HHS Public Access

Author manuscript

Stat Med. Author manuscript; available in PMC 2017 July 10.

Published in final edited form as:

Stat Med. 2016 July 10; 35(15): 2516–2524. doi:10.1002/sim.6886.

THE RAPID ENROLLMENT DESIGN FOR PHASE I CLINICAL TRIALS

Anastasia Ivanova[1,*], **Yunfei Wang**[2], and **Matthew C Foster**[3]

[1]Department of Biostatistics, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-7420, U.S.A.

[2]Departments of Pediatrics and Biostatistics, George Washington University and Children's National Medical Center, Washington DC 20010, U.S.A.

[3]University of North Carolina, Lineberger Comprehensive Cancer Center, Chapel Hill, North Carolina, 27599, U.S.A.

SUMMARY

We propose a dose-finding design for Phase I oncology trials where each new patient is assigned to the dose most likely to be the target dose given observed data. The main model assumption is that the dose-toxicity curve is non-decreasing. This method is beneficial when it is desirable to assign a patient to a dose as soon as the patient is enrolled into a study. To prevent assignments to doses with limited toxicity information in fast accruing trials we propose a conservative rule that assigns temporary fractional toxicities to patients still in follow-up. We also recommend using a safety rule with any dose-finding method.

Keywords

Phase I trial; mTPI; RED; CRM; TITE-CRM

1. INTRODUCTION

The main objective of dose-finding trials in oncology is to learn about the dose-toxicity relationship of a drug and to estimate the maximum tolerated dose (MTD). The MTD is usually defined as the dose with a certain probability of dose limiting toxicity (DLT). We will use DLT and toxicity interchangeably. There is a long history of oncology dose-finding methods for estimating a dose with a certain mean response when outcome is binary [1, 2, 3]. Recent reviews of dose-finding methods can be found in [4, 5]. Often it takes a long time to observe toxicity compared to accrual rate and a number of methods have been developed for dose-finding with delayed outcome [6, 7, 8]. In some trials, when the pace of disease is so rapid that a patient may succumb to the disease while waiting, it is desirable to enroll a patient into the trial and start treating as soon as possible. The method of Bekele et al. [8] solves the problem of fast accrual by prescribing when a trial should be halted as toxicity

*Correspondence to: Anastasia Ivanova, Department of Biostatistics, CB #7420, The University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-7420, U.S.A. ; Email: aivanova@bios.unc.edu, Tel (919) 8438086, Fax (919) 9663804

rate at the current dose becomes uncertain. In TITE-CRM [6] the accrual is not halted and a new patient is assigned to the current MTD estimated based on all data including partial data from patients still in follow-up. As noted in [6] rapid accrual, especially in the beginning of the trial, may lead to assigning patients to doses with high toxicity rate when TITE-CRM is used. None of the existing dose-finding methods developed for delayed toxicity provide a conservative approach to enrollment in trials where a patient is assigned to a dose as soon as the patient is enrolled.

In most dose-finding designs that do not rely on an parsimonious model for the dose-toxicity relationship [1, 7, 9, 10, 11] the next assignment is determined based on the data at the current dose. We propose assigning patients at lower doses while the toxicity profile at higher doses is being explored. When assignments are being made to lower doses as well as the higher dose, it is no longer clear what dose is the current dose and therefore none of the above cited rules can be used for assignment. In this paper we describe a dose-finding method where the dose-assignment strategy is not tied to a current dose and the next assignment is made to the dose that is most likely to be the MTD. We refer to the new method as the Rapid Enrollment Design (RED). We describe this dose-assignment design in Section 2. In Section 3 we propose how to mitigate the uncertainty from patients still in follow-up when a new assignment is made. We give an example in Section 4. We compare designs in Section 5 and discuss the findings in Section 6.

2. DOSE-FINDING METHOD

Let $D = \{d_1,\ldots,d_K\}$ denote the set of ordered dose levels selected for a trial. Let T be the length of follow-up for toxicity. A subject's response at d_k has a Bernoulli distribution with parameter q_k, where $(q_1,\ldots,q_K)$ is the vector of true toxicity rates at the K doses. We assume that $q_1,\ldots,q_K$ are non-decreasing, $q_1 \le \ldots \le q_K$. Observations from different subjects are independent. The goal is to find dose $d_m \in D$ such that $q_m = \Gamma$. If there is no such dose, the goal is to find the dose d_m with mean response closest to Γ. We refer to Γ as the target DLT rate and d_m as the target dose or the MTD.

First we consider the case where T is short compared to accrual rate and therefore toxicity responses from all patients are known when a new patient is enrolled. In the proposed design, each new patient is assigned to the dose with the highest probability of being the target dose given data available so far. Let $\boldsymbol{n} = (n_1,\ldots,n_K)$ be the number of subjects assigned to each dose and $\boldsymbol{m} = (m_1,\ldots,m_K)$ be the number of DLTs observed at each dose. Then the DLT rates at each dose can be estimated by simple proportions $\hat{q_j} = m_j / n_j$, $j = 1,\ldots,K$. Since $\hat{q_1},\ldots,\hat{q_K}$ might not be monotone we compute isotonic estimates $(\hat{q}_1^*,\ldots,\hat{q}_K^*)$ of DLT rates using the pool adjacent violator algorithm (PAVA) [12]. PAVA is described as follows. First we set $\hat{q}_j^* = \hat{q}_j$, $j = 1,\ldots, K$. If $\hat{q_1} \le \ldots \le \hat{q_K}$, then $\hat{q}_1^* \le \ldots \le \hat{q}_K^*$ and nothing else is needed. Otherwise, find the smallest k such that $m_k / n_k > m_{k+1} / n_{k+1}$ and replace these two estimates with $\hat{q}_k^* = \hat{q}_{k+1}^* = (m_{k-1}+m_k)/(n_{k-1}+n_k)$. We denote the estimate based on pooling data from doses d_k and d_{k+1} by $\hat{q_{k,k+1}}$ with "dose" $d_{k,k+1}$ and denominator $n_{k,k+1} = n_k + n_{k+1}$. Look again for a dose where the working isotonic estimate is greater than the working estimate at the next dose, and repeat the averaging and concatenation process until

