Proceedings of the 2017 Winter Simulation Conference
W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, eds.

SEQUENTIAL PROBABILITY RATIO TEST FOR MULTIPLE-OBJECTIVE RANKING AND SELECTION

Wenyu Wang
Hong Wan

Department of Industrial Engineering
Purdue University
315 N Grant St
West Lafayette, IN 47906, USA

ABSTRACT

In this paper, we introduce a sequential procedure for the Multi-Objective Ranking and Selection (MOR&S) problems that identifies the Pareto front with a guaranteed probability of correct selection (PCS). In particular, the proposed procedure is fully sequential using the test statistics built upon the generalized sequential probability ratio test (GSPRT). The main features of the new proposed procedure are: 1) a unified framework, the new procedure treats the multi-objective problems in the same way as the single-objective problems; 2) an indifference-zone-free formulation, the new procedure eliminates the necessity of indifference-zone parameter; 3) asymptotically optimality, the *GSPRT* achieves asymptotically the shortest expected sample size among all sequential tests; 4) general distribution, the procedure uses the empirical likelihood for generally distributed observation. A numerical evaluation demonstrates the efficiency of the new procedure.

1 INTRODUCTION

The Ranking and Selection problems (R&S) are to select the best system(s) among a finite set of alternatives. This paper considers the multi-objective R&S (MOR&S) problems. In the context of multi-objectives, the "best" systems are characterized by the *Pareto front* that systems within the Pareto front are not *dominated* by any other alternatives. According to Pasupathy and Ghosh (2013), most procedures on R&S fall into two categories: (a) class $\mathscr{P}$ procedures that guarantee a lower bound for the *probability of correct selection* (PCS), (b) class $\mathscr{L}$ procedures that efficiently allocate a finite computing budget. In this paper, we propose a class $\mathscr{P}$ procedure for the MOR&S.

Most previous studies on MOR&S are class $\mathscr{L}$ procedures. These procedures include the multi-objective optimal computing budget allocation (Lee, Chew, and Teng 2007; Lee et al. 2010), which proposes a computing allocation framework for MOR&S under the assumption of normally distributed outputs and independent objectives; optimal sampling laws proposed by Feldman, Hunter, and Pasupathy (2015) that allocates sampling budget accounting for dependence between the objective vector estimates for a single system in the context of bi-objective; myopic strategy M-MOBA (Branke and Zhang 2015; Branke, Zhang, and Tao 2016) that targets to minimize the difference in hyper-volume between the observed means of the perceived Pareto front and the true Pareto front.

To the best of the authors' knowledge, this work is the first class $\mathscr{P}$ procedure on MOR&S. There are many established procedures that control PCS for the single-objective R&S problems, i.e., KN procedure (Kim and Nelson 2001), KN++ procedure (Kim and Nelson 2006), BIZ procedure (Frazier 2014), and IZ-free procedure (Fan, Hong, and Nelson 2016). These procedures, however, cannot be easily extended to MOR&S since there could be more than one best systems in the Pareto front. Therefore, their common strategy that compares an alternative with the single perceived best system fails.

To solve this problem, we directly test on whether an alternative is in the Pareto front or not. We call such a test the *candidacy test* and approach it using the generalized sequential probability ratio test (GSPRT) (Kiefer and Weiss 1957; Li, Liu, and Ying 2014). To achieve the desired level of error probabilities, the GSPRT continues to gather samples until the likelihood ratio statistic crosses one of two decision boundaries. The proposed procedure has following characteristics.

1. **An indifference-zone-free formulation**. This formulation eliminates the necessity to specify the indifference-zone parameter.
2. **A GLR-based procedure for multi-objective R&S problems**. We propose to test directly on whether an alternative is "best", i.e., in the Pareto front. We call such a test the *candidacy test* and approach it using the sequential probability ratio test (SPRT) (Wald 1945, Wald 1945, Kiefer and Weiss 1957, Pavlov 1988, Pavlov 1991). To achieve the desired level of error probabilities, the SPRT continues to gather samples until the test statistic crosses one of the decision boundaries.
3. **Adaptive implementation for various distributions**. Most classic R&S procedures are valid only when observations are from normal distributions. In many cases, normality assumption is reasonable. Nonetheless, there are special situations when this is not the case. The proposed procedures are flexible to be implemented under various assumptions of distributions.

The rest of the paper is organized as follows. We introduce the notations and definitions in section 2. In Section 3, we present and justify the sequential procedure to identify the Pareto front for MOR&S under the normality assumption. In section 4, we relax the normality assumption and propose a sequential procedure based on empirical likelihood. A special case of single-objective R&S is considered in Section 5. Section 6 presents numerical evaluations.

2 PROBLEM STATEMENT

Multi-objective R&S (MOR&S) is a natural extension of single-objective R&S. Let us consider a MOR&S problem with K independent candidate systems and d objectives for each system. Denote $\Omega = \{1,...,K\}$ as the candidate set. The mean vector of the alternative i is $\boldsymbol{\mu}_i = [\mu_{i1}, \mu_{i2}, ..., \mu_{iD}]' \in \mathbb{R}^D$. We assume that $\boldsymbol{\mu}_i$ is unknown and can only be observed with noise. When the alternative i is sampled, a random vector $\mathbf{X}_i \in \mathbb{R}^D$ with $\mathbb{E}[\mathbf{X}_i] = \boldsymbol{\mu}_i$ is observed. Throughout this thesis, we assume following condition hold:

A1 Systems are unique in all objectives, i.e., $\mu_{id} \neq \mu_{jd}$ for all $i, j \in \{1,2,3,...,K\}$ and $d \in \{1,2,3,...,D\}$.

