A Bayesian Adaptive Design in Cancer Phase I Trials using Dose Combinations in the Presence of a Baseline Covariate

Sungjin Kim[*], Marcio Diniz[*] and Mourad Tighiouart[*†]

Abstract

We describe a Bayesian adaptive design for estimating the maximum tolerated dose curve as a function of a baseline covariate using two cytotoxic agents. Parametric models are used to describe the relationship between the doses, baseline covariate, and the probability of dose limiting toxicity (DLT). Trial design proceeds by treating cohorts of two patients simultaneously using escalation with overdose control, where at each stage of the trial, we seek a dose of one agent using the current posterior distribution of the MTD of this agent given the current dose of the other agent and the next patients baseline covariate value. At the end of the trial, we estimate MTD curves as functions of Bayes estimates of the model parameters. We evaluate design operating characteristics in terms of safety of the trial and percent of dose recommendation at dose combination neighborhoods around the true MTD by comparing the design that uses the covariate to the one that ignores the baseline characteristic.

Cancer Phase I trials, Maximum tolerated dose, Escalation with overdose control, Drug combination, Dose limiting toxicity, Continuous dose, Baseline covariate

1. Introduction

The combination of several cytotoxic and biologic agents in drug development and cancer treatment can help reduce tumor resistance to chemotherapy by targeting different signaling pathways simultaneously and improve tumor response when using additive or synergistic drugs [1]. Although the majority of cancer phase I trials use drug combinations of several agents, most of them are designed to estimate the MTD of one drug for fixed dose levels of the other drugs. This approach may provide a single tolerable dose for the combination but it may be suboptimal in terms of therapeutic effects in subsequent efficacy studies.

Trials where the dose levels of at least two agents are allowed to vary yield more than one MTD, or even an infinite number of MTDs in the case of continuous dose levels. Estimating the resulting set of MTDs by designing a safe trial is the main goal of phase I trials with dose combinations of several agents. Statistical methodologies for designing such trials have been studied extensively in the past decade [2]-[13]. These methods assume that the patient population is homogeneous in terms of treatment tolerance and every patient should be treated at a dose combination corresponding to a pre-defined target probability of DLT. Therefore, individual patients differences in susceptibility to treatment are not adapted. For single agent trials, strategies of drug allocation that accommodate individual patient needs based on pharmacokinetics and the genetics of drug metabolism have been used in [14]-[16]. Statistical designs permitting individualized MTD determination in single agent cancer phase I trials have also been proposed and implemented in real trials by a number of authors [17]-[21]. For drug combinations, an additional layer of complexity in specifying the dose-toxicity relationship given a baseline covariate is needed. Using the notation in [11], the general problem can be stated as follows. Let A_i, $i = 1, \ldots, k$ be k drugs and $S_i \subset R^+$ be the set of all possible doses of drug A_i. Denote by $x = (x_1, \ldots, x_k)$

[*]Biostatistics and Bioinformatics Research Center, Cedars-Sinai Medical Center 8700 Beverly Blvd, Los Angeles, CA 90048

[†]Corresponding author

a dose combination of k drugs and $S = S_1 \times \cdots \times S_k$. Consider a dose-toxicity model

$$\text{Prob}(\text{DLT}|\text{dose} = x, Z = z) = F_z(x, \xi), \tag{1.1}$$

where Z is the baseline covariate taking a value of 0 or 1, F is a known link function, and $\xi \in R^d$ is an unknown parameter. The MTD for a patient with baseline covariate z is defined as the set C_z of dose combinations x such that the probability of DLT for a patient with baseline covariate z given dose x equals to a target probability of DLT θ:

$$C_z = \{x \in S : F_z(x, \xi) = \theta\}, \tag{1.2}$$

In this work, we extend the design described by Tighiouart et al. [11] using escalation with overdose control (EWOC) principle [21]-[23], by treating cohorts of two patients simultaneously and accounting for patients baseline binary covariate. We use a reparameterization that allows MTD curve to lie anywhere within the Cartesian plane determined according to the range of continuous doses of the two drugs. Furthermore, a simplified form of model (1.1) is considered by assuming that patients with different covariate values will have parallel MTD curves. This assumption is mathematically convenient and allows us to use parsimonious models due to the small sample size constraints in cancer phase I trials.

This paper is organized as follows. Section 2 will describe the dose-toxicity model and trial design for continuous dose levels. In Section 3, we evaluate the performance of the proposed method by assessing the safety of the trial design and the efficiency of the estimate of the MTD curve. Discussions will be in Section 4.

2. Model

2.1 Dose-Toxicity Model

Consider the problem of identifying tolerable dose combinations of two cytotoxic agents A and B given a patient with a binary baseline covariate value of z. We consider the dose-toxicity model of the form

$$\text{Prob}(\delta = 1|x, y, z) = F(\beta_0 + \beta_1 x + \beta_2 y + \beta_3 z + \beta_4 xy), \tag{2.1}$$

where δ is the indicator of DLT, $\delta = 1$ if a patient given the dose combination (x, y) exhibits DLT within one cycle of therapy and $\delta = 0$ otherwise, $x \in [X_{min}, X_{max}]$ is the dose level of agent A, $y \in [Y_{min}, Y_{max}]$ is the dose level of agent B, z is a binary baseline covariate, and F is a known cumulative distribution function. Suppose that the doses of agents A and B are continuous and standardized to be in the interval $[0, 1]$. We will assume that the probability of DLT increases with the dose of any one of the agents when the other one is held constant for $z = 0, 1$. A necessary and sufficient condition for this property to hold is to assume $\beta_i > 0$, $i = 1, 2$ and the interaction term β_4 is nonnegative. The MTD C_z is defined as a set of combinations (x^*, y^*) such that