isotonicity is actually obtained, that is, until $\hat{q}_1^* \leq \ldots \leq \hat{q}_K^*$. We note that the isotonic estimates obtained by PAVA are the maximum likelihood estimates assuming $q_1 \leq \ldots \leq q_K$. In the dose-assignment design described in the next paragraph when estimates $(\hat{q}_k^*, \ldots, \hat{q}_{k+j}^*)$ are computed based on pooled data $\hat{q}_k^* = \ldots = \hat{q}_{k+j}^* = (m_k + \ldots + m_{k+j})/(n_k + \ldots + n_{k+j})$, we use the highest dose on the estimated plateau, d_{k+j}, to represent the pooled doses if $\hat{q}_k^* \leq \Gamma$ and the lowest dose on the estimated plateau, d_k, if $\hat{q}_k^* > \Gamma$. To facilitate the Bayesian decision rule described below we assign the average number of DLTs $(m_k + \ldots + m_{k+j}) / (j + 1)$ and the average sample size $(n_k + \ldots + n_{k+j}) / (j + 1)$ of the pooled doses to that dose.

To compute the probability of a dose being the MTD we use Bayesian computations and assume Beta(α,β) prior on q_j, $j = 1, 2, \ldots, K$, the posterior distribution of q_j conditional on outcome data is

$$q_j | \boldsymbol{m}, \boldsymbol{n} \sim \text{Beta}(\alpha + m_j, \beta + n_j - m_j), j = 1, 2, \ldots, K.$$

The prior Beta(α,β) reflects the belief that there exists data from α +β patients, α patients with a DLT and β patients without a DLT. We use Jeffrey's prior α = β = 0.5. This prior leads to decisions similar to the 3+3 design for small sample sizes when the target probability of DLT is around 0.2 and 0.25.

The design is based on the following rules. Let $\{d_1, \ldots, d_k\}$, $k \leq K$, be the set of doses with at least one patient assigned.

1. *Initial escalation*. Do not change the dose unless at least s patients are assigned to the current highest dose, k, and their toxicity outcomes observed (excep when the safety rule is invoked). The number s depends on the target rate Γ [13]. Cohort size $s = 4$ is recommended if Γ = 0.1, $s = 3$ if Γ = 0.25, and $s = 1$ if Γ = 0.5.
2. If $\hat{q}_k^* < \Gamma$ and at least s patients are assigned to d_k, the next patient is assigned to d_{k+1} (or d_K if $k = K$).
3. If $\hat{q}_k^* \geq \Gamma$ and if there is a dose d_j such that $\hat{q}_j^* = \Gamma$, the next patient is assigned to d_j. Otherwise, let j, $j \leq k-1$, be such that $\hat{q}_j^* < \Gamma$ and $\hat{q}_{j+1}^* > \Gamma$. Let probability π_j reflect how close is the DLT rate at d_j to Γ, $\pi_j = \Pr\{\Gamma - \varepsilon < q_j < \Gamma + \varepsilon\}$. The next patient is assigned to the dose corresponding to the value π_j or π_{j+1} that is higher.
4. *Safety rule*. Do not assign patients to a dose with $\Pr\{q_j > \Gamma\} > 0.95$, $j = 1, 2, \ldots, k$. If $\Pr\{q_1 > \Gamma\} > 0.95$, the trial is stopped because the lowest dose is too toxic.
5. At the end of the trial the dose that would have been recommended for the next patient is selected as the estimated MTD. No dose is selected if the lowest dose is deemed too toxic by the stopping rule in 4).

In the rules above, ε, ε ≤ min(Γ,1-Γ), is a design parameter. We recommend ε = 0.05 Robustness of this choice is discussed in Section 5.

Decision rules for two candidate doses, one with estimated DLT rate lower than Γ and one with rate higher than Γ, are shown in Table 1 for Γ = 0.25. The first column of the table

contains data yielding DLT proportion less than $\Gamma = 0.25$, and the second column DLT proportion higher than $\Gamma = 0.25$. The decision rule for each pair of the first and second column data is in column 3. For example, if 3 patients were enrolled at dose 1 with no DLTs, and 6 patients were enrolled at dose 2 with 2 DLTs, the decision rule for these data is to assign to the higher dose (seventh line from the bottom in Table 1). This is because the data at the two doses, (0/3, 2/6), yield $\pi_1 = 0.100$ and $\pi_2 = 0.210$.

