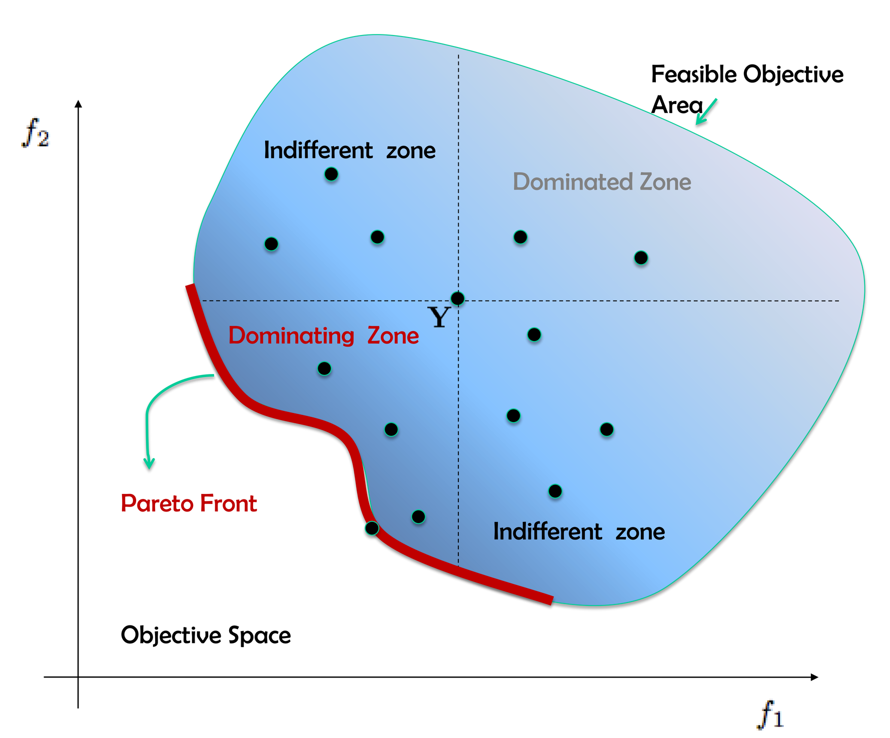

Multi-Objective Stochastic Optimization

We consider black-box problems where we do not have analytic forms for the objective functions. The objective values can only be estimated through expensive simulations. The output of these simulations is subject to stochastic variation. Let ξ be a random vector. We focus on multi-objective stochastic problems with the form:

$$\min_{x \in X} F(x) = E[f_1(x,\xi), f_2(x,\xi), \dots, f_q(x,\xi)]$$

We use the sample average approximation (SAA) approach [1,2] to approximate the above problem. Essentially, we replace the expected value function by a sample average, and use deterministic optimization to solve the ensuing problem. The SAA approximation to is given by

$$\min_{x \in X} \hat{F}_N(x) = \left[\frac{1}{N}\sum_{i=1}^{N} f_1(x,\xi_i), \frac{1}{N}\sum_{i=1}^{N} f_2(x,\xi_i) \dots, \frac{1}{N}\sum_{i=1}^{N} f_q(x,\xi_i)\right]$$

Pareto Approximation

We iteratively apply a trust-region method [3] to the SAA problem and find a set of non-dominated points. While we would like to identify solutions close to the Pareto front, we also want to generate well-spread solutions in order to approximate as much of the Pareto front as possible. To this end, we select the most isolated point among the points that have thus far been determined to be non-dominated. A trust region centered at the point is determined to maintain the uniformity of the optimal solution set by exploring non-visited areas.

