

Dose-finding in phase I clinical trials based on toxicity probability intervals

Yuan Ji¹, Yisheng Li², B. Nebiyu Bekele²

1 Department of Bioinformatics and Computational Biology

2 Department of Biostatistics

University of Texas M. D. Anderson Cancer Center, Houston, TX 77030

Email: *yuanji@mdanderson.org*

ABSTRACT

Background: Most phase I clinical trials conducted at the M. D. Anderson Cancer Center use the algorithmic 3+3 design, despite the availability of more advanced model-based designs such as the continual reassessment method.

Purpose: Through simple statistical modeling and computing, we develop a dose-finding design that can be easily understood and implemented by nonstatisticians.

Methods: We propose a beta/binomial Bayesian model and a probabilistic up-and-down rule that allow all possible dose-assignment actions to be tabulated in a spreadsheet. We have developed an Excel macro (available at <http://odin.mdacc.tmc.edu/~yuanj/tm.xls>) that generates trial monitoring tables, which contain the dose-assignment actions corresponding to various toxicity outcomes.

Results: The new design outperforms the 3+3 design and performs comparably to other model-based methods in the literature.

Limitations: The proposed method assumes that the observed toxicity is a binary variable and that toxicity increases with dose level.

Conclusions: The new dose-finding design enables physicians to readily determine dose assignments for new patients by referencing a trial monitoring table.

KEY WORDS: Dose escalation; Dose de-escalation; Maximum tolerated dose; Trial monitoring.

1 Introduction

Phase I clinical trials in oncology are designed to locate the maximum tolerated dose (MTD): the dose for which the probability of toxicity is closest to a prespecified target toxicity rate p_T . A rich literature on dose-finding has been developed over the past two decades and many advanced methods have been proposed for designing phase I oncology trials. These methods include the continual reassessment method (CRM) by O'Quigley, Pepe, and Fisher (1990) and its various extensions; the efficient dose escalation with overdose control (EWOC) by Babb, Rogatko, and Zacks (1998); the curve-free method (CFM) by Gasparini and Eisele (2000); the biased-coin design (BCD) with isotonic regression estimator by Stylianou and Flournoy (2002); and the decision theoretic designs by Haines, Perevozskaya, and Rosenberger (2003), among many others.

This research project would be unnecessary if most of the phase I oncology trials were designed using one of the methods discussed above. However, between April and May of 2006, 20 of 22 phase I clinical trials submitted for Institutional Review Board (IRB) approval at the M. D. Anderson Cancer Center used the 3+3 design. In fact, most ongoing phase I trials conducted at M. D. Anderson are based on the 3+3 design, despite its well-documented undesirable performance (e.g., Thall and Lee, 2003).

This observation motivated us to examine why most investigators still prefer the 3+3 design. Our research revealed two primary reasons for this preference. First, the 3+3 design is very easy to understand and implement. In fact, it is so simple that investigators can design their trials without any knowledge of statistics. In contrast, investigators usually do not understand the statistical methodology behind most advanced designs such as the CRM and therefore do not want to implement such methods for their trials. Second, and more specific to our own experience, the large volume of phase I trials at M. D. Anderson makes it

impossible to provide every phase I trial with a statistician who can implement sophisticated designs. Although the Department of Biostatistics at M. D. Anderson has provided investigators a graphical user interface (GUI) to monitor trials using the CRM, many investigators and their research administrators remain reluctant to use the CRM because of the perceived black-box nature of the method. Consequently, they avoid using this and other model-based methods in their trials.

Our goal is to develop a method that can be easily understood and implemented by clinical investigators, is relatively easy to explain to these investigators and their research teams, and performs well. By combining the up-and-down rule embedded in the 3+3 method with a simple probability model, we develop a method in which all possible dose-assignment actions are tabulated via a spreadsheet, thus allowing investigators to monitor their dose-finding trials without any subsequent statistical input. In addition, as noted by the Associate Editor, a serious flaw with the 3+3 design is that investigators cannot choose a target toxicity probability. Our design allows investigators to specify target toxicity probabilities that are clinically relevant for the disease under study.

The proposed method consists of two components: a beta/binomial model and a dose-assignment rule based on posterior toxicity probabilities. First, we adapt a conjugate beta/binomial Bayesian model that gives us closed-form beta posterior distributions for the dose toxicity probabilities. Second, using the beta posteriors, we compute posterior probabilities of three toxicity intervals that are associated with high, acceptable, and low toxicity for the dose at which patients are being treated; we associate each interval with a different dose-assignment action for treating future patients. For example, if at a dose the interval associated with high toxicity has the largest posterior probability, then the dose is likely to be highly toxic and we take a de-escalation action, treating future patients at a lower

dose level. By combining the beta/binomial model and the proposed dose-assignment rule based on toxicity probability intervals, we can compute all possible dose-assignment actions and summarize them in a spreadsheet. This allows investigators to see exactly how the rule behaves for all possible outcomes at any given dose.

To facilitate the implementation of our design, we have developed an Excel macro available at <http://odin.mdacc.tmc.edu/~yuanj/tm.xls>. The macro allows the user to define the target toxicity probability and the maximum sample size for a trial. Once these two values are provided, the macro generates a corresponding spreadsheet, which can be used to monitor the trial and determine the dose assignment for each patient enrolled. In addition, An R program for performing simulations is available at <http://odin.mdacc.tmc.edu/~yuanj/ppi.R>, which will provide operating characteristics of our method for different toxicity scenarios.

The remainder of the article is organized as follows. Section 2 introduces the probability model including the beta prior and the binomial likelihood. Section 3 describes the proposed dose-assignment rule. Section 4 explains how to use the proposed rule to monitor a phase I clinical trial. Section 5 presents simulation results comparing the operating characteristics of the new method with those of four other methods, the 3+3, the BCD, the CFM and the CRM. Also in this section, we provide an example in which the proposed method is implemented for a practical trial. Section 6 ends the article with conclusions and discussions.