To implement the RED we developed web-based software available at http://cancer.unc.edu/biostatistics/program/ivanova/. The input is the number of DLTs at each dose, $\boldsymbol{m} = (m_1, \ldots, m_k)$, and the number of subjects at each dose $\boldsymbol{n} = (n_1, \ldots, n_k)$. The program identifies two candidate doses d_j and d_{j+1} using isotonic estimates and computes π_j and π_{j+1}. It also computes the probabilities that the DLT rates at the two doses are higher than Γ needed for the safety rule.

3. MITIGATING UNCERTAINTY FROM PATIENTS STILL IN FOLLOW-UP

In many Phase I oncology trials follow-up time is long compared to accrual rate. As mentioned in Section 2, when $\Gamma = 0.25$ we require that 3 patients complete follow-up at a previously untried dose before more patients can be assigned to that or a higher dose level. This however, does not fully resolve uncertainty about safety of future assignments. For example, if one out of the three patients had a DLT at the lowest dose, can we assign, say, 6 more patients to that dose at once or is it too risky? We propose a simple way to mitigate this risk. A patient without a DLT who has been in the follow-up for time u and therefore has completed a fraction u/T of the total follow-up with $1-u/T$ of the follow-up still remaining is counted as a patient with $1-u/T$ of a temporary DLT. The total DLT count at a dose is the number of actual DLTs, m, plus the sum of temporary DLTs $1-u_i/T$, where the sum is over all the patients assigned at that dose that are still in follow-up. The denominator is the total number of patients at that dose irrespective of their follow-up time. These data are used to determine the next assignment according to the rules in Section 2.

For example, in a trial with the target $\Gamma = 0.25$ three patients have completed follow up at d_1 with one DLT and several patients are available to enroll. With 1 DLT out of 3, $\Pr\{q_1 > \Gamma\} = 0.67$, therefore, since this probability is less than 0.95, at least one more patient can be enrolled. If one patient is enrolled, to mitigate potential DLT outcome from this patient, the data are augmented with 1 DLT at d_1 yielding 2 DLTs out of 4 at d_1 and $\Pr\{q_1 > \Gamma\} = 0.87$. Since this probability is less than 0.95 we can enroll one more patient. This assignment is less conservative than in the 3+3 design. The probability to observe 2 DLTs out of 4 patients if the true DLT rate is 0.25 is 0.21, that is, it is quite likely to see 2 DLTs in 4 patients. After augmenting 1/3 with 2/2, for two newly enrolled patients with $u = 0$ follow-up each, the DLT data at d_1 are 3/5 yielding $\Pr\{q_1 > \Gamma\} = 0.96$, therefore we cannot assign more patients to d_1 at this point. After the two newly enrolled patients have been followed for, say, half the total follow-up time without a DLT, the data at d_1 are 2/5, obtained by adding 1/3, 0.5/1 and 0.5/1, yielding $\Pr\{q_1 > \Gamma\} = 0.79$ and more patients can be enrolled. Patients are assigned to d_1 until the estimated DLT rate at d_1 becomes lower than $\Gamma = 0.25$, at which point patients are assigned to d_2.

In another example, there are 3 patients enrolled at d_1 with no DLTs, and 3 patients enrolled at d_2 with 1 DLT observed in these 3 patients. The data at the two doses are (0/3, 1/3), corresponding to $\pi_1 = 0.10$ and π_2, therefore the next patient, patient seven, is assigned to d_2. If patient eight is available at the time when patient seven is assigned, the augmented data are (0/3, 2/4), corresponding to $\pi_1 = 1.10$ and $\pi_2 = 0.11$, therefore, after checking the safety rule at d_2, $\Pr\{q_2 > \Gamma\} = 0.87$, patient eight should be assigned to d_2.

Note that for applying the stopping rule, $\Pr\{q_1 > \Gamma\} > 0.95$, at the lowest dose that might lead to terminating the trial, we only use real toxicities from patients who would have completed the follow-up by the time the analysis is performed and patients without toxicities who have completed the follow-up.

4. ACUTE MYELOID LEUKEMIA TRIAL EXAMPLE

Dose-finding trials in acute leukemia typically require a long follow-up for toxicity. This is because it is difficult to distinguish undue hematologic drug toxicity from the bone marrow effects of the disease itself. Often this requires follow-up for toxicity from an individual cycle of therapy that lasts 4–6 weeks rather than 3–4 weeks typical in solid tumors. In addition it is desirable to offer continuous enrollment in acute leukemia trials, since the pace of these leukemias is so rapid that the patient may succumb to the disease while waiting. Our new method was implemented in a dose-finding study of a new derivative of thalidomide for older adults with acute myeloid leukemia. Since the trial is ongoing, we used data from a recently completed Phase I trial [14] to illustrate the method. The trial investigated clofarabine in combination with gemtuzumab ozogamicin in relapsed or refractory acute myeloid leukemia patients. The DLT was defined as grade 3 or greater treatment-related toxicity lasting greater than 2 weeks or delay in hematologic recovery beyond 35 days from initiation of induction and not related to persistent or recurrent leukemia. Therefore the observation period for toxicity was 5 weeks from the start of therapy. The goal of the trial was to find a dose with the DLT rate of 0.26. The trial used the time-to-event cumulative cohort design [7] with an ad-hoc modification that allowed rapid enrollment with immediate assignment. We illustrate how the proposed dose-assignment algorithm would have worked if used in the gemtuzumab trial. We use patient enrollment times, their DLT outcomes and the time when a DLT has occurred. There were three patients who progressed or died between days 32 and 35 without a DLT, and therefore these patients were permanently censored for DLT before $T = 35$. In this illustration these patients are counted as patients with full follow-up of 35 days and no DLT. Dose assignments for all 20 patients are presented in Table 2. To make assignments in Table 2 coincide with those in the gemtuzumab trial we used the safety rule with 0.85 instead of 0.95. For example, after observing 2 DLTs in the first 4 patients at d_2 in the gemtuzumab trial the next assignment was to d_1. For these data, the posterior probability that the DLT rate is higher than the target of 0.26 is 0.86. Therefore if we set the cut-off for the safety rule at 0.85, the next assignment will be to d_1. Similarly, after 5 DLTs were observed in 13 patients at d_1 in the gemtuzumab trial, subsequent patients were assigned to d_{-1}. The probability that the DLT rate at d_1 was higher than the target of 0.26 was 0.85. Table 2 illustrates that our proposed design and the mitigation rule yielded rational dose assignments in a trial with rapid enrollment such as the gemtuzumab trial.