$$\text{Prob}(\delta = 1|x^*, y^*, z) = \theta. \tag{2.2}$$

The target probability of DLT, θ is set relatively high when the DLT is a reversible or non-fatal condition, and low when it is life threatening. Using (2.1) and (2.2), the MTD C_z is

$$C_z = \left\{(x^*, y^*) \in [0, 1]^2 : y^* = \frac{F^{-1}(\theta) - \beta_0 - \beta_1 x^* - \beta_3 z}{\beta_2 + \beta_4 x^*}\right\}. \tag{2.3}$$

We reparametrize model (2.1) in terms of the following parameters: ρ_{00}, the probability of DLT at the minimum available doses of agents A and B for a patient with covariate value $z = 0$; ρ_{100}, the probability of DLT when the level of drug A is X_{max}, the level of drug B is Y_{min} and $z = 0$; ρ_{101}, the probability of DLT when the level of drug A is X_{max}, the level of drug B is Y_{min} and $z = 1$; ρ_{010}, the probability of DLT when the level of drug A is X_{min}, the level of drug B is Y_{max} and $z = 0$; and the interaction parameter β_4. We will assume that $0 < \rho_{100}, \rho_{101}$, and $\rho_{010} < 1$. It follows that

$$\begin{aligned} \beta_0 &= F^{-1}(\rho_{00}) \\ \beta_1 &= F^{-1}(\rho_{100}) - F^{-1}(\rho_{00}) \\ \beta_2 &= F^{-1}(\rho_{010}) - F^{-1}(\rho_{00}) \\ \beta_3 &= F^{-1}(\rho_{101}) - F^{-1}(\rho_{100}). \end{aligned} \tag{2.4}$$

The MTD given in (2.3) can be expressed in terms of these new parameters as

$$C_z = \left\{ (x^*, y^*) \in [0,1]^2 : y^* = \frac{\begin{array}{r} F^{-1}(\theta) - F^{-1}(\rho_{00}) - (F^{-1}(\rho_{100}) - F^{-1}(\rho_{00}))x^* \\ - (F^{-1}(\rho_{101}) - F^{-1}(\rho_{100}))z \end{array}}{F^{-1}(\rho_{010}) - F^{-1}(\rho_{00}) + \beta_4 x^*} \right\}. \tag{2.5}$$

Let $D_n = (x_i, y_i, z_i, \delta_i), i = 1, \ldots, n$ be the data after enrolling n patients in the trial. The likelihood function under the reparametrization is

$$\begin{aligned} L(\rho_{00}, \rho_{100}, \rho_{101}, \rho_{010}, \beta_4 | D_n) = \prod_{i=1}^{n} &(H(\rho_{00}, \rho_{100}, \rho_{101}, \rho_{010}, \beta_4; x_i, y_i, z_i))^{\delta_i} \\ &\times (1 - H(\rho_{00}, \rho_{100}, \rho_{101}, \rho_{010}, \beta_4; x_i, y_i, z_i))^{1-\delta_i} \end{aligned} \tag{2.6}$$

where

$$\begin{aligned} &H(\rho_{00}, \rho_{100}, \rho_{101}, \rho_{010}, \beta_4; x_i, y_i, z_i) = \\ &\quad F(F^{-1}(\rho_{00}) + (F^{-1}(\rho_{100}) - F^{-1}(\rho_{00}))x_i + (F^{-1}(\rho_{010}) - F^{-1}(\rho_{00}))y_i \\ &\quad + (F^{-1}(\rho_{101}) - F^{-1}(\rho_{100}))z_i + \beta_4 x_i y_i) \end{aligned} \tag{2.7}$$

2.2 Prior and Posterior Distributions

Equations (2.4) imply that $0 < \rho_{00} < \theta$ since $\beta_i > 0$, $i = 1, 2$. We consider the priors $\rho_{00}/min(\rho_{100}, \rho_{010}) \sim \beta(a_0, b_0)$, $\rho_{100} \sim \beta(a_1, b_1)$, $\rho_{101} \sim \beta(a_2, b_2)$, $\rho_{010} \sim \beta(a_3, b_3)$, $\beta_4 \sim \gamma(a, b)$ with mean $E(\beta_4) = a/b$ and variance $Var(\beta_4) = a/b^2$. Vague priors for these parameters are achieved by taking $a_j = b_j = 1$, $j = 0, 1, 2, 3$. A vague prior for β_4 is achieved by selecting a large variance. As described in [11], a vague prior for β_4 is similarly selected with mean of 21 and variance of 542. Using Bayes rule, the posterior distribution of the model parameters is proportional to the product of the likelihood and prior distribution

$$\begin{aligned} \pi(\rho_{00}, \rho_{100}, \rho_{101}, \rho_{010}, \beta_4 | D_n) \propto \prod_{i=1}^{n} &(H(\rho_{00}, \rho_{100}, \rho_{101}, \rho_{010}, \beta_4; x_i, y_i, z_i))^{\delta_i} \\ &\times (1 - H(\rho_{00}, \rho_{100}, \rho_{101}, \rho_{010}, \beta_4; x_i, y_i, z_i))^{1-\delta_i} \\ &\times \pi(\rho_{00} | \rho_{101}, \rho_{010}) \pi(\rho_{100}) \pi(\rho_{101}) \pi(\rho_{010}) \pi(\beta_4) \end{aligned} \tag{2.8}$$

Features of this posterior distribution are estimated using WinBUGS [24] and JAGS.

