Published in final edited form as:

Stat Med. 2011 July 30; 30(17): 2098–2108. doi:10.1002/sim.4164.

Bayesian hybrid dose-finding design in phase I oncology clinical trials

Ying Yuan^{a,*†} and Guosheng Yin^b

^aDepartment of Biostatistics, The University of Texas MD Anderson Cancer Center, Houston, TX 77030, U.S.A

^bDepartment of Statistics and Actuarial Science, The University of Hong Kong, Pokfulam Road, Hong Kong

Abstract

In oncology, dose escalation is often carried out to search for the maximum tolerated dose (MTD) in phase I clinical trials. We propose a Bayesian hybrid dose-finding method that inherits the robustness of model-free methods and the efficiency of model-based methods. In the Bayesian hypothesis testing framework, we compute the Bayes factor and adaptively assign a dose to each cohort of patients based on the adequacy of the dose–toxicity information that has been collected thus far. If the data observed at the current treatment dose are adequately informative about the toxicity probability of this dose (e.g. whether this dose is below or above the MTD), we make the decision of dose assignment (e.g. either to escalate or to de-escalate the dose) directly without assuming a parametric dose–toxicity curve. If the observed data at the current dose are not sufficient to deliver such a definitive decision, we resort to a parametric dose–toxicity curve, such as that of the continual reassessment method (CRM), in order to borrow strength across all the doses under study to guide dose assignment. We examine the properties of the hybrid design through extensive simulation studies, and also compare the new method with the CRM and the ‘3+3’ design. The simulation results show that our design is more robust than parametric model-based methods and more efficient than nonparametric model-free methods.

Keywords

model-free; model-based; Bayes factor; hypothesis testing; robust

1. Introduction

The goal of phase I trials in oncology is to find the maximum tolerated dose (MTD) of the investigational drug [1, 2]. Patients are sequentially enrolled in phase I trials and are treated in cohorts at a series of prespecified ascending doses. The first cohort of patients is often treated at the lowest dose. Based on the observed dose-limiting toxicities (DLTs) from patients who have been treated, a decision is made according to a set of prespecified rules as to which dose will be given to the next cohort of patients. Eventually, the MTD is identified based on the observed DLTs after the trial is complete.

Phase I dose-finding methods can be generally classified into algorithm-based (dose–toxicity curve-free or model-free) and dose–toxicity model-based approaches. Typically, algorithm-based designs do not assume any dose–toxicity curve except that toxicity monotonically increases with the dose. This family of designs conducts dose escalation and de-escalation strictly according to the prespecified algorithm. The most commonly used algorithm-based method is the so-called ‘3+3’ design, in which the decision of dose assignment for the next cohort of patients depends solely on the observed data at the most recently administered dose level [3]: for example, if 0 out of 3 patients experiences DLT, we escalate the dose; if 1 out of 3 patients experiences DLT, we treat the next three patients at the current dose level; and if 2 or more out of 3 patients experience DLT, we de-escalate the dose. These rules rigorously specify the action under each circumstance in the trial. Other algorithm-based methods include the random walk rule [4], the improved up-and-down design [5] and the biased coin design [6]. Algorithm-based methods are robust because they do not rely on any parametric model structures. However, these designs, in particular the ‘3+3’ design, has been criticized for inefficiency because the data information across different dose levels is not borrowed or fully utilized, i.e. the decision of dose assignment for the next cohort is based only on the toxicity outcomes observed at the most recently administered dose.

Model-based dose-finding methods have been proposed to improve the performance of dose finding by pooling all the information together across different doses. Model-based methods often assume a parametric dose–toxicity model indexed by one or two parameters, such as a probability power function or a logistic regression model. During the trial, the model parameter is continuously updated based on all the observed data and dose-escalation decisions are made based on the estimated toxicity probabilities using the dose–toxicity model. Among the model-based approaches, an influential phase I trial design is the continual reassessment method (CRM) [7]. Assuming a one-parameter power function form of the dose–toxicity curve, the CRM utilizes all the observed toxicity information to direct the dose assignment. By modeling all the observed data, the CRM has been shown to be superior to the algorithm-based ‘3+3’ design. A variety of extensions of the CRM have been proposed to improve its practical implementation and operating characteristics [8–12]. Garrett-Mayer [13] provided a comprehensive introduction to the CRM and information about its practical use in phase I clinical trials.

When using the CRM, because all the decisions on dose assignment center around the assumed dose–toxicity model, correct specification of the dose–toxicity model is crucial for the performance of the design. Cheung and Chappell [14] investigated the model sensitivity in the CRM. Shen and O’Quigley [15] showed that under certain conditions the CRM is consistent in the sense that the MTD identified by the CRM generally converges to the true MTD, even when the working model is misspecified. However, the sample size in a typical phase I trial is very small, often as few as 20–40 patients. Under such small sample sizes, Yin and Yuan [16] showed that the performance of the CRM may be compromised if the dose–toxicity model is misspecified. In addition, as shown in our simulation studies (described in Section 3), when the assumed dose–toxicity curve is substantially deviated from the truth, the performance of the CRM can be downgraded because of some inappropriate actions, such as continuing to treat patients at the current dose when 3 out of 3 patients have experienced toxicity, or de-escalating the dose when 0 out of 6 patients has experienced toxicity. This issue has raised considerable concerns about the application of the CRM in clinical practice [17].