5. COMPARISONS WITH OTHER DOSE-FINDING METHODS

5.1 Comparison with mTPI, the *t*-statistic designs and the CRM in trials with a short follow-up time

First we considered the case where DLT outcome is known right away. We compared the RED with the modified toxicity probability interval (mTPI) method [11], the *t*-statistics design [10] and the CRM [2].

The mTPI is based on computing Bayesian posterior probabilities of the DLT rate being in certain intervals. Let d_j be the current dose, that is, the dose the last patient was assigned to. Calculate three probabilities $E = \Pr\{0 < q_j < \Gamma - \varepsilon_1\}/(\Gamma - \varepsilon_1)$, $S = \Pr\{\Gamma - \varepsilon_1 < q_j < \Gamma + \varepsilon_2\}/(\varepsilon_2 + \varepsilon_1)$ and $D = \Pr\{\Gamma + \varepsilon_2 < q_j < 1\}/(1 - \Gamma - \varepsilon_2)$. The next patient is assigned to d_{j+1} if E is the largest, to d_j if S is the largest, and to d_{j-1} if D is the largest. If $\Pr\{q_j > \Gamma\} > 0.95$, $j > 1$, patients are assigned to lower doses. If $\Pr\{q_1 > \Gamma\} > 0.95$ the trial is stopped. The estimated MTD is the dose with the estimated DLT rate closest to Γ.

The *t*-statistics method is a dose-finding design in which the *t*-statistic, t, to test the hypothesis that the mean at the current dose is equal to the target, is computed at each step. The next patient is assigned to the current dose if $-\Delta < t < \Delta$, otherwise the dose is reduced or increased depending on the sign of t. Here Δ is a design parameter and $\Delta = 1$ is recommended. In the *t*-statistic design to estimate the target dose after the trial for the new design, we first obtained the isotonic estimates of DLT rates. The dose with the estimated DLT rate closest to Γ is the estimated MTD. If there are two or more such doses, the highest dose with the estimated DLT rate below Γ is chosen. If all the estimated rates at these doses are higher than Γ, the lowest of these doses is chosen.

To describe the CRM, let x_i be the dose level received by subject i, $x_i \in D$, and $y_i = 1$ if the ith subject had a DLT and 0 otherwise. In the CRM [2], the calculation of posterior mean of parameter θ at the time when $(n+1)$th subject enters the trial is based on the likelihood:

$$L_n(\theta) = \prod_{i=1}^{n} F(x_i, \theta)^{y_i} \{1 - F(x_i, \theta)\}^{1-y_i},$$

where $F(x_i, \theta) = b_{x_i}^{\theta}$ and $b_1 < \ldots < b_K$ is a set of positive constants. We used the prior and $(b_1, \ldots, b_K) = (0.05, 0.1, 0.2, 0.3, 0.5, 0.7)$ as in [11] with our six-dose scenarios.

To compare designs we used scenarios from [6] and [11]. Here we show results for four of the five scenarios from [6] and one scenario from [11] that we reduced to a six-dose scenario by removing the two highest doses. That is, scenario 2 from [6] was replaced with a more interesting scenario from [11]. Results for the remaining scenarios are available from the authors. Simulations were performed in R and comparison was made based on 4000 simulation runs for each scenario. The target DLT rate was $\Gamma = 0.2$. Patients were assigned in cohorts of 3 in all designs. Note that all designs can be used with any number of patients per cohort. In the RED we set $\alpha = \beta = 0.5$ and $\varepsilon = 0.05$. In mTPI design the prior parameters were $\alpha = \beta = 1$ as in [11]. To make all the designs comparable we used the safety rule, rule

4) from Section 2, with the t-statistic design and the CRM (CRM+S) not only with the mTPI method and the RED. We also simulated the CRM without the safety rule. Simulation results are displayed in Table 3 and summary measures of these results in Table 5.

Overall mTPI, t-statistic, the RED and the CRM+S perform very similarly as far as average targeting precision and the average overdose (Table 3 and Table 5). In scenario 2 adopted from [11], the DLT probability at d_2 is 0.05 and the DLT probability at d_3 is 0.50. The CRM recommended d_3 in 34% of all trials. This is the consequence of relying solely on a parsimonious model. To correct that, we simulated the CRM with the safety rule (CRM+S). The addition of the safety rule reduced the proportion of trials where the unsafe dose d_3 was selected to 10%, however, it also worsened the CRM performance as far as the true MTD selection across scenarios, from 51% on average to 43% (Table 5). Since safety is important the CRM with the safety rule might be a good alternative to the CRM.

