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.

BAYESIAN PROCEDURES FOR SELECTING SUBSETS OF ACCEPTABLE ALTERNATIVES

Jinbo Zhao[1]

[1]Department of Industrial and Systems Engineering, Texas A&M University, College Station, TX, USA

ABSTRACT

We develop a Bayesian framework for subset selection in ranking-and-selection problems that goes beyond the classical focus on identifying a single best alternative. The framework accommodates broad notions of acceptability, including delta-optimality, stochastically constrained optimality, and Pareto optimality. The central challenge is that evaluating posterior quantities involves integrating over high-dimensional and potentially non-convex acceptance regions and must be repeatedly computed for many candidate subsets. We circumvent this challenge by employing sample-average approximation to reformulate optimization problems that involve these integrals into mixed-integer programs. We also propose new myopic sequential sample allocation policies.

1 INTRODUCTION

Ranking-and-selection (R&S) concerns the problem of identifying acceptable alternatives from a finite set of alternatives, denoted by $\mathscr{K} = \{1,\dots,k\}$, where the multi-valued performance of each alternative $i \in \mathscr{K}$ is represented by a d-dimensional vector $W_i \in \mathbb{R}^d$ and all alternatives' performances collectively are represented by a matrix $\mathbf{W} = (W_1,\dots,W_k) \in \mathbb{R}^{k\times d}$. We assume $\mathbf{W}$ is unknown but can be inferred based on independent and identically distributed (i.i.d.) simulation outputs $X_{i1}, X_{i2},\dots$ drawn from alternative i for all $i \in \mathscr{K}$.

Classical R&S research has largely focused on the case where the decision maker aims to identify the single best alternative, often the one with the smallest expected performance. In practice, however, the definition of what constitutes an acceptable alternative is more varied. We let $\mathscr{A}\colon \mathbb{R}^{k\times d} \to 2^{\mathscr{K}}$ denote an *acceptability rule* that maps a performance matrix $\mathbf{W}$ to the set of acceptable alternatives $\mathscr{A}(\mathbf{W})$. Common notions of acceptability include feasibility and Pareto optimality. Unlike the single-best setting, the number of acceptable alternatives is unknown. This naturally motivates a *subset-selection* approach in which the goal is to return a subset $\mathscr{S} \subseteq \mathscr{K}$ that, with high probability, contains at least one, all, or at least a specified proportion of the acceptable alternatives.

We adopt a Bayesian perspective, treating $\mathbf{W}$ as random with a posterior distribution conditional on the evidence $\mathscr{E}$, which combines prior beliefs with observed simulation outputs. Posterior quantities derived from this distribution can be used to evaluate the quality of a candidate subset $\mathscr{S}$ and, importantly, remain statistically valid as additional samples are collected based on previous observations. Posterior quantities of interest can be expressed as integrals of the posterior distribution over combinations of the *acceptable regions* of the alternatives, defined as $\mathbb{A}_i = \{\mathbf{W} \in \mathbb{R}^{k\times d}\colon i \in \mathscr{A}(\mathbf{W})\}$ for $i \in \mathscr{K}$. In particular, we focus on three metrics for evaluating the quality of a subset $\mathscr{S}$ that capture different desiderata of subset selection:

- **Probability of acceptable inclusion**: $\mathrm{PAI}(\mathscr{S}) = \mathbb{P}(|\mathscr{A}(\mathbf{W}) \cap \mathscr{S}| \geq 1 \mid \mathscr{E}) = \mathbb{P}\big(\mathbf{W} \in \bigcup_{j\in\mathscr{S}} \mathbb{A}_j \mid \mathscr{E}\big)$.
- **Probability of acceptable subset selection**: $\mathrm{PASS}(\mathscr{S}) = \mathbb{P}(\mathscr{A}(\mathbf{W}) \subseteq \mathscr{S} \mid \mathscr{E}) = \mathbb{P}\big(\mathbf{W} \in \bigcap_{j\in\mathscr{S}} \mathbb{A}_j \mid \mathscr{E}\big)$.
- **Expected acceptable inclusion rate**: $\mathrm{EAIR}(\mathscr{S}) = \mathbb{E}\left[\frac{|\mathscr{A}(\mathbf{W})\cap\mathscr{S}|}{|\mathscr{A}(\mathbf{W})|} \,\middle|\, \mathscr{E}\right]$, where $0/0$ is defined to be 1 when $\mathscr{A}(\mathbf{W}) = \emptyset$.

Each of these subset-quality measures can be written as $Q(\mathscr{S};\mathscr{E}) = \mathbb{E}[q(\mathscr{A}(\mathbf{W}),\mathscr{S}) \mid \mathscr{E}]$ for some mapping $q\colon 2^{\mathscr{K}} \times 2^{\mathscr{K}} \to [0,1]$ that evaluates the quality of a fixed subset $\mathscr{S} \subseteq \mathscr{K}$ with respect to the (random) acceptable set $\mathscr{A}(\mathbf{W})$. Proposition 1 shows how these three measures are related; a frequentist analog is presented in Zhao et al. (2023).