To bridge the algorithm-based and model-based methods, we propose a hybrid dose-finding approach to phase I clinical trials. Unlike the ‘3+3’ design, which never borrows information across different doses, and the CRM, which always enforces a parametric dose–toxicity

curve to fully borrow information across the doses, our method borrows information adaptively in the sense that it borrows information only when it is necessary. That is, if the data observed at the latest dose level administered are informative enough about the toxicity probability of this dose (e.g. whether this dose is below or above the MTD), we make the decision to escalate or de-escalate the dose without borrowing information from other doses. This is similar to the '3+3' design. If the observed data at the latest dose level administered do not contain enough information to make such a definitive decision, we use the CRM model to borrow strength across all the dose levels in order to determine the dose assignment for the next cohort of patients. This adaptation gives our approach the robustness of algorithm-based methods and also takes advantage of the efficiency of model-based approaches. Our design is built upon the Bayesian hypothesis testing framework, in which the strength of the information in the observed data is measured by the Bayes factor. Given two hypotheses under examination, the Bayes factor represents the ratio of the posterior odds to the prior odds, which also has an explicit probability interpretation regarding the evidence contained in the data to support one hypothesis [18]. Based on frequentist hypothesis testing, Cheung [19] considered a stepwise sequential procedure for finding the MTD.

The remainder of this article is organized as follows. In Section 2, we briefly review the CRM and propose the Bayesian hybrid dose-finding method based on the Bayes factor. In Section 3, we present simulation studies to compare the operating characteristics of the new hybrid design with those of the standard '3+3' and CRM designs. We conclude with a brief discussion in Section 4.

2. Method

2.1. Continual reassessment method

In a typical dose-finding trial, the DLT is recorded as a binary outcome and the true dose-toxicity is assumed to monotonically increase with respect to the dose level. The CRM specifies a prior dose-toxicity curve and then continuously updates this curve given the observed toxicity outcomes from the cumulating patients in the trial. Based on the updated dose-toxicity curve, a new cohort of patients is assigned to the dose with an estimated toxicity probability closest to the prespecified target. Let $(d_1, \dots, d_J)$ denote a set of J prespecified increasing doses for the drug under investigation. In the CRM, we assume a working dose-toxicity model, such as

$$\text{pr}(\text{toxicity at } d_j) = \pi_j(\alpha) = p_j^{\exp(\alpha)} \quad (1)$$

for $j = 1, \dots, J$, where α is an unknown parameter, and $(p_1, \dots, p_J)$ are prespecified constants satisfying $0 < p_1 < \dots < p_J < 1$.

Let φ be the target toxicity rate specified by physicians. The first cohort of patients usually receives the lowest dose d_1 . Let $\mathbf{y}$ denote the observed data, $\mathbf{y} = \{y_j; j = 1, \dots, J\}$, i.e. among n_j patients treated at dose level j , y_j patients have experienced DLT. Based on the binomial distribution for the toxicity outcome, the likelihood function is

$$L(\mathbf{y}|\alpha) \propto \prod_{j=1}^J \{p_j^{\exp(\alpha)}\}^{y_j} \{1 - p_j^{\exp(\alpha)}\}^{n_j - y_j}.$$

We estimate the toxicity probabilities by the corresponding posterior means of $\pi_j(\alpha)$, which is given by

$$\hat{\pi}_j = \int p_j^{\exp(\alpha)} \frac{L(\mathbf{y}|\alpha)f(\alpha)}{\int L(\mathbf{y}|\alpha)f(\alpha) d\alpha} d\alpha,$$

where $f(\alpha)$ is a prior distribution for the parameter α . We take a zero-mean normal prior distribution, $N(0,2)$, for α . After updating the posterior estimates of the toxicity probabilities at all of the doses considered, the recommended dose level j^* for the next cohort of patients is the one that has a toxicity probability closest to the target φ :

$$j^* = \operatorname{argmin}_{j \in \{1, \dots, J\}} |\hat{\pi}_j - \varphi|.$$

The trial continues until the exhaustion of the total sample size, and then the dose with a posterior toxicity probability closest to φ is selected as the MTD. For safety, in practice, we often impose a stopping rule during the trial, such as

$$\text{if } \Pr(\pi_1 > \varphi | \mathbf{y}) > 0.9 \quad \text{terminate the trial.} \quad (2)$$

2.2. Bayesian hybrid design

The intuition behind our hybrid dose-finding method is that if the data observed at the current dose level provide overwhelming evidence that this dose is below or above the MTD, there is no need to enforce a parametric dose-toxicity model to borrow information across different doses. We can directly make the decision of dose assignment for the next cohort of patients based on nonparametric methods as long as the data information at the current dose is strong enough (e.g. 0 out of 6 or 3 out of 3 patients experienced DLTs). On the other hand, if the data at the current dose level contain only limited information on the toxicity of the current dose that is not strong enough for decision making, we turn to the model-based approach to borrow strength from all the other doses to facilitate the decision making on the next dose assignment. For this purpose, a parsimonious model is more suitable because the data are often sparse in phase I trials. Through combining the model-free and model-based approaches coherently using the Bayes factor, our design can efficiently eliminate occurrences of some extreme cases of dose escalation or de-escalation that may be considered inappropriate in clinical practice: for example, continuing to treat patients at the same dose when 3 out of 3 patients have experienced DLTs, escalating the dose when 2 DLTs are observed out of 3 patients, or de-escalating the dose when 0 DLT is observed out of 6 patients.