2.3 Trial Design

The algorithm for dose escalation is similar to the ones described in [11], [25]. It uses the EWOC principle [21]-[23] where at each stage of the trial, we seek a dose of one agent using the current posterior distribution of the MTD of the agent given the current dose of the other agent and the next patients baseline covariate value. For instance, if agent A is held constant at level x, the dose of agent B is y such that the posterior probability that y exceeds the MTD of agent B given the dose of agent $A = x$ and covariate value z is bounded by a feasibility bound α. Cohorts of two patients are enrolled simultaneously receiving different dose combinations. Specifically, the design proceeds as follows.

1. Let $D_2 = \{(x_1, y_1, z_1, \delta_1), (x_2, y_2, z_2, \delta_2)\}$ the first cohort of two patients such that each patient receives the same dose combination $(x_i, y_i) = (X_{minA}, X_{minB}) = (0, 0)$ for $i = 1, 2$.

2. In the second cohort of two patients, patient 3 receives dose (x_1, y_3) and patient 4 receives dose (x_4, y_2), where y_3 is the α-th percentile of $\pi(\Gamma_{B|A=x_1,Z=z_3}|D_2)$ and x_4 is the α-th percentile of $\pi(\Gamma_{A|B=y_2,Z=z_4}|D_2)$. Here, $\pi(\Gamma_{B|A=x_1,Z=z_3}|D_2)$ is the posterior distribution of the MTD of agent B given that the level of agent A is x_1 and the baseline covariate value of patient 3 is z_3, given the data D_2. $\pi(\Gamma_{A|B=y_2,Z=z_4}|D_2)$ is defined similarly. $\Gamma_{B|A=x}$ and $\Gamma_{A|B=y}$ can be expressed in terms of $\rho_{00}, \rho_{100}, \rho_{101}$, and ρ_{010}.

3. In the i-th cohort of two patients,

   (a) If i is even, patient $2i - 1$ receives dose (x_{2i-3}, y_{2i-1}) and patient $2i$ receives dose (x_{2i}, y_{2i-2}), where $y_{2i-1} = \Pi^{-1}_{\Gamma_{B|A=x_{2i-3},Z=z_{2i-1}}}(\alpha|D_{2i-2})$ and $x_{2i} = \Pi^{-1}_{\Gamma_{A|B=y_{2i-2},Z=z_{2i}}}(\alpha|D_{2i-2})$.

   (b) If i is odd, then patient $2i1$ receives dose (x_{2i-1}, y_{2i-3}) and patient $2i$ receives dose (x_{2i-2}, y_{2i}), where $x_{2i-1} = \Pi^{-1}_{\Gamma_{A|B=y_{2i-3},Z=z_{2i-1}}}(\alpha|D_{2i-2})$ and $y_{2i} = \Pi^{-1}_{\Gamma_{B|A=x_{2i-2},Z=z_{2i}}}(\alpha|D_{2i-2})$.

4. Repeat step 3 until n patients are enrolled to the trial subject to the following stopping rule.

Stopping rule We stop enrollment to the trial if $P(P(\text{DLT}|(x, y) = (0, 0), z) \geq \theta + \delta_1|\text{data}) > \delta_2$, i.e. if the posterior probability that the probability of DLT at the minimum available dose combination in the trial exceeds the target probability of DLT is high for $z = 0, 1$ where δ_1 and δ_2 are design parameters chosen to achieve desirable model operating characteristics. At the end of the trial, we estimate the MTD curve using Bayes estimates of the parameters defining this curve. For example, an estimate of the MTD curve for $z = 0, 1$ is obtained using (2.5) as

$$\hat{C}_z = \left\{ (x^*, y^*) \in [0, 1]^2 : y^* = \frac{\begin{array}{r} F^{-1}(\theta) - F^{-1}(\hat{\rho}_{00}) - (F^{-1}(\hat{\rho}_{100}) \\ - F^{-1}(\hat{\rho}_{00}))x^* - (F^{-1}(\hat{\rho}_{101}) - F^{-1}(\hat{\rho}_{100}))z \end{array}}{F^{-1}(\hat{\rho}_{010}) - F^{-1}(\hat{\rho}_{00}) + \hat{\beta}_4 x^*} \right\}. \tag{2.9}$$

where $\hat{\rho}_{00}, \hat{\rho}_{100}, \hat{\rho}_{101}, \hat{\rho}_{010}$, and $\hat{\beta}_4$ are the posterior medians given the data D_n.

3. Simulation Studies

3.1 Simulation Set-up and Scenarios

We evaluate design operating characteristics by assuming a logistic link function $F(u) = (1 + e^{-u})^{-1}$ for the working and true models. In all simulations, the target probability of DLT is fixed at $\theta = 0.33$, the trial sample size is $n = 40$ patients with 20 patients in each group, $\delta_1 = 0.05$ and $\delta_2 = 0.8$.