Condition A1 is critical for the finite-time termination of the proposed procedure as it avoids the situations when two or more systems have the same mean performances.

The R&S problems are to identify the "best" system(s). In the context of multi-objective, the "best" systems are alternatives that are *not-dominated*, which are defined as below

Definition 1 An alternative i is **dominated** by an alternative j (denoted by $\boldsymbol{\mu}_i \preceq \boldsymbol{\mu}_j$), if $\mu_{id} \leq \mu_{jd}$ for all objectives $d = 1,2,...,D$.

An alternative i is **not dominated** by an alternative j (denoted by $\boldsymbol{\mu}_i \npreceq \boldsymbol{\mu}_j$), if $\mu_{id} > \mu_{jd}$ for at least one objective.

An alternative is called **Pareto optimal** if it is not dominated by all other alternatives, and the set of Pareto optimal alternatives is called the **Pareto front**.

Denote the **Pareto front** as $S_1 \triangleq \{i | \forall j \in \Omega, i \npreceq j\}$ and the **non-Pareto-front** as $S_0 \triangleq \{i | \exists j \in \Omega, i \preceq j\}$.

Remark In multi-objective optimization, the definition of *dominated* above is also called the *weakly dominated*, as opposed to the concept of strongly dominated that requires $\mu_{id} < \mu_{jd}$ for all objectives. This distinction, in stochastic optimization, however, is impractical as a sharp distinction is not possible due to sampling error. Therefore, we do not distinguish these two concepts.

The objective of the proposed procedures is to identify the Pareto front, S_1, correctly with a guaranteed *probability of correct selection*. Denote the perceived Pareto front as $\tilde{S}_1$, the *PCS* can be written as $PCS = P(\tilde{S}_1 = S_1) \geq 1-\alpha$, where α is a pre-specified error probability. The proposed procedures deliver the guarantee of *PCS* through conduct two types of hypothesis tests, namely, *pairwise test* and *candidacy test*.

Definition 2 The **pairwise test** between i and j determines whether system i is dominated by j.

$$H_0 \ : \ \boldsymbol{\mu}_i \preceq \boldsymbol{\mu}_j \text{ and } H_1 \ : \ \boldsymbol{\mu}_i \npreceq \boldsymbol{\mu}_j \tag{1}$$

The **candidacy test** of i determines whether an alternative i is in the Pareto front.

$$H_0 \ : \ i \in S_0 \text{ and } H_1 \ : \ i \in S_1 \tag{2}$$

Remark The *pairwise test* defined above differs from pairwise comparison in single-objective R&S procedures (i.e. KN, KN++, and IZ-free) in two aspects. First, those procedures only compare an alternative with the practical best alternative; whereas the pairwise test in our procedures consider all pairs. Second, pairwise comparisons in those procedures are reflexive: $\mu_i > \mu_j$ is equivalent to $\mu_j > \mu_i$; whereas $\boldsymbol{\mu}_i \npreceq \boldsymbol{\mu}_j$ doesn't imply $\boldsymbol{\mu}_j \preceq \boldsymbol{\mu}_i$ in our procedures.

3 METHODOLOGY

A common modeling strategy in other leading frequentist approaches is to approximate the sum of differences between two alternatives as a Brownian motion with a drift term, where the drift term is determined by the true difference between their means (Kim and Nelson 2001; Kim and Nelson 2006; Fan and Hong 2014; and Fan, Hong, and Nelson 2016). These procedures construct the continuation region and make an elimination decision when the BM process exist the region. This strategy, however, is an indirect approach to compare systems, and thus conservative.

In this thesis, we propose a direct modeling method using the Sequential Probability Ratio Test (SPRT). The classic *Sequential Probability Ratio Test* (SPRT) was proposed by Wald and others (Wald 1945, Wald and Wolfowitz 1948) in response to the need of efficient testing tools. The SPRT is a sequential procedure that does not fix the sample size. Let $X_1, X_2, \ldots$ be i.i.d. random variables with the same distribution f. The SPRT tests a simple hypothesis $H_0 : f = f_0$ versus a simple alternative hypothesis $H_1 : f = f_1$. It tracks the likelihood ratio (LR) statistic $\Lambda_n = \prod_i^n (f_1(X_i)/f_0(X_i))$, and terminates at stopping time $\tau = \inf\{n : \Lambda_n \geq B \text{ or } \Lambda_n \leq A\}$, where A and B are decision boundaries satisfying $0 < A < 1 < B < \infty$. At termination, the null hypothesis H_0 is rejected if $\Lambda_\tau \geq B$ or is accepted if otherwise. Let α_1 be the type I error probability and α_2 be the type II error probability, Wald (1945) shows the explicit solution of A and B in terms of α_1 and α_2 are approximately: $A = \frac{\alpha_2}{1-\alpha_1}$ and $B = \frac{1-\alpha_2}{\alpha_1}$. It is proved that SPRT is optimal in terms of minimizing the average sample size among all tests under both H_0 and H_1 at a given level of type I and type II error probabilities (Wald and Wolfowitz 1948).