We investigated the robustness of the new design with respect to parameter ε. We performed simulations with values of ε in the range (0.01, 0.1) and obtained results very similar to ε = 0.05.

5.2 Comparison with TITE-CRM when the follow-up for DLT is long

Cheung and Chappell [6] proposed a time-to-event modification of the CRM, TITE-CRM, for dose-finding trials where follow-up for DLT is long. In clinical trials that require long follow-up times, the toxicity rate at dose x_i is defined as the probability of observing toxicity at x_i during a time period of length T after initiation of therapy. Data for the ith subject, i = 1,…, r, when (r+1)th subject is assigned to a treatment, consists of dose x_i, toxicity indicator y_i and the time u_i that has elapsed from the time of the ith subject's treatment assignment to the time of the (r+1)th subject's treatment assignment. For TITE-CRM, Cheung and Chappell [6] suggested the weighted likelihood

$$\tilde{L}_r(\theta)=\prod_{i=1}^{r}\{w_i F(x_i,\theta)\}^{y_i}\{1-w_i F(x_i,\theta)\}^{1-y_i},$$

where w_i is the weight assigned to the ith observation prior to the entry of the (r+1)th subject. For example, setting $w_i = \min(u_i / T,1)$ reflects an assumption that the density of time to toxicity is flat in $(0, T)$.

We compared the performance of the TITE-CRM and the RED for delayed toxicity outcome via simulations. In the RED the uncertainty from patients still in follow-up was mitigated as described in Section 3 (RED with mitigation). In the TITE-CRM, a patient who has been followed for a fraction w of full follow-up time T and has not had a DLT yet is contributing roughly as a fractional patient (w of a patient) without a DLT. On the other hand, in our approach such a patient is contributing as 1 patient with $(1 - w)$ DLTs. We simulated the original CRM with our mitigation rule from Section 3 (CRM with mitigation) as well.

Table 4 displays comparisons of the designs in a trial with a 35-day follow up for DLT with Table 5 displaying summary results. A new patient is enrolled on Monday every two weeks. For patients with DLT, time to DLT was generated as a uniform random variable on (0,35).

Simulating time to DLT using exponential distribution yielded very similar selection probabilities for all designs. The performance of the TITE-CRM in the five scenarios does not differ much from the CRM where our conservative toxicity mitigation rule was used to "impute" toxicity of patients still in follow-up (Table 4 and Table 5). The TITE-CRM selected unsafe doses in 16% of the trials on average and the CRM with our mitigation rule selected unsafe doses in 15% of the trials (Table 5). Similar performance of the two methods is due to the fact that the accrual is that very rapid compared to the follow-up length and, therefore, the way how DLTs are "imputed" is affecting the escalation considerably only in the beginning of the trial when not much DLT information is available. In contrast, an addition of the safety rule to the CRM with mitigation reduced the selection of unsafe doses to 6% (Table 5). The RED with mitigation performed similarly to the CRM with the safety rule and mitigation as far as average targeting precision and the average overdose (Table 5).

Interestingly, the CRM with mitigation performed much better in scenario 5 compared to the TITE-CRM. This scenario is an example of a scenario where CRM might not converge to the true MTD, d_5, and converges to d_4 instead in a large proportion of trials [15]. For the CRM with mitigation the probability of correctly selecting the true MTD in scenario 5 was 0.50 compared to 0.33 for the TITE-CRM (Table 4). This might be because the CRM with mitigation yields low estimates of the DLT probability at d_4 in the beginning of the trial allowing the CRM to escalate to the true MTD at d_5.

The length of a trial with rapid enrollment is at most 65 weeks. For comparison in a trial with the same accrual pattern where we wait until each cohort of 3 patients is fully followed before we enrolled the first patient from the next cohort, the trial length is 109 weeks. The results of the latter are the same as those presented in Table 3. Simulation summary in Table 5 shows that the RED and the CRM+SCRM+S with toxicity mitigation provide good estimation of the MTD without exposing patients to doses with high DLT rates in trials with long follow-up for toxicity when patients are enrolled at any time.

6. DISCUSSION

We proposed a new method to assign patients to doses in a Phase I oncology trial and a method to mitigate uncertainty of unobserved toxicities from patients still in follow-up. Ji and Wang [16] recently compared the mTPI and the 3+3 design [17] and showed that mTPI can identify the correct MTD better and also assigns fewer patients to the toxic doses above the MTD. Our new method performs similar to mTPI but also has the flexibility to accommodate rapid enrollment in trials with long follow-up for toxicity.

Our simulations of the TITE-CRM showed that in trials with delayed toxicity outcome the TITE-CRM selects unsafe doses relatively often. The classical CRM with our toxicity mitigation rule and the safety rule performs much better compared to the TITE-CRM in trials with delayed toxicity outcome. Another advantage is that while the TITE-CRM requires the TITE-CRM specific software, the CRM with toxicity mitigation can be implemented using the CRM software (though one will need to add a safety rule). Of note, we did not study the convergence of the CRM with addition of the mitigation and safety rules.