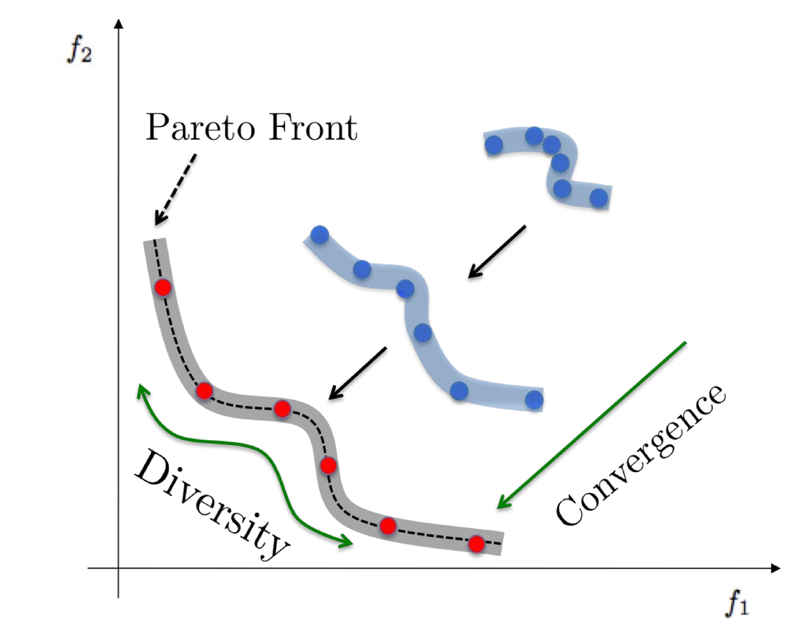

Algorithm for Bi-Objective Problem

Step 1 (Next iterate selection): Select the most isolated point $x_c^{(k)}$ from the non-dominated points at the end of iterate (k-1).

Step 2 (Regression model): Sample points in a trust region and construct a fully linear model for f_j, j=1,2,3 using sample average values, where

$$f_3(x) = -\prod_{j=1,2}\left\{\left(f_j(x_c^{(k)}) - f_j(x)\right)_+\right\}^2$$

Step 3 (Minimize approximate functions) : Find the minimizer for each single objective function.

Step 4 (Trust region radius update): Compute the reduction ratio at each minimizer with the associated approximate objective function. If the approximate function is poor, reduce the trust region radius.

Step 5 (Approximate Pareto front): Update the set of non-dominated points using all evaluated points.