Suppose at the most recently administered dose level j , we observe that y_j out of n_j patients have experienced toxicity. To gauge the distance between the toxicity probability of dose level j and the target φ , we define three complementary hypotheses:

$$\begin{aligned} H_1: \pi_j &< \varphi - \delta, \\ H_2: \varphi - \delta &\leq \pi_j \leq \varphi + \delta, \\ H_3: \pi_j &> \varphi + \delta, \end{aligned}$$

where π_j is the true toxicity probability of dose level j , and δ is the tolerable margin prespecified by physicians. The hypotheses H_1 , H_2 and H_3 represent the situations in which

dose level j is below, approximately equal to, and above the MTD, respectively. We formulate H_2 as an interval hypothesis $\varphi - \delta \leq \pi_j \leq \varphi + \delta$ rather than a point hypothesis $\pi_j = \varphi$ because in clinical practice, as long as the toxicity probability of a dose is adequately close to φ , this dose can be chosen as the MTD. For instance, with $\varphi = 0.3$ and $\delta = 0.03$, a dose with a toxicity probability lying within $(0.27, 0.33)$ would be accepted as the MTD.

Given the data observed at dose level j , $\{n_j, y_j\}$, we evaluate the evidence of supporting each hypothesis by calculating their posterior probabilities. We assign the toxicity probabilities π_j a uniform prior distribution under each hypothesis:

$$\begin{aligned} p(\pi_j|H_1) &= \text{Unif}(0, \varphi - \delta), \\ p(\pi_j|H_2) &= \text{Unif}(\varphi - \delta, \varphi + \delta), \\ p(\pi_j|H_3) &= \text{Unif}(\varphi + \delta, 1). \end{aligned} \quad (3)$$

It then follows that the marginal distribution of y_j under H_1 is given by

$$\begin{aligned} p(y_j|H_1) &= \int_0^{\varphi - \delta} \binom{n_j}{y_j} \pi_j^{y_j} (1 - \pi_j)^{n_j - y_j} \frac{1}{\varphi - \delta} d\pi_j \\ &= \frac{F_{\text{beta}}(\varphi - \delta; y_j + 1, n_j - y_j + 1)}{(\varphi - \delta)(n_j + 1)}, \end{aligned}$$

where $F_{\text{beta}}(c; a, b)$ is the cumulative distribution function of a beta distribution with the shape and scale parameters a and b , evaluated at the value c . Similarly, the marginal distributions of y_j under H_2 and H_3 are given by

$$p(y_j|H_2) = \frac{F_{\text{beta}}(\varphi + \delta; y_j + 1, n_j - y_j + 1) - F_{\text{beta}}(\varphi - \delta; y_j + 1, n_j - y_j + 1)}{2\delta(n_j + 1)}$$

and

$$p(y_j|H_3) = \frac{1 - F_{\text{beta}}(\varphi + \delta; y_j + 1, n_j - y_j + 1)}{(1 - \varphi - \delta)(n_j + 1)},$$

respectively. Therefore, at dose level j , the posterior probability of H_k ($k = 1, 2, 3$) is given by

$$p(H_k|y_j) = \frac{p(H_k)p(y_j|H_k)}{p(H_1)p(y_j|H_1) + p(H_2)p(y_j|H_2) + p(H_3)p(y_j|H_3)},$$

which can also be expressed equivalently using the Bayes factors

$$p(H_k|y_j) = \frac{p(H_k)}{p(H_1)\text{BF}_{1k} + p(H_2)\text{BF}_{2k} + p(H_3)\text{BF}_{3k}},$$

where $BF_{ik} = p(y_j | H_i) / p(y_j | H_k)$ ($i, k = 1, 2, 3$), is the Bayes factor of H_i against H_k . With no evidence favoring any of the hypotheses over the others *a priori*, we take $p(H_1) = p(H_2) = p(H_3) = 1/3$, and then

$$p(H_k | y_j) = \frac{1}{BF_{1k} + BF_{2k} + BF_{3k}}. \quad (4)$$

To gauge the strength of the evidence contained in the data in favor of each hypothesis, Jeffreys [20] suggested interpreting the Bayes factor in the unit of 1/2 on the $\log_{10}$ scale: if $\log_{10}(BF_{12}) > 1/2$, this indicates that the data contain substantial evidence in favor of H_1 against H_2 ; if $\log_{10}(BF_{12}) > 1$, such evidence is strong in the data; and if $\log_{10}(BF_{12}) > 2$, then the evidence appears to be decisive. In our case, if $\log_{10}(BF_{12}) > 1/2$ and $\log_{10}(BF_{13}) > 1/2$, or equivalently $p(H_1 | y_j) > 0.61$, there is substantial evidence in favor of H_1 against both H_2 and H_3 , suggesting that dose level j is far below the MTD. Therefore, we should directly escalate the dose to level $j+1$, without the need to borrow any information from other doses. Similarly, if $p(H_3 | y_j) > 0.61$, we should de-escalate the dose to level $j-1$ as there is substantial evidence indicating that dose level j is far above the MTD. Finally, if $p(H_2 | y_j) > 0.61$, i.e. there is substantial evidence that dose level j is close to the MTD, the next dose should then stay at the same level.

