Probably Approximately Correct (PAC) Selection in Simulation/Best-Arm Problems

David Eckman
Cornell University, ORIE
dje88@cornell.edu

Shane Henderson
Cornell University, ORIE
sgh9@cornell.edu

INFORMS Annual Meeting
October 22, 2017

Problem Setting

- Finite number of alternatives, i.e., arms.
- Optimize a scalar performance measure of interest.
- An alternative's performance is observed with simulation noise.

Examples:

Alternative	Performance Measure
hospital bed allocation	*expected* diversion costs
ambulance base location	*expected* call response time
MDP policy	*expected* discounted total cost

Assumptions

Alternative 1	X_{11}	X_{12}	$\cdots$	i.i.d., $\sim F_1$ with mean μ_1
Alternative 2	X_{21}	X_{22}	$\cdots$	i.i.d., $\sim F_2$ with mean μ_2
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
Alternative k	X_{k1}	X_{k2}	$\cdots$	i.i.d., $\sim F_k$ with mean μ_k

Assume $\mu_1 \leq \mu_2 \leq \cdots \leq \mu_k$, where the order is unknown.

Observations across alternatives are independent.

- Unless CRN used for variance reduction.

Marginal distributions F_i:

- R&S: Normal distribution
- MAB: Bounded support or sub-Gaussian distribution

Selection Procedures

Typical Procedure

1. Obtain observations to estimate alternatives' performances.
   - Calculate estimators $Y_1, \ldots, Y_k$ of $\mu_1, \ldots, \mu_k$.
2. Select the alternative with the best estimated performance.
   - Select alternative $K := \arg\max Y_i$.

Would like to take as few samples as possible.

Most efficient procedures use *screening* to eliminate inferior systems.

Objective

PAC Selection Guarantee

A type of *fixed-confidence* guarantee on the performance of the chosen alternative *relative* to the other alternatives.

Probably	**Approximately**	**Correct**
w.p. $1 - \alpha$	within δ of	the best

"Close enough is good enough."

- Frequentist ranking and selection (R&S) $\rightarrow$ known as PGS.
- Multi-armed bandits (MAB) in full exploration.

Proving PAC Selection Guarantees

MAB

- Concentration inequalities, e.g., Hoeffding, Chernoff.

R&S

- Multiple comparisons with the best (MCB).
- Often hard to prove directly for sequential procedures.
  - Session MB57 – “An Efficient Fully Sequential Procedure Guaranteeing Probably Approximately Correct Selection”
- A more common guarantee deals with *correct selection.*

Indifference-Zone Formulation

Bechhofer (1954) developed the idea of an **indifference zone** (IZ).

IZ parameter $\delta > 0$ is often described as the *smallest difference in performance worth detecting.*

- **Preference Zone**: PZ$(\delta) = \{\mu : \mu_k - \mu_{k-1} \geq \delta\}$
  "The best alternative is at least δ better than all the others."
- **Indifference Zone**: IZ$(\delta) = \{\mu : \mu_k - \mu_{k-1} < \delta\}$
  "There are close competitors to the best alternative."

Space of Configurations

E.g., for $F_i := \mathcal{N}(\mu_i, \sigma_i^2)$:

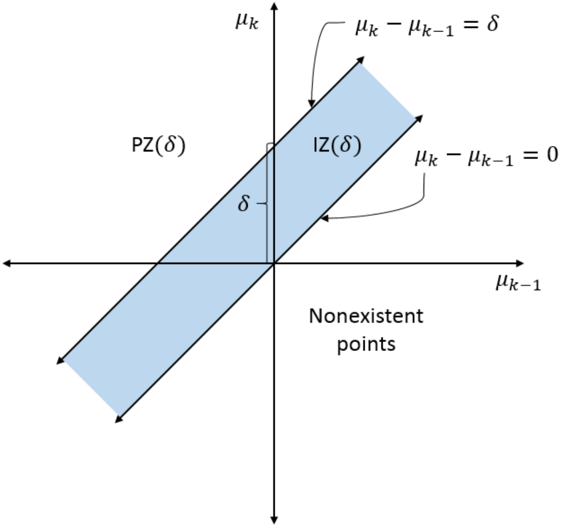

Goals of Selection Procedures

Two Frequentist Guarantees

Let K be the index of the chosen alternative. For specified confidence level $1-\alpha \in (1/k, 1)$ and $\delta > 0$, guarantee

$$\mathbb{P}_\mu(\mu_K > \mu_k - \delta) \geq 1-\alpha \quad \text{for all } \mu, \qquad \text{(Goal PACS)}$$

$$\mathbb{P}_\mu(\mu_K = \mu_k) \geq 1-\alpha \quad \text{for all } \mu \in \text{PZ}(\delta). \qquad \text{(Goal PCS-PZ)}$$

Goal PACS $\implies$ Goal PCS-PZ.

Goal PCS-PZ is the standard in the frequentist R&S community, but doesn't appear in the MAB literature.

Goal PCS-PZ vs Goal PACS

"Goal PCS-PZ is weaker, but is that so bad?"