2 Probability model

Define p_T to be the target toxicity probability of the MTD (e.g., $p_T = 0.3$). The goal of phase I clinical trials is to find the highest dose with a toxicity probability closest to p_T . Let p_i denote the toxicity probability for dose $i = 1, \dots, d$, where d is the total number of candidate doses in the trial. The observed data include the n_i patients treated at dose i and

the corresponding x_i experiencing toxicity. Let $\mathbf{p} = (p_1, \dots, p_d)'$. The likelihood function is a product of binomial densities,

$$l(\mathbf{p}) \propto \prod_{i=1}^d p_i^{x_i} (1 - p_i)^{n_i - x_i}. \quad (1)$$

To correctly locate the MTD by assigning patients to appropriate doses in the dose-finding process, the proposed statistical model needs to be able to accurately estimate the probability of toxicity at the current tried dose, the dose currently used to treat patients. When little information is known about the toxicity of the candidate doses, we propose to use models with vague priors for p_i so that the shape of the resulting posterior distributions will be decided mainly by the shape of the likelihood based on the observed data. Models of this type include the priors used in Yin, Li, and Ji (2006) and a set of vague independent beta priors. Because the latter is much easier to implement, we use it in our method. Assume that the priors of p_i are i.i.d. $B(0.005, 0.005)$. Combined with the likelihood in (1), the posterior of p_i follows independent $B(0.005 + x_i, 0.005 + n_i - x_i)$, for $i = 1, \dots, d$. The beta prior $B(0.005, 0.005)$ has been shown to be noninformative with respect to the Bernoulli family (Zhu and Lu, 2004). More specifically, given the U shape of the prior $B(0.005, 0.005)$, the posterior estimate of p_i is very close to the observed toxicity proportion at dose i when $0 < x_i < n_i$; the posterior estimate of p_i is close to 0 or 1 when $x_i = 0$ or $x_i = n_i$, respectively. We exploit this property in the next section when we introduce the proposed dose-assignment rules.

However, when strong prior information on the toxicity of the candidate doses is available, informative beta priors can replace the vague priors given above. For example, suppose an initial clinical trial explored various doses and found none to be toxic. The highest dose had various patients enrolled, none of whom experienced toxicity. A subsequent phase I trial is conducted in which the first and second doses are the second highest and highest doses from

the first trial. In this example, we can incorporate the information on the two doses from the first trial into their priors in the second trial.

3 Methodology

3.1 Dose-assignment rule

We consider the following up-and-down idea: the decision to escalate to a higher dose, stay at the same dose, or de-escalate to a lower dose is based on the estimated probability of toxicity at the current tried dose. In general, if the current dose is well tolerated, the next cohort is treated at a higher dose; if the current dose is too toxic, the next cohort is treated at a lower dose; otherwise, the next cohort is treated at the current dose.

Specifically, suppose only a change of one dose level is allowed in the trial. If patients are being treated at dose i , there are only three possible actions: de-escalate (D) to dose $(i - 1)$; stay (S) at dose i ; or escalate (E) to dose $(i + 1)$. Consider a partition of the unit interval $(0,1)$ $\Delta = \{(0, p_T - K_1\sigma_i), [p_T - K_1\sigma_i, p_T + K_2\sigma_i], (p_T + K_2\sigma_i, 1)\}$, where σ_i is the posterior standard deviation of p_i and K_1 and K_2 are some small positive constants such that $0 < p_T - K_1\sigma_i < p_T + K_2\sigma_i < 1$. The three intervals represent the situations in which dose i has low, acceptable, and high toxicity, respectively. If the posterior distribution of p_i puts most of the mass at the first interval $(0, p_T - K_1\sigma_i)$, then p_i is most likely smaller than p_T , which implies that the dose should be escalated. Likewise, if the posterior distribution puts most of the mass at the third interval $(p_T + K_2\sigma_i, 1)$, then p_i is most likely greater than p_T , which implies that the dose should be de-escalated. Finally, if the posterior distribution puts most of the mass at the middle interval $[p_T - K_1\sigma_i, p_T + K_2\sigma_i]$, then p_i is close to p_T , which implies that dose i is near the true MTD and should not be changed. Based on the proposed

U-shape prior, the posterior estimates of p_i will be close to 0 or 1 if no or all patients experience toxicity, respectively. Therefore, when we observe that no patients experience toxicity at dose i , the posterior distribution of p_i puts most of the mass at the first interval, leading to dose escalation. Likewise, when we observe that all patients experience toxicity, the posterior distribution of p_i puts most of the mass at the third interval, leading to dose de-escalation. Finally, when we observe that some but not all patients experience toxicity, the posterior distribution is dominated by the likelihood, allowing the dose assignments to be decided by the observed binomial data.

We translate the above ideas into a formal mathematical decision formula. Given any prior distribution for \mathbf{p} with a proper density function $\pi(\mathbf{p})$, we define the posterior probabilities of the three intervals in partition Δ as

$$q(D, i) = P(p_i - p_T > K_1\sigma_i | \text{data}), \quad (2)$$

$$q(S, i) = P(-K_2\sigma_i \leq p_i - p_T \leq K_1\sigma_i | \text{data}), \quad (3)$$

$$q(E, i) = P(p_i - p_T < -K_2\sigma_i | \text{data}). \quad (4)$$

A dose-assignment rule \mathcal{B}_i based on these three probabilities is given by

$$\mathcal{B}_i = \arg \max_{m \in \{D, S, E\}} q(m, i), \quad (5)$$

which chooses the action corresponding to the interval with the largest posterior probability. However, due to ethical considerations, the above rule is incomplete. Specifically, \mathcal{B}_i would allow an escalation action as long as the current dose is not toxic, i.e., $q(E, i)$ is the largest among the three q 's. Such escalation is not safe and should not be allowed if the data suggest that the next higher dose is very likely to be highly toxic. Consequently, we modify the dose-assignment rule \mathcal{B}_i to account for this. We call this modification a toxicity exclusion rule