In the case in which none of the posterior probabilities of the hypotheses is larger than 0.61, i.e. $p(H_k | y_j) \leq 0.61$ for all k , then there is not adequate information at dose level j alone to support any action. As a consequence, we invoke the CRM dose-toxicity model to pool the information together from all the dose levels to direct the dose assignment for new patients. In other words, if the toxicity information at the currently administered dose is strong enough, we base the decision upon the Bayes factors obtained in (4); otherwise we resort to the model-based approach to maximize the borrowing of information across different dose levels. Following this route, we adopt the CRM model in (1) for such a switch from model-free to model-based schemes. Under the three hypotheses, the uniform priors for π_j given in (3) can be transformed into the following prior distributions for α :

$$\begin{aligned} p(\alpha | H_1) &= \frac{-\log(p_j)}{\varphi - \delta} p_j^{\exp(\alpha)} \exp(\alpha) \quad \text{for } \alpha > \log \left\{ \frac{\log(\varphi - \delta)}{\log(p_j)} \right\}, \\ p(\alpha | H_2) &= \frac{-\log(p_j)}{2\delta} p_j^{\exp(\alpha)} \exp(\alpha) \quad \text{for } \log \left\{ \frac{\log(\varphi + \delta)}{\log(p_j)} \right\} < \alpha < \log \left\{ \frac{\log(\varphi - \delta)}{\log(p_j)} \right\}, \\ p(\alpha | H_3) &= \frac{-\log(p_j)}{1 - \varphi - \delta} p_j^{\exp(\alpha)} \exp(\alpha) \quad \text{for } \alpha < \log \left\{ \frac{\log(\varphi + \delta)}{\log(p_j)} \right\}. \end{aligned}$$

Thus, the marginal distribution of the observed data at all of the doses $y = \{y_j; j = 1, \dots, J\}$ under H_k ($k = 1, 2, 3$), is given by

$$p(y | H_k) = \int p(\alpha | H_k) \prod_{i=1}^J \binom{n_i}{y_i} \{p_i^{\exp(\alpha)}\}^{y_i} \{1 - p_i^{\exp(\alpha)}\}^{n_i - y_i} d\alpha.$$

The posterior distribution of H_k takes the same form as (4), based on which Jeffereys' rule can be used for decision making. If still none of the decisions can be reached based on all the observed data under the CRM model, we treat the next cohort of patients at the current dose level to accumulate more information.

2.3. Dose-finding algorithm

Let φ be the physician-specified toxicity target. Patients are treated in cohorts, for example, with a cohort size of three. To be conservative and protect patients' safety, we restrict dose escalation or de-escalation by one dose level of change at a time. Assuming that $p(H_1) = p(H_2) = p(H_3) = 1/3$, the Bayesian hybrid dose-finding algorithm is described as follows:

1. Patients in the first cohort are treated at the lowest dose d_1 , or the physician-specified dose.
2. At the current dose level j^{curr} with the observed data $y_{j^{\text{curr}}}$, we evaluate $p(H_1|y_{j^{\text{curr}}})$, $p(H_2|y_{j^{\text{curr}}})$ and $p(H_3|y_{j^{\text{curr}}})$. If $p(H_1|y_{j^{\text{curr}}}) > 0.61$, we escalate the dose level to $j^{\text{curr}} + 1$; if $p(H_3|y_{j^{\text{curr}}}) > 0.61$, we de-escalate the dose level to $j^{\text{curr}} - 1$; and if $p(H_2|y_{j^{\text{curr}}}) > 0.61$, the dose stays at the same level as j^{curr} for the next cohort of patients.
3. Otherwise, we switch to the CRM model to evaluate $p(H_1|y)$, $p(H_2|y)$ and $p(H_3|y)$ based on all the observed data $y = \{y_j; j = 1, \dots, J\}$, and apply similar decision rules as in step 2 for dose assignment.
4. Once the maximum sample size is reached, we select dose j^* as the MTD, such that at dose level j^* the toxicity probability estimate $\hat{\pi}_j^*$ is closest to φ , where $\hat{\pi}_j^*$ is the isotonically transformed proportion y_j^*/n_j^* using the pooled-adjacent-violators algorithm [21].

The isotonic transformation is used to ensure that $\hat{\pi}_j = y_j/n_j$ monotonically increases with the dose levels. For safety, we impose the early stopping rule given in (2) to account for the possibility that all the doses considered are overly toxic.

3. Simulation studies

We investigated the operating characteristics of the proposed Bayesian hybrid dose-finding design through simulation studies under eight toxicity scenarios listed in Table I. We considered six dose levels with the target toxicity probability $\varphi = 0.3$, and assumed that toxicity monotonically increased with respect to the dose. The maximum sample size was 24 and patients were treated in cohorts of size 3. We compared the proposed hybrid design with the '3+3' design and two versions of the CRM design: one with an arbitrary dose-toxicity model having $(p_1, \dots, p_6) = (0.14, 0.20, 0.25, 0.30, 0.35, 0.40)$ which is also used in the hybrid design for the parametric model, and the other with a correctly specified dose-toxicity model obtained by setting the p_j 's at the true toxicity probabilities of each scenario (denoted as CRM_{True}). Although the CRM_{True} is typically not available in practice because the underlying true toxicity profile is unknown; we use the CRM_{True} as a benchmark for comparison because it represents the optimal case. Among various existing algorithm-based designs, such as the improved up-and-down design [5] or the biased coin design [6], we chose the standard '3+3' design, due to its practical popularity, as an illustrative example for comparison with the proposed design. For the CRMs, we used the modified version [8], in which the first cohort is treated at the lowest dose and dose escalation/de-escalation is limited to one dose level of change at a time. In the proposed hybrid design, we took the tolerable margin $\delta = 0.03$, and under each scenario, we simulated 10 000 trials. The simulation code was written in R, and is available for downloading at <http://odin.mdacc.tmc.edu/~yyuan/>.

