Computational Statistics and Data Analysis 113 (2017) 136–153

Contents lists available at ScienceDirect

Computational Statistics and Data Analysis

journal homepage: www.elsevier.com/locate/csda

A Bayesian adaptive design for clinical trials in rare diseases

S. Faye Williamson [a,*], Peter Jacko [b], Sofía S. Villar [c], Thomas Jaki [a]

[a] Department of Mathematics and Statistics, Lancaster University, UK
[b] Department of Management Science, Lancaster University, UK
[c] MRC Biostatistics Unit, Cambridge, UK

ARTICLE INFO

Article history:
Received 30 January 2016
Received in revised form 9 September 2016
Accepted 10 September 2016
Available online 28 September 2016

Keywords:
Clinical trials
Rare diseases
Bayesian adaptive designs
Sequential allocation
Bandit models
Dynamic programming

ABSTRACT

Development of treatments for rare diseases is challenging due to the limited number of patients available for participation. Learning about treatment effectiveness with a view to treat patients in the larger outside population, as in the traditional fixed randomised design, may not be a plausible goal. An alternative goal is to treat the patients within the trial as effectively as possible. Using the framework of finite-horizon Markov decision processes and dynamic programming (DP), a novel randomised response-adaptive design is proposed which maximises the total number of patient successes in the trial and penalises if a minimum number of patients are not recruited to each treatment arm. Several performance measures of the proposed design are evaluated and compared to alternative designs through extensive simulation studies using a recently published trial as motivation. For simplicity, a two-armed trial with binary endpoints and immediate responses is considered. 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 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. Furthermore, this design is fully randomised which is an advantage from a practical point of view because it protects the trial against various sources of bias. As such, the proposed design addresses some of the key issues that have been suggested as preventing so-called bandit models from being implemented in clinical practice.

1. Introduction

Before any new medical treatment is made available to the public, clinical trials must be undertaken to ensure that the treatment is safe and efficacious. Development of treatments for rare diseases is particularly challenging due to the limited number of patients available for experimentation.

The current gold standard design is the randomised controlled trial, in which patients are randomised to either the experimental or control treatment in a pre-fixed proportion. Its main goal is to learn about treatment effectiveness with a view to prioritising future patients outside of the trial. Although this design can detect a significant treatment difference with a high probability, i.e. it maximises the statistical power, which is of benefit to future patients, it lacks the flexibility to incorporate other desirable criteria, such as the trial participant's well-being. As such, a large number of patients within the trial receive the inferior treatment. This is particularly concerning for rare disease trials in which a substantial proportion

* Corresponding author.
E-mail address: s.f.williamson@lancaster.ac.uk (S.F. Williamson).

http://dx.doi.org/10.1016/j.csda.2016.09.006

of all patients with the disease may be included in the trial. Moreover, there will be fewer patients available outside of the trial to benefit from the learning. Therefore, in this case, the priority should be on treating those patients within the trial as effectively as possible.

This motivates the use of response-adaptive designs for clinical trials involving rare diseases in which the accruing data on patient responses are used to skew the allocation towards the superior treatments, thus reducing patient exposure to inferior treatments. Although it does not fully eliminate the ethical problem of randomising patients to the inferior treatment, it certainly mitigates it by reducing the probability of allocation to the inferior treatment, if it exists.

Berry and Eick (1995) compare the performance of the traditional design, in which half of the participants receive treatment A and the other half receive treatment B, to four response-adaptive designs. They conclude that if the condition being treated is rare, then response-adaptive methods can perform substantially better and might be a more suitable alternative.

Despite the long history in clinical trials methodology, very few response-adaptive designs have actually occurred in practice and applications thus far have been disappointing (Rosenberger, 1999). This is largely attributable to the extracorporeal membrane oxygenation (ECMO) trial by Bartlett et al. (1985) which employed the randomised play-the-winner rule, a response-adaptive design described briefly in Section 2.

The problem of designing a clinical trial which aims to identify the superior treatment (exploration or learning) whilst treating the trial participants as effectively as possible (exploitation or earning) is a natural application area for bandit models, a type of response-adaptive design. Bandit models seek to balance the exploration versus exploitation trade-off in order to obtain an optimal allocation policy which maximises the expected number of patient successes over a finite number of patients. As such, they present an appealing alternative to the traditional approach used in clinical trials. Across the bandit literature, the use of bandit models to optimally design a clinical trial is often referred to as the primary motivation for their study (Gittins, 1979). However, to the best of our knowledge, they have never been implemented in real clinical practice for reasons including lack of randomisation and biased treatment effect estimates. Moreover, in contrast to the traditional approach taken in clinical trials, bandit models exhibit very low power since it is not possible to maximise both power and patient successes simultaneously. For a discussion of the benefits and challenges of bandit models in clinical trial practice, see Villar et al. (2015a).

In this paper, we propose a novel bandit-based design which provides a very appealing compromise between these two conflicting objectives and addresses some of the key issues that have prevented bandit models from being implemented in clinical trial practice. We modify the optimal design, which aims to maximise the expected number of patient successes, in such a way that we overcome its limitations without having a significant negative impact on the patient benefit.

The modifications involve incorporating randomisation into a currently deterministic design, which was considered by Cheng and Berry (2007), and adding a constraint which forces a minimum number of patients on each treatment. These are described in Sections 2.2 and 2.3, respectively, building on the standard dynamic programming approach presented in Section 2.1. In Section 4, we compare our design to alternative designs via extensive simulations in several scenarios in the context of a recently published Phase II clinical trial of isotonic fluid resuscitation in children with severe malnutrition and hypovolaemia (Akech et al., 2010). We evaluate each design's performance according to the measures set out in Section 3. We summarise the main conclusions in Section 5 and highlight areas for future research.

2. Methods

In this section, we introduce different methods for allocating patients to treatments in a clinical trial. For simplicity of exposition, we consider a two-armed clinical trial with a binary endpoint and a finite number of patients within the trial, n. Patients enter the trial sequentially over time, one-by-one, and each patient is allocated to either treatment A or B on arrival. We assume that n is fixed but that the sample sizes for treatment groups A and B, denoted by N_A and N_B respectively, are random, where $N_A + N_B = n$. Let X and Y denote the patient's response (either a success or failure) from treatments A and B respectively, which we model as independent Bernoulli random variables. That is,

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

where θ_A and θ_B are the unknown success probabilities of treatments A and B respectively. Further, assume that each patient's response from the allocated treatment becomes immediately available.

The *fixed randomised* design randomises patients to either treatment A or B with an equal, fixed probability, i.e. 50% in a two-armed trial. This will act as a reference to which each of the response-adaptive designs described below will be compared against.

One of the most well-known response-adaptive designs is the *randomised play-the-winner* (RPW) rule, a type of urn model, proposed by Wei and Durham (1978). This design is very intuitive and applies specifically to clinical trials comparing two treatments with binary responses. Initially, an urn contains u balls of type A and u balls of type B. When a patient is recruited, a ball is drawn randomly from the urn with replacement; if it is a type A ball, the patient receives treatment A and if it is a type B ball, the patient receives treatment B. After each patient's outcome is observed, a decision about the urn composition is made depending on the observed result. Thus, a success on treatment A, or a failure on treatment B, generates an additional β type A balls and α type B balls in the urn. Similarly, a success on treatment B, or a failure on treatment A, will generate an additional β type B balls and α type A balls in the urn, where $0 \leq \alpha \leq \beta$ are integers. In this way, the