Contemporary Clinical Trials Communications 8 (2017) 227–233

Contents lists available at ScienceDirect

Contemporary Clinical Trials Communications

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

Bayesian adaptive clinical trials of combination treatments

Kristian Thorlund[a,*], Shirin Golchi[b], Edward Mills[a]

[a] Department of Health Research Methods, Evidence, and Impact, McMaster University, Hamilton, ON, Canada
[b] Department of Statistics and Actuarial Science, Simon Fraser University, Vancouver, BC, Canada

ARTICLE INFO

Keywords:
Adaptive trials
Additivity
Bayesian
Combination treatment

ABSTRACT

Randomized clinical trials (RCT) increasingly investigate combination therapies. Strong biological rationale or early clinical evidence commonly suggest that the effect of the combination treatment is importantly greater than the maximum effect of any of the individual treatments. While these relationships are commonly well-accepted, RCTs do not incorporate them into the design or analysis plans. We therefore propose a simple Bayesian framework for incorporating the known relationships that the effectiveness of a combination treatment exceeds that of any individual treatment, but does not necessarily exceed the sum of individual effects. We term the collation of these two relationships 'fractional additivity'. We performed a binary outcome simulation study of a response adaptive randomized three-arm clinical trial with treatment arms A, B, and A&B that allowed for dropping an inferior treatment arm and terminating the trial early for superiority during any of 4 interim analyses. We compared the Bayesian fractional additivity model to a conventional analysis by measuring the expected proportion of failures, sample size at trial termination, time to termination, and root mean squared error of final estimates. We also compared the fractional additivity model to a 'full additivity' model where the effect of A&B was assumed to be the sum of the effect of A and B. In simulation scenarios where important fractional additivity or full additivity existed, the Bayesian fractional additivity model yielded a 3–4% relative reduction in expected number of failures, and a 30%–50% relative reduction in sample size at trial termination compared to a conventional analysis. These results held true even when the Bayesian fractional additivity model employed a biased prior. The full additivity model had slightly higher gains, but too frequently terminated the trial at the first interim look. In scenarios where no or weak fractional additivity exists, the expected sample size and time to termination were similar for the Bayesian fractional additivity model with a moderately optimistic bias about fractional additivity and the conventional model. Lastly, the fractional additivity model generally yielded similar or lower root mean squared error compared to the other models. In conclusion, our proposed Bayesian *fractional additivity* model provides an efficient approach for estimating effects of combination treatments in clinical trials. The approach is not only highly applicable in adaptive clinical trials, but also provides added power in a conventional RCT.

1. Background

Several clinical trials investigate combinations of interventions that have already been demonstrated to be individually effective. Historically, the superiority of combination therapies (vs single agent therapies) have been demonstrated medical areas such as in cardio-vascular diseases (e.g., the poly-pill) and respiratory diseases [3,4]. Recently, superiority of combination therapies have been definitely demonstrated in phase III randomized clinical trials (RCT) in areas such as immuno-oncology and type II diabetes (see example Box 1 for detailed description) [1,2,5,6]. In these combination therapy RCTs, the effectiveness of the individual interventions is typically well known, and there are typically substantial biological rationale, early clinical evidence, or evidence from related disease areas to suggest that the combination of interventions will work markedly better than any of the interventions alone [3,4,7]. However, RCTs of combination treatments commonly analyse a combination therapy arm as if it is a separate individual intervention. For example, this is generally the case in 2 × 2 factorial trials across all areas of medicine. Thus, no advantage is taken of prior knowledge and assumptions about the combination therapy in the conduct and analysis of the clinical trial.

While true additivity (i.e. the property that the effect of the combination of two intervention equals the sum of the two individual treatment effects) is rare, in many cases it is plausible to assume that treatment combinations investigated in clinical trials will exhibit markedly better effects than each of the individual treatments alone. In

* Corresponding author.
E-mail address: thorluk@mcmaster.ca (K. Thorlund).

https://doi.org/10.1016/j.conctc.2017.11.001
Received 16 June 2017; Received in revised form 17 October 2017; Accepted 1 November 2017
Available online 03 November 2017

other words, it is plausible to assume clinically meaningful *fractionally additive* effect of the combination treatment investigated (also previously referred to as *antagonistic additivity* [3,7]). Thus, structurally incorporating *fractional additivity* into the statistical analysis of a clinical trial should theoretically suffice to optimize the trial design in terms of mitigating sample size requirements and trial duration. Yet to our knowledge, no clinical trial has previously capitalized on prior knowledge about additivity or fractional additivity.

We therefore propose a Bayesian framework for incorporating *fractional additivity* in the statistical model and analysis of clinical trial of combination treatments. The proposed model expresses the effect of a combination treatment, A&B, as the maximum of A and B, plus a *fractional additivity parameter* times the minimum of A and B. A weakly informative prior distribution is assigned to the *fractional additivity* parameter to reflect plausible ranges of *fractional additivity*, yet does not rule out equipoise nor full additivity (also previously referred to as *synergistic additivity*). Due to the hypothesized efficiency gain we apply the proposed Bayesian *fractional additivity* model within an adaptive clinical trial setting. We test the performance of the proposed model against a conventional approach using simulations.