and base it on a random variable

$$\mathcal{T}_i = 1 \{P(p_i > p_T | \text{data}) > \xi\},$$

where $1\{\cdot\}$ is the indicator function and $\xi \in (0, 1)$ is a cutoff value (e.g., $\xi = 0.95$). For a large value of ξ , $\mathcal{T}_i = 1$ implies that dose i is very likely to be highly toxic and escalation to this dose should be permanently prohibited. Suppose the current tried dose is i ; if $\mathcal{T}_{i+1} = 1$, then escalation from dose i to $(i + 1)$ must not be allowed. We incorporate \mathcal{T}_{i+1} into the proposed dose-assignment rule \mathcal{B}_i as follows: Let $q(\tilde{E}, i) = q(E, i)(1 - \mathcal{T}_{i+1})$ and define the new dose-assignment rule with toxicity exclusion to be

$$\mathcal{B}_i^{(e)} = \arg \max_{m \in \{D, S, \tilde{E}\}} q(m, i).$$

Therefore, if $\mathcal{T}_{i+1} = 1$, the quantity $q_{\tilde{E}, i}$ equals 0 and the assignment rule $\mathcal{B}_i^{(e)}$ can be only the action D , to de-escalate, or the action S , to stay, whichever has the larger probability. Under the proposed beta/binomial model, it is easy to show that the action will always be S . A special case is when $\mathcal{T}_1 = 1$, which indicates that the first dose and, by implication, all higher doses are highly toxic and that the trial should be terminated immediately. There is one exception for trials with cohorts of size 1. Under the proposed prior $B(0.005, 0.005)$, we do not apply the exclusion rule until two patients have been treated at a dose. This is a limitation of the vague prior since it will exclude a dose when only one patient is treated and the patient experienced toxicity. In practice, excluding a dose with only one out of one toxicity is deemed unreasonable. Therefore, we impose in our dose-finding algorithm (introduced below) that the exclusion rule only applies when toxicity outcomes of at least two patients have been observed at a dose.

When applying the exclusion rule, it is important to check if an untried dose will be excluded. Under our model, when a dose is untried, one computes \mathcal{T}_{i+1} according to the

proposed beta prior $B(0.005, 0.005)$. For a large ξ (say > 0.5), the resulting \mathcal{T}_{i+1} will always equal 0. Therefore, our model will not exclude an untried dose. If a different prior is used, as long as that dose ($i + 1$) is not considered highly toxic *a priori*, i.e., $P(p_{i+1} > p_T) \leq \xi$ (where the probability is computed with respect to the prior distribution of p_{i+1}), the stopping rule \mathcal{T}_{i+1} will equal 0, and therefore not exclude the untried dose ($i + 1$). If the prior is constructed in such a way that dose ($i + 1$) is highly toxic *a priori*, then either this dose should not be included in the trial or the prior is not appropriate and should be modified.

Our escalation/de-escalation rules only consider the toxicity of the current tried dose, which changes as the trial proceeds. Suppose the current tried dose is the j th dose. If dose j is very toxic, we de-escalate to dose ($j - 1$), which becomes our current tried dose for the next decision. If dose j is not toxic, we escalate to dose ($j + 1$), which becomes our current tried dose for the next decision. If dose j has similar toxicity to p_T , we continue to treat it as our current tried dose. Under these rules, if the current tried dose is j (assuming j is not the first or last dose), then it must be true that dose ($j - 1$) is safe or dose ($j + 1$) is too toxic if it has been tried; because to be at dose j , one must have escalated from dose ($j - 1$) or de-escalated from dose ($j + 1$). Based on this rationale, we can see that although our escalation/de-escalation rules only use the data at the most recently tried dose, they implicitly and sequentially use toxicity information from other doses. When these escalation and de-escalation decision rules are performed sequentially, it is impossible to have cases in which a lower dose level is highly toxic and a higher dose level is highly nontoxic because to reach the higher dose, all lower doses must be safe to warrant the escalation decisions. We acknowledge that other rules (such as the minimum distance rule in the CRM) in the literature perform very well in most scenarios as will be shown in the simulation study discussed in Section 5.1. However, one advantage of our rules is that they are based

on very simple models and can be tabulated before the trial starts, which is attractive to nonstatisticians.

One should consider K_1 and K_2 as tuning parameters that allow for flexible dose-finding strategies. Based on the partition Δ , it is easy to see that $K_2 > K_1$ leads to a conservative design that favors de-escalation over escalation. This is desirable when patient safety is critical and is recommended for phase I trials in oncology. In contrast, $K_1 > K_2$ leads to an aggressive design that favors escalation, which is suitable for trials in which most toxicity events are acceptable and fast escalation to the MTD is preferred. In our simulation studies, we will use $K_1 = 1$ and $K_2 = 1.5$ which in general give reasonable results for the scenarios studied. Therefore, we recommend these two values as the default choices for practical trials. However, if one wants to calibrate K_1 and K_2 , we recommend to try values in the interval $(1, 2)$ based on our sensitivity analysis (results not shown). In general, a different set of values will usually improve the performance of our method for some scenarios while worsen it for others. The K_1 and K_2 which best correspond to the preferences of the investigator should then be used.

The partition Δ accounts for the posterior variability by using the posterior standard deviation σ_i of p_i in the three intervals. This allows the length of the three intervals to vary with respect to the posterior variability in p_i , which is critical to the success of the proposed rule. In contrast, a partition such as $\{(0, p_T - 0.1), [p_T - 0.1, p_T + 0.1], (p_T + 0.1, 1)\}$ may not be desirable when only a few patients are treated at dose i . Suppose the posterior is centered at p_T , but the posterior standard deviation is 0.4. Then the redefined $q(S, i)$ based on the above partition could be smaller than $q(E, i)$ or $q(D, i)$, even when the posterior of p_i is centered at p_T . Consequently, the rule \mathcal{B}_i would wrongly choose to escalate or de-escalate.

3.2 The dose-finding algorithm

With the proposed dose-assignment rules and the available Excel macro, physicians can easily determine the appropriate dose assignment for each new patient throughout the trial. At the end of the trial when the toxicity outcomes of all the enrolled patients are observed, a dose will be selected as the estimated MTD for subsequent studies. We propose to select the MTD by performing an isotonic regression procedure that borrows strength across doses. Specifically, we first compute the posterior mean \hat{p}_i under the beta posterior distribution and then perform the pooled adjacent violators algorithm (PAVA) (Robertson, Wright, and Dykstra, 1988) on \hat{p}_i so that the resulting transformed values \hat{p}_i^* increase with the dose levels. That is, $\hat{p}_j^* \geq \hat{p}_i^*$ for $j > i$. Note that the computation using the PAVA is straightforward and many software packages such as R have build-in functions for implementation. In addition, this computation only needs to be performed when the trial is finished. Therefore, it does not affect our proposed dose-assignment rules.