We present four scenarios for the true MTD curve. The first scenario is a case where the two true MTD curves for two groups are parallel and close to the minimum doses as shown in Figure 1(a). The corresponding true parameters are $\rho_{00} = 0.01$, $\rho_{100} = 0.4$, $\rho_{101} = 0.8$, $\rho_{010} = 0.4$, and $\beta_4 = 10$. The second scenario is a case where the two true MTD curves for two groups are parallel but very close to each other. The third scenario is a case where two true MTD curves for two groups are not parallel. The last scenario is a case where the two true MTD curves are parallel but lie far away from each other and close to the maximum doses.

In all scenarios, $a_i = b_i = 1$, $i = 0, 1, 2, 3$, which correspond to non-informative priors for $\rho_{00}/min(\rho_{100}, \rho_{010}), \rho_{100}, \rho_{101}$, and ρ_{010}, and a vague prior for β_4 is selected with mean of 21 and variance of 540 as discussed in [25]. For each scenario, 1000 trials were simulated with the logistic link function for the working and true models.

A variable feasibility bound α was started from 0.25 and increased by 0.05 each time when we compute the dose for the next patient until α was reached to 0.5 [21], [25]. A dose escalation is restricted to be no more than 20% of the dose range of the corresponding agent.

3.2 Design Operating Characteristics

In order to assess the performance of this method when designing a prospective trial, we evaluate its operating characteristics by comparing the following three designs.

- Design using a covariate (WC); patients are accrued to the trial sequentially and the dose combinations given to the next cohort of patients is calculated assuming model (2.1).
- Design ignoring the covariate (IC); patients are accrued to the trial sequentially and the dose combinations given to the next cohort of patients is calculated assuming model (2.1) without the covariate, i.e., as in [11, 25].
- Design using parallel trials (S); in each group, patients are accrued to the trial sequentially and model (2.1) without the covariate is implemented in each group.

We evaluate the performance of the methods by assessing the safety of the trial designs and the efficiency of the estimate of the MTD curve.

3.2.1 Trial Safety

We assess trial safety by reporting the average percent of DLTs across all $m = 1000$ trials and the percent of trials that have a DLT rate exceeding $\theta + 0.1$, which is usually considered to be an indication of an excessive DLT rate.

3.2.2 Trial Efficiency

We present an estimate of the MTD curve using the average posterior medians of the model parameters. The estimate for $z = 0, 1$ is

$$\bar{C}_z = \left\{ (x^*, y^*) \in [0,1]^2 : y^* = \frac{\begin{array}{r} F^{-1}(\theta) - F^{-1}(\bar{\rho}_{00}) - (F^{-1}(\bar{\rho}_{100}) - F^{-1}(\bar{\rho}_{00}))x^* \\ - (F^{-1}(\bar{\rho}_{101}) - F^{-1}(\bar{\rho}_{100}))z \end{array}}{F^{-1}(\bar{\rho}_{010}) - F^{-1}(\bar{\rho}_{00}) + \bar{\beta}_4 x^*} \right\}. \tag{3.1}$$

where $F(\cdot)$ is the logistic function and $\bar{\rho}_{00}, \bar{\rho}_{100}, \bar{\rho}_{101}, \bar{\rho}_{010}$, and $\bar{\beta}_4$ are the average posterior medians of the parameters $\rho_{00}, \rho_{100}, \rho_{101}, \rho_{010}$, and β_4 from all $m = 1000$ trials, respectively. The next measure of efficiency is the pointwise average relative minimum distance from the true MTD curve to the estimated MTD curve. For $i = 1, \ldots, m$, let C_i be the estimated MTD curve and C_{true} be the true MTD curve. For every point $(x, y) \in C_{true}$, let

$$d^{(i)}_{(x,y)} = sign(y' - y) \times min_{\{(x^*,y^*):(x^*,y^*)\in C_i\}} \sqrt{(x - x^*)^2 + (y - y^*)^2} \tag{3.2}$$

where y' is such that $(x, y') \in C_i$. This is the minimum relative distance of the point (x, y) on the true MTD curve to the estimated MTD curve C_i. If the point (x, y) is below C_i, then $d^{(i)}_{(x,y)}$ is positive. Otherwise, it is negative. Let

$$d_{(x,y)} = m^{-1} \sum_{i=1}^{m} d^{(i)}_{(x,y)} \tag{3.3}$$

This is the pointwise average relative minimum distance from the true MTD curve to the estimated MTD curve and can be interpreted as the pointwise average bias in estimating the MTD. Let $\Delta(x, y)$ be the Euclidian distance between the minimum dose combination $(0, 0)$ and the point (x, y) on the true MTD curve and $0 < p < 1$. The last measure of efficiency we consider is

$$P_{(x,y)} = m^{-1} \sum_{i=1}^{m} I\left(|d^{(i)}_{(x,y)}| \leq p\Delta(x, y) \right) \tag{3.4}$$

This is the pointwise percent of trials for which the minimum distance of the point (x, y) on the true MTD curve to the estimated MTD curve C_i is no more than $(100 \times p)\%$ of the true MTD. This statistic is equivalent to drawing a circle with center (x, y) on the true MTD curve and radius $p\Delta(x, y)$ and calculating the percent of trials with MTD curve estimate C_i falling inside the circle. This will give us the percent of trials with MTD recommendation within $(100 \times p)\%$ of the true MTD for a given tolerance p.

