Adaptive designs with response-adaptive randomization

MedianaDesigner package

1. Introduction

This document provides a description of the statistical methodology used in the adaptive design module that supports response-adaptive randomization (ADRand function).

For more information on the MedianaDesigner package, visit the following web pages at

http://www.mediana.us/medianadesigner

http://medianasoft.github.io/MedianaDesigner

2. Adaptive designs with response-adaptive randomization

2.1. Trial design

Response-adaptive randomization is an attractive option for exploratory clinical trials, including dose-finding Phase II trials. Adaptive randomization designs enable trial sponsors update the randomization scheme in a data-driven way. Key advantages of this approach include

- Ability to treat the patients in the trial as effectively as possible by assigning most patients to the most effective doses or regimens.
- Improved statistical inferences/precision for the most effective doses or regimens.

To introduce adaptive randomization designs, consider a Phase II trial that will be conducted to investigate the efficacy and safety of multiple doses of an experimental treatment compared to placebo. The primary efficacy endpoint in this trial is assumed to be a continuous endpoint that follows a normal distribution (the current version of the ADRand function only supports continuous endpoints). A multi-stage design will be employed in the trial. In the first stage, patients will be randomly assigned to receive placebo or a dose of the experimental treatment using an equal randomization approach. An unblinded interim analysis will be conducted at the end of this stage, as well as the other stages in the trial, and the available efficacy data will be examined to model the dose-response function and evaluate the efficacy profile of each dose compared to placebo. The information on the strength of the efficacy signal in each dosing arm will be used to update the randomization probabilities in the next stage with the ultimate goal of treating most patients at the most promising doses. The dose-response data will be evaluated at

the final analysis to assess the statistical significance of the dose-response relationship in the trial.

It needs to be pointed out that the current (experimental) version of the ADRand function implements one approach to designing dose-finding trials with response-adaptive randomization. Multiple alternative frequentist and Bayesian designs that support response-adaptive features have been proposed in the literature. For a detailed overview of the theory of response-adaptive designs, the reader is referred to Hu and Rosenberger (2008). Examples of multi-arm clinical trials that employed adaptive randomization designs can be found in Grieve and Krams (2005) and Lee, Gu and Liu (2010).

The adaptive design methodology is introduced in Section 2.2 and a case study to illustrate adaptive randomization designs in Phase II trials is presented in Section 3.

2.2. Adaptive design methodology

Consider a dose-finding Phase II trial with an adaptive randomization design and a normally distributed primary endpoint. The data at the end of each trial stage will be used to examine the dose-response relationship in the trial and modify the randomization probabilities for the individual trial arms based on the accumulating efficacy data.

A multi-stage design with K stages will be employed in the trial. Equal randomization will be used in the first stage and an interim analysis will be performed at the end of the kth stage, $k = 1, \dots, K-1$, to support a decision to update the randomization scheme in the next stage.

The following algorithm will be applied at the end of the kth stage (kth interim analysis). Let $d_0, \dots, d_m$ denote the dose levels in the trial with $d_0 = 0$ corresponding to the placebo arm. The dose-response data will be modeled using the MCPMod (Multiple Comparisons and Modeling) method originally introduced in Bretz et al. (2005) based on a set of pre-defined dose-response models. These dose-response models will be denoted by

$$f_l(d_i, \theta_l),\ i = 0, \dots, m,\ l = 1, \dots, L,$$

where θ_l is the vector of model-specific parameters and L is the number of candidate dose-response models. The common variance of the responses across the trial arms is denoted by σ^2.

The dose-response models will be fitted to the data at the kth interim analysis and the model-specific predicted mean effects will be computed across the trial arms, i.e.,

$$\hat{\mu}_{il} = f_l(d_i, \theta_l),\ i = 0, \dots, m,\ l = 1, \dots, L.$$

To simplify the notation, the dependence on the kth interim analysis is suppressed.