The Generalized Sequential Probability Ratio Test (GSPRT) is an extension of SPRT is to test two or more composite hypotheses. Let Θ be the parameter space of the probability density function f. In this paper, GSPRT tests two composite hypotheses

$$H_0 \ : \ \theta \in \Theta_0 \ \text{ and } H_1 \ : \theta \in \Theta_1 \tag{3}$$

where Θ_1 and Θ_2 are disjoint subsets of Θ.

The GSPRT uses the generalized likelihood ratio (GLR) statistics instead of simple likelihood ratios. In this paper, we use the GSPRT of the following form,

$$\Lambda_n = \frac{\sup_{\theta_1 \in \Theta_1} \prod_{i=1}^n f_{\theta_1}(X_i)}{\sup_{\theta_0 \in \Theta_0} \prod_{i=1}^n f_{\theta_0}(X_i)} \tag{4}$$

Essentially, the GLR statistic "self-tunes" to information about the true θ over the course of the test. The GSPRT stops sampling after the nth observations if the GLR crosses either of the two boundaries $\Lambda_n \geq B_n$ or $\Lambda_n \leq A_n$. The null hypothesis is rejected if $\Lambda_n \geq B_n$ and is accepted if $\Lambda_n \leq A_n$. For a pair of A_n and B_n, denote the stopping time as

$$\tau = \inf\{n : \Lambda_n \geq B_n \text{ or } \Lambda_n \leq A_n\} \tag{5}$$

The classification error at the stopping time τ is determined by A_n and B_n together.

$$\max\{\sup_{\theta\in\Theta_1} \mathbf{P}_\theta(\Lambda_\tau \leq A_\tau), \sup_{\theta\in\Theta_0} \mathbf{P}_\theta(\Lambda_\tau \geq B_\tau)\} \tag{6}$$

where $\mathbf{P}_\theta$ represents the probability measure for a given θ.

For the validity of the GSPRT in the context of the R&S problem, we introduce following technical conditions:

A2 The sample spaces of X is compact;
A3 The density $f_\theta(x)$ is continuous in (θ, x);
A4 $\rho(\theta_0, \theta_1) > 0$ for all $\theta_0 \neq \theta_1$, where $\rho(\theta_0, \theta_1)$ is the Kullback-Leibler distance from θ_0 to θ_1:

$$\rho(\theta_0, \theta_1) = \mathbb{E}_{\theta_0}\{\log f_{\theta_0}(X) - \log f_{\theta_1}(X)\} \tag{7}$$

where $\mathbb{E}_\theta$ represents taking the expectation with respect to $f(\theta)$.

Under these conditions, we have the following theorem to show that the GSPRT with the time-invariant testing boundaries $(A, B) = (\alpha, \frac{1}{\alpha})$ controls the type I and type II error probabilities at the level of α.

Theorem 1 (Theorem 2.1 in Li, Liu, and Ying (2014)) Consider the composite null hypothesis against composite alternative hypothesis given as in (3). With testing boundaries $A = \alpha$ and $B = \frac{1}{\alpha}$, the GSPRT admits stopping time (5) and the GLR statistic (4). Under Conditions A1-5, the maximal error probability is asymptotically bounded above by

$$\max\{\sup_{\theta\in\Theta_1} \mathbf{P}_\theta(\Lambda_\tau \leq \alpha), \sup_{\theta\in\Theta_0} \mathbf{P}_\theta(\Lambda_\tau \geq \frac{1}{\alpha})\} \cong \alpha \tag{8}$$

where $\mathbf{P}_\theta$ represents the probability measure for a given θ.

Remark In the classic SPRT method, the thresholds $A = \frac{\alpha_2}{1-\alpha_1}$ and $B = \frac{1-\alpha_2}{\alpha_1}$, where α_1 and α_2 are the pre-specified type I and type II error probabilities, have the same decay rate as in 8.

The choice of testing boundaries $(A, B) = (\alpha, \frac{1}{\alpha})$ is not only valid to control the error probability at α level, but also asymptotically optimal in the sense that it achieves the shortest expected sample size as α tends to zero. The next theorem states the asymptotic approximation of the expected stopping time.

Theorem 2 (Theorem 2.3 in Li, Liu, and Ying (2014)) Under the setting and the conditions of Theorem 1, the expected stopping time of the GSPRT with $(A, B) = (\alpha, \frac{1}{\alpha})$ admits the following asymptotic approximation

$$\mathbb{E}_{\theta_0}(\tau) \leq \frac{\log \alpha^{-1}}{\inf_{\theta_1\in\Theta_1} \rho(\theta_0, \theta_1)}(1 + o(1)) \text{ and } \mathbb{E}_{\theta_1}(\tau) \leq \frac{\log \alpha^{-1}}{\inf_{\theta_0\in\Theta_0} \rho(\theta_1, \theta_0)}(1 + o(1)) \tag{9}$$

as $\alpha \to 0$ for all $\theta_0 \in \Theta_0$ and $\theta_1 \in \Theta_1$.

The asymptotic optimality of the GSPRT is proved by showing that the expected stopping time stated in Theorem 2 asymptotically achieves the lower bound of the stopping time of any sequential procedures with guaranteed error probabilities.