3.3 Results

3.3.1 Trial Safety

Table 1 shows that the average percent of DLTs varies between 3% and 46% across all designs under the four scenarios (a)-(d). In general, the average DLT rate tends to be lower when the true MTD curve is farther away from the minimum dose combination. Furthermore, the percent of trials with an excessive rate of DLT as defined by a DLT rate exceeding $\theta + 0.1$ is 0.0% for all designs under the all four scenarios. This rate does not exceed 0.02% across all the cases where the maximum rate occurs with the design using parallel trials under the first scenario with true parameters of $\rho_{00} = 0.01$, $\rho_{100} = 0.4$, $\rho_{101} = 0.8$, $\rho_{010} = 0.4$, and $\beta_4 = 10$. Based on these findings, we conclude that the methodology is safe in general.

Table 1: Operating characteristics summarizing trial safety for three designs under four scenarios

Scenario	Design	Average % DLT			% Trials:
$(\rho_{00}, \rho_{100}, \rho_{101}, \rho_{010}, \beta_4)$		Overall	Z = 0	Z = 1	DLT rate $> \theta + 0.1$
(0.01, 0.4, 0.8, 0.4, 10)	WC	29.073	20.680	37.465	0.004
	P	25.748	21.95	29.545	0.018
	IC	30.098	14.05	46.145	0.004
(0.005, 0.1, 0.2, 0.1, 10)	WC	23.828	19.280	28.375	0.000
	P	17.953	16.865	19.040	0.000
	IC	25.718	19.070	32.365	0.000
(0.005, 0.2, 0.7, 0.01, 10)	WC	23.070	13.470	32.670	0.000
	P	18.493	14.775	22.210	0.001
	IC	24.900	8.595	41.205	0.000
(0.0001, 0.001, 0.05, 0.001, 10)	WC	16.143	5.360	26.925	0.000
	P	10.485	6.430	14.54	0.000
	IC	21.550	2.885	40.215	0.000

WC = With covariate, P = Parallel and IC = Ignoring covariates

3.3.2 Trial Efficiency

Figure 1 shows the true and estimated MTD curves for each group of patients under the four scenarios (a)-(d) when using the proposed design with a baseline covariate, parallel trials, and a design ignoring the baseline covariate. The estimated MTD curves were obtained using (3.1) and DLT responses were simulated from the true logistic model. Figure 1 shows that the estimated MTD curves are closer to the true MTD curves when accounting for a significant baseline covariate using the proposed design and parallel trials. When ignoring the covariate, the estimated MTD curve tends to be in between the true MTD curves. This shows that when the two MTD curves are well separated, not accounting for a baseline covariate results in suboptimal MTD curve estimation for one group of patients and a too toxic MTD curve recommendation for the other group.

Figure 2 displays the pointwise average relative minimum distance from the true MTD curve to the estimated MTD curve under the four scenarios (a)-(d) as defined by (3.3). This is a measure of pointwise average bias of the estimate of the MTD. In the first scenario, Figure 2(a) shows that the average bias varies between -0.10 and 0.10 where the higher values are observed with the design ignoring the covariate. The pointwise average bias is quite similar when using the design with covariate and parallel trials and the highest absolute value is observed with the design ignoring the covariate. Similar findings are observed under the third (Figure 2(c)) and the last scenarios (Figure 2(d)). In the second scenario where the true MTD curves for $z = 0$ and $z = 1$ are very close (Figure 2(b)), the highest absolute value is achieved at the edges of the true MTD curves for both groups of $z = 0$ and $z = 1$ when using the design with covariate.

Figure 3 shows the pointwise percent of trials for which the minimum distance from the true MTD curve to the estimated MTD curve is no more than $(100 \times p)\%$ of the true MTD for $p = 0.2$ as defined by (3.4). This can be interpreted as the percent of MTD recommendation for a given tolerance p. Under the first scenario, Figure 3(a) shows that the percent of trials with correct MTD recommendation within 20% of the true value of the MTD varies between 61% and 99% with the design accounting for a baseline covariate while it varies more widely between 29% and 100% with the design ignoring the covariate.