To obtain a single set of predictions, a model-averaging approach will be applied as follows. Let $c_1, \ldots, c_L$ denote the values of the model selection criterion (Akaike information criterion) for the selected L dose-response models. The average predicted means are given by

$$\hat{\mu}_i = \sum_{l=1}^{L} w_l \hat{\mu}_{il}, \, i = 0, \ldots, m,$$

where $w_l = \frac{\exp(-c_l/2)}{\sum_{j=1}^{L} \exp(-c_j/2)}$.

The resulting predicted means will be used to compute the posterior distributions of the true mean effects in the placebo and dosing arms ($\mu_0, \ldots, \mu_m$). To enable these calculations, non-informative conjugate priors will be assumed for $\mu_0, \ldots, \mu_m$ and σ^2, i.e., it will be assumed that σ^2 follows a scaled inverse chi-square distribution and, conditional on σ^2, each mean follows a normal distribution. Under these assumptions, closed-form expressions for the posterior distributions of $\mu_0, \ldots, \mu_m$ and σ^2 are easily derived (Gelman et al., 2013).

Based on the posterior distributions, the posterior probability of the target efficacy profile, i.e., the probability of providing a pre-specified amount improvement over placebo, will be computed for each dose. The threshold corresponding to a clinically meaningful improvement over placebo is denoted by δ. Assuming that a larger value of the endpoint indicates a beneficial effect, the posterior probability of the target efficacy profile for the ith dose is given by

$$P_i = P(\mu_i \geq \mu_0 + \delta | D_k), \, i = 1, \ldots, m,$$

where D_k denotes the dose-response data at the kth interim analysis.

Finally, the randomization scheme for the next stage will be updated as follows. The randomization probability for the placebo arm will be fixed at a pre-defined value denoted by r_0 ($0 < r_0 < 1$). The randomization probabilities for the dosing arms will be computed by normalizing the posterior probabilities of target efficacy, i.e.,

$$r_i = (1 - r_0) \frac{P_i^{\rho}}{\sum_{j=1}^{m} P_j^{\rho}}, \, i = 1, \ldots, m,$$

where ρ is the parameter controlling the degree of balance for adaptive randomization. If $\rho = 0$, a perfectly balanced patient allocation is achieved and, with larger values of this parameter, the allocation will be increasingly imbalanced.

The adaptive randomization approach will be applied to Stages 2 through K. At the final analysis, the significance of the dose-response will be evaluated using the MCPMod approach with the same set of candidate dose-response models.

3. Case study

The adaptive randomization framework described in Section 2 will be applied to a case study based on the Phase II trial for the treatment of cognitive impairment in patients with schizophrenia. The trial will be conducted to assess the efficacy of four doses of an experimental treatment (20 mg, 50 mg, 100 mg or 250 mg) compared to placebo. The primary efficacy assessment will be based on the change from baseline to 3 months in the MCCB (MATRICS Consensus Cognitive Battery) overall composite score. This continuous endpoint is normally distributed and a greater positive change from baseline indicates a more favorable effect.

This setting was used to evaluate the performance of a multi-stage design with response-adaptive randomization and compare it to that of a traditional design with equal randomization. The total sample size in the trial was set to 320 patients and a four-stage design with 80 enrolled patients per stage was used in the adaptive trial. As indicated in Section 2, a balanced randomization approach was used in the first stage of the adaptive trial and the randomization probabilities for the four dosing groups were updated at the end of the first, second and third stages. The randomization ratio for the placebo group was fixed at 20% throughout the trial.

At each interim analysis, the dose-response data at 3 months were modeled using the MCPMod method to predict the mean value of the efficacy endpoint in each trial arm. The predictions relied on a model-averaging approach based on four dose-response models (linear, Emax, exponential and logistic models). These dose-response models were specified using the following values of the non-linear parameters:

- Linear: No non-linear parameters.
- Exponential: Delta was set to 100.
- Emax: ED50 was set to 200.
- Logistic: ED50 and delta were set to 125 and 25, respectively.

The resulting dose-response relationships are depicted in Figure 1.

The posterior probability of the target efficacy profile at each dose level was computed assuming that the threshold for a clinically meaningful improvement over placebo was equal to 1.5. The randomization probabilities for the four doses were updated to be proportional to the resulting posterior probabilities. The balance parameter was set to 2, which means that a more imbalanced patient allocation was considered desirable and thus more patients would be assigned to the most promising doses of the experimental treatment.