The main purpose of the mitigation rule we proposed is to guide dose assignments when there is little or no information available at a dose. When many patients have been enrolled at a dose, data from these patients overwhelms temporary toxicities assigned by the mitigation rule to patients still in follow-up and the mitigation rule does not play much of a role in dose assignment. We took a conservative approach to mitigate uncertainty of unobserved toxicities as this is often a choice of Phase I investigators. The new mitigation approach allows eliminating a set of rules we sometimes see in Phase I protocols. These rules prescribe when to wait to enroll a patient, when to enroll a patient and to what dose the patient should be assigned. Such rules can be cumbersome, especially if enrolling patients is allowed (usually at lower doses) before an initial cohort of patients at each new dose is fully followed. Our new dose assignment rule and the method of assigning temporary toxicity to patients still in follow-up eliminates the need for these rules. After temporary toxicities are assigned to patients still in follow-up, the assignment follows the primary design as if there is no delay in observing the DLT outcome. A web-based program is available to identify a dose for the next assignment given trial data.

We also developed a dose-assignment design similar to the RED that can be used for dose-finding with continuous outcome or where toxicity score or continuous biomarker is used in place of binary toxicity outcome [18].

Acknowledgments

This work was supported in part by NIH grant P01 CA142538. The authors thank Anna Snavely for helpful discussions and comments. The authors thank an associate editor and anonymous reviewers for their helpful comments.

REFERENCES

1. Wetherill GB. Sequential estimation of quantal response curves. Journal of the Royal Statistical Society B. 1963; 25:1–48.
2. O'Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase 1 clinical trials in cancer. Biometrics. 1990; 46:33–48. [PubMed: 2350571]
3. Babb J, Rogatko A, Zacks S. Cancer phase I clinical trials: efficient dose escalation with overdose control. Statistics in Medicine. 1998; 17:1103–1120. [PubMed: 9618772]
4. Tighiouart M, Liu Y, Rogatko A. Escalation with overdose control using time to toxicity for cancer phase I clinical trials. PLoS One. 2014 Mar 24.9(3):e93070. [PubMed: 24663812]
5. Sverdlov O, Wong WK, Ryeznik Y. Adaptive clinical trial designs for phase I cancer studies. Statistics Surveys. 2014; 8(2014):2–44.
6. Cheung YK, Chappell R. Sequential designs for phase I clinical trials with late-onset toxicities. Biometrics. 2000; 56:1177–1182. [PubMed: 11129476]
7. Ivanova I, Flournoy N, Chung Y. Cumulative cohort design for dose-finding. Journal of Statistical Planning and Inference. 2007; 137:2316–2327.
8. Bekele BN, Ji Y, Shen Y, Thall PF. Monitoring late-onset toxicities in phase I trials using predicted risks. Biostatistics. 2008; 9(3):442–457. [PubMed: 18084008]
9. Yuan Z, Chappell R. Isotonic designs for Phase I cancer clinical trials with multiple risk groups. Clinical Trials. 2004; 1:499–508. [PubMed: 16279290]
10. Ivanova A, Kim SH. Dose finding for continuous and ordinal outcomes with a monotone objective function: a unified approach. Biometrics. 2009; 65:307–315. [PubMed: 18479486]
11. Ji Y, Liu P, Li Y, Bekele BN. A modified toxicity probability interval method for dose-finding trials. Clinical trials. 2010; 7:653–663. [PubMed: 20935021]

12. Barlow, RE.; Bartholomew, DJ.; Bremner, JM.; Brunk, HD., editors. Statistical Inference under Order Restrictions. London, New York: John Wiley & Sons; 1972.
13. Ivanova A, Haghighi AM, Mohanti SG, Durham SD. Improved up-and-down designs for phase I trials. Statistics in Medicine. 2003; 22:69–82. [PubMed: 12486752]
14. Foster MC, Amin C, Voorhees PM, van Deventer HW, Richards KL, Ivanova A, Whitman J, Chiu WM, Barr ND, Shea T. A phase I dose escalation study of clofarabine in combination with fractionated gemtuzumab ozogamicin in patients with refractory or relapsed acute myeloid leukemia. Leukemia and Lymphoma. 2012; 53(7):1331–1337. [PubMed: 22149206]
15. Cheung YK, Chappell R. A simple technique to evaluate model sensitivity in the continual reassessment method. Biometrics. 2002; 58:671–674. [PubMed: 12230003]
16. Ji Y, Wang S. Modified toxicity probability interval design: a safer and more reliable method than the 3+3 design for practical phase I trials. Journal of Clinical Oncology. 2013; 31:1785–1791. [PubMed: 23569307]
17. Storer BE. Design and analysis of phase I clinical trials. Biometrics. 1989; 45:795–798.
18. Wang Y, Ivanova A. Dose-finding with continuous outcome in Phase I oncology trials. Pharmaceutical Statistics. 2015; 14:102–107. [PubMed: 25408518]

Table 1

Dose allocation decision based on the posterior probability at two candidate doses d_j and d_{j+1}. Target probability is $\Gamma = 0.25$. The observed toxicity rate at d_j is less than Γ and the observed toxicity rate at d_{j+1} is higher than Γ. Data at each dose is the number of DLTs over the number of patients assigned to the dose. The entry 0/1:3, for example, indicates one of the three possibilities: 0 DLTs were observed out of 1 or 2 or 3 patients.