Table I shows the dose selection probability, the percentage of inconclusive trials (denoted by 'None'), the average number of patients treated at each dose level, the average number of patients treated at doses above the MTD ($N_{\text{MTD}+}$), and the average number of DLTs (N_{DLT}). In scenario 1, for which dose 3 is the MTD, the '3+3' design had the lowest selection probability, 41.2 per cent, and the CRM had a slightly higher selection percentage of 42.2

per cent. The hybrid design outperformed these two methods with a selection probability of 53.8 per cent. In addition, the hybrid design was safer than the CRM, with 3.6 versus 6.8 being the respective number of patients treated at doses above the MTD, and 5.6 versus 6.9 being the respective observed number of DLTs. As expected, the CRM_{True} had the best performance with the highest selection probability and the smallest number of patients treated at doses above the MTD. This suggests that when the dose–toxicity model is correctly specified, the CRM indeed performs very well in terms of both the MTD selection and patient safety. In scenario 2, where the MTD is at dose level 4, the hybrid design had a slightly better selection probability for the MTD and was substantially safer than the CRM. The number of patients treated at doses above the MTD under the hybrid design was less than a half of that under the CRM. Although the ‘3+3’ design was relatively safe, with 3.9 patients treated at doses above the MTD, its selection probability was slightly inferior to those of the other designs. In scenario 3, which has the fifth dose as the MTD, the selection probability of the hybrid design was 9 per cent higher than that of the CRM, and the number of patients treated at doses above the MTD when using the hybrid design was smaller than that when using the CRM (2.7 versus 5.0). Similar results were observed in scenario 4, in which the MTD is the second dose. In scenario 5, there is a sudden toxicity jump from dose level 4 to 5. In this case, the CRM did not behave as well as the others. The selection probability of the MTD using the hybrid design was 27.7 per cent higher than that of the CRM. Nevertheless, if the dose–toxicity model was correctly specified, the CRM_{True} performed the best with a selection probability of 75.7 per cent. When the MTD is the last dose, as in scenario 6, the CRM outperformed the hybrid design with a much higher selection probability of the MTD; and more interestingly, the CRM even yielded a higher MTD selection percentage than the CRM_{True}. When the MTD is the first dose (scenario 7), both the CRM and the hybrid design surprisingly outperformed the CRM_{True}. These results suggest that, under scenarios 6 and 7, the assumed dose–toxicity model might underestimate toxicity probabilities and consequently the CRM and the hybrid design were more likely to select high dose levels as the MTD than the CRM_{True}. Scenario 8 is designed to examine whether the proposed hybrid design can terminate the trial properly when all the doses considered are overly toxic, and clearly all of the designs were able to terminate the trial early with high probabilities.

By adaptively using the parametric dose–toxicity model to borrow information across different doses, the hybrid design efficiently limits the impact of model misspecification and thus substantially improves the performance of dose finding and trial safety. Comparatively, the CRM is more vulnerable to model misspecification because it always enforces the assumed parametric dose–toxicity model. As a result, the CRM may lead to some practically inappropriate actions of dose escalation or de-escalation. For example, the CRM may escalate the dose although the observed toxicity rate at the current dose level is high (e.g. 2 out of 3 patients have experienced toxicity, for which the standard ‘3+3’ design would result in immediate dose de-escalation). In addition, the CRM may de-escalate the dose when the observed toxicity rate is already low (e.g. only 1 out of 9 patients has experienced toxicity), or continue treating patients at the same dose when the observed toxicity rate is very high (e.g. 3 out of 3 patients have experienced toxicity). In Table II we present the percentage of trials for which the CRM directs such practically inappropriate actions. More specifically, the following dose assignments are considered to be inappropriate: at dose level j , escalate the dose when the observed toxicity rate is $y_j/n_j = 2/3$; de-escalate the dose when $y_j/n_j = (0/6, 1/6, 1/9)$; and retain the same dose level when $y_j/n_j = (3/3, 5/6, 0/6, 0/9, 1/9)$. We can see that the percentages of CRM-directed trials with inappropriately dose assignments varied across different scenarios. For example, in scenario 1, the CRM tended to be overly conservative with 6.5 per cent of the trials continuing to treat patients at the same dose although none of the 6 patients experienced toxicity at that dose. Also in scenario 1, 5.1 per cent of the trials de-escalated the dose when only 1 out of 6 patients experienced toxicity. In scenario 5, the

CRM was particularly aggressive as 12.7 per cent of the trials escalated the dose when 2 out of 3 patients experienced toxicity at the current dose. Although the chance of such inappropriate dose escalation or de-escalation is generally low, the occurrence of these improper actions may discourage the widespread use of the CRM in practice. In contrast, the proposed hybrid design did not conduct any of these inappropriate actions in the simulated trials. Therefore, the hybrid design provides a potential solution to avoid the inappropriate and potentially harmful actions for the CRM.

The CRM used in our simulation study is a modified version with some practical safety rules, i.e. the trial starts at the lowest dose and limits a change in the dose to one level at a time. One consequence of imposing such safety rules is that the modified CRM may lose the group-coherent property of the original CRM [22], as demonstrated by the aforementioned inappropriate dose escalation and de-escalation. The group-coherent property means that a dose-finding design does not escalate the dose if the observed toxicity rate in the current cohort is equal to or higher than the target toxicity rate. Figure 1 provides a numerical example to illustrate the violation of the group-coherent property under the modified CRM. In this example, we considered scenario 3 and assumed an ideal case in which the dose-toxicity model was correctly specified with the true toxicity probabilities, i.e. $(p_1, \dots, p_6) = (0.06, 0.08, 0.1, 0.15, 0.3, 0.45)$. We treated the first cohort at dose level one. After no toxicity was observed, we escalated the dose to level 2 to treat the second cohort, in which 1 out of 3 patients experienced toxicity. Based on the observed data, the estimated toxicity rates at six dose levels were $(0.17, 0.20, 0.23, 0.29, 0.43, 0.56)$, respectively, leading to a dose escalation to level 3. In other words, although the observed toxicity rate of 1/3 in the second cohort was higher than the target toxicity probability of 0.3, the CRM still escalated the dose and thus violated the group-coherent property.