Based on the dose-assignment rule $\mathcal{B}_i^{(e)}$ and the transformed and order-restricted estimates \hat{p}_i^* , we propose a dose-finding algorithm as follows:

- Suppose that the current tried dose is i , $i \in \{1, \dots, d\}$. After the toxicity outcomes of the last cohort are observed, select the dose for treating the next cohort among $\{(i-1), i, (i+1)\}$ based on the assignment rule $\mathcal{B}_i^{(e)}$. There are two exceptions: if $i = 1$, the next available doses are $\{i, (i+1)\}$; if $i = d$, the next available doses are $\{(i-1), i\}$.
- Suppose that dose 1 has been tried previously. If $\mathcal{T}_1 = 1$, terminate the trial due to excessive toxicity. Otherwise, terminate the trial when the maximum sample size is reached. In the special case of cohorts of size 1, do not apply the exclusion rule \mathcal{T}_i until

two or more patients have been evaluated at a dose.

- At the end of the trial, select the dose as the estimated MTD with the smallest difference $|\hat{p}_i^* - p_T|$ among all the tried doses i for which $\mathcal{T}_i = 0$. If two or more doses tie for the smallest difference, perform the following rule. Let p^* denote the transformed posterior mean \hat{p}_i^* of the tied doses.
 - If $p^* < p_T$, choose the highest dose among the tied doses.
 - If $p^* > p_T$, choose the lowest dose among the tied doses.

4 Trial monitoring

Given the values of p_T , n_i , and x_i , determining the assignment rule $\mathcal{B}_i^{(e)}$ only requires computing and comparing the three posterior probabilities $q(D, i)$, $q(S, i)$, and $q(E, i)$. These three probabilities can be easily obtained based on the posterior distribution of p_i , $B(0.005 + x_i, 0.005 + n_i - x_i)$. Therefore, one can compute $\mathcal{B}_i^{(e)}$ for any values of n_i and x_i and tabulate the results. Moreover, these computations can be done before the trial begins and are the same for any dose i . Therefore, we developed an Excel macro that provides users a table of dose-assignment actions. For example, we used our Excel macro with $p_T = 0.3$ and sample size 12 to generate a trial monitoring table, part of which is presented in Table 1 containing the dose-assignment actions when $n_i = 3, 6, 9$, or 12 patients have been treated at a given dose. If dose i is the current tried dose, then one can use the values of (x_i, n_i) to determine the action to be taken for treating the next cohort of patients. For instance, when $x_i = 2$ and $n_i = 2$, the corresponding table entry is the action S , which indicates that the next cohort of patients should be treated at dose i . The label DU for dose i means two things: D indicates that the dose-assignment action is to de-escalate to $(i - 1)$, and U

indicates that dose i is unacceptably toxic, i.e., $\mathcal{T}_i = 1$, and should not be used again in the trial. Therefore, label U accommodates the toxicity exclusion rule.

To further illustrate the use of our algorithm and the trial monitoring table, consider a hypothetical trial with a target toxicity probability 0.3 and cohorts of size 3. In this example, the trial starts by treating three patients at dose 1 and no toxicity is observed. The posterior distribution of p_1 is $B(0.005, 3.005)$. By comparing the posterior probabilities $q(m, 1)$, $m \in \{D, S, E\}$, we find that the probability $q(E, 1)$ is the largest among the three. Consequently, the action is $\mathcal{B}_1 = E$, to escalate, which corresponds to the action in Table 1 with 0 toxicities out of 3 patients. Therefore the trial proceeds by treating three more patients at dose 2. Suppose by chance that no toxicity is observed out of these three patients. Through the same type of posterior computation, we find that the action is E to escalate again and treat three more patients at dose 3. Subsequently, we put three patients at dose 3 and observe two toxicities. We compute the posterior probabilities $q(D, 3)$, $q(S, 3)$, and $q(E, 3)$ under the posterior of p_3 , which is $B(2.005, 1.005)$, and find that $q(D, 3)$ is the largest. Therefore, the action is D to de-escalate to dose 2, which corresponds to the entry in Table 1 for 2 toxicities out of 3 patients. Following this action, we treat three additional patients at dose 2 and observe one toxicity. Combining these patients with the initial three patients treated at dose 2, we have one toxicity out of a total of six patients at this dose. Therefore, the posterior of q_2 is $B(1.005, 5.005)$. Based on the posterior computation, the action is S , corresponding to the entry in Table 1 for 1 toxicity out of 6 patients, i.e., we need to stay at dose 2 and put more patients on this dose. The trial proceeds by treating six more patients at dose 2 and eventually selects dose 2 as the recommended MTD based on our method.

Note that if three toxicities out of three patients had been observed at dose 3, the posterior computation would have resulted in $\mathcal{T}_3 = 1$; Table 1 would have indicated the action DU , in

which D implies to de-escalate to dose 2 and U implies that dose 3 is unacceptably toxic and no more patients should be treated at this dose for the remainder of the trial. Consequently, we would not have allowed escalation from dose 2 to dose 3 again. In general, if dose i is the current tried dose and the trial monitoring table indicates the action E , one must check whether the table entry contains the action U for x_{i+1} toxicities out of n_{i+1} patients. If it is, then the action E at dose i must be changed to the action S because dose $(i + 1)$ is too toxic.

As the above hypothetical example shows, all the dose-assignments are obtained through posterior computations based on our beta/binomial model. However, instead of performing the computations to obtain appropriate actions each time a cohort has been evaluated, we can obtain these actions from Table 1 directly. Through the above example, we explained how one can use our trial monitoring table (e.g., Table 1) to perform dose assignment actions without additional computations. Compared to the 3+3 method which does not require computations either in deciding dose-assignments, our algorithm does not terminate a trial if two out three toxicities are observed (it will de-escalate) or two out of six toxicities are observed (it will stay at the current dose). In both cases, the 3+3 method would terminate the trial.