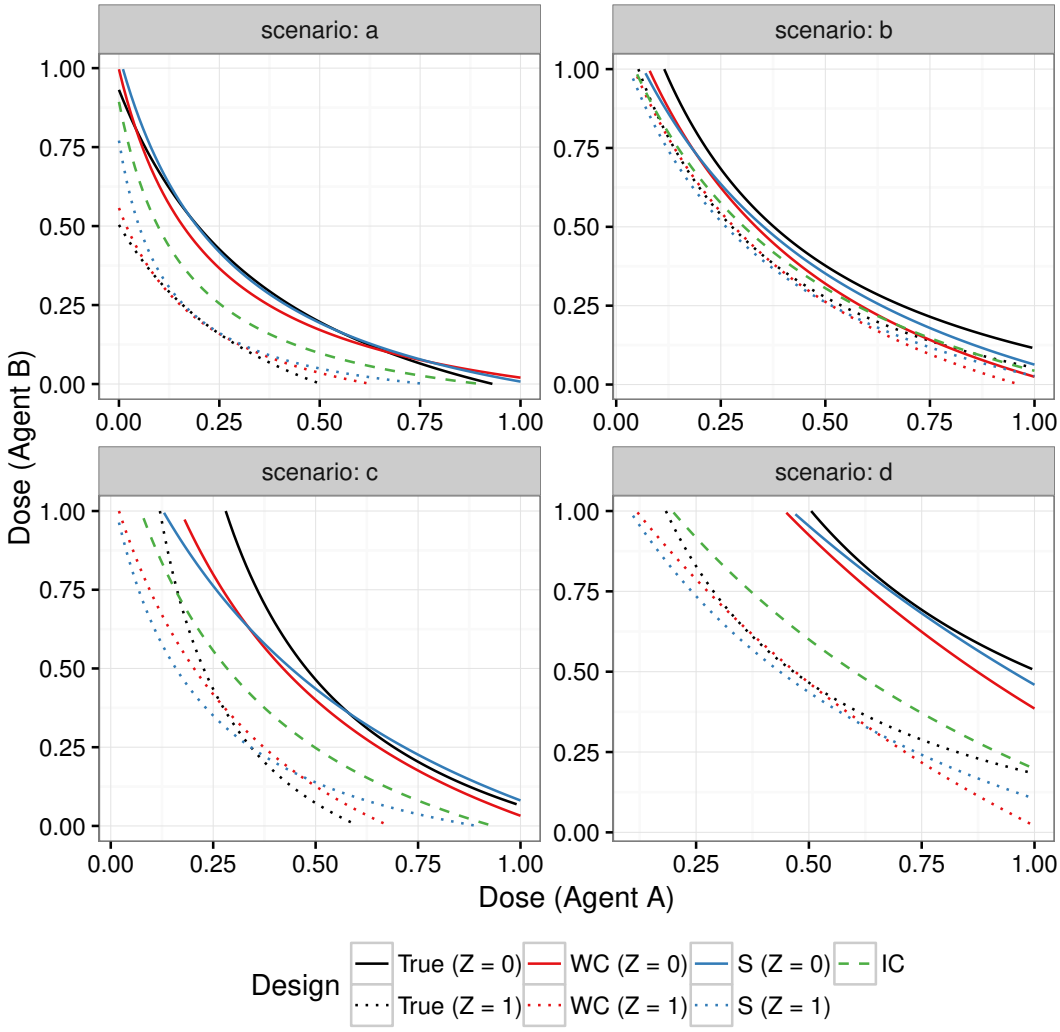

Figure 1: True and estimated MTD curves from m = 1000 simulated trials under scenarios (a)-(d)

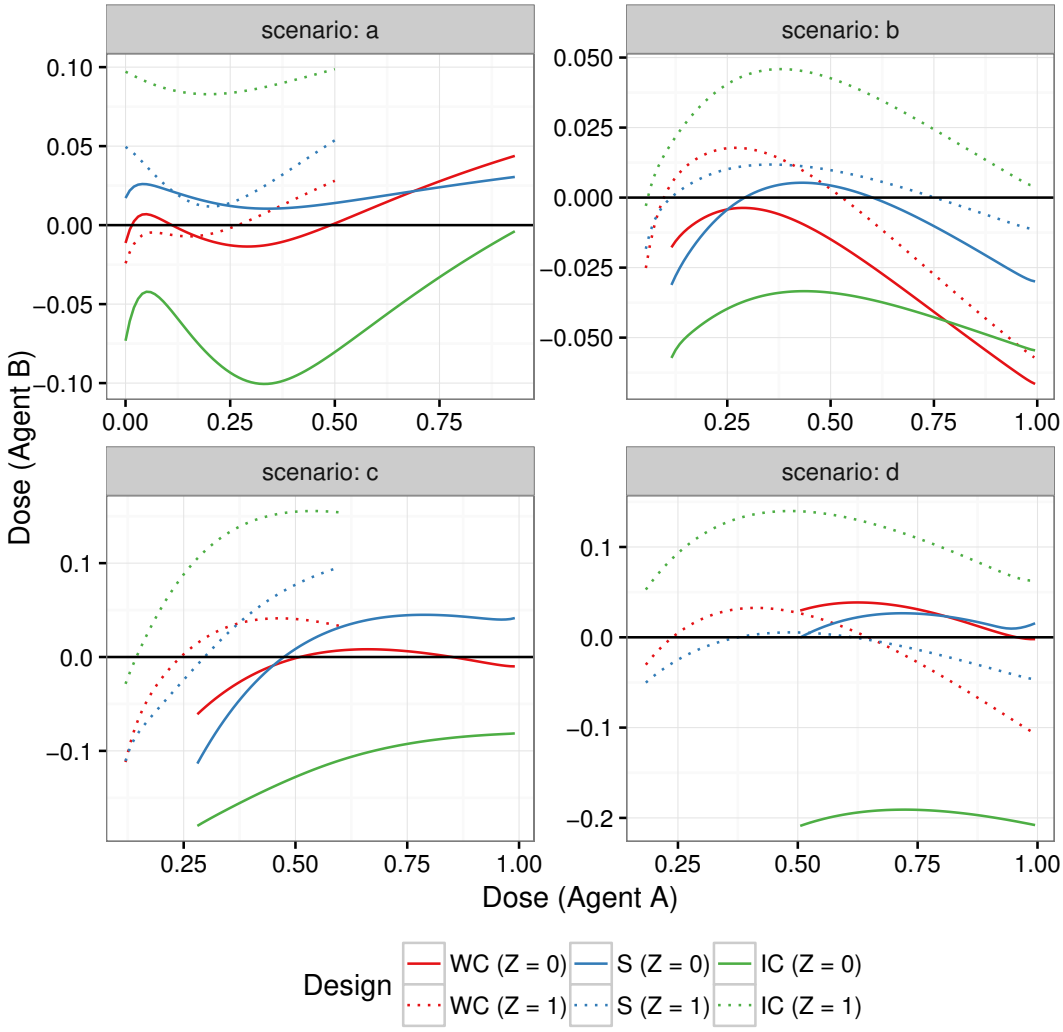

Figure 2: Pointwise average relative minimum distance from the true MTD curve to the estimated MTD curve under scenarios (a)-(d)

