Contemporary Clinical Trials Communications 7 (2017) 11–17

Contents lists available at ScienceDirect

Contemporary Clinical Trials Communications

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

A placebo-controlled Bayesian dose finding design based on continuous reassessment method with application to stroke research

Chunyan Cai[a,b,*], Mohammad H. Rahbar[a,b,c], Md Monir Hossain[d], Ying Yuan[e], Nicole R. Gonzales[f]

[a] Division of Clinical and Translational Sciences, Department of Internal Medicine, McGovern Medical School, The University of Texas Health Science Center at Houston, Houston, TX 77030, USA
[b] Biostatistics/Epidemiology/Research Design Component, Center for Clinical and Translational Sciences, The University of Texas Health Science Center at Houston, Houston, TX 77030, USA
[c] Division of Epidemiology, Human Genetics, and Environmental Sciences, University of Texas School of Public Health at Houston, Houston, TX 77030, USA
[d] Division of Biostatistics and Epidemiology, Cincinnati Children's Hospital Medical Center, Cincinnati, OH 45229, USA
[e] Department of Biostatistics, Division of Quantitative Sciences, The University of Texas MD Anderson, Houston, TX 77030, USA
[f] Department of Neurology, McGovern Medical School, The University of Texas Health Science Center at Houston, Houston, TX 77030, USA

ARTICLE INFO

Keywords:
Continuous reassessment method
Dose finding
Placebo-controlled study
Stroke research

ABSTRACT

Traditional dose-finding designs do not require assignment of patients to a control group. Motivated by SHRINC (Safety of Pioglitazone for hematoma resolution in intracerebral hemorrhage), we developed a placebo-controlled dose-finding study to identify the maximum tolerated dose for pioglitazone in stroke patients with spontaneous intracerebral hemorrhage. We designed an extension of the continuous reassessment method that allowed to incorporate information from the control group (i.e., the standard of care), and utilized it to determine the maximum tolerated dose in the SHRINC trial. We evaluated the operating characteristics of our design by conducting extensive simulation studies. Our findings from the simulation studies demonstrate that our proposed design is robust and performs well. By estimating the toxicity rate in the control group, we were able to obtain more accurate information about the natural history of the disease and identify appropriate dose for the next phase of this study. The proposed design provides a tool to incorporate the information from the control group into the dose-finding framework for trials with similar objectives.

1. Introduction

Clinical trials are considered as the most reliable method for evaluation of safety and efficacy of new drugs and other clinical interventions. Phase I dose-finding trials are designed to identify the maximum tolerated dose (MTD) of a new drug, defined as the highest dose within a tolerable dose-limiting toxicity (DLT). After the MTD is determined, the drug will be carried forward for subsequent assessments through phase II and III trials. Inaccurate evaluation of the MTD can lead to waste of resources if an inappropriate dose is moved to subsequent phases. Therefore, it is important to design efficient dose-finding trials to determine the most appropriate dose before the drug is tested in future phases of drug development.

Broadly, dose-finding trial designs are classified into two types, algorithm-based designs, and model-based designs. The algorithm-based designs, also known as up-and-down designs, are used often in practice due to their simplicity in implementation. The most popular algorithm-based design is the "3 + 3" design [1]. Although it is popular in practice, the reported shortcomings of this design include unreliable estimation of the MTD [2], a significant large proportion of patients treated at subtherapeutic dose levels [3], and the restricted choice of the target DLT rate [4]. Several investigators attempted to develop improved up-and-down designs to identify the MTD, including the accelerated titration design [5], the biased coin design [6] and its extension with isotonic regression [7], the k-in-a-row design [8], the up-and-down design based on isotonic regression [9], the modified toxicity probability interval design [10], and the Bayesian optimal interval design [11]. Comprehensive reviews of up-and-down designs are provided by Ivanova [12] and Liu et al. [4]. For model-based designs, the most popular one is the continual reassessment method (CRM) [13]. In contrast to 3 + 3 design, CRM design provides a more accurate estimation of toxicity probability of the MTD and a more flexible setup of target DLT rate. Due to the popularity of CRM, a variety of extensions have been proposed to improve its practical