Finally, to understand the performance of an adaptive randomization design in a real-life setting, it is important to model the patient enrollment and dropout processes. It was assumed that patients would be enrolled over the period of 24 months with the time to patient enrollment following a truncated normal distribution. The median enrollment time, i.e., the time point by

which 50% of the patients are expected to be enrolled into the trial, was set to 18 months. The dropout rate was equal to 10%, i.e., 10% of the enrolled patients were expected to be lost to follow-up by the end of the 3-month treatment period.

Figure 1. Candidate dose-response models used in the MCPMod method

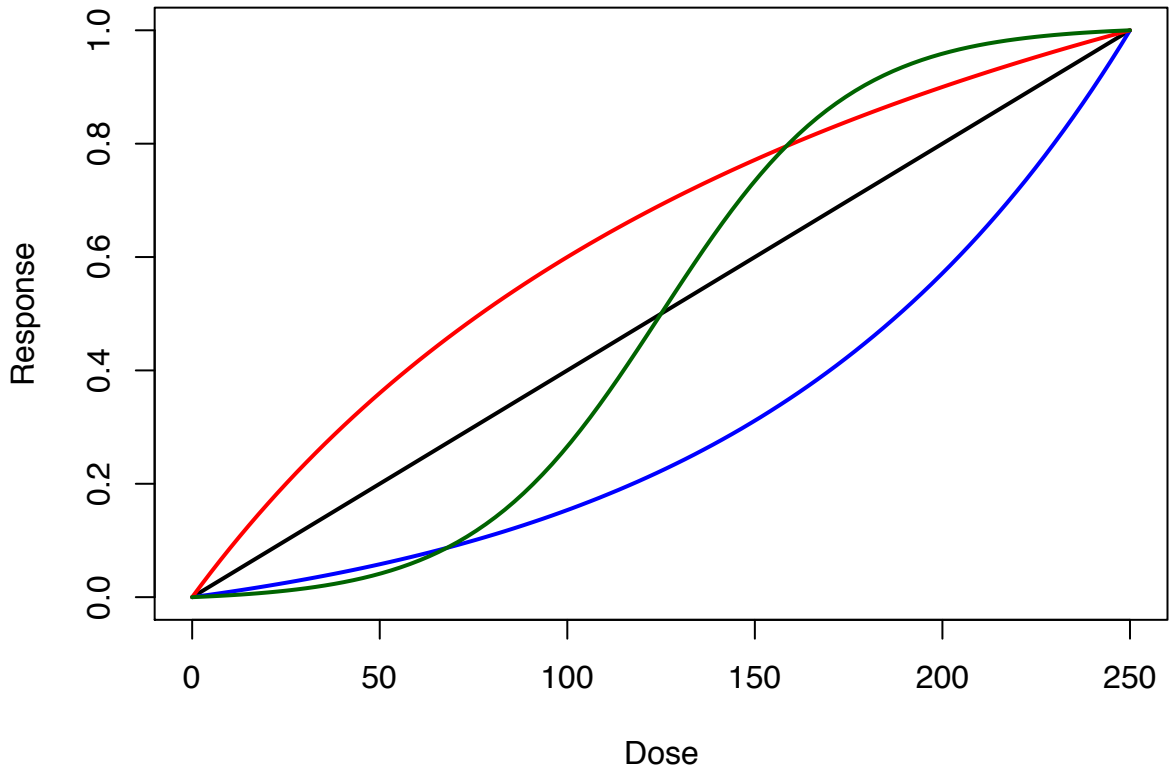

Black curve: Linear model, Blue curve: Exponential model, Red curve: Emax model, Green curve: Logistic model.

Operating characteristics of the traditional and adaptive designs were evaluated using a single set of treatment effect assumptions shown in Table 1. This table indicates that a linear dose-response relationship was assumed in the trial.

