Asymptotically Optimal Adaptive A/B Tests for Average Treatment Effect

Achal Bassamboo
Northwestern University

Abstract:

Typically, in A/B testing, an experiment designer sequentially assigns treatment A or B to arriving individuals to identify the better treatment (known as the best treatment identification (BTI) problem), that is, the one with better mean performance, with the probability of error restricted to a pre-specified δ. We focus on a related, equally important, but more informative problem of estimating the difference between the two means or the average treatment effect (ATE). For computational efficiency, we restrict accuracy to a confidence interval (CI) of width at most ϵ, where again the probability of CI not containing ATE is restricted to at most δ. The objective of the experiment designer is to estimate a CI of ATE with the minimum sample size, i.e., minimize the total number of individuals getting treatment A or B. We first establish a lower bound on the expected sample size of the A/B test (experiment) needed for any adaptive experimental policy, which constructs a CI of ATE with desired properties as the solution to a max-min optimization problem for small δ. Using the insights provided by the max-min optimization problem, we construct an adaptive policy that is asymptotically optimal, i.e., matches the lower bound on the expected sample size for small δ. To reduce the computational burden of our policy, we propose another adaptive policy that is asymptotically optimal for small ϵ and δ. We find that, for small ϵ and δ, the asymptotically optimal fraction of treatment assignment for A and B is inversely proportional to the square root of Fisher's information of the outcome distributions of treatments A and B, respectively. Further, we compare the popular randomized controlled policy with any asymptotically optimal adaptive policies and show that there are meaningful gains from any asymptotically optimal adaptive policy in terms of the length of the experiment. Finally, we present a comparative analysis between our ATE problem and the BTI problem revealing marked differences in the asymptotically optimal assignment of treatments in both ATE and BTI problems. (Joint work with Vikas Deep and Sandeep Juneja.).

Bio: Achal Bassamboo is the Charles E. Morrison Professor at Kellogg School of Management, Northwestern University. He is also the co-director of the MMM program, which is a dual degree program between Kellogg and Segal Design at McCormick School. Professor Bassamboo joined the Kellogg School of Management faculty in 2005 after completing his Ph.D. in Operations, Information, and Technology at the Stanford Graduate School of Business. His research interests lie in the areas of service systems, statistical learning, revenue management, and information sharing. His current research focuses on service systems with dependency and statistical learning in experimentation.