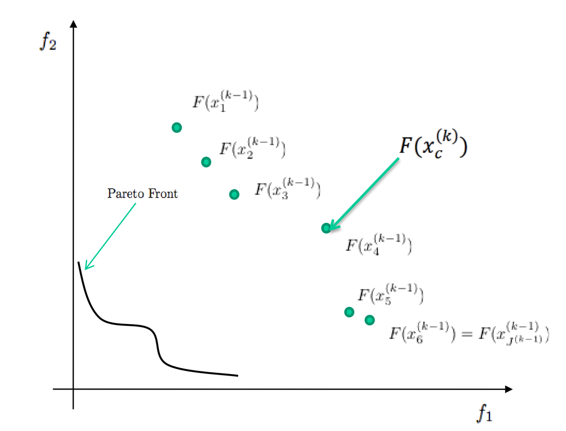

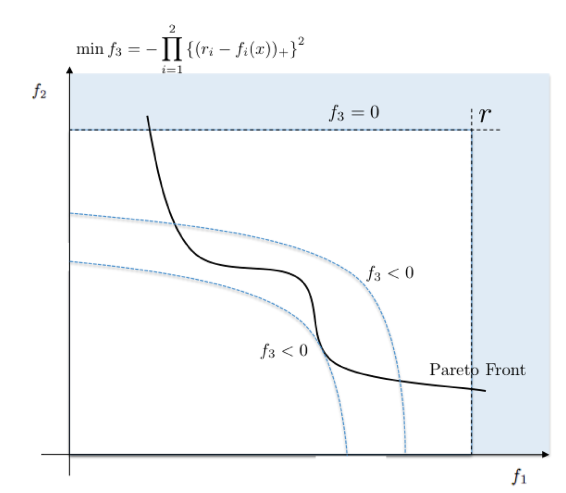

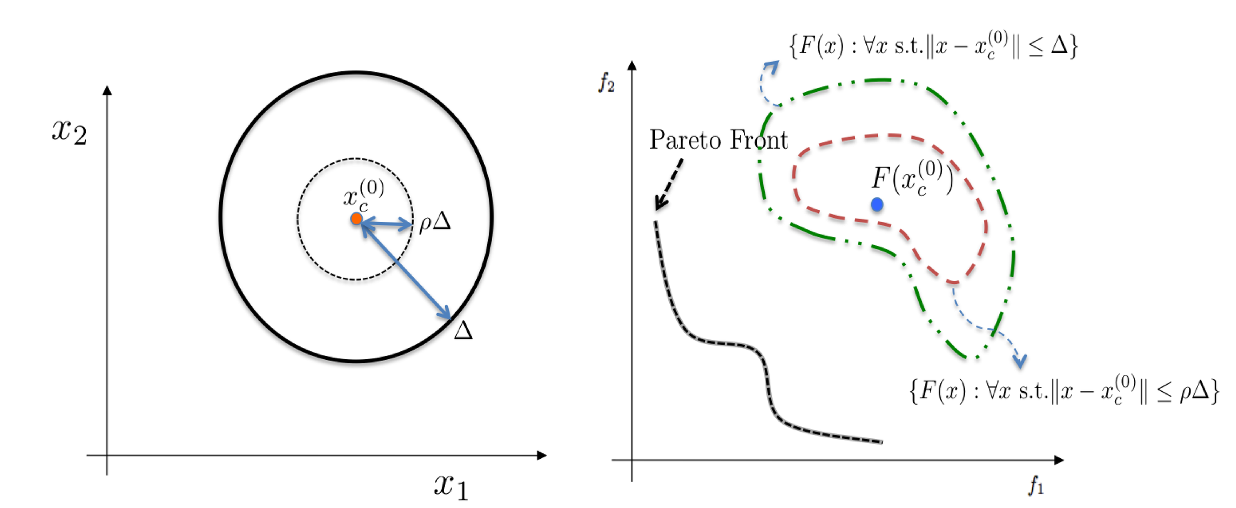

Numerical Results

We tested an unconstrained bi-objective problem with a convex Pareto front.

$$\min_{x=(x_{(1)},x_{(2)})\in\mathbb{R}^2} E[f_1(x,\xi), f_2(x,\xi)]$$

$$f_1(x,\xi) = \left(x_{(1)} - 2\xi_{(1)}\right)^2 + \left(x_{(2)} - \xi_{(2)}\right)^2$$

$$f_2(x,\xi) = \left(x_{(1)}\right)^2 + \left(x_{(2)} - 6\xi_{(3)}\right)^2$$

$$\xi = \left(\xi_{(1)}, \xi_{(2)}, \xi_{(3)}\right), \xi_{(i)} \overset{i.i.d.}{\sim} \chi^2$$

We find a set of solutions H of around 2,500 non-dominated solutions that are uniformly-spaced. To evaluate our method, we use the generational distance (GD) criterion. Let $H = \{x_1, \dots, x_e\}$ be the set of solutions. The GD is computed by

$$GD = \frac{\sqrt{\sum_{j=1}^{e}\left\{\min_{x_j^* \in H}\left\|F(x_j) - F(x_j^*)\right\|\right\}^2}}{e}$$

This measures the average distance between the objective value at the obtained solution and the true Pareto front. We can observe that with 5,000 function evaluations, N = 10 performs the best, and with 20,000 function valuations, N = 50 performs significantly better than others. This implies that the sample size should be carefully determined taking into account the computational budget.