* Corresponding author. Division of Clinical and Translational Sciences, Department of Internal Medicine, McGovern Medical School, The University of Texas Health Science Center at Houston, Houston, TX 77030, USA.
E-mail address: chunyan.cai@uth.tmc.edu (C. Cai).

http://dx.doi.org/10.1016/j.conctc.2017.05.002
Received 20 September 2016; Received in revised form 6 April 2017; Accepted 3 May 2017
Available online 06 May 2017

implementation and operating characteristics [14–24]. Other model-based designs include designs based on Bayesian decision-theoretic approach [25], the escalation with overdose control [26], Bayesian nonparametric approach [27], and stochastic approximation [28]. Comprehensive reviews for dose-finding designs can be found in Rosenberger and Haines [29], Chevret [30], Ting [31], and Le Tourneau et al. [3].

In stroke trials, DLT usually is defined as a clinical event with substantial morbidity and mortality such as a cerebral hemorrhage [32]. Therefore, the acceptable target rate of patients undergoing such events is quite small. With the development of aforementioned dose-finding designs, CRM is the most promising detection method to address this challenge. There has been some effective utilization of dose-finding designs in stroke trials, especially the application of CRM [33–35]. In addition to the utilization of CRM design, other dose-finding designs were also applied to stroke trials. For example, Krams et al. [36] used a Bayesian adaptive dose finding design based on normal dynamic linear model in an acute ischemic stroke study. Whelan et al. [37] described the utilization of the Bayesian phase I/II design proposed by Thall and Cook [2] to find an optimal dose for treatment of ischemic stroke in children.

All the aforementioned dose-finding designs do not require assignment of patients to a control group. However, for an acute intracerebral hemorrhage (ICH) trial, SHRINC (Safety of Pioglitazone for hematoma resolution in intracerebral hemorrhage), the investigators were interested in finding a dose with the target toxicity rate dependent upon the rate in the concurrent control group. That means, the target toxicity rate in the SHRINC study is unknown before the initiation of the trial and needs to be determined based on the toxicity data collected from the concurrent control group. Motivated by SHRINC, we developed a placebo-controlled dose-finding study based on CRM to identify the MTD. The inclusion of a control group allowed us to study the natural history of the disease and co-morbidities. Therefore, our proposed design can lead to more meaningful MTD identification than traditional designs. The purpose of this paper is to describe the unique features of our proposed design, and share our experience on its application to the SHRINC study. Using a stroke study as a motivating example, our proposed design provides a tool to incorporate the information from the control group into the dose-finding framework for trials in other diseases with similar objectives. We evaluate the operating characteristics of the proposed design by conducting extensive simulation studies.

2. Methods

2.1. Rational of the inclusion of a control group

SHRINC study was designed to assess the safety of PIO in spontaneous ICH compared with the standard of care. It is a prospective, randomized, blinded, placebo-controlled, and dose-escalation safety trial in which patients are randomly allocated to control or treatment groups. The primary objective of the SHRINC trial is to determine the MTD of PIO, a dose with DLT rate closest to a target rate. More detailed information including inclusion/exclusion criteria, informed consent, safety outcomes, and clinical and radiographic outcomes can be found in Gonzales et al. [38]. The SHRINC study is registered at http://www.clinicaltrials.gov/(registration number NCT00827892).