Issues with Goal PCS-PZ

- Says nothing about performance in IZ(δ).
- Configurations in PZ(δ) may be unlikely in practice.
  - Large number of alternatives.
  - Alternatives found from search.
- Choice of δ restricts the problem.
- May require Bayesian belief about μ.

Goal PACS has none of these issues!

Equivalence of Goals

When does Goal PCS-PZ $\Longrightarrow$ Goal PACS?

Intuition: More good alternatives, more likely to pick a good alternative.

Scattered results dating back to Fabian (1962), though none in the past 20 years.

Reasons for studying this:

- Show that R&S procedures meet Goal PACS.
- Determine how MAB procedures might be designed for Goal PCS-PZ, as a means to achieve Goal PACS.

Main Equivalence Results: Condition 1

Condition 1 (Guiard 1996)

For all subsets $A \subset \{1, \ldots, k\}$, the joint distribution of the *estimators* Y_i for $i \in A$ does not depend on μ_j for $j \notin A$.

“Changing the mean of an alternative doesn’t change the distribution of other alternatives’ estimators.”

Limitation: Can only be applied to procedures without screening.

- Normal (i.i.d.): Bechhofer (1954), Dudewicz and Dalal (1975), Rinott (1978)
- Normal (CRN): Clark and Yang (1986), Nelson and Matejcik (1995)
- Bernoulli: Sobel and Huyett (1957)
- Support $[a, b]$: Naive Algorithm of Even-Dar et al. (2006)

Main Equivalence Results: Condition 2

Condition 2 (Hayter 1994)

For all alternatives $i = 1, \ldots, k$,

$$\mathbb{P}_{\mu}(\text{Select alternative } i)$$

is non-increasing in μ_j for every $j \neq i$.

"Improving an alternative doesn't help any other alternative get selected."

Limitation: Deriving an expression for $\mathbb{P}_{\mu}(\text{Select alternative } i)$ is hard.

Main Equivalence Results: Condition 2

Procedure not satisfying Condition 2

1. Take n_0 samples of each alternative.
2. Eliminate all but the two alternatives with the highest means.
3. Take n_1 additional samples for the two surviving alternatives.
4. Select the surviving alternative with the highest overall mean.

Consider the three-alternative case: $\mu_1 < \mu_2 < \mu_3$.

- Track $\mathbb{P}_\mu(\text{Select alternative } 2)$ as μ_1 increases up to μ_2.
- Consider $n_1 = 0$ and $n_1 = \infty$ as extreme cases.

Main Equivalence Results: Condition 3

Condition 3

For all alternatives $i = 1, \ldots, k$,

$$\mathbb{P}_\mu(\text{Select alternative } j, \text{ for some } j < i)$$

is non-increasing in μ_i.

"Improving an alternative doesn't help inferior alternatives get selected."

Condition 2 $\Rightarrow$ Condition 3.

Conclusions

Main take-aways

- Goal PACS is superior to Goal PCS-PZ.
- Goal PACS can follow immediately from Goal PCS-PZ.
- Condition 3 has the potential to hold for many procedures, *if only it could be verified.*

Do modern sequential selection procedures achieve Goal PACS?

- KN of Kim and Nelson (2001)
- BIZ of Frazier (2014)

Can MAB procedures be designed for Goal PCS-PZ while also satisfying one of these conditions?

Questions

Questions

Acknowledgments

This material is based upon work supported by the National Science Foundation under grants DGE–1144153 and CMMI–1537394. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- Robert E. Bechhofer. 1954. A single-sample multiple decision procedure for ranking means of normal populations with known variances. *The Annals of Mathematical Statistics* 25, 1 (1954), 16–39.
- Vaclav Fabian. 1962. On multiple decision methods for ranking population means. *The Annals of Mathematical Statistics* 33, 1 (1962), 248–254.
- V. Guiard. 1996. Different definitions of Δ-correct selection for the indifference zone formulation. *Journal of Statistical Planning and Inference* 54, 2 (1996), 175–199.
- Anthony J. Hayter. 1994. On the selection probabilities of two-stage decision procedures. *Journal of Statistical Planning and Inference* 38, 2 (1994), 223–236.
- Edward J. Dudewicz and Siddhartha R. Dalal. 1975. Allocation of observations in ranking and selection with unequal variances. *The Indian Journal of Statistics* 37, 1 (1975), 28–78.
- Yosef Rinott. 1978. On two-stage selection procedures and related probability inequalities. *Communications in Statistics – Theory and Methods* 7, 8 (1978), 799–811.

References

- Milton Sobel and Marilyn J. Huyett. 1957. Selecting the best one of several binomial populations. *Bell Labs Technical Journal* 36, 2 (1957), 537–576.
- Seong-Hee Kim and Barry L. Nelson. 2001. A fully sequential procedure for indifference-zone selection in simulation. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 11, 3 (2001), 251–273.
- Peter Frazier. 2014. A fully sequential elimination procedure for indifference-zone ranking and selection with tight bounds on probability of correct selection. *Operations Research* 62, 4 (2014), 926–942.