Observed data at d_j	Observed data at d_{j+1}	Decision
1/4:9, 2/8:12, 3/12	1/3, 2/5	Lower Dose
1/4:10, 2/8:12, 3/12	3/7, 5/12	Lower Dose
2/8:12, 3/12	2/7, 3/9, 4/11	Lower Dose
1/4:7, 2/8:12, 3/12	3/8	Lower Dose
1/5:6, 2/8:12, 3/12	2/6	Lower Dose
2/9:10, 3/12	3/10	Lower Dose
3/12	3/11	Lower Dose
1/4:8, 2/8:12, 3/12	4/10	Lower Dose
2/8:11, 3/12	4/12	Lower Dose
1/4:11, 2/8:12, 3/12	4/9	Lower Dose
0/1:2, 1/4:12, 2/8:12, 3/12	1/2, 2/4, 5/11	Lower Dose
0/1:3, 1/4:12, 2/8:12, 3/12	3/6	Lower Dose
0/1:4, 1/4:12, 2/8:12, 3/12	4/8	Lower Dose
0/1:5, 1/4:12, 2/8:12, 3/12	5/10, 6/12	Lower Dose
0/1:8, 1/4:12, 2/8:12, 3/12	6/11	Lower Dose
0/3:12	1/2, 2/4, 5/11	Higher Dose
0/4:12	3/6	Higher Dose
0/5:12	4/8	Higher Dose
0/2:12, 1/10:12	1/3, 2/5	Higher Dose
0/2:12, 1/11:12	3/7, 5/12	Higher Dose
0/2:12, 1/4:12	2/6:7, 3/9, 4/11	Higher Dose
0/1:12, 1/8:12	3/8, 4/10	Higher Dose
0/1:12, 1/4:12, 2/8:12	3/10:11	Higher Dose
0/1:12, 1/4:12, 2/12	4/12	Higher Dose
0/1:12, 1/12	4/9	Higher Dose
0/6:12	5/10, 6/12	Higher Dose
0/9:12	6/11	Higher Dose

Table 2

Dose assignments for patients in the gemtuzumab trial. The target DLT rate is 0.26. The length of follow-up is 35 days. DLTs were observed in patients 4, 9, 11, 12, 14, 16 and 17 (in bold).

Pt	Day of enrollment	Data from patients with full follow-up at the time of enrollment		Additional temporary DLTs from patients still in follow-up at the time of enrollment		Estimated DLT rate		Posterior probability that DLT rate is in (0.21, 0.31)		Dose Assignment
		d_1	d_2	d_1	d_2	d_1	d_2	d_1	d_2	
1	1	---	---	---	---	---	---	---	---	1
2	77	0/1	---	0/0	---	0.00	---	0.108	---	1
3	77	0/1	---	1/1	---	0.50	---	0.111	---	1
4	172	0/3	---	0/0	---	0.00	---	0.095	---	**2**
5	194	0/3	0/0	0/0	0.37/1	0.00	0.37	0.095	0.112	2
6	327	0/3	1/2	0/0	0/0	0.00	0.50	0.095	0.111	2
7	327	0/3	1/2	0/0	1/1	0.00	0.67	0.095	0.058	1
8	348	0/3	1/2	0.4/1	0.4/1	0.10	0.47	0.149	0.127	1
9	369	0/4	1/3	0.4/1	0/0	0.08	0.33	0.135	0.165	**2**
10	437	0/5	2/4	0/0	0/0	0.00	0.50 *	0.067	0.114	1
11	448	0/5	2/4	0.69/1	0/0	0.12	0.50	0.165	0.114	**1**
12	508	1/7	2/4	0/0	0/0	0.14	0.50	0.193	0.114	**1**
13	516	1/7	2/4	0.77/1	0/0	0.22	0.50	0.257	0.114	1
14	565	2/9	2/4	0/0	0/0	0.22	0.50	0.270	0.114	**1**
15	636	3/10	2/4	0/0	0/0	0.30	0.50	0.278	0.114	1
16	671	3/11	2/4	0/0	0/0	0.27	0.50	0.302	0.114	**1**
17	676	3/11	2/4	0.86/1	0/0	0.32	0.50	0.281	0.114	**1**
18	801	5/13	2/4	0/0	0/0	0.38 **	0.50	0.204	0.114	−1
19	815	5/13	2/4	0/0	0/0	0.38 **	0.50	0.204	0.114	−1
20	850	5/13	2/4	0/0	0/0	0.38 **	0.50	0.204	0.114	−1

* The probability that the DLT rate at d_2 exceeds 0.26 is 0.86 hence the dose is reduced

** The probability that the DLT rate at d_1 exceeds 0.26 is 0.85 hence the dose is reduced

Table 3

Proportion of trials each dose was recommended. All designs except the CRM were simulated with addition of the safety rule. The CRM with the safety rule is denoted by the CRM+S. The target DLT probability is 0.20 and the total sample size is $n = 30$. Numbers at the true MTD are in bold.