4 THE SEQUENTIAL PROCEDURE FOR MOR&S

In this section, we present the GLR-based framework that controls *PCS* for R&S . This framework is valid for both the single-objective and the multi-objective situations. The framework makes no assumption on the underlying distributions and thus could be implemented under various assumptions on distributions. Section 4.1 presents the framework, and Section 4.2 focuses on the calculation of the statistics. Statistical validities are given in Section 4.3.

4.1 The Framework

The framework is illustrated in Algorithm 1. The first step in the framework is divide the *PCS* using the Bonferroni method. Recall that the objective is to identify the Pareto front S_1 with a guaranteed *PCS*. Denote the perceived "best" systems as $\tilde{S}_1$ and the perceived "inferior" alternatives $\tilde{S}_0$. Thus, we can decompose the *PCS* as:

$$\begin{aligned} PCS &= P(\tilde{S}_1 = S_1) \\ &= 1 - P\{\tilde{S}_1 \neq S_1\} \\ &= 1 - P\{[\cup_{i\in S_1} i \in \tilde{S}_0] \cup [\cup_{i\in S_0} i \in \tilde{S}_1]\} \\ &\geq 1 - \sum_{i\in S_1} P(i \in \tilde{S}_0) - \sum_{i\in S_0} P(x_i \in \tilde{S}_1) \\ &= 1 - \sum_{i\in\Omega} P(i \text{ is selected incorrectly}) \end{aligned}$$

Algorithm 1 The GSPRT Framework for Ranking and Selection

Inputs:

$PCS = 1-\alpha$ ▷ Desired PCS level

$\Omega = \{1,...,K\}$ ▷ Pool of candidate systems

Procedure:

set $\alpha_i = \frac{\alpha}{K}$
set $\tilde{S}_0 = \emptyset$ and $\tilde{S}_1 = \Omega$;
set $A_i = \alpha_i$ and $B_i = 1/\alpha_i$; ▷ A and B are the GSPRT test boundaries
generate n_0 samples for each system in $\tilde{S}_1$;
set $n = n_0$;
repeat
    generate 1 more sample for all systems in $\tilde{S}_1$;
    set $n = n+1$;
    for $i \in \tilde{S}_1$ **do**
        calculate the GLR statistic Λ_n^i for i;
        if the Λ_n^i crosses the lower boundary **then**
            set $\tilde{S}_1 = \tilde{S}_1/\{i\}$ and $\tilde{S}_0 = \tilde{S}_0 \cup \{i\}$
        end if
    end for
until Λ_n^i are above upper boundary for all $i \in \tilde{S}_1$
return $\tilde{S}_1$

Let us denote the error probability of the candidacy test on the alternative i is α_i, the *PCS* is bounded below by $1 - \sum \alpha_i$. Therefore, to deliver the $PCS = 1-\alpha$ guarantee, we use $\alpha_i = \frac{\alpha}{K}$. Note that $\alpha_i =$

$P(x_i \text{ is selected incorrectly})$ is both the maximal type I and type II error probabilities of the candidacy test of x_i. Note that this decomposition is valid without the independent assumptions. Thus, the framework is robust against any dependency or topological structures.

According to Theorem 1, the framework sets the testing boundaries as $A = \alpha_i$ and $B = 1/\alpha$. Then, the framework conducts the candidacy tests on the remaining systems in $\tilde{S}_1$ and eliminates inferior systems sequentially. It terminates when all of the remaining systems are found as the "best" systems. Next subsection shows how to calculate the GLR statistic for the candidacy test.

4.2 The GLR Statistic

The GLR statistic of the candidacy test is the cornerstone of the framework. Recall that the candidacy test of the alternative i is on whether an alternative is in the Pareto front S_1 or not, i.e., $H_0 : i \in S_0$ and $H_1 : i \in S_1$. Let us denote Λ_n^i as the generalized likelihood ratio (GLR) statistic of the *candidacy test* on the alternative i after accumulating n observations, and state Λ_n^i as

$$\Lambda_n^i = \frac{\mathscr{L}(i \in S_1)}{\mathscr{L}(i \in S_0)} \tag{10}$$

where $\mathscr{L}(i \in S_1)$ ($\mathscr{L}(i \in S_0)$) is the maximized likelihoods under the condition that the alternative i is (not) in the Pareto front.

$$\begin{aligned}
\mathscr{L}(i \in S_1) &= \sup_{\{\mu_1,\ldots,\mu_K\} \in \{i \in S_1\}} L(\mu_1, \ldots, \mu_K | \text{all observations accumulated to the } n\text{th round}) \\
\mathscr{L}(i \in S_0) &= \sup_{\{\mu_1,\ldots,\mu_K\} \in \{i \in S_0\}} L(\mu_1, \ldots, \mu_K | \text{all observations accumulated to the } n\text{th round})
\end{aligned}$$

The closed-form expression of Λ_n^i is hard, if possible, to obtain analytically, since the event that the alternative i is not in the Pareto front contains all possible scenarios that alternative i is dominated by another system. Instead, we consider a conservative approximation of Λ_n^i using the fact that $\{i \in S_1\} = \{i \not\preceq j \text{ for all } j \neq i\}$ and $\{i \in S_0\} = \{i \preceq j \text{ for at least one } j \neq i\}$.