5 Examples

The proposed dose-finding algorithm is being implemented in clinical trials at M. D. Anderson, including a phase I trial studying the toxicity of a combination of two agents on patients with non-small cell lung cancer and another investigating the toxicity of a combination of a chemotherapy and a new agent on patients with lymphoma. However, these trials involve at most six doses. For generality, we compare our method with other existing methods in a

setting with more doses.

5.1 Simulation study

Using a clinical trial described by Goodman, Zahurak, and Piantadosi (1995), we perform simulations to compare the proposed method with two algorithmic methods, the 3+3 and the BCD, and two model-based methods, the CFM and the CRM.

The 3+3 method is the most popular method in practice. Since our target toxicity probability $p_T = 0.3$ in this simulation study, we used a version of 3+3 in which two out of six toxicities observed at a dose result in selection of that dose as the MTD (Thall and Lee, 2003).

The BCD is a sequential design (i.e., cohorts of size one) based on the following dose-assignment algorithm: when $p_T < .5$, which is the case considered in this study, the BCD steps down a dose if toxicity is observed in the previous patient, and randomizes with probability $h(p_T) = p_T/(1 - p_T)$ to the next higher dose and $\{1 - h(p_T)\}$ to the same dose if no toxicity is observed in the previous patient. The BCD uses the isotonic regression estimator with linear interpolation to estimate toxicity probabilities. For detailed discussions of this version of the BCD, see Stylianou and Flournoy (2002).

Based on a transformation that induces monotonicity among the toxicity probabilities, the CFM assumes unimodal beta prior distributions on the transformed parameters. Closed-form solutions are obtained for the posterior mean probabilities of toxicity \hat{p}_i , with which the CFM allocates the next cohort to the dose that minimizes the absolute difference $|\hat{p}_i - p_T|$.

The CRM skeleton $\boldsymbol{\phi} = (\phi_1, \dots, \phi_d)$ is a set of prespecified and fixed toxicity probabilities with the constraint $\phi_1 < \phi_2 < \dots < \phi_d$. According to Shen and O'Quigley (1996), the CRM models the probability of toxicity at the i th dose as $p_i = \phi_i^{\exp(-\beta)}$; we assume that β follows

a normal prior distribution with mean 0 and standard deviation 2. Note that there are many versions of the CRM (e.g., Piantadosi, Fisher, and Grossman, 1998) and it may be that some of these models perform better than the one-parameter model used here. Like the CFM, the CRM allocates the next cohort to the dose that minimizes $|\hat{p}_i - p_T|$. To obtain \hat{p}_i , the CRM integrates $\phi_i^{\exp(-\beta)}$ with respect to the posterior distribution of β . We evaluated several different CRM skeletons ϕ and found that the operating characteristics of the CRM were sensitive to the choice of ϕ (results not shown). We present the simulation results for the skeleton that provided the highest probabilities of selecting the MTD, which corresponds to $\phi_i = i * 0.05$ for $i = 1, \dots, d$. We also found that the hyperparameter values in the beta prior affect the performance of the CFM, and the results presented here are the best we obtained after trying several different values.

For our design, we set $\xi = 0.95$, $K_1 = 1$, and $K_2 = 1.5$. The value of ξ can be decreased if one wants to exclude a dose more easily. For example, we used $\xi = 0.70$ in a trial at M. D. Anderson so that a dose would be excluded if the posterior probability that its toxicity rate is greater than p_T is greater than 0.7. In general, a smaller ξ leads to a more conservative design using our method.

For comparison purposes, we modified the standard BCD, CFM, and CRM by using an additional safety stopping rule equivalent to calculating \mathcal{T}_1 under our method. To apply this additional rule for the three methods, we assumed that the number of toxicity events for dose 1 follows a binomial distribution with $B(0.005, 0.005)$ prior. We then computed \mathcal{T}_1 and terminated the trial if $\mathcal{T}_1 = 1$. Note that other than this slight modification, no other aspects of these three methods were changed. In addition, we only allow the BCD to stop the trial if two or more patients' outcomes have been observed at dose 1, which is the same rule as what we imposed for our dose-finding algorithm. There are eight doses ($d = 8$) available at

levels of 50, 100, 200, 300, 400, 500, 650, and 800 mg/day. The cohort size is three (except for the BCD) and $p_T = 25\%$. The first cohort is treated at the lowest dose level, 50 mg/day. The maximum number of patients allowed in the trial is 30. For all the methods except the 3+3, the trial is stopped when the maximum sample size is reached or when $\mathcal{T}_1 = 1$.

We simulated 1,000 trials. Table 2 summarizes the results from six scenarios and is organized by scenario sections. Each scenario section contains 11 rows. The first row of each scenario section is the true toxicity probability for each dose, from which we generated the trial data. The next 10 rows are the percentages of times the i th dose was selected as the MTD in 1,000 simulations, and the average number of patients treated at dose i using the proposed Bayesian design (denoted as “Bayes”), the 3+3, the BCD, the CFM, and the CRM, respectively. The column “none” contains the percentages of not selecting any of the doses as the MTD due to excessive toxicity in a given scenario. The last two columns give the overall toxicity percentage of the 1,000 simulations and the average number of patients used for one trial. Scenarios 1 and 2 assume that the true MTDs are doses 2 and 7, respectively. The probability of correctly selecting the true MTD decreases for all five methods over these two scenarios. It seems that the proposed method, the BCD, the CFM, and the CRM performed similarly well in these scenarios; the 3+3 performed much worse. Scenarios 3 and 4 are special but important cases. In Scenario 3, dose 2 is very nontoxic and dose 3 is very toxic. Surprisingly, the 3+3 method performed very well in this scenario. Because of the proposed exclusion rule \mathcal{T}_{i+1} (which applies to dose 3 in this scenario), the proposed method also performed well. In Scenario 4, the first dose is too toxic and the trial should be terminated as early as possible without declaring any dose as the MTD; here, all the methods performed reasonably well. In Scenarios 5 and 6, the toxicity probabilities increase gradually with the dose levels. The CRM performed well for these two scenarios since the CRM skeletons used

here were close to the true probabilities of toxicity. The proposed method, the BCD and the CFM performed comparably with the CRM; the 3+3 performed a little worse.