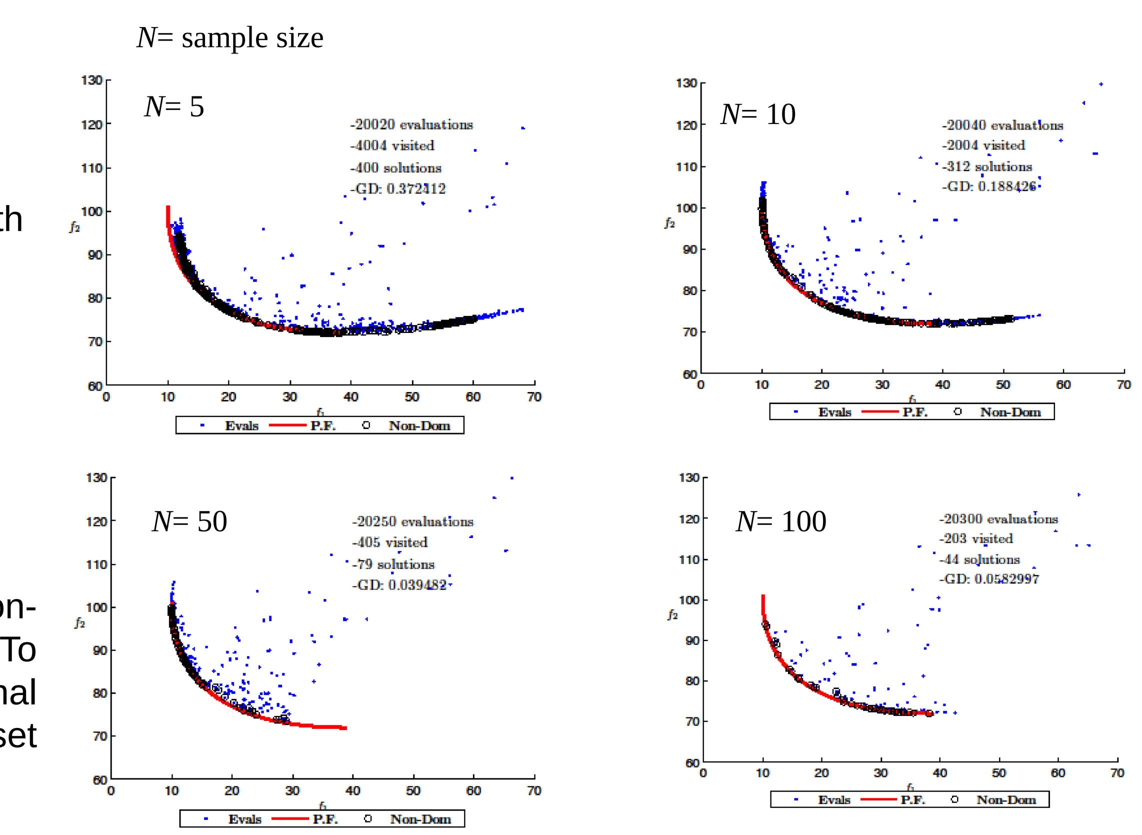

Performance comparisons with 20,000 function evaluations

N	5,000 evaluations			20,000 evaluations		
	GD < 0.1	GD < 0.5	GD < 1	GD < 0.1	GD < 0.5	GD < 1
5	23	81	90	47	83	90
10	47	86	91	50	86	97
50	6	67	88	63	94	98
100	0	0	0	56	92	98

The number of runs with GD less than 0.1, 0.5, and 1, with 100 independent runs

Conclusion

We developed the framework of the SAA method for MOP. The convergence of the SAA method can be obtained under a set of fairly general regularity conditions [4]. We applied an iterative algorithm for bi-objective stochastic optimization problems, based on the trust region method, to the SAA problems. The algorithm does not require any strong modeling assumptions, and has great potential to work well in various real-world settings. The numerical results show that the our proposed method is feasible, and can perform robustly with a large enough size N. To improve the finite time performance of the algorithm, the sample size should be carefully determined with consideration for the trade-off between sampling and optimization errors. The difference between the solutions obtained from SAA and the solutions to the true problems can be reduced by taking a larger sample size. On the other hand, as the number of iterations grows, the distance between solutions from each iteration and the Pareto front decreases. We consider an algorithm to solve a sequence of SAA problems with increasing sample size as a future work.

References

[1] Shapiro, A. 2003. "Monte Carlo Sampling Methods". In Stochastic Programming, Edited by A. Ruszczynski and A. Shapiro, Volume 10 of Handbooks in Operations Research and Management Science, 353 –425. Elsevier.
[2] Shapiro, A., D. Dentcheva, and A. Ruszczynski. 2009. Lectures on Stochastic Programming: Modeling and Theory. MPS-SIAM Series on Optimization.
[3] Ryu, J., and S. Kim. 2011. "A Trust Region Method for Multiobjective Optimization Problems". Working paper.
[4] Kim, S. & Ryu, J.-H. 2011. "A sample average approximation method for multi-objective stochastic optimization." Proceedings of the 2011 Winter Simulation Conference (eds: S. Jain, R.R. Creasey, J. Himmelspach, K.P. White, M. Fu). To appear.