Recent Progress on Bayesian Decision-Theoretic Clinical Trial Designs

Peter Jacko
Lancaster University, UK

Abstract

Statistical testing based on randomized equal allocation is a widespread state-of-the-art approach in the design of experiments today, known as *randomized controlled trial* in biostatistics, *between-group design* in social sciences, and *A/B testing* in Internet marketing. Already in the 1930s, William R. Thompson, a biostatistician from Yale University, considering the case of Bernoulli outcomes, proposed an outcome-adaptive algorithm which would in expectation yield higher outcomes than (alternating) equal allocation, using the following words "...there can be no objection to the use of data, however meagre, as a guide to action required before more can be collected ... Indeed, the fact that such objection can never be eliminated entirely-no matter how great the number of observations-suggested the possible value of seeking other modes of operation than that of taking a large number of observations before analysis or any attempt to direct our course... This would be important in cases where either the rate of accumulation of data is slow or the individuals treated are valuable, or both." This was done in a Bayesian framework, using the Beta distribution for maintaining the belief about the success probability, and it has become known as the first formulation of the *bandit problem* in academic literature; the two-armed case in 1933, extended to the multi-armed case in 1935 by the same author. Formulation of the problem using the Bayesian decision-theoretic framework allows for Bayes-optimality, i.e., maximising the benefit for both the in-trial and after-trial patients. The practical application of this Bayesian approach has however been long hindered by its computational complexity. A variety of approximations (including Thompson's Bayesian posterior sampling and non-Bayesian algorithms based on a stay-with-the-winner property or on confidence bounds) have been developed and studied across several disciplines in order to overcome this issue, but failed to come sufficiently close to Bayes-optimality for finite horizon problems, which are the problems relevant to clinical trials, as there is a finite number of in-trial patients. Recently a few novel bandit-based designs have been developed and proposed especially for rare and/or life-threatening diseases, and are being implemented in a growing number of trials, mainly in cancers, where patients are stratified into smaller groups based on a number of biomarkers. We will show that, computationally, much larger problems can be solved to Bayes-optimality, i.e., much larger trials can be designed, than what is commonly believed. In particular, using an efficient code in Julia programming language run on a standard laptop, a completely enumerated solution of the two-armed problem with a horizon of around $1,000$ periods can be computed in a few minutes and stored for an offline use, while even longer horizons can be solved by calculating optimal allocations in an online fashion.