Denote Λ_n^{ij} as the GLR statistic of the pairwise test between the alternative i and the alternative system j:

$$\Lambda_n^{ij} = \frac{\mathscr{L}(\boldsymbol{\mu}_i \not\preceq \boldsymbol{\mu}_j)}{\mathscr{L}(\boldsymbol{\mu}_i \preceq \boldsymbol{\mu}_j)} \tag{11}$$

where $\mathscr{L}(\boldsymbol{\mu}_i \preceq \boldsymbol{\mu}_j)$ ($\mathscr{L}(\boldsymbol{\mu}_i \not\preceq \boldsymbol{\mu}_j)$) is the maximized likelihood under the condition that the alternative x is (not-)dominated by x.

$$\begin{aligned}
\mathscr{L}(i \in S_1) &= \sup_{\{\mu_1,\ldots,\mu_K\} \in \{i \in S_1\}} L(\mu_1, \ldots, \mu_K | \text{all observations accumulated to the } n\text{th round}) \\
\mathscr{L}(i \in S_0) &= \sup_{\{\mu_1,\ldots,\mu_K\} \in \{i \in S_0\}} L(\mu_1, \ldots, \mu_K | \text{all observations accumulated to the } n\text{th round})
\end{aligned}$$

Define $\Lambda_n^{i\star}$ as

$$\Lambda_n^{i\star} \triangleq \min_{j \in \Omega/i} \Lambda_n^{ij} \tag{12}$$

$\Lambda_n^{i\star}$ is the minimum of Λ_n^{ij} over all $j \neq i$ and provides a good estimation of Λ_n^i. We claim that $\Lambda_n^{i\star}$ is a conservative approximation of Λ_n^i in the sense that decision made upon $\Lambda_n^{i\star}$ admits less error rate than Λ_n^i.

Theorem 3 Under the definition of $\Lambda_n^{i\star}$ as in 12, Λ_n^i and $\Lambda_n^{i\star}$ are in favor of the same decision, and the absolute value of the $\log \Lambda_n^{i\star}$ is bounded above by the absolute value of $\log \Lambda_n^i$.

Remark $|\log \Lambda_n^i| \geq |\log \Lambda_n^{i,\star}|$ indicates Λ_n^i overshoot when $\Lambda_n^{i,\star}$ crosses the boundaries. Therefore, $\Lambda_n^{i\star}$ is a less powerful GLR statistic than Λ_n^i. Although conservative, this approximation is found very tight through simulations (See numerical evaluation section).

The GLR statistic of the pairwise test is much easier to calculate than its counterpart of the candidacy test. The calculation, however, requires the assumptions of the underlying distributions. We will come back to the detailed calculations in the next two sections.

4.3 The Statistical Validity

Combining Theorem 1 and Theorem 3, we establish the statistical validity of Framework 1 and summarize it as the following theorem.

Theorem 4 Under the settings of Theorem 1 and Theorem 3 and Assumptions A1 and A2, the Procedure 2 terminates in finite time with probability 1 and selects the Pareto front with probability at least $1-\alpha$ in asymptotically sense, i.e.,

$$\limsup_{\alpha \to 0^+} \frac{1-P(CS)}{\alpha} \geq 1 \tag{13}$$

The above theorem states that the Framework 1 selects the entire Pareto front at the stopping time τ. Nonetheless, if the procedure stops before τ, it is still guaranteed that the Pareto front is retained in the candidate pool $\tilde{S}_1$. The result is summarized as the following Theorem.

Theorem 5 Under the settings of Theorem 1 and Theorem 3 and Assumption A1 and A2, if the Framework 1 terminates at n before τ, then the candidate pool $\tilde{S}_1$ retains the Pareto front with S_1 with probability at least $1-\alpha$, in asymptotically sense, i.e.,

$$\limsup_{\alpha \to 0^+} \frac{S_1 \in \tilde{S}_1}{\alpha} \geq 1 \tag{14}$$

4.4 The Implementation under Normality Assumption

The proposed framework is valid for a variety of distributions. For demonstration purposes, we consider a MOR&S problem with observations from multivariate normal distributions.

Let us denote $\{X_{i,1},...,X_{i,n}\}$ as the i.i.d. random observations obtained from the alternative i. Assume $\{X_{i,n}\}$ are from multivariate normal distribution $N(\boldsymbol{\mu}_i, \Sigma_i)$ with unknown mean and unknown covariance. Hence, the differences between $\{X_{i,n}\}$ and $\{X_{j,n}\}$ are also normally distributed. Denote $Z_{ij,n} = X_{i,n} - X_{j,n}$. Therefore, $Z_{ij,n} \sim N(\boldsymbol{\mu}_{ij} = \boldsymbol{\mu}_i - \boldsymbol{\mu}_j, \Sigma_{ij} = \Sigma_i + \Sigma_j)$. Let $\bar{Z}_{ij} \triangleq \frac{1}{N}\sum_{n=1}^N Z_{ij,n}$ be the sample mean of $Z_{ij,\cdot}$ and $S_{ij} \triangleq \sum_{n=1}^N (Z_{ij,n} - \bar{Z}_{ij})(Z_{ij,n} - \bar{Z}_{ij})'$ be the sample covariance. We have following result to compute Λ_n^{ij}.