2. Methods

We propose a Bayesian *fractional additivity* modelling framework to optimize estimation of additive effects in clinical trials. Under the conjecture that the proposed model adds considerable efficiency compared to the conventional framework, we conduct a simulation study to assess its performance in an adaptive clinical trial setting. In addition, we illustrate the evolution of posterior probabilities informing trial adaptation in 3 simulated clinical trials.

For simplicity, we only consider a binary outcome clinical trial setting in this paper. However, the proposed model can easily be generalized to other types of outcomes (e.g., continuous or time-to-event data).

2.1. The Bayesian fractional additivity model

Under the proposed *fractional additivity* model, we make two seminal assumptions:

1) The effect of A&B is likely larger than the maximum of A and B;
2) The effect of A&B is likely smaller than the sum of the effects of A and B.

Letting θ_A, and θ_B denote the log odds of the treatment responses for A and B, respectively, and letting $\theta_{A\&B}$ denote the log odds treatment response of A&B, we can express $\theta_{A\&B}$ as a function of θ_A, and θ_B as follows:

$$\theta_{A\&B} = \max(\theta_A, \theta_B) + f\min(\theta_A, \theta_B), \quad (1)$$

where f is likely a number between 0 and 1 that denotes the fraction of additivity that the combination treatment exhibits (note that '$\theta_{A\&B} = \theta_A + \theta_B$' is what is conventionally referred to as 'full additivity').

This model is easily set up in a Bayesian framework that places non-informative priors on the effect sizes (i.e., the log odds) of individual treatment effects, θ_A and θ_B, and a weakly-informative prior on the fractional additivity parameter f. The model is fit using RStan version 2.14.1 (Stan is a probabilistic programming language that implements Hamiltonian Monte Carlo and RStan is an R interface to Stan) [8]. The Stan implementation of the model is provided in the supplementary material.

2.1.1. Prior choice for fractional additivity parameter

The parameter f represents the fractional additivity. Under the assumptions for model (1), f should lie between 0 and 1, and so a first natural choice would be a beta distribution. However, strictly constraining f to the (0,1) interval, by the choice of prior, implicitly violates the assumption of equipoise in RCTs. Thus, hard constraints should be avoided to allow deviations from the fractional additivity assumption. We instead propose to use a normal distribution as a prior for f. Expert belief can be used to determine the prior mean, while a variance of 0.16 is supposed to introduce sufficient uncertainty under model (1) to allow the accumulating data to shape the inference, while still being sufficiently informative to stabilize and strengthen estimation (see Table 1 for 95% confidence intervals for group responses under this choice of variance). In practice, f is not known, but good biological rationale or early clinical evidence is typically available to inform f's distribution. In this paper, we specifically test scenarios where the mean prior distribution for f is unbiased (i.e., the truth in the simulation), and where f is biased positively or negatively by a 25% (see section 4.2 for further details).

Table 1
Overview of simulation scenarios.

Simulation Parameter	Fixed values by scenario
Response probabilities for Tx A and Tx B	1 Pr(response with A) = 35%, Pr (response with B) = 40% 2 Pr(response with A) = 40%, Pr(response with B) = 40%
Fractional additivity	1 f = 0.50 2 f = 0.75 3 f = 1.00
Prognostic factor variability	σ^2 = 0.16 corresponding to 95%CI of group response of: 1 6.8%–80% when Pr(response with A) = 35% 2 8.4%–83% when Pr(response with A) = 40%

2.2. Adaptive design

Due to the anticipated efficiencies of the proposed fractional additivity model as well as the Bayesian nature of the model, we propose applying the model within an adaptive trial design framework. For completeness, however, we confirmed the superior power of the model in a conventional parallel design framework (see Fig. S.1 in supplementary material). We consider a three-arm response adaptive randomized (RAR) clinical trial design that allows for 1) continually adapting the allocation ratios by the updated probabilities of superiority for any treatment; 2) dropping of an inferior treatment arm; and 3) early stopping for superiority. At the beginning of the trial patients are assigned to each of the three intervention arms with equal probabilities (1:1:1). Adaptations can in principle be applied anytime new outcome data becomes available. However, for simplicity and computational feasibility we consider 4 interim analyses at which adaptations can be made. The four interim looks are spaced equally between the first patient enrollment and reaching a fixed parallel design sample size requirement between A&B and the maximum of A and B (e.g., 80% power and 5% type I error to detect a 20% difference). Thus, the first interim analysis occurs when outcome data on 20% of this required sample size has been accrued, the second at 40%, and so forth. Trial adaptations are based on the interim posterior probabilities that A, B, and A&B, respectively are better than the two other interventions. Let $p_{A\ best}$, $p_{B\ best}$, and $P_{A\&B\ best}$ denote these three probabilities. At each interim look, the allocation proportions are updated to the ratio between the square roots of these three probabilities (i.e., $\sqrt{p_{A\ best}}$: $\sqrt{p_{B\ best}}$: $\sqrt{P_{A\&B\ best}}$). The use of square roots rather than crude probabilities avoids too rapid adaptation and has become common place in adaptive trials [9]. We also allow for dropping an inferior treatment arm if the square root probability of superiority falls below 0.01, as well as early termination of the trial for