In the proposed design, the tolerable margin δ is prespecified by physicians based on clinical knowledge, i.e. the largest deviation from the target toxicity rate for a dose that physicians are comfortable to choose as the MTD. We assessed the sensitivity of the proposed method to the choice of δ using a simulation study based on two other values of δ : 5 and 15 per cent of the target toxicity rate, i.e. $\delta = 0.015$ and 0.045 , respectively. As shown in Table III, our method was not sensitive to the value of δ . Both the dose selection probabilities and the number of patients treated at each dose level were generally close to those listed in Table I.

4. Conclusion

We have proposed a new Bayesian hybrid dose-finding design based on the Bayes factors, which combines the advantages of model-free and model-based approaches. If the data observed at the current treatment dose are adequately informative about the toxicity probability of this dose (e.g. whether this dose is below or above the MTD), we make the decision of dose assignment (e.g. either to escalate or to de-escalate the dose) directly based on a model-free approach. However, if the toxicity information observed at the current dose is not strong enough to inform the decision making on dose assignment, a parametric model is used to collect information from all the dose levels under investigation in order to facilitate the next dose assignment.

In our approach, three hypotheses are formulated to represent that the current dose is below, adequately close to, or above the MTD. We measure the strength of the toxicity evidence contained in the data using the Bayes factor, and apply Jeffreys' rule to make a decision on dose escalation or de-escalation. We take the original CRM as the model-based aspect of our approach. The hybrid design offers straightforward implementation and easy computation based on the Gaussian quadrature approximation. It represents a nice compromise of the

model-free and model-based dose-finding methods, and thus inherits advantages from these two classes of designs.

Acknowledgments

We thank the referees, the Associate Editor and the Editor for very helpful comments that substantially improved this paper. The research of Ying Yuan was partially supported by the National Cancer Institute (U.S.A.) grant R01CA154591-01A1, and the research of Guosheng Yin was partially supported by a grant from the Research Grants Council of Hong Kong.

References

1. Chevret, S. Statistical Methods for Dose-finding Experiments. Wiley; England: 2006.
2. Ting, N. Dose Finding in Drug Development. Springer; Cambridge: 2006.
3. Storer BE. Design and analysis of phase I clinical trials. *Biometrics*. 1989; 45:925–937. [PubMed: 2790129]
4. Durham SD, Flournoy N, Rosenberger WF. A random walk rule for phase I clinical trials. *Biometrics*. 1997; 53:745–760. [PubMed: 9192462]
5. Leung D, Wang YG. An improved up-and-down design for phase I trials. *Controlled Clinical Trials*. 2001; 22:126–138. [PubMed: 11306151]
6. Stylianou M, Flournoy N. Dose finding using the biased coin up-and-down design and isotonic regression. *Biometrics*. 2002; 58:171–177. [PubMed: 11890313]
7. O’Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics*. 1990; 46:33–48. [PubMed: 2350571]
8. Goodman SN, Zahurak ML, Piantadosi S. Some practical improvements in the continual reassessment method for phase I studies. *Statistics in Medicine*. 1995; 14:1149–1161. [PubMed: 7667557]
9. Møller S. An extension of the continual reassessment methods using a preliminary up-and-down design in a dose finding study in cancer patients, in order to investigate a greater range of doses. *Statistics in Medicine*. 1995; 14:911–922. [PubMed: 7569510]
10. Piantadosi S, Fisher J, Grossman S. Practical implementation of a modified continual reassessment method for dose finding trials. *Cancer Chemotherapy and Pharmacology*. 1998; 41:429–436. [PubMed: 9554585]
11. Heyd JM, Carlin PB. Adaptive design improvements in the continual reassessment method for phase I studies. *Statistics in Medicine*. 1999; 18:1307–1321. [PubMed: 10399198]
12. Braun TM. The bivariate CRM: extending the CRM to phase I trials of two competing outcomes. *Controlled Clinical Trials*. 2002; 23:240–255. [PubMed: 12057877]
13. Garrett-Mayer E. The continual reassessment method for dose-finding studies: a tutorial. *Clinical Trials*. 2006; 3:57–71. [PubMed: 16539090]
14. Cheung YK, Chappell R. A simple technique to evaluate model sensitivity in the continual reassessment method. *Biometrics*. 2002; 58:671–674. [PubMed: 12230003]
15. Shen L, O’Quigley J. Consistency of continual reassessment method under model misspecification. *Biometrika*. 1996; 83:395–405.
16. Yin G, Yuan Y. Bayesian model averaging continual reassessment method in phase I clinical trials. *Journal of the American Statistical Association*. 2009; 104:954–968.
17. Mathew P, Thall PF, Jones D, Perez C, Bucana C, Troncoso P, Kim S, Fidler I, Logothetis C. Platelet-derived growth factor receptor inhibitor imatinib mesylate and docetaxel: a modular phase I trial in androgen-independent prostate cancer. *Journal of Clinical Oncology*. 2004; 16:3323–3329. [PubMed: 15310776]
18. Goodman SN. Toward evidence-based medical statistics. Part 2: the Bayes factor. *Annals of Internal Medicine*. 1999; 130:1005–1013. [PubMed: 10383350]
19. Cheung YK. Sequential implementation of stepwise procedures for identifying the maximum tolerated dose. *Journal of the American Statistical Association*. 2007; 102:1448–1461.
20. Jeffreys, H. *Theory of Probability*. 3. Oxford University Press; London: 1961.