Lemma 1 (Wang and Wan 2014) Under the normality assumption, the likelihood ratio in Equation 11 is given by

$$\Lambda_n^{ij} \triangleq \frac{\mathscr{L}(\boldsymbol{\mu}_i \not\preceq \boldsymbol{\mu}_j)}{\mathscr{L}(\boldsymbol{\mu}_i \preceq \boldsymbol{\mu}_j)} = \frac{\max_{\hat{\boldsymbol{\mu}}_{ij} \not\preceq \mathbf{0}} \left| \frac{S_{ij}}{N} + (\bar{Z}_{ij} - \hat{\boldsymbol{\mu}}_{ij})(\bar{Z}_{ij} - \hat{\boldsymbol{\mu}}_{ij})' \right|^{-\frac{N}{2}}}{\max_{\boldsymbol{\mu}_{ij} \preceq -\mathbf{0}} \left| \frac{S_{ij}}{N} + (\bar{Z}_{ij} - \hat{\boldsymbol{\mu}}_{ij})(\bar{Z}_{ij} - \hat{\boldsymbol{\mu}}_{ij})' \right|^{-\frac{N}{2}}} \tag{15}$$

where $\mathbf{0}$ is vector of 0.

Using Lemma 1, Λ_n^{ij} can be easily calculated and updated using quadratic programming techniques. With these formulations, we state the procedure in Algorithm 2.

5 NUMERICAL STUDIES

In this section, we run numerical experiments to test the performance of the proposed procedure.

Algorithm 2 The Sequential Procedure with Normally Distributed Observations

Inputs:

$PCS = 1-\alpha$ ▷ Desired PCS level
$\Omega = \{x_1, ..., x_K\}$ ▷ Set of candidate systems

Procedure:

set $\tilde{S}_0 = \emptyset$ and $\tilde{S}_1 = \Omega$;
set $\alpha_i = \frac{\alpha}{K}$, $A_i = \alpha_i$, and $B_i = \frac{1}{\alpha_i}$;
set $n = n_0$;
generate n_0 samples for each system;
repeat
    generate 1 more sample for all systems in $\tilde{S}_1$;
    set $n = n+1$;
    for $i \in \tilde{S}_1$ **do**
        for $j \in \tilde{S}_1$ **do**
            calculate Λ_n^{ij} using 15; ▷ GLR statistic for the pairwise test
        end for
    end for
    for $i \in \tilde{S}_1$ **do**
        approximate Λ_n^i using 12; ▷ GLR statistic for the candidacy test
        if $\Lambda_n^i \leq A_i$ **then**
            set $\tilde{S}_1 = \tilde{S}_1 / i$ and $\tilde{S}_0 = \tilde{S}_0 \cup \{i\}$
        end if
    end for
until $\Lambda_n^i > B_i$ for all $x_i \in \tilde{S}_1$
return $\tilde{S}_1$

5.1 Numerical Study on Single-objective R&S

The first experiment uses the same setting as the experiment in Fan and Hong (2014) and Fan, Hong, and Nelson (2016). Consider a single-objective R&S problem with various number of alternatives, with $\mu_i = 1.5 - 0.5i$ and $\sigma_i^2 = 10$ for all $i = 1, ...K$. For $K = 20,\ 50,\ 100,\ 500$, we report the estimated *PCS* and average sample sizes over 1,000 independent evaluations in Table 1.

Although the proposed procedure is indifference-zone free, it has the flexibility to take the IZ parameter as an input. If a valid IZ parameter is provided, it will account for the minimum distance between two systems in pairwise test. In the numerical study, we denote the IZ parameter as δ and vary it from 0 to 0.5, the gap between the best system and the other systems.

From Table 1, we obtain the following three findings. Firstly, the average sample size increases as the δ increases. The same pattern is also found in KN++. Secondly, under the same IZ setting, the new procedure consistently outperforms both the IZ-free procedure and KN++ procedure in terms of using less samples. Thirdly, the relative difference between the average sample sizes tends to be narrower as the number of alternatives K increases. Note that the third finding is due to newly added alternatives are likely to be screened out very fast.

5.2 Numerical Study on MOR&S

The second experiment is on MOR&S. We consider bi-objective R&S problems with 20, 50, 100 candidates. True means are generated randomly and depicted in Figure 1. We consider three variance configurations: (1)

Table 1: Comparison of IZ-free procedure, KN++ procedure, and the proposed sequential procedure for single-objective R&S to select the largest mean from K normal variables with $(\mu_1, \mu_2, ...\mu_K) = (0.5, 0, -0.5, ..., 1.5 - 0.5K)$, $\sigma_i^2 = 10$, $PCS = 0.95$, and $IZ = \delta$. (In each cell, the first line is the estimated PCS, the second line is the average sample size N_{avg}.)

K		KN++ procedure				IZ-free	Sequential Procedure				
		$\delta = \frac{1}{2}$	$\delta = \frac{1}{4}$	$\delta = \frac{1}{8}$	$\delta = \frac{1}{16}$		$\delta = \frac{1}{2}$	$\delta = \frac{1}{4}$	$\delta = \frac{1}{8}$	$\delta = \frac{1}{16}$	$\delta = 0$
20	PCS	0.99	1.00	1.00	1.00	0.99	0.994	0.998	0.999	0.998	1.00
	N_{avg}	1371	3247	7045	14700	2816	871	1272	1659	1962	2396
50	PCS	1.00	1.00	1.00	1.00	0.99	0.999	0.996	0.999	1.00	0.999
	N_{avg}	2014	4454	9779	20700	3588	1352	1814	2331	2615	3120
100	PCS	0.99	1.00	1.00	1.00	1.00	1.00	0.999	1.00	1.00	1.00
	N_{avg}	2710	5506	11790	25140	4388	1997	2516	2954	3403	3947
500	PCS	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	N_{avg}	7217	10770	18770	36280	9138	6611	7276	7919	8469	9058

Independent, $\Sigma_1 = \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}$, (2) Positive-Correlated, $\Sigma_2 = \begin{bmatrix} 4 & 1 \\ 1 & 4 \end{bmatrix}$, (3) Negative-Correlated, $\Sigma_3 = \begin{bmatrix} 4 & -1 \\ -1 & 4 \end{bmatrix}$. We conduct 1,000 independent macroreplications for each scenario and report the estimated PCS and the average sample size in Table 2. The results indicate the proposed procedure is valid for various types of covariance structure.

