Proceedings of the 2016 Winter Simulation Conference
T. M. K. Roeder, P. I. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. E. Chick, eds.

CHALLENGES IN APPLYING RANKING AND SELECTION AFTER SEARCH

David J. Eckman
Shane G. Henderson

School of Operations Research
and Information Engineering
Cornell University
206 Rhodes Hall
Ithaca, NY 14853, USA

ABSTRACT

It is tempting to reuse simulation replications taken during a simulation optimization search as input to a ranking-and-selection procedure, especially when generating replications is computationally expensive. Yet when a search identifies new systems based on the observed performance of explored systems, the resulting search replications are conditionally dependent given the sequence of returned systems. This dependence can mislead the selection decisions of ranking-and-selection procedures that reuse search data, in some cases to the point of violating guarantees on the probabilities of correct and good selection.

1 INTRODUCTION

Ranking-and-selection (R&S) procedures typically provide a statistical guarantee of selecting or maintaining the best of a finite set of systems whose identities are fixed a priori. For simulation optimization problems, R&S procedures have been applied on the (random) set of systems returned by a search heuristic (Boesel, Nelson, and Kim 2003). To make decisions more efficiently, these R&S procedures reuse rather than discard the replications taken during search.

In the setting of R&S applied after search, we consider the overall guarantee on the probability of correct selection (PCS):

$$\mathbb{P}(\text{Correct Selection after Search} \mid \mu(S) \in \mathrm{PZ}(\delta)) \geq 1-\alpha, \tag{1}$$

where $\mu(S)$ is the configuration of the set of systems S returned by the search and $\mathrm{PZ}(\delta)$ is the preference zone specified by the indifference-zone parameter δ. In words, Guarantee (1) states that when a given R&S procedure is run after a search and the set of systems returned by the search is in the preference zone, a correct selection is made with probability $\geq 1-\alpha$. Guarantee (1) is difficult to prove directly because the probability measure on the left-hand side is with respect to the outcomes of the search and the replications taken by the R&S procedure.

Alternatively, we consider a stronger PCS guarantee, conditioned on the set of systems returned by the search:

$$\mathbb{P}(\text{Correct Selection after Search} \mid S) \geq 1-\alpha \quad \text{for all } S \text{ s.t. } \mu(S) \in PZ(\delta), \tag{2}$$

Guarantee (2) is easier to prove directly because the probability measure is now only with respect to the replications taken by the R&S procedure. Similar guarantees for good selection—selecting or maintaining a system whose performance is within δ of the best system—can also be formulated.

We study whether Guarantees (1) and (2) hold for R&S procedures that reuse replications taken during search. For searches that identify new systems based on the observed performance of previously returned systems, the search replications are conditionally dependent on the sequence of returned systems. By means of an adversarial search heuristic, we show that this dependence can reduce PCS and for some R&S procedures invalidate Guarantees (1) and (2).

2 ADVERSARIAL SEARCH

We construct a search heuristic named Adversarial Search (AS) that weakens the selection decisions of R&S procedures by inducing an unfavorable dependence in the search replications. AS achieves this by returning a δ-better system when the best system thus far has the highest observed performance, and a δ-worse system otherwise. In this manner, AS returns sets of systems in which the best system often does not look best.

Although AS relies on complete knowledge of the performance of systems, we believe it can serve as a useful benchmark for testing the robustness of R&S guarantees. Furthermore, it is possible to construct particular instances of optimization problems and searches that behave exactly the same as AS, without relying on this oracle-like knowledge. Every set of systems returned by AS has a unique best system that is at least δ better than all of the others, i.e., $\mu(S) \in PZ(\delta)$ for all S. Thus in our experiments with AS, guarantees on the probabilities of correct and good selection are equivalent.

3 EXPERIMENTAL RESULTS

For our experiments, we fixed $1-\alpha = 0.95$, $\delta = 1$, an initial sample size of $n_0 = 10$, and a common variance of $\sigma^2 = 1$ and varied the number of systems returned by AS from 10 to 1000. We tested selection and subset-selection procedures applied after AS, reusing the search replications. For each setting, we estimated the overall PCS of Guarantee (1) by the empirical PCS based on 10 000 macroreplications.

We tested the selection procedures of Bechhofer (Bechhofer 1954) and Rinott (Rinott 1978), designed for common, known variances and uncommon, unknown variances, respectively. For the Bechhofer procedure, the empirical PCS dropped below the guaranteed level as the number of systems increased, falling to as low as 0.75. For the Rinott procedure, the empirical PCS decreased to 0.97 at around 100 systems before increasing again, suggesting that the conservativeness of Rinott offsets the unfavorable dependence induced by AS.

We also tested two subset-selection procedures, both variations of the Gupta procedure (Gupta 1965), designed to provide PCS guarantees in the preference zone. The first, Modified Gupta, assumes common, known variances, while the second, Screen-to-the-Best (Nelson, Swann, Goldsman, and Song 2001), handles unknown and uncommon variances. For Modified Gupta, empirical PCS dropped well below the guaranteed level as the number of systems increased, falling to as low as 0.5. The empirical PCS of Screen-to-the-Best decreased to 0.955 for around 150 systems before increasing again, suggesting that the PCS guarantee of Screen-to-the-Best might be violated for certain problems.

We expect that similar results would be observed for more general selection procedures reusing search data, such as those for multi-armed bandits in full exploration that provide probably approximately correct (PAC) guarantees.

REFERENCES

Bechhofer, R. E. 1954. "A single-sample multiple decision procedure for ranking means of normal populations with known variances". *The Annals of Mathematical Statistics* 25 (1): 16–39.

Boesel, J., B. L. Nelson, and S.-H. Kim. 2003. "Using ranking and selection to 'clean up' after simulation optimization". *Operations Research* 51 (5): 814–825.

Gupta, S. S. 1965. "On some multiple decision (selection and ranking) rules". *Technometrics* 7 (2): 225–245.

Nelson, B. L., J. Swann, D. Goldsman, and W. Song. 2001. "Simple procedures for selecting the best simulated system when the number of alternatives is large". *Operations Research* 49 (6): 950–963.

Rinott, Y. 1978. "On two-stage selection procedures and related probability inequalities". *Communications in Statistics – Theory and Methods* 7 (8): 799–811.