The inclusion of a control group is not a characteristic of the traditional dose finding trials. In SHRINC, we included a control group for the following reasons: first, many expected adverse events (AEs) and serious AEs (SAEs) in the PIO group were also expected in the ICH population as part of the natural history of the disease and co-morbidities. These AEs and SAEs were not captured in our prospective stroke registry [39] and are difficult to collect accurately in a retrospective manner. Therefore, we planned to collect this information prospectively and compared the rates of AE/SAEs between the PIO-treated and control groups to obtain a more accurate measure of safety with the use of PIO in our patient population. Second, in our preclinical work, more rapid hematoma resolution was correlated with improved neurologic recovery. This finding has not been demonstrated in the clinical setting. In addition, the control group data related to the rate of hematoma resolution would help to determine an optimal duration of PIO. Third, retrospective studies usually overestimate the expected benefit of treatment. Therefore, the inclusion of a concurrent control group would improve our knowledge of toxicity and efficacy profile of the study population and guide us to choose a meaningful target DLT for MTD identification. Thus, instead of specifying a fixed target DLT rate before the initiation of the trial, we consider the target rate dependent up on the rate in the concurrent control group. As a result, many existing methods, such as the CRM, cannot be directly applied. Toward this goal, we propose a modified CRM dose-finding design with a control group for our study.

2.2. Dose-finding method

In this section, we propose an extension of the CRM that allows incorporating information from the control group for dose finding. Suppose K dose levels, denoted as $d_1, d_2, \ldots, d_K$, have been chosen for the investigation with the true toxicity probability P_k for dose d_k, $k = 1, \ldots, K$. Let Y be a binary variable to denote whether a patient has experienced the prespecified DLT event, with 1 denoting an event and 0 otherwise, and let x denote the dose level for this patient. Usually, a one-parameter model $\psi(x,\alpha)$ is proposed to model the relationship between dose level x and its toxicity probability $P(Y = 1|x)$ with the unknown parameter α. There are three popular dose-toxicity models of $\psi(x,\alpha)$ proposed for CRM: logistic model, power model, and hyperbolic tangent model. In our proposed design, we utilize the following one-parameter logistic model,

$$P(Y = 1|x) = \psi(x, \alpha) = \frac{\exp(c + \alpha x)}{1 + \exp(c + \alpha x)}, \quad (1)$$

where c is a constant and recommended to be 3 [40]. As part of the design, we also need a "skeleton" for the CRM, which is the investigator's prior estimates of DLT at each dose level, denoted as P_k^0 for dose level d_k. By plugging the toxicity probability in equation (1) with each value of the "skeleton", we obtain the standardized dose $d_k' = \frac{logit(P_k^0) - c}{\hat{\alpha}}$, $k = 1, \ldots, K$, where $\hat{\alpha}$ is the mean or median of the prior estimate of the parameter and $logit(p) = \log\left(\frac{p}{1-p}\right)$.

During the course of the trial, if n patients have been enrolled into the study and assigned to different dose levels, we denote the standardized dose for ith patient, $i = 1, \ldots, n$ as x_i with his/her observed outcome y_i, where $x_i \in \{d_1', d_2', \ldots, d_K'\}$. Denoting the observed data $\{x_i, y_i, i = 1, \ldots, n\}$ as D, the likelihood function of observed data D can be written as $L(D|\alpha) \propto \prod_{i=1}^{n} \psi(x_i, \alpha)^{y_i}(1 - \psi(x_i, \alpha))^{1-y_i}$. Denoting the prior density of α as $\pi(\alpha)$, the posterior density of α can be written as $f(\alpha|D) \propto L(D|\alpha)\pi(\alpha)$. The prior distribution for α is set to unit exponential distribution. The other choice can be uniform distribution with $0 < \alpha < 3$, and the log-normal distribution. The Gibbs sampler [41] can be used to obtain the posterior samples for α. After that, the posterior mean for P_i at dose d_i, $i = 1, \ldots, K$, can be calculated as $\hat{P}_i = \int \psi(d_i', \alpha) f(\alpha|D) d\alpha$.

In the traditional CRM, if the target DLT rate is set at a fixed value of θ, then the MTD d^* is taken to be one of the specified dose in the set of $d_1, d_2, \ldots, d_k$ which satisfies the following criteria $d^* = argmin|\hat{P}_i - \theta|$, $i \in 1, \ldots, k$. Without losing the generality of assumption, we consider a design to identify the MTD with the DLT rate closest to a target rate, defined as a rate higher than the rate in the control group by a constant magnitude of δ. We denote the toxicity rate in the control group as P_0. To reduce the variability of the toxicity rate in the control group among different cohorts, we estimate P_0 during the