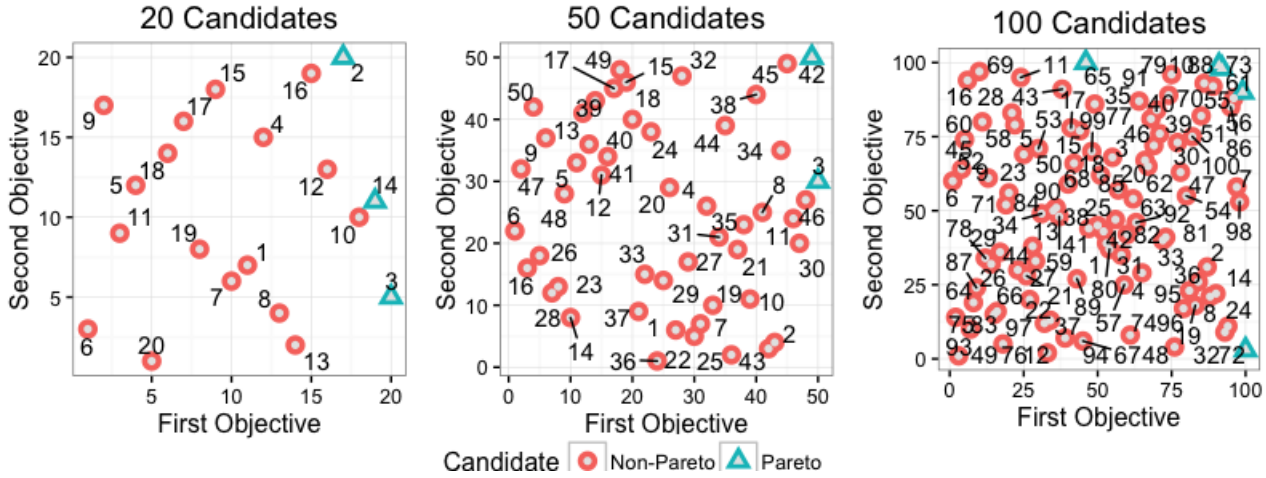

Figure 1: True means of 2-Dimensional candidates. Values along each objective are generated by randomly shuffling $[1, 2..., K]$. For $K = 20$, the Pareto front are $S_1 = \{x_2, x_3, x_{14}\}$; for $K = 50$, $S_1 = \{x_3, x_{42}\}$; and for $K = 100$, $S_1 = \{x_{61}, x_{65}, x_{72}, x_{73}, x_{88}\}$.

We continue the numerical study on MOR&S by considering R&S problems with 20 or 50 candidates and 2 to 5 objectives. For simplicity, We take the independent variance configuration $\Sigma = I_K$. The true means are randomly generated. We conduct 1,000 macroreplications and report the estimated PCS and average sample size in Table 3.

These two studies test the validity of the proposed procedure from two important aspects: (1) covariance structure, (2) number of objectives. From Table 2 and Table 3, we have two interesting findings. Firstly, independent covariance structure uses the less samples than dependent covariance structure (Note that the proposed procedure does not assume a particular covariance structure). It makes sense since dependent

Table 2: The proposed sequential procedure for Bi-Objective R&S to select the Pareto front from K normal variables with $PCS = 0.95$ and three variance structures. (In each cell, the first line is the estimated PCS, the second line is the average sample size with 95% confidence interval over 1,000 macroreplications.)

K	Sequential Procedure ($D = 2$)		
	$\Sigma_1 = \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}$	$\Sigma_2 = \begin{bmatrix} 4 & 1 \\ 1 & 4 \end{bmatrix}$	$\Sigma_3 = \begin{bmatrix} 4 & -1 \\ -1 & 4 \end{bmatrix}$
20	1.000 283.3 ± 3.6	1.000 381.5 ± 6.1	1.000 397.9 ± 5.9
50	1.000 368.7 ± 2.7	1.000 422.6 ± 3.9	1.000 423.6 ± 4.5
100	1.000 816.6 ± 4.6	0.999 998.8 ± 8.2	1.000 993.5 ± 7.9

samples provide redundancy due to correlation and thus, to achieve the same amount of information, the procedure needs less independent samples. Secondly, the sample size increases as the number of objective increases.

Table 3: The proposed sequential procedure for MOR&S to select the Pareto front from K normal variables with $PCS = 0.95$ and three variance structures. (In each cell, the first line is the estimated PCS, the second line is the average sample size with 95% confidence interval over 1,000 macroreplications.)