21. Robertson, T.; Wright, F.; Dykstra, R. Order Restricted Statistical Inference. Wiley; New York: 1988.
22. Cheung YK. Coherence principles in dose-finding studies. *Biometrika*. 2005; 92:863–873.

Figure 1. Group-incoherent dose escalation under the modified CRM. The open circles denote patients without toxicity and the solid circles denote patients with toxicity.

Table 1

Dose selection probability, average number of patients treated at each dose level, average number of patients treated at the doses above the MTD ($N_{\text{MTD}+}$), and average number of toxicities (N_{DLT}) under the ‘3+3’, CRM, and the proposed Bayesian hybrid designs.

Design	Dose level							None	$N_{\text{MTD}+}$	N_{DLT}
	1	2	3	4	5	6				
Scenario 1										
3+3	True toxicity	0.10	0.12	0.30	0.50	0.60	0.65			
	Selection (per cent)	6.0	33.1	41.2	13.6	2.2	0.0	3.9		
	Patients #	3.9	4.4	4.7	2.8	0.8	0.1		3.7	3.1
CRM	Selection (per cent)	3.9	24.5	42.2	22.7	4.9	1.5	0.3		
	Patients #	4.7	5.2	7.3	4.5	1.9	0.38		6.8	6.9
Hybrid	Selection (per cent)	6.0	18.2	53.8	19.3	2.3	0.3	0.2		
	Patients #	6.1	6.2	8.1	2.8	0.7	0.1		3.6	5.6
CRM _{True}	Selection (per cent)	0.5	13.6	71.4	13.3	0.6	0.2	0.3		
	Patients #	4.5	5.3	11.1	2.5	0.4	0.1		3.0	6.0
Scenario 2										
3+3	True toxicity	0.08	0.12	0.20	0.30	0.40	0.50			
	Selection (per cent)	5.2	14.3	26.1	27.3	17.9	0.0	9.2		
	Patients #	3.7	4.1	4.4	3.9	2.7	1.3		3.9	3.0
CRM	Selection (per cent)	0.7	6.1	18.6	29.6	26.0	18.9	0.1		
	Patients #	4.2	3.9	5.1	4.6	3.9	2.2		6.2	5.9
Hybrid	Selection (per cent)	3.1	11.8	29.0	33.3	16.8	5.9	0.1		
	Patients #	5.3	6.1	5.9	3.6	2.2	0.8		3.0	4.7
CRM _{True}	Selection (per cent)	0.3	3.6	25.2	40.6	23.4	6.8	0.1		
	Patients #	4.1	4.0	6.1	5.9	2.9	0.9		3.8	5.4
Scenario 3										
3+3	True toxicity	0.06	0.08	0.10	0.15	0.30	0.45			
	Selection (per cent)	2.1	3.3	8.7	29.4	35.2	20.0	1.3		
	Patients #	3.5	3.6	3.8	4.3	4.3	2.8		2.8	2.6
CRM	Selection (per cent)	0.1	0.5	3.3	15.9	32.1	48.1	0.0		
	Patients #	3.8	3.3	3.5	3.8	4.7	5.0		5.0	5.1

Design		Dose level							N _{MDD+}	N _{DLT}
		1	2	3	4	5	6	None		
Hybrid	Selection (per cent)	1.1	4.0	8.2	24.9	41.1	20.6	0.1		
	Patients #	0.5	4.9	4.3	3.7	3.9	2.7		2.7	4.0
CRM _{True}	Selection (per cent)	0.0	0.2	0.9	18.2	56.6	24.0	0.0		
	Patients #	3.7	3.1	3.2	5.1	6.5	2.4		2.4	4.6
<i>Scenario 4</i>										
3+3	True toxicity	0.2	0.3	0.4	0.5	0.6	0.7			
	Selection (per cent)	27.4	29.9	18.8	6.4	1.1	0.1	16.3		
CRM	Patients #	4.7	4.2	2.9	1.3	0.4	0.1		4.7	3.1
	Selection (per cent)	27.5	35.7	21.5	8.0	1.9	0.3	5.1		
Hybrid	Patients #	9.7	6.6	4.5	1.7	0.6	0.1		7.0	7.1
	Selection (per cent)	32.4	40.8	17.5	4.1	0.5	0.1	4.7		
CRM _{True}	Patients #	11.2	8.4	3.0	0.6	0.2	0.0		3.7	6.3
	Selection (per cent)	25.6	40.8	22.7	5.2	0.5	0.0	5.1		
<i>Scenario 5</i>										
3+3	True toxicity	0.02	0.03	0.04	0.05	0.3	0.5			
	Selection (per cent)	0.3	0.4	1.0	35.7	44.7	17.8	0.1		
CRM	Patients #	3.2	3.3	3.4	4.3	5.0	3.1		3.1	2.5
	Selection (per cent)	0.0	0.0	0.8	10.8	36.0	52.5	0.0		
Hybrid	Patients #	3.2	3.0	3.2	3.4	4.8	6.5		6.5	5.2
	Selection (per cent)	0.1	0.2	0.7	11.3	63.7	24.0	0.0		
CRM _{True}	Patients #	3.4	3.6	3.5	3.5	5.4	4.6		4.6	4.5
	Selection (per cent)	0.0	0.0	0.0	4.7	75.7	19.6	0.0		
<i>Scenario 6</i>										
3+3	True toxicity	0.02	0.05	0.08	0.1	0.14	0.3			
	Selection (per cent)	0.9	2.4	3.7	7.4	32.2	53.3	0.1		
CRM	Patients #	3.2	3.5	3.7	3.8	4.3	4.6		0.0	1.6
	Selection (per cent)	0.0	0.0	0.2	1.9	10.2	87.7	0.0		
	Patients #	3.2	3.1	3.2	3.1	3.6	7.7		0.0	3.6

