A Bayesian Adaptive Design for Clinical Trials in Rare Diseases

S. Faye Williamson *Department of Mathematics and Statistics, Lancaster University, UK*
Peter Jacko *Department of Management Science, Lancaster University, UK*
Sofía S. Villar *MRC Biostatistics Unit, Cambridge, UK*
Thomas Jaki *Department of Mathematics and Statistics, Lancaster University, UK*

1 Introduction

The main goal of the current gold standard design for clinical trials, the *fixed randomised*, is to learn about treatment effectiveness with a view to treat future patients outside of the trial. Its drawbacks for trials involving rare diseases motivate the use of *response-adaptive designs* in which the accruing data on patient responses are used to skew the allocation towards the superior treatments, with an alternative goal of treating the patients within the trial as effectively as possible. The problem of designing a trial which aims to identify the superior treatment (exploration/learning) whilst treating the trial participants as effectively as possible (exploitation/earning) is a natural application area for *bandit models*, which seek to balance this trade-off in order to obtain an optimal allocation policy which maximises the expected number of patient successes during the given time horizon. We use a bandit model set in the framework of finite-horizon Markov decision processes, where dynamic programming (DP) can be used to develop a Bayesian response-adaptive design. Although the use of bandit models to optimally design a trial is often referred to as the primary motivation for their study, they have never been implemented in real clinical practice for reasons including lack of randomisation, low power, and biased treatment effect estimates [1].
We propose a novel bandit-based design which addresses these key issues in a very appealing way. We incorporate randomisation and add a constraint which penalises if a minimum number of patients are not recruited to each treatment arm. Simulation results for the proposed design show that: (i) the percentage of patients allocated to the superior arm is much higher than in the traditional fixed randomised design; (ii) relative to the optimal (non-rndomised and non-constrained) DP design, the power is largely improved upon and (iii) it exhibits only a very small bias and mean squared error of the treatment effect estimator.

2 The method

We consider a two-armed trial with binary endpoints, immediate responses and a finite number of patients. Patients enter sequentially over time, one-by-one, and each patient is allocated to either treatment A or B. Let X and Y denote the patient's response (either a success 1 or failure 0) from treatments A and B respectively, which we model as independent Bernoulli random variables,

$$X \sim \text{Bernoulli}(1,\ \theta_A) \ \text{ and } \ Y \sim \text{Bernoulli}(1,\ \theta_B), \ \text{ for } \ 0 \leq \theta_A,\ \theta_B \leq 1,$$

where θ_A (θ_B) is the unknown success probability of treatment A (B).
In Section 2.3 of [2] we develop the optimal design using *Constrained Randomised Dynamic Programming* (CRDP). We force actions to be randomised by assigning a probability so that each treatment has a probability of at least $1 - p$ of being allocated, where $0.5 \leq p \leq 1$, and will be referred to as the *degree of randomisation*. Note that $p = 0.5$ corresponds to fixed equal randomisation design. We further add a constraint to ensure that we always obtain at least ℓ observations from each treatment arm, where ℓ is a fixed predefined value and will be referred to as the *degree of constraining*. For details of how this design was implemented in `R`, refer to the online supplementary material of [2].

3 Overall Performance

Through extensive simulation studies we compare our proposed CRDP design with other designs, including Fixed (equal randomisation), RPW (randomised play-the-winner), DP (non-randomised and non-constrained), WI (the Whittle index approximation of DP) and RDP (randomised but non-constrained). Figure 1 summarises the key features of each design showing that our proposed CRDP design performs well with respect to all of the performance measures.

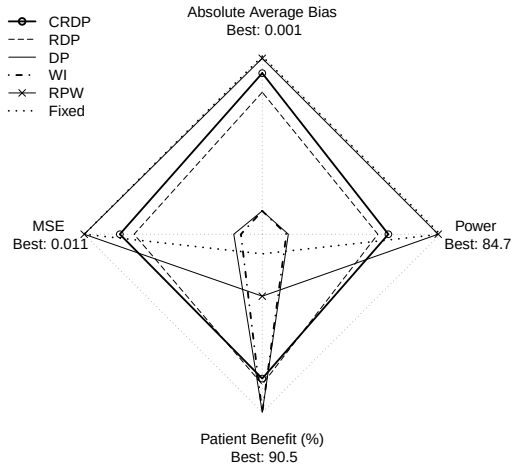

Figure 1: Star plot showing the performance of each design with respect to power, patient benefit, absolute average bias of the treatment effect estimator and MSE in a trial with 75 patients when $\theta_A = 0.5$ and $\theta_B = 0.2$. The best achieved values for each performance measure are depicted at the outer edge.

References

[1] Villar, S. S., Bowden, J., Wason, J. (2015). Multi-armed bandit models for the optimal design of clinical trials: Benefits and challenges, *Statistical Science* 30 (2), pp. 199–215.

[2] Williamson, S. F., Jacko, P., Villar, S. S., Jaki, T. (2016). A Bayesian adaptive design for clinical trials in rare diseases, *Computational Statistics and Data Analysis*, DOI: 10.1016/j.csda.2016.09.006.