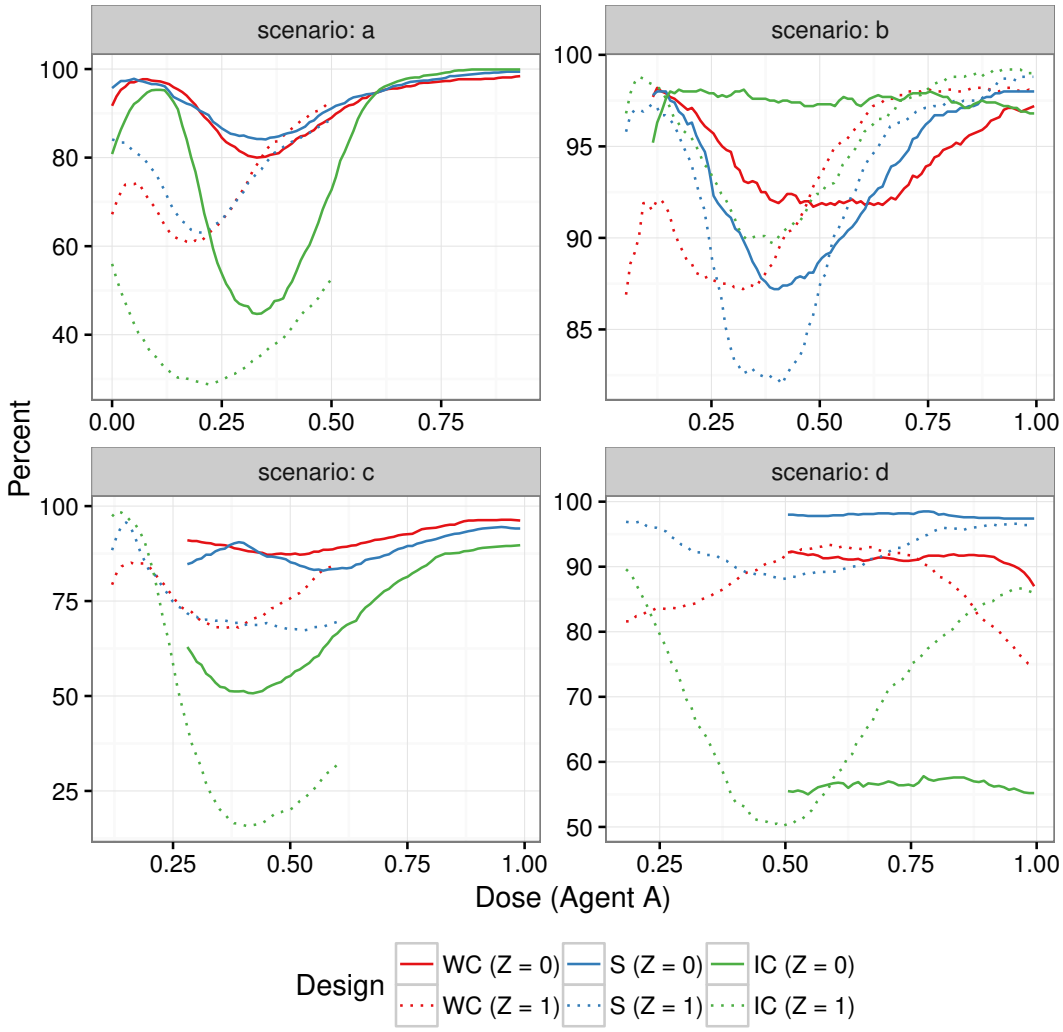

Figure 3: Pointwise percent of MTD recommendation for p = 0.2 under scenarios (a)-(d)

Similar pattern is found under the third (Figure 3(c)) and the fourth (Figure 3(d)) scenarios. Under the second scenario (Figure 3(b)), the percent recommendation is similar between all designs varying between 87% and 98% with the design with a covariate , $82\% - 99\%$ with the parallel trials, and $90\% - 99\%$ with the design without the baseline covariate. The percent recommendation increases as we move away from the minimum available dose combination. In general, ignoring a practically important baseline covariate results in a lower MTD recommendation rate relative to a design accounting for this covariate. These findings support that the design is efficient in recommending the MTD curve estimates.

4. Conclusion

We described Bayesian adaptive designs for cancer phase I clinical trials using two drugs with continuous dose levels in the presence of a binary baseline covariate. The goal is to estimate the MTD curve in the two-dimensional Cartesian plane for a patients specific baseline covariate value. The methodology extends the single agent trial design with a baseline covariate and two agents design without a covariate. In each case, vague priors were used

to quantify the toxicity profile of each agent a priori. We used an algorithm for dose escalation where cohorts of two patients are enrolled simultaneously and the patients receive different dose combinations. We studied design operating characteristics of the method under four practical scenarios by comparing this method with the design that ignores the baseline covariate and design using parallel trials. In all simulations, we used a sample size of $n = 40$ patients, 20 patients in each group. We found that in general, the methodology is safe in terms of the probability that a prospective trial will result in an excessively high number of DLTs when accounting for a significant covariate. We used several measures to assess the efficiency of the estimate of the MTD. In the presence of a practically significant baseline covariate, the design with a covariate had a smaller pointwise average bias and a higher percent of MTD recommendation relative to a design which ignores the covariate or when using parallel trials. When the two true MTD curves are very close, including a baseline covariate in the model results in a slightly higher but still negligible bias and a small reduction in percent of MTD recommendation relative to a design that ignores this covariate. Therefore, we stand to lose little if we include a practically not important covariate in the model. While we are currently investigating the approach for pre-specified discrete dose combination of the two agents, other extensions include accommodating late onset toxicity and efficacy studies.