5.2 Illustrative trial

As an illustrative example of our algorithm, we use a practical trial at M. D. Anderson for patients with non-small cell lung cancer. In this trial, a new agent in a family of nonreceptor tyrosine kinases is found to be promising in treating patients with non-small cell lung cancer. It is believed that a combination of the new agent and a second and third line standard therapy will have synergistic effects on the patients, i.e., the effect of the combination will be greater than the added effects of the two treatments. A phase I dose-finding clinical trial is being planned to establish the safety of the combination. Specifically, the dosage of the standard therapy is fixed at 150 mg/day, and four doses are prepared for the new agent at the levels of $\{0.25, 0.5, 1.0, 1.5\}$ units. Starting at the lowest dose, the trial will enroll patients in cohorts of size three. The maximum sample size is 24 and the MTD has a probability of toxicity $p_T = 0.3$. Table 3 presents the operating characteristics of the proposed design for this trial under six toxicity scenarios. The operating characteristics using the proposed design are reasonable.

6 Discussion

We have proposed a new dose-finding algorithm based on a simple statistical model and a new dose-assignment rule. The new method performed well in our simulation studies, and its implementation and computation were simple. This simplicity should appeal to physicians, who ultimately decide which dose-finding method to use in their trials. To implement the

proposed dose-assignment decisions and perform simulations using our method, one simply needs to specify the value of p_T , the number of doses, the maximum sample size, and the true dose probabilities of toxicity. For trial monitoring, given observed interim data, the dose-assignment actions can be directly obtained from the provided software.

The new dose-assignment rule could be easily extended to many other applications. For example, one can apply a similar dose-assignment rule to trials in which toxicity is considered an ordinal response. Bekele and Thall (2004) proposed a complicated model for this type of trial, but one could use a simple model, e.g., a vague Dirichlet prior coupled with a similar dose-assignment rule, to solve the problem. Other potential applications include finding the best schedule, as in Braun, Yuan, and Thall (2005) and dose-finding based on toxicity and efficacy. When these extensions are realized, one gains the ability to solve these problems and obtain comparable results through simple statistical modeling and computing, which investigators will easily accept and use.

References

- Babb, J., Rogatko, A. and Zacks, S. (1998), “Cancer phase I clinical trials: efficient dose escalation with over dose control”, *Statistics in Medicine* **17**, 1103–1120.
- Bekele, B. and Thall, P. (2004), “Dose-finding based on multiple toxicities in a soft tissue sarcoma trial.”, *Journal of the American Statistical Association* **60**, 684–693.
- Braun, T., Yuan, Z. and Thall, P. (2005), “Determining a maximum-tolerated schedule of a cytotoxic agent.”, *Biometrics* **61**, 335–343.
- Gasparini, M. and Eisele, J. (2000), “A curve-free method for phase I clinical trials”, *Biometrics* **56**, 609–615.

- Goodman, S., Zahurak, M. and Piantadosi, S. (1995), “Some practical improvements in the continual reassessment method for phase I studies”, *Statistics in Medicine* **14**, 1149–1161.
- Haines, L., Perevozskaya, I. and Rosenberger, W. (2003), “Bayesian optimal designs for phase I clinical trials”, *Biometrics* **59**, 591–600.
- O’Quigley, J., Pepe, M. and Fisher, L. (1990), “Continual reassessment method: A practical design for phase 1 clinical trials in cancer”, *Biometrics* **46**, 33–48.
- Piantadosi, S., Fisher, J. and Grossman, S. (1998), “Practical implementation of a modified continual reassessment method for dose finding trials”, *Cancer Chemotherapy and Pharmacology* **41(6)**, 429–436.
- Robertson, T., Wright, F. and Dykstra, R. (1988), *Order Restricted Statistical Inference*, Wiley, New York.
- Rosenberger, W. and Haines, L. (2002), “Competing designs for phase I clinical trials: a review”, *Statistics in Medicine* **21**, 2757–2770.
- Shen, L. and O’Quigley, J. (1996), “Consistency of continual reassessment method under model misspecification”, *Biometrika* **83**, 395–405.
- Stylianou, M. and Flournoy, N. (2002), “Dose finding using the biased coin up-and-down design and isotonic regression”, *Biometrics* **58**, 171–177.
- Thall, P. and Lee, S. (2003), “Practical model-based dose-finding in phase I clinical trials: methods based on toxicity”, *International Journal of Gynecological Cancer* **13(3)**, 251–261.

Yin, G., Li, Y. and Ji, Y. (2006), “Bayesian dose-finding in phase I/II clinical trials using toxicity and efficacy odds ratios”, *Biometrics* **62**, 777–797.

Zhu, M. and Lu, A. (2004), “The counter-intuitive non-informative prior for the Bernoulli family”, *Journal of Statistics Education* **12(2)**, 1–10.

Table 1: Dose-assignment actions based on the number of patients treated and the number of patients who experienced toxicity at dose i . The target toxicity probability of the MTD $p_T = 0.3$. Actions D , S , and E represent de-escalation to dose $(i - 1)$, staying at dose i , and escalation to dose $(i + 1)$. U implies that dose i is unacceptable and should be excluded from the trial. The table entries determine the dose for treating the next cohort of patients. One exception is that if the action is E for the current tried dose i and if the trial data at dose $(i + 1)$ indicate that dose $(i + 1)$ is unacceptable (i.e., U) according to this table, the action must be changed to S , to stay at dose i , because dose $(i + 1)$ is too toxic.

| Number of patients | Dose-assignment actions | | | |
|----------------------|-------------------------|------|------|------|
| | 3 | 6 | 9 | 12 |
| Number of toxicities | | | | |
| 0 | E | E | E | E |
| 1 | S | S | E | E |
| 2 | D | S | S | E |
| 3 | DU | D | S | S |
| 4 | | DU | S | S |
| 5 | | DU | D | S |
| 6 | | DU | DU | D |
| 7 | | | DU | DU |
| 8 | | | DU | DU |
| 9 | | | DU | DU |
| 10 | | | | DU |
| 11 | | | | DU |
| 12 | | | | DU |