Proposition 1 When $\mathbb{P}(\mathscr{A}(\mathbf{W}) = \emptyset \mid \mathscr{E}) = 0$, the following chain of inequalities holds for all $\mathscr{S} \subseteq 2^{\mathscr{K}}$:

$$\text{PASS}(\mathscr{S}) \leq \text{EAIR}(\mathscr{S}) \leq \text{PAI}(\mathscr{S}).$$

2 SAMPLE-AVERAGE APPROXIMATION

To avoid evaluating integrals for all candidate subsets, we employ a *sample-average approximation* (SAA) framework. The framework operates under minimal assumptions: we do not restrict the form of the acceptability rule $\mathscr{A}(\cdot)$ or the posterior distribution of $\mathbf{W}$, but assume only that posterior samples can be generated. More specifically, we propose drawing posterior samples $\mathbf{W}^1, \ldots, \mathbf{W}^N \overset{\text{i.i.d.}}{\sim} p(\mathbf{W} \mid \mathscr{E})$ and approximating $Q(\mathscr{S}; \mathscr{E})$ by

$$\widehat{Q}_N(\mathscr{S}; \mathscr{E}) = \frac{1}{N} \sum_{n=1}^{N} q(\mathscr{A}(\mathbf{W}^n), \mathscr{S}) \text{ for all } \mathscr{S} \in 2^{\mathscr{K}}.$$

Let m be a user-specified upper bound on the subset size, representing the maximum number of alternatives the decision maker is willing to consider. Consider the SAA subproblem of identifying the subset of size at most m with the largest subset quality:

$$\max_{\mathscr{S} \subseteq \mathscr{K}} \quad \widehat{Q}_N(\mathscr{S}; \mathscr{E}) \quad \text{s.t.} \quad |\mathscr{S}| \leq m. \tag{FS-SAA}$$

(FS-SAA) can be formulated as a mixed-integer program. If the resulting estimated optimal quality is lower than some user-specified threshold $1 - \alpha$, then this suggests that the current evidence is insufficient and additional simulation replications are needed.

We develop a general-purpose myopic sample allocation policy that, like the SAA formulation, requires only the ability to generate posterior samples. At iteration t, let $\mathscr{E}_t$ denote the evidence collected so far and $p(\mathbf{W} \mid \mathscr{E}_t)$ denote the posterior distribution. Solving (FS-SAA) yields a candidate subset $\widehat{\mathscr{S}_t^*}$ and its estimated quality $\widehat{Q}_N(\widehat{\mathscr{S}_t^*}; \mathscr{E}_t)$. To choose the next alternative, we consider the hypothetical effect of obtaining one additional simulation output from one alternative. Let $\tilde{X}_i$ represent a prospective observation of alternative i, and let $\mathscr{E}_t \cup \tilde{X}_i$ represent the augmented evidence if $\tilde{X}_i$ were observed. We define the one-step myopic gain as

$$\Delta_i(\mathscr{E}_t) = Q(\widehat{\mathscr{S}_t^*}; \mathscr{E}_t \cup \tilde{X}_i) - Q(\widehat{\mathscr{S}_t^*}; \mathscr{E}_t),$$

and consider $i_t^* \in \arg\max_{i \in \mathscr{K}} \Delta_i(\mathscr{E}_t)$ as the next alternative to simulate.

To avoid intractable integration, we approximate this selection rule by Monte Carlo:

$$\widehat{i_t^*} = \arg\max_{i \in \mathscr{K}} \widehat{\Delta}_i(\mathscr{E}_t) = \arg\max_{i \in \mathscr{K}} \widehat{Q}_{N'}\left(\widehat{\mathscr{S}_t^*}; \mathscr{E}_t \cup \tilde{X}_i\right),$$

where

$$\widehat{Q}_{N'}\left(\widehat{\mathscr{S}_t^*}; \mathscr{E}_t \cup \tilde{X}_i\right) = \frac{1}{N'} \sum_{n=1}^{N'} q\left(\mathscr{A}(\mathbf{W}^{n,i}), \widehat{\mathscr{S}_t^*}\right) \text{ and } \mathbf{W}^{n,i} \overset{\text{i.i.d.}}{\sim} p(\mathbf{W} \mid \mathscr{E}_t \cup \tilde{X}_i) \text{ for all } i \in \mathscr{K} \text{ and } n = 1, \ldots, N'.$$

Alternative $\widehat{i_t^*}$ is then sampled, the posterior distribution updated, and the procedure repeated until a stopping rule is satisfied, such as when the optimal value of (FS-SAA) exceeds $1 - \alpha$.

We are currently investigating methods for generating prospective simulation outputs $\tilde{X}_i$ and studying how $\widehat{\mathscr{S}_t^*}$ performs in terms of the aforementioned quality metrics. Numerical experiments indicate that sampling $\tilde{X}_i$ from the empirical CDF can perform poorly with respect to these Bayesian metrics, sometimes even underperforming equal allocation, yet still performs adequately under frequentist criteria.

REFERENCES

Zhao, J., J. Gatica, and D. J. Eckman. 2023. “Screening Simulated Systems for Optimization”. In *2023 Winter Simulation Conference (WSC)* https://doi.org/10.1109/WSC60868.2023.10408036.