References

[1] Frey E, III., Karon M, Levin RH, Freireich EJ, Taylor RJ, Hananian J, Selawry O, Holland JF, Hoogstraten B, Wolman IJ, Abir E, Sawitsky A, Lee S, Mills SD, Burgert EO, Spurr CL, Patterson RB, Ebaugh FG, James GWI, Moon JH. The Effectiveness of Combinations of Antileukemic Agents in Inducing and Maintaining Remission in Children with Acute Leukemia. *Blood* 1965; 26: 642-656.

[2] Thall PF, Millikan RE, Mueller P, Lee SJ. Dose-finding with two agents in phase I oncology trials. *Biometrics* 2003; 59: 487-496.

[3] Wang K, Ivanova A. Two-dimensional dose finding in discrete dose space. *Biometrics* 2005; 61: 217-222.

[4] Yin GS, Yuan Y. A Latent Contingency Table Approach to Dose Finding for Combinations of Two Agents. *Biometrics* 2009; 65: 866-875.

[5] Yin GS, Yuan Y. Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society Series C-Applied Statistics* 2009; 58: 211-224.

[6] Braun TM, Wang SF. A Hierarchical Bayesian Design for Phase I Trials of Novel Combinations of Cancer Therapeutic Agents. *Biometrics* 2010; 66: 805-812.

[7] Wages NA, Conaway MR, O'Quigley J. Continual Reassessment Method for Partial Ordering. *Biometrics* 2011; 67: 1555-1563.

[8] Wages NA, Conaway MR, O'Quigley J. Dose-finding design for multi-drug combinations. *Clinical Trials* 2011; 8: 380-389.

[9] Sweeting MJ, Mander AP. Escalation strategies for combination therapy Phase I trials. *Pharm Stat* 2012; 11: 258-266.

[10] Shi Y, Yin G. Escalation with overdose control for phase I drug-combination trials. *Stat Med 2013*: in press.

[11] Tighiouart M, Piantadosi S, Rogatko A. Dose finding with drug combinations in cancer phase I clinical trials using conditional escalation with overdose control. *Stat Med* 2014; 33: 3815-3829.

[12] Riviere MK, Yuan Y, Dubois F, Zohar S. A Bayesian dose-finding design for drug combination clinical trials based on the logistic model. *Pharm Stat* 2014; 13: 247-257.

[13] Mander AP, Sweeting MJ. A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Stat Med* 2015; 34: 1261-1276.

[14] Decoster G, Stein G, Holdener EE. Responses and Toxic Deaths in Phase-I Clinical-Trials. *Annals of Oncology* 1990; 1: 175-181.

[15] Ratain MJ, Mick R, Schilsky RL, Siegler M. Statistical and Ethical Issues in the Design and Conduct of Phase-I and Phase-Ii Clinical-Trials of New Anticancer Agents. *Journal of the National Cancer Institute* 1993; 85: 1637-1643.

[16] Ratain MJ, Mick R, Janisch L, Berezin F, Schilsky RL, Vogelzang NJ, Kut M. Individualized dosing of amonafide based on a pharmacodynamic model incorporating acetylator phenotype and gender. *Pharmacogenetics* 1996; 6: 93-101.

[17] O'Quigley J, Shen LZ, Gamst A. Two-sample continual reassessment method. *J Biopharm Stat* 1999; 9: 17-44.

[18] Babb JS, Rogatko A. Patient specific dosing in a cancer phase I clinical trial. *Stat Med* 2001; 20: 2079-2090.

[19] O'Quigley J, Paoletti X. Continual reassessment method for ordered groups. *Biometrics* 2003; 59: 430-440.

[20] Tighiouart M, Cook-Wiens G, Rogatko A. Incorporating a Patient Dichotomous Characteristic in Cancer Phase I Clinical Trials Using Escalation with Overdose Control. *Journal of Probability and Statistics* 2012; 2012: 10.

[21] Tighiouart M, Rogatko A. Dose Finding with Escalation with Overdose Control (EWOC) in Cancer Clinical Trials. *Statistical Science* 2010; 25: 217-226.

[22] Babb J, Rogatko A, Zacks S. Cancer Phase I clinical Trials: efficient dose escalation with overdose control. *Stat Med*1998; 17: 1103-1120.

[23] Tighiouart M, Rogatko A, Babb JS. Flexible Bayesian methods for cancer phase I clinical trials. Dose escalation with overdose control. *Stat Med* 2005; 24: 2183-2196.

[24] Lunn DJ, Thomas A, Best N, Spiegelhalter D. WinBUGS - A Bayesian modelling framework: Concepts, structure, and extensibility. *Statistics and Computing* 2000; 10: 325-337.

[25] Tighiouart M, Piantadosi S, Rogatko A. Dose finding for drug combination in early cancer phase I trials using conditional escalation with overdose control. *Stat Med* 2014; 33:3815-3829.

[26] Chu PL, Lin Y, Shih WJ. Unifying CRM and EWOC designs for phase I cancer clinical trials. *Journal of Statistical Planning and Inference* 2009; 139:11461163.