Table 2: Simulation results comparing the proposed method, the 3+3, the BCD, the CFM and the CRM. The selection percentages for the true MTDs are in **bold face**.

| | | Recommendation percentage at dose level | | | | | | | | Toxicity | Average | |
|------------|-------|---|-----------|-----|-----|-----|-----|-----------|-----|-------------|-----------|----|
| | | $p_T = 0.25$ | | | | | | | | percentage* | number of | |
| Dose | | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | | patients | |
| Scenario 1 | | 5 | 25 | 50 | 60 | 70 | 80 | 90 | 95 | none | | |
| Bayes | % MTD | 13 | 79 | 8 | 0 | 0 | 0 | 0 | 0 | 0 | 25 | 30 |
| | # Pts | 7.7 | 16.1 | 5.8 | 0.5 | 0 | 0 | 0 | 0 | | | |
| 3+3 | % MTD | 24 | 58 | 16 | 2 | 0 | 0 | 0 | 0 | 0 | 25 | 12 |
| | # Pts | 4.0 | 5.0 | 2.6 | 0.4 | 0 | 0 | 0 | 0 | | | |
| BCD | % MTD | 10 | 78 | 11 | 1 | 0 | 0 | 0 | 0 | 0 | 24 | 30 |
| | # Pts | 11.4 | 11.5 | 5.2 | 1.4 | 0.3 | 0.1 | 0 | 0 | | | |
| CFM | % MTD | 6 | 80 | 14 | 0 | 0 | 0 | 0 | 0 | 0 | 29 | 30 |
| | # Pts | 5.2 | 16.3 | 7.5 | 0.9 | 0 | 0 | 0 | 0 | | | |
| CRM | % MTD | 6 | 83 | 11 | 0 | 0 | 0 | 0 | 0 | 0 | 27 | 30 |
| | # Pts | 5.7 | 18.6 | 4.9 | 1.0 | 0 | 0 | 0 | 0 | | | |
| Scenario 2 | | 1 | 2 | 3 | 4 | 5 | 6 | 25 | 50 | none | | |
| Bayes | % MTD | 0 | 0 | 0 | 0 | 2 | 22 | 62 | 14 | 0 | 12 | 30 |
| | # pt | 3.2 | 3.2 | 3.4 | 3.5 | 3.7 | 4.5 | 5.9 | 2.6 | | | |
| 3+3 | % MTD | 0 | 0 | 0 | 2 | 3 | 21 | 46 | 8 | 0 | 11 | 27 |
| | # pt | 3.1 | 3.2 | 3.3 | 3.4 | 3.3 | 3.7 | 4.5 | 2.2 | | | |
| BCD | % MTD | 0 | 0 | 1 | 2 | 7 | 24 | 56 | 10 | 0 | 10 | 30 |
| | # pt | 3.2 | 3.6 | 3.6 | 3.5 | 3.8 | 5.4 | 4.8 | 2.1 | | | |
| CFM | % MTD | 0 | 0 | 0 | 0 | 1 | 22 | 61 | 16 | 0 | 12 | 30 |
| | # pt | 3.1 | 3.0 | 3.1 | 3.5 | 3.7 | 5.1 | 6.3 | 2.1 | | | |
| CRM | % MTD | 0 | 0 | 1 | 1 | 5 | 22 | 50 | 21 | 0 | 13 | 30 |
| | # pt | 3.1 | 3.4 | 3.3 | 3.7 | 3.6 | 4.4 | 5.1 | 3.4 | | | |

Table 2 (continued)

| | | Recommendation percentage at dose level | | | | | | | | Toxicity percentage* | Average number of patients | |
|--------------|-------|---|-----------|------|-----|-----|----|----|----|-------------------------|----------------------------------|----|
| | | $p_T = 0.25$ | | | | | | | | | | |
| | Dose | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | | | |
| Scenario 3 | | 1 | 5 | 50 | 60 | 70 | 80 | 90 | 95 | none | | |
| Bayes | % MTD | 0 | 79 | 21 | 0 | 0 | 0 | 0 | 0 | 0 | 22 | 30 |
| | # pt | 5.5 | 13.2 | 10.2 | 1.0 | 0 | 0 | 0 | 0 | | | |
| 3+3 | % MTD | 0 | 70 | 28 | 2 | 0 | 0 | 0 | 0 | 0 | 22 | 13 |
| | # pt | 3.1 | 5.2 | 4.4 | 0.7 | 0.1 | 0 | 0 | 0 | | | |
| BCD | % MTD | 0 | 60 | 39 | 1 | 0 | 0 | 0 | 0 | 0 | 22 | 30 |
| | # pt | 4.9 | 14.3 | 8.2 | 2.2 | 0.4 | 0 | 0 | 0 | | | |
| CFM | % MTD | 0 | 56 | 44 | 0 | 0 | 0 | 0 | 0 | 0 | 28 | 30 |
| | # pt | 3.1 | 11.7 | 13.1 | 2.0 | 0.1 | 0 | 0 | 0 | | | |
| CRM | % MTD | 0 | 49 | 51 | 0 | 0 | 0 | 0 | 0 | 0 | 26 | 30 |
| | # pt | 3.1 | 13.0 | 12.0 | 1.8 | 0 | 0 | 0 | 0 | | | |
| Scenario 4** | | 40 | 50 | 60 | 70 | 80 | 90 | 95 | 99 | none | | |
| Bayes | % MTD | 31 | 2 | 0 | 0 | 0 | 0 | 0 | 0 | 67 | 41 | 19 |
| | # pt | 16.8 | 1.8 | 0.2 | 0 | 0 | 0 | 0 | 0 | | | |
| 3+3 | % MTD | 38 | 9 | 1 | 0 | 0 | 0 | 0 | 0 | 52 | 43 | 6 |
| | # pt | 4.7 | 0.5 | 0.6 | 0.7 | 0 | 0 | 0 | 0 | | | |
| BCD | % MTD | 38 | 2 | 0 | 0 | 0 | 0 | 0 | 0 | 60 | 45 | 18 |
| | # pt | 12.6 | 4.6 | 1.2 | 0.2 | 0 | 0 | 0 | 0 | | | |
| CFM | % MTD | 38 | 3 | 1 | 0 | 0 | 0 | 0 | 0 | 58 | 42 | 14 |
| | # pt | 11.7 | 1.9 | 0.5 | 0.1 | 0 | 0 | 0 | 0 | | | |
| CRM | % MTD | 47 | 2 | 0 | 0 | 0 | 0 | 0 | 0 | 51 | 42 | 23 |
| | # pt | 20.2 | 2.5 | 0.2 | 0 | 0 | 0 | 0 | 0 | | | |