Table 1. Treatment effect assumptions

Trial arm	Parameter	Value
Placebo	Mean	1.5
	SD	4
Dose 1 (20 mg)	Mean	2
	SD	4
Dose 2 (50 mg)	Mean	2.5
	SD	4
Dose 3 (100 mg)	Mean	3
	SD	4
Dose 4 (250 mg)	Mean	3.5
	SD	4

The key operating characteristics of the two trial designs are presented in Tables 2, 3 and 4.

Tables 2 and 3 present summaries of patient allocation by stage and trial arm. Beginning with Table 2, standard descriptive statistics were computed for the number of patients enrolled in each trial stage. It is important to note that the actual number of patients per stage was quite different from the planned number of patients (80 patients per stage) due to the so-called pipeline patients, i.e., patients who are enrolled in the trial between the end of a stage and the time the patients in this stage compete the 3-month treatment period. Due to the pipeline patients, the average number of patients in the first stage approached 120 and then the stage size decreased with virtually no patients expected to be enrolled in the last stage. To address this obvious imbalance of the number of patients per stage, it would be recommended to decrease the stage size in the earlier stages and enroll more patients in the later stages. A simulation-based assessment of the adaptive design would help ensure a more balanced allocation of patients across the stages.

Continuing to Table 3, it is instructive to examine the expected patient allocation in the trial. As stated earlier, the balance parameter that controls the degree of balance with respect to the patient allocation was set to 2. A larger value of this parameter would be expected to result in a more imbalanced patient allocation, which is indeed the case. It follows from Table 3 that, as expected, about 20% of the patients would receive placebo in this trial, the average number of placebo patients was 59. A lower proportion of patients (on average 52 patients) would receive the least effective dose (20 mg) and a higher proportion of patients (on average 67 patients) would receive the most effective dose (250 mg). This table clearly illustrates a key advantage of adaptive randomization designs, i.e., an option to adjust the randomization scheme in a data-driven way to ultimately ensure that most patients will be treated at the most effective doses of an experimental treatment.

Table 4 provides a power comparison of the adaptive and traditional designs. As stated in Section 2, the final analysis focuses on the statistical significance of the dose-response relationship. This assessment relies on the MCPMod method with the same four dose-response models. As shown in Table 4, the adaptive design delivers the same probability of success as the traditional design; however, unlike the traditional approach, the adaptive approach results in a much higher likelihood of treating the patients in the trial at the most effective doses.

References

Bretz, F., Pinheiro, J.C., Branson, M. (2005). Combining multiple comparisons and modeling techniques in dose response studies. Biometrics. 61, 738-748.

Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A., Rubin, D. (2013). Bayesian Data Analysis. Third Edition. New York: Chapman and Hall/CRC Press.

Grieve, A.P., Krams, M. (2005). ASTIN: A Bayesian adaptive dose-response trial in acute stroke. Clinical Trials. 4, 340-351.

Hu, F., Rosenberger, W.F. (2008). The Theory of Response-Adaptive Randomization in Clinical Trials. New York: Wiley.

Lee, J.J., Gu, X., Liu, S. (2010). Bayesian adaptive randomization designs for targeted agent development. Clinical Trials. 5, 584-596.

Table 2. Sample size by stage

Stage	Statistic	Sample size
Stage 1	Min	102
	Median	119
	Mean	119.5
	Max	152
Stage 2	Min	68
	Median	109
	Mean	108.4
	Max	146
Stage 3	Min	60
	Median	92
	Mean	91.8
	Max	121
Stage 4	Min	0
	Median	0
	Mean	0.3
	Max	19

Table 3. Sample size by trial arm

Dose	Statistic	Sample size
Placebo	Min	57
	Median	59
	Mean	58.6
	Max	60
Dose 1 (20 mg)	Min	43
	Median	52
	Mean	52.4
	Max	107
Dose 2 (50 mg)	Min	35
	Median	56
	Mean	55.3
	Max	68
Dose 3 (100 mg)	Min	43
	Median	60

Dose	Statistic	Sample size
	Mean	59.8
	Max	69
Dose 4 (250 mg)	Min	42
	Median	67
	Mean	66.9
	Max	90

Table 4. Comparison of traditional and adaptive designs

Design	Power (%)
Traditional	81.3
Adaptive	81.8