Design		Dose level							N _{MDD+}	N _{DLT}
		1	2	3	4	5	6	None		
Hybrid	Selection (per cent)	0.0	0.7	3.3	7.8	27.8	60.4	0.0		
	Patients #	3.4	4.0	4.2	3.4	3.6	5.3		0.0	3.1
CRM _{True}	Selection (per cent)	0.0	0.0	0.1	0.7	18.6	80.6	0.0		
	Patients #	3.2	3.1	3.1	3.2	4.4	7.1		0.0	3.5
<i>Scenario 7</i>										
3+3	True toxicity	0.3	0.45	0.5	0.6	0.7	0.8			
	Selection (per cent)	39.5	17.6	5.8	1.0	0.0	0.0	36.0		
CRM	Patients #	5.0	3.2	1.2	0.3	0.0	0.0		4.8	2.9
	Selection (per cent)	53.6	17.7	3.9	1.0	0.1	0.0	23.7		
Hybrid	Patients #	13.9	4.7	1.8	0.4	0.1	0.0		7.0	7.5
	Selection (per cent)	52.6	20.4	3.0	0.4	0.0	0.0	23.5		
CRM _{True}	Patients #	14.7	5.3	0.7	0.1	0.0	0.0		6.2	7.2
	Selection (per cent)	47.4	15.4	2.9	0.4	0.0	0.0	33.8		
<i>Scenario 8</i>	Patients #	12.2	4.2	1.5	0.3	0.0	0.0		6.0	6.5
	True toxicity	0.50	0.60	0.70	0.75	0.78	0.80			
3+3	Selection (per cent)	21.9	3.2	0.2	0.0	0	0	74.7		
	Patients #	4.5	1.2	0.2	0.0	0.0	0.0		5.9	2.6
CRM	Selection (per cent)	15.2	0.5	0.1	0.0	0.0	0.0	84.3		
	Patients #	10.7	1.0	0.2	0.0	0.0	0.0		12.0	6.2
Hybrid	Selection (per cent)	14.2	1.5	0.0	0.0	0.0	0.0	84.3		
	Patients #	10.8	1.2	0.0	0.0	0.0	0.0		12.1	6.2
CRM _{True}	Selection (per cent)	9.6	0.4	0.0	0.0	0.0	0.0	90.0		
	Patients #	8.1	0.9	0.1	0.0	0.0	0.0		9.1	4.7

CRM_{True} denotes the CRM design with the true dose-toxicity model. The target dose is in boldface.

Table II

Percentages of the CRM-guided trials that make inappropriate decisions of dose assignment.

Dose action	Tox/pat	Scenario					
		1	2	3	4	5	6
Escalation	2/3	7.4	12.9	8.0	0.0	12.7	0.0
De-escalation	0/6	0.6	0.2	0.1	0.0	0.0	0.0
	1/6	5.1	2.8	0.8	0.2	0.4	0.2
Staying the same	1/9	0.1	0.1	0.0	0.0	0.0	0.0
	3/3	2.5	3.2	1.8	0.0	4.6	0.0
	5/6	0.0	0.3	0.5	0.0	1.9	0.0
	0/6	6.5	4.0	3.0	0.7	2.7	0.7
	0/9	0.5	0.2	0.3	0.1	0.2	0.1
	1/9	4.0	2.0	2.3	1.1	1.5	1.1

Sensitivity analysis of the proposed Bayesian hybrid design based on different values of the tolerable margin δ under scenarios 1–3.

Table III

Tolerable margin	Dose level							$N_{\text{MDT+}}$	N_{DLT}	
	1	2	3	4	5	6	None			
<i>Scenario 1</i>										
$\delta = 0.015$	Selection (per cent)	8.8	22.4	49.4	16.3	2.7	0.3	0.2		
	Patients #	9.7	5.3	5.4	3.0	0.6	0.1		3.6	5.1
$\delta = 0.045$	Selection (per cent)	6.0	18.2	53.8	19.3	2.3	0.3	0.2		
	Patients #	6.2	6.4	7.9	2.7	0.8	0.1		3.6	5.6
<i>Scenario 2</i>										
$\delta = 0.015$	Selection (per cent)	4.1	16.4	28.9	29.1	16.4	5.0	0.1		
	Patients #	8.6	5.3	4.2	3.6	1.8	0.5		2.3	4.2
$\delta = 0.045$	Selection (per cent)	2.7	12.2	29.2	33.3	17.6	4.9	0.1		
	Patients #	5.3	6.0	6.1	3.6	2.2	0.8		3.0	4.7
<i>Scenario 3</i>										
$\delta = 0.015$	Selection (per cent)	1.8	7.8	8.7	24.8	38.2	18.6	0.0		
	Patients #	7.8	4.4	3.4	3.5	3.1	1.8		1.8	3.4
$\delta = 0.045$	Selection (per cent)	1.3	4.4	8.3	24.0	41.5	20.5	0.0		
	Patients #	4.6	4.9	4.3	3.6	3.9	2.7		2.7	4.0

The target dose is in boldface.