Table 2 (continued)

| | | Recommendation percentage at dose level | | | | | | | | Toxicity | Average | |
|------------|-------|---|-----------|-----------|-----|-----|-----|-----|----|-------------|-----------|----|
| | | $p_T = 0.25$ | | | | | | | | percentage* | number of | |
| Dose | | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | | patients | |
| Scenario 5 | | 15 | 25 | 35 | 45 | 55 | 65 | 75 | 85 | none | | |
| Bayes | % MTD | 31 | 41 | 21 | 7 | 0 | 0 | 0 | 0 | 0 | 24 | 30 |
| | # pt | 12.4 | 9.5 | 5.5 | 1.9 | 0.3 | 0 | 0 | 0 | | | |
| 3+3 | % MTD | 29 | 37 | 20 | 7 | 1 | 0 | 0 | 0 | 8 | 26 | 12 |
| | # pt | 4.4 | 3.9 | 2.4 | 0.9 | 0.2 | 0 | 0 | 0 | | | |
| BCD | % MTD | 21 | 46 | 22 | 6 | 1 | 0 | 0 | 0 | 5 | 26 | 29 |
| | # pt | 10.6 | 9.2 | 5.7 | 2.5 | 0.8 | 0.1 | 0 | 0 | | | |
| CFM | % MTD | 15 | 44 | 32 | 7 | 0 | 0 | 0 | 0 | 0 | 27 | 30 |
| | # pt | 8.0 | 10.6 | 8.0 | 2.6 | 0.4 | 0 | 0 | 0 | | | |
| CRM | % MTD | 36 | 47 | 14 | 2 | 0 | 0 | 0 | 0 | 0 | 23 | 30 |
| | # pt | 13.8 | 11.4 | 3.6 | 0.9 | 0.2 | 0 | 0 | 0 | | | |
| Scenario 6 | | 5 | 15 | 25 | 35 | 45 | 55 | 65 | 75 | none | | |
| Bayes | % MTD | 2 | 24 | 42 | 24 | 7 | 0 | 0 | 0 | 0 | 22 | 30 |
| | # pt | 5.1 | 8.2 | 9.2 | 5.7 | 1.6 | 0.3 | 0 | 0 | | | |
| 3+3 | % MTD | 9 | 28 | 34 | 22 | 5 | 0 | 0 | 0 | 0 | 21 | 15 |
| | # pt | 3.6 | 4.3 | 3.8 | 2.3 | 0.8 | 0.2 | 0 | 0 | | | |
| BCD | % MTD | 1 | 29 | 44 | 19 | 6 | 1 | 0 | 0 | 0 | 21 | 30 |
| | # pt | 6.8 | 8.7 | 7.5 | 4.4 | 1.9 | 0.6 | 0.1 | 0 | | | |
| CFM | % MTD | 0 | 14 | 49 | 29 | 6 | 0 | 0 | 0 | 0 | 24 | 30 |
| | # pt | 3.9 | 6.2 | 10.7 | 7.1 | 1.8 | 0.3 | 0 | 0 | | | |
| CRM | % MTD | 4 | 37 | 45 | 12 | 2 | 0 | 0 | 0 | 0 | 20 | 30 |
| | # pt | 5.5 | 11.5 | 8.9 | 3.4 | 0.7 | 0.1 | 0 | 0 | | | |

* Overall % toxicity out of all the simulated trials.

** The same exclusion rule as the one we proposed for our design is implemented for the BCD and CFM so that the trial will stop early upon observation of excessive toxicity at dose 1.

Table 3: Operating characteristics of the proposed method for the non-small cell lung cancer trial. The selection percentages for the true MTD are in **bold** face.

| Dose | Recommendation percentage | | | | | Toxicity percentage* | Average number of patients |
|------------|---------------------------|-----------|-----------|------------|-----------|----------------------|----------------------------|
| | $p_T = 0.30$ | | | | | | |
| | 1 | 2 | 3 | 4 | none | | |
| Scenario 1 | 30 | 45 | 55 | 60 | | | |
| % MTD | 67 | 17 | 2 | 0 | 13 | 34 | 22 |
| # pt | 17.1 | 4.1 | 0.8 | 0.1 | | | |
| Scenario 2 | 10 | 30 | 45 | 55 | | | |
| % MTD | 18 | 58 | 21 | 3 | 0 | 27 | 24 |
| # pt | 8.4 | 10.3 | 4.4 | 0.8 | | | |
| Scenario 3 | 1 | 10 | 30 | 50 | | | |
| % MTD | 0 | 18 | 61 | 20 | 0 | 24 | 24 |
| # pt | 3.5 | 6.6 | 9.6 | 4.3 | | | |
| Scenario 4 | 1 | 5 | 10 | 30 | | | |
| % MTD | 0 | 0 | 16 | 84 | 0 | 17 | 24 |
| # pt | 3.3 | 3.8 | 5.6 | 11.3 | | | |
| Scenario 5 | 50 | 65 | 75 | 80 | | | |
| % MTD | 28 | 0 | 0 | 0 | 72 | 50 | 14 |
| # pt | 13.7 | 0.7 | 0 | 0 | | | |
| Scenario 6 | 1 | 2 | 3 | 5 | | | |
| % MTD | 0 | 0 | 0 | 100 | 0 | 4 | 24 |
| # pt | 3.2 | 3.2 | 3.3 | 14.3 | | | |

* Overall % toxicity out of 1,000 simulations