K	Sequential Procedure ($\Sigma = I_K$)			
	$D = 2$	$D = 3$	$D = 4$	$D = 5$
20	1.000 266.0 ± 4.0	0.994 376.8 ± 4.6	1.000 478.4 ± 6.4	0.996 494.4 ± 8.0
50	1.000 636.6 ± 6.4	0.996 1061.5 ± 12.4	0.999 1812.3 ± 23.6	1.000 1939.2 ± 30.9

6 CONCLUSION

In this paper, we propose a new sequential testing procedure for MOR&S with a guaranteed PCS. To our best knowledge, this is the first class $\mathscr{P}$ procedure for MOR&S. For future research, we identify two opening questions.

1. Bonferroni-free procedure. To account for multiple comparison, most Class $\mathscr{P}$ procedures (except BIZ procedure (Frazier 2014)) adjust the error rates for multiple alternatives. This is conservative for large-scale problems. We are experimenting with the $MSPRT$, which is a sequential testing procedure that considers more than two arms simultaneously.
2. Truncated procedure. Although the stopping time in our proposed procedure is finite with probability 1, it is still possible that the procedure takes too long to stop. We are developing a truncated procedure that guarantees to terminate at a predefined stopping time while maintains the PCS.

REFERENCES

Branke, J., and W. Zhang. 2015. "A New Myopic Sequential Sampling Algorithm For Multi-Objective Problems". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 3589–3598. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Branke, J., W. Zhang, and Y. Tao. 2016. "Multiobjective Ranking and Selection Based on Hypervolume". In *Proceedings of the 2016 Winter Simulation Conference*, edited by T. M. K. Roeder, P. I. Frazier, R. Szechtman, and E. Zhou, 859–870. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Fan, W., and L. J. Hong. 2014. "A Frequentist Selection-of-the-Best Procedure Without Indifference Zone". In *Proceedings of the 2014 Winter Simulation Conference*, edited by A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 3737–3748. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Fan, W., L. J. Hong, and B. L. Nelson. 2016. "Indifference-Zone-Free Selection of the Best". *Operations Research* 64 (6): 1499–1514.

Feldman, G., S. R. Hunter, and R. Pasupathy. 2015. "Multi-objective simulation optimization on finite sets: Optimal allocation via scalarization". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 3610–3621. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Frazier, P. I. 2014. "A Fully Sequential Elimination Procedure for Indifference-Zone Ranking and Selection with Tight Bounds on Probability of Correct Selection". *Operations Research* 62 (4): 926–942.

Kiefer, J., and L. Weiss. 1957. "Some Properties of Generalized Sequential Probability Ratio Tests". *The Annals of Mathematical Statistics* 28 (1): 57–74.

Kim, S.-H., and B. L. Nelson. 2001. "A Fully Sequential Procedure for Indifference-zone Selection in Simulation". *ACM Transactions on Modeling and Computer Simulation* 11 (3): 251–273.

Kim, S.-H., and B. L. Nelson. 2006. "On the Asymptotic Validity of Fully Sequential Selection Procedures for Steady-State Simulation". *Operations Research* 54 (3): 475–488.

Lee, L. H., E. P. Chew, and S. Teng. 2007. "Finding the Pareto Set for Multi-Objective Simulation Models by Minimization of Expected Opportunity Cost". In *Proceedings of the 2007 Winter Simulation Conference*, edited by S. G. Henderson, B. Biller, M. Hseih, J. Shortle, J. D. Tew, and R. R. Barton, 513–521. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Lee, L. H., E. P. Chew, S. Teng, and D. Goldsman. 2010. "Finding the Non-Dominated Pareto Set for Multi-Objective Simulation Models". *IIE Transactions* 42 (9): 656–674.

Li, X., J. Liu, and Z. Ying. 2014. "Generalized Sequential Probability Ratio Test for Separate Families of Hypotheses". *Sequential Analysis* 33 (4): 539–563.

Pasupathy, R., and S. Ghosh. 2013. "Simulation Optimization: A Concise Overview and Implementation Guide". In *Theory Driven by Influential Applications*, INFORMS Tutorials in Operations Research, 122–150. INFORMS.

Pavlov, I. 1988. "A Sequential Procedure for Testing Many Composite Hypotheses". *Theory of Probability & Its Applications* 32 (1): 138–142.

Pavlov, I. 1991. "Sequential Procedure of Testing Composite Hypotheses with Applications to the Kiefer-Weiss Problem". *Theory of Probability & Its Applications* 35 (2): 280–292.

Wald, A. 1945. "Sequential Method of Sampling for Deciding between Two Courses of Action". *Journal of the American Statistical Association* 40 (231): 277–306.

Wald, A., and J. Wolfowitz. 1948. "Optimum Character of the Sequential Probability Ratio Test". *The Annals of Mathematical Statistics* 19 (3): 326–339.

Wang, W., and H. Wan. 2014. "Sequential Procedures for Multiple Responses Factor Screening". In *Proceedings of the 2014 Winter Simulation Conference*, edited by A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 745–756. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

Wenyu Wang is a Ph.D. student in the School of Industrial Engineering at Purdue University. His research interests include stochastic optimization. His email address is wang1708@purdue.edu.

Hong Wan is Associate Professor in the School of Industrial Engineering at Purdue University. Her research interests include design and analysis of simulation experiments, simulation optimization, applied statistics, quality management, and healthcare engineering. She is a member of INFORMS and ASA. Her email address is hwan@purdue.edu.