Proceedings of the 2024 Winter Simulation Conference
H. Lam, E. Azar, D. Batur, S. Gao, W. Xie, S. R. Hunter, and M. D. Rossetti, eds.

RANKING AND CONTEXTUAL SELECTION FOR DATA-DRIVEN DECISION MAKING

Gregory Keslin

Dept. of Industrial Engineering & Management Sciences, Northwestern University, Evanston, IL, USA

ABSTRACT

This extended abstract is an overview of ranking and contextual selection (R&CS), a new procedure for ranking and selection with covariates. R&CS runs individual ranking-and-selection experiments at each covariate in an experiment design sampled from the covariate distribution. The systems selected by the ranking-and-selection procedures and the design itself form a classifier for selecting the best system at any future covariate value. Associated with the classifier is an assessment of its accuracy that is proven to satisfy a finite-sample coverage guarantee.

1 INTRODUCTION

In operations research there is significant interest in optimal decision making among systems with stochastic outputs. Specifically, a set of alternatives, $\{1,2,\ldots,p\}$ are associated with a measure of performance, $\mu_i, i \in \{1,2,\ldots,p\}$, which is the expectation of a stochastic output from system i. The optimal decision, $j^\star$, is the system that attains the maximum mean performance across the set of alternatives: $j^\star \in \text{argmax}_{i\in\{1,2,\ldots,p\}}\mu_i$. The setting of simulation optimization (SO) further assumes that each system can be simulated to generate synthetic replications which mimic the behavior of the true system to estimate μ_i. Such decision-making problems are typically referred to as ranking-and-selection (R&S) problems. Decades of literature has focused on designing procedures to solve R&S problems. For a more in-depth overview, see Nelson and Pei (2021).

More recently a context-sensitive version of R&S has been studied. Rather than selecting a static single best system, the goal is to select the best conditional on additional covariate information, $\boldsymbol{x} \in \mathscr{X}$. The covariate is modeled as a random vector with a distribution $\mathscr{F}$. Associated with each system is the mean performance function conditional on the value of the covariate, $\mu_i(\cdot), i \in \{1,2,\ldots,p\}$.

A simple approach to accommodate the inclusion of the covariate is to run a R&S procedure after the user observes the realization of the covariate, $\boldsymbol{X} = \boldsymbol{x}$. The problem with this approach is that all simulation must be done "online," so only after the covariate is realized can inference about the performance of any system be gleaned. If the computational effort required to simulate each system is large, or the allowable time to make a decision after the covariate is realized is short, then running a full blown R&S procedure may be impractical or infeasible. Instead, approaches have been developed that simulate "offline" prior to the realization of the covariate.

For the purposes of the extended abstract, we call these "offline" procedures ranking and selection with covariates (R&S+C). To obtain strong finite-sample guarantees, existing R&S+C procedures assume known structure about each $\mu_i(\cdot)$. For example, Shen et al. (2021) assume $\mu_i(\cdot)$ is linear in the covariate. Further, R&S+C procedures do not exploit the wealth of existing and highly efficient R&S procedures, which makes implementation more difficult.

To rectify these problems, Keslin et al. (2022) and Keslin et al. (2024) introduced ranking and contextual selection (R&CS). Further, Keslin et al. (2024) provides finite-sample guarantees for the R&CS procedure that require no knowledge about the structure of each $\mu_i(\cdot)$.

2 R&CS OVERVIEW

We assume that there exists a covariate $\boldsymbol{X} \sim \mathscr{F}$ obtaining values in $\mathscr{X}$. For each system j and covariate $\boldsymbol{x} \in \mathscr{X}$, replications can be simulated, $Y_{\ell,j}(\boldsymbol{x}),\ \ell = 1,2,\ldots$ with $\mathbb{E}\left(Y_{\ell,j}\right) = \mu_j(\boldsymbol{x})$. These replications are used to fit a nearest-neighbor classifier, $\widehat{J}(\cdot)$, which selects a system for each covariate value. Because the replications are random, $\widehat{J}(\boldsymbol{x})$ is a random function of $\boldsymbol{x}$.

R&CS requires a design consisting of m covariates, $\mathscr{D}_m = \{\boldsymbol{X}_1 \boldsymbol{X}_2, \ldots, \boldsymbol{X}_m\}$. For each design covariate, $\boldsymbol{X}_i$, a R&S procedure is run and a system is selected, $R(\boldsymbol{X}_i)$. The corresponding set of R&S choices is $\mathscr{R}_m = \{R(\boldsymbol{X}_1), R(\boldsymbol{X}_2), \ldots, R(\boldsymbol{X}_m)\}$. The pairing of $\mathscr{D}_m$ and $\mathscr{R}_m$ forms a database used to fit a nearest-neighbor classifier function. Keslin et al. (2024) provide a detailed justification for using a single neighbor.