	d_1	d_2	d_3	d_4	d_5	d_6
Scenario 1	0.05	0.10	**0.20**	0.30	0.50	0.70
mTPI	0.05	0.26	**0.41**	0.24	0.03	0.00
RED	0.05	0.20	**0.39**	0.33	0.04	0.00
t-statistics	0.05	0.26	**0.44**	0.22	0.02	0.00
CRM+S	0.05	0.28	**0.46**	0.20	0.00	0.00
CRM	0.01	0.19	**0.52**	0.28	0.01	0.00
Scenario 2	0.01	**0.05**	0.50	0.60	0.70	0.80
mTPI	0.01	**0.93**	0.06	0.00	0.00	0.00
RED	0.01	**0.94**	0.05	0.00	0.00	0.00
t-statistics	0.01	**0.93**	0.06	0.00	0.00	0.00
CRM+S	0.01	**0.87**	0.10	0.02	0.00	0.00
CRM	0.00	**0.66**	0.34	0.00	0.00	0.00
Scenario 3	0.05	0.06	0.08	0.11	**0.19**	0.34
mTPI	0.02	0.05	0.10	0.22	**0.42**	0.19
RED	0.02	0.02	0.05	0.19	**0.44**	0.27
t-statistics	0.02	0.05	0.13	0.28	**0.37**	0.14
CRM+S	0.01	0.04	0.12	0.33	**0.46**	0.04
CRM	0.00	0.01	0.06	0.31	**0.56**	0.05
Scenario 4	0.06	0.08	0.12	**0.18**	0.40	0.71
mTPI	0.04	0.10	0.22	**0.46**	0.15	0.00
RED	0.03	0.06	0.16	**0.55**	0.19	0.00
t-statistics	0.05	0.11	0.27	**0.47**	0.11	0.00
CRM+S	0.02	0.10	0.28	**0.54**	0.06	0.00
CRM	0.00	0.04	0.21	**0.64**	0.10	0.00
Scenario 5	0.00	0.00	0.03	0.05	0.11	**0.22**
mTPI	0.00	0.00	0.01	0.06	0.32	**0.60**
RED	0.00	0.00	0.01	0.04	0.28	**0.66**

	d_1	d_2	d_3	d_4	d_5	d_6
t-statistics	0.00	0.00	0.01	0.09	0.37	**0.53**
CRM+S	0.00	0.00	0.02	0.14	0.58	**0.27**
CRM	0.00	0.00	0.00	0.05	0.63	**0.32**

Table 4

Proportion of trials each dose was recommended as the MTD. The CRM with the safety rule added is denoted by CRM+S. The length of follow-up for DLT is 35 days, one patients is enrolled every two weeks. The target DLT rate is 0.20 and the total sample size is $n = 30$. Numbers at the true MTD are in bold.

	d_1	d_2	d_3	d_4	d_5	d_6
Scenario 1	0.05	0.10	**0.20**	0.30	0.50	0.70
TITE-CRM	0.01	0.19	**0.51**	0.28	0.01	0.00
CRM with mitigation	0.02	0.27	**0.52**	0.18	0.01	0.00
CRM+S with mitigation	0.05	0.30	**0.48**	0.17	0.01	0.00
RED with mitigation	0.09	0.28	**0.41**	0.20	0.02	0.00
Scenario 2	0.01	**0.05**	0.50	0.60	0.70	0.80
TITE-CRM	0.00	**0.65**	0.34	0.00	0.00	0.00
CRM with mitigation	0.00	**0.67**	0.33	0.00	0.00	0.00
CRM+S with mitigation	0.01	**0.90**	0.09	0.00	0.00	0.00
RED with mitigation	0.02	**0.87**	0.12	0.00	0.00	0.00
Scenario 3	0.05	0.06	0.08	0.11	**0.19**	0.34
TITE-CRM	0.00	0.02	0.08	0.32	**0.54**	0.05
CRM with mitigation	0.01	0.06	0.19	0.30	**0.34**	0.09
CRM+S with mitigation	0.01	0.07	0.20	0.30	**0.32**	0.09
RED with mitigation	0.06	0.05	0.12	0.31	**0.34**	0.12
Scenario 4	0.06	0.08	0.12	**0.18**	0.40	0.71
TITE-CRM	0.01	0.06	0.22	**0.60**	0.11	0.00
CRM with mitigation	0.02	0.14	0.35	**0.40**	0.10	0.00
CRM+S with mitigation	0.03	0.14	0.34	**0.41**	0.07	0.00
RED with mitigation	0.09	0.11	0.26	**0.45**	0.09	0.00
Scenario 5	0.00	0.00	0.03	0.05	0.11	**0.22**
TITE-CRM	0.00	0.00	0.00	0.05	0.62	**0.33**
CRM with mitigation	0.00	0.00	0.00	0.07	0.43	**0.50**
CRM+S with mitigation	0.00	0.00	0.01	0.09	0.42	**0.48**
RED with mitigation	0.00	0.00	0.02	0.12	0.40	**0.46**

Table 5

Summary measures for Tables 4 and 5. The average MTD targeting precision (Average Targeting Precision) is the average of the probabilities of correctly recommending the true MTD in scenarios 1, 3, 4 and 5. The average overdose measure (Average Overdose) is the probability of selecting a dose with the DLT rate of 0.4 or higher averaged over scenarios with such doses, scenarios 1, 2 and 4.

	Average Targeting Precision	Average Overdose
No Delay in Outcome		
mTPI	0.47	0.08
RED	**0.51**	0.09
t-statistics	0.45	**0.06**
CRM+S	0.43	**0.06**
CRM	**0.51**	0.15
Delayed Outcome		
TITE-CRM	**0.49**	0.16
CRM with mitigation	0.44	0.15
CRM+S with mitigation	0.42	**0.06**
RED with mitigation	0.41	0.07