Along with the classifier, R&CS provides an associated performance assessment $\widehat{\Delta}_m^{1-\alpha}$. Details about $\widehat{\Delta}_m^{1-\alpha}$ can be found in Keslin et al. (2024); here we present the finite-sample guarantee provided by $\widehat{\Delta}_m^{1-\alpha}$.

We define the optimality gap for covariate $\boldsymbol{x}$ as the difference between the mean of the optimal and selected system at $\boldsymbol{x}$: $\Delta(\boldsymbol{x}) = \mu_{j^*(\boldsymbol{x})}(\boldsymbol{x}) - \mu_{\widehat{J}(\boldsymbol{x})}(\boldsymbol{x})$. Controlling this optimality gap across the covariate space is the primary goal of an effective classifier, $\widehat{J}(\cdot)$. If this optimality gap is "small" across the covariate space, a user can be confident that the systems selected by $\widehat{J}(\cdot)$ are "good."

Because $\Delta(\boldsymbol{x})$ may be small at some values of $\boldsymbol{x}$ and large at other values, an assessment with guarantees on the optimality gap across the entire covariate space must be considered. Therefore, R&CS requires a user-specified confidence level, $0 < 1-\alpha < 1$, for the assessment $\widehat{\Delta}_m^{1-\alpha}$. Theorem 1 proves that the probability that the optimality gap at a random covariate will be below $\widehat{\Delta}_m^{1-\alpha}$ is above $1-\alpha$. Thus, $\widehat{\Delta}_m^{1-\alpha}$ assesses the performance of $\widehat{J}(\cdot)$. If $\widehat{\Delta}_m^{1-\alpha}$ is below a user-specified acceptable gap, δ, then the classifier has sufficient performance. On the other hand, if $\widehat{\Delta}_m^{1-\alpha} > \delta$, then there exists a non-negligible probability that the optimality gap of the covariate will be above δ.

Theorem 1 assumes the following (i) Each $\boldsymbol{X}_i \in \mathscr{D}_m$ is i.i.d. with $\boldsymbol{X}_i \sim \mathscr{F}$; (ii) $\widehat{J}(\cdot)$ is a nearest-neighbor classifier that breaks ties by choosing the design point with smallest index in $\mathscr{D}_m$; and (iii) $\Delta(\cdot)$ is a known function at the design points, $\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_m$.

Theorem 1 Under Assumptions (i)–(iii),

$$\mathbb{P}\left(\Delta(\boldsymbol{X}) \leq \widehat{\Delta}_m^{1-\alpha}\right) \geq 1-\alpha.$$

Notice that the probability in Theorem 1 is for the random covariate, $\boldsymbol{X}$, and the random classifier function, $\widehat{J}(\cdot)$. The guarantee given by Theorem 1 places no assumptions on each $\mu_i(\cdot)$ or the R&S procedure used to generate the database. However, Assumption (iii) from Theorem 1 is unlikely to be satisfied in practice. Therefore, Keslin et al. (2024) provide asymptotic inference when using plug-in estimators of $\Delta(\boldsymbol{X}_i)$ instead of $\Delta(\boldsymbol{X}_i)$ for each design covariate $\boldsymbol{X}_i$.

Extensive empirical results in Keslin et al. (2022) and Keslin et al. (2024) demostrate not only that the desired optimality-gap coverage is achieved, but also that R&CS makes effective use of offline simulation by exploiting highly efficient, off-the-shelf R&S procedures.

REFERENCES

Keslin, G., B. L. Nelson, B. K. Pagnoncelli, M. Plumlee and H. Rahimian. 2024. "Ranking and Contextual Selection". *Operations Research*. Forthcoming.

Keslin, G., B. L. Nelson, M. Plumlee, B. K. Pagnoncelli and H. Rahimian. 2022. "A Classification Method for Ranking and Selection with Covariates". In *2022 Winter Simulation Conference (WSC)*, 1–12. IEEE.

Nelson, B. L. and L. Pei. 2021. *Foundations and Methods of Stochastic Simulation*. 2nd ed. New York: Springer.

Shen, H., L. J. Hong, and X. Zhang. 2021. "Ranking and Selection with Covariates for Personalized Decision Making". *INFORMS Journal on Computing* 33(4):1500–1519.