Published in final edited form as:

Stat Med. 2017 January 30; 36(2): 280–290. doi:10.1002/sim.6961.

A Bayesian Adaptive Design for Estimating the Maximum Tolerated Dose Curve using Drug Combinations in Cancer Phase I Clinical Trials

Mourad Tighiouart*, Quanlin Li, and André Rogatko

Samuel Oschin Comprehensive Cancer Institute, 8700 Beverly Blvd., Los Angeles, CA 90048

Abstract

We present a cancer phase I clinical trial design of a combination of two drugs with the goal of estimating the maximum tolerated dose curve in the two-dimensional Cartesian plane. A parametric model is used to describe the relationship between the doses of the two agents and the probability of dose limiting toxicity. The model is re-parameterized in terms of the probabilities of toxicities at dose combinations corresponding to the minimum and maximum doses available in the trial and the interaction parameter. Trial design proceeds using cohorts of two patients receiving doses according to univariate escalation with overdose control (EWOC), 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. The maximum tolerated dose curve is estimated as a function of Bayes estimates of the model parameters. Performance of the trial is studied by evaluating its design operating characteristics in terms of safety of the trial and percent of dose recommendation at dose combination neighborhoods around the true MTD curve and under model misspecifications for the true dose-toxicity relationship. The method is further extended to accommodate discrete dose combinations and compared with previous approaches under several scenarios.

Keywords

Cancer Phase I trials; Maximum tolerated dose curve; Escalation with overdose control; Drug combination; Dose limiting toxicity; Continuous dose; Discrete dose

1 Introduction

The majority of cancer phase I trials use drug combinations of several cytotoxic/biologic agents. This allows targeting different signaling pathways simultaneously and helps reduce potential tumor resistance to chemotherapy. However, most of these trials are designed to estimate the maximum tolerated dose (MTD) of a single agent for fixed dose levels of the other agents. This approach may provide a single safe dose for the combination but it may be suboptimal in terms of therapeutic effects. Statistical designs that allow the two agents to vary during the trial have been studied extensively in the last decade [1-12]. Some of these

* mourad.tighiouart@cshs.org.

designs are aimed at identifying a single MTD ([4, 5, 7-9, 11]) whereas others can recommend more than one MTD combination ([1-3, 6, 12]) and even an infinite number of MTDs [10]. Most of these methods use a parametric model for the dose-toxicity relationship

$$Prob(Z=1|\text{dose}=\mathbf{x})=F(\mathbf{x},\xi), \quad (1.1)$$

where $\mathbf{x} = (x_1, \dots, x_k)$ denotes the dose combination of k drugs, F is a known link function, Z is the indicator of dose limiting toxicity (DLT), and $\xi \in \mathcal{R}^d$ is an unknown parameter. Let S be the set of all dose combinations available in the trials. The MTD is defined as the set C of dose combinations $\mathbf{x}$ such that the probability of DLT for a patient given dose combination $\mathbf{x}$ equals to a target probability of DLT θ

$$C = \{\mathbf{x} \in S : F(\mathbf{x}, \xi) = \theta\}. \quad (1.2)$$

An alternative definition of the MTD is the set of dose combinations $\mathbf{x}$ that satisfy $|F(\mathbf{x}, \xi) - \theta| \leq \delta$ since the set C in (1.2) may be empty. This can happen for example when S is finite and the MTD is not part of the dose combinations available in the trial. The threshold parameter δ , $0 < \delta < 1$, referred to as “100× δ -point window” in [6] must be pre-specified by the clinician. Except for the methods in [1, 10], the above designs do not apply to the case of continuous dose levels and it is not clear how they perform in the presence of a large number of dose combinations relative to the trial sample size. For instance, when the total number of dose combinations is 64 or more, the methods in [3, 9] use over 100 patients, which is impractical in cancer phase I clinical trials.

In this article, we extend the method of Tighiouart et al. [10] by relaxing the assumption that the MTD of each drug is within the range of available doses of the corresponding agent. We use a new reparameterization that can be easily interpreted by the clinicians, permits the use of vague priors, and allow the MTD curve to lie anywhere within the Cartesian plane defined by the range of the continuous doses of the two agents. Furthermore, the algorithm described in [10] is extended by treating cohorts of two patients simultaneously and receiving different dose combinations according to the escalation with overdose control (EWOC) scheme [13-15]. Finally, we show how the method can be adapted to the setting of discrete dose combinations and compare it with the two-dimensional design of Wang and Ivanova [2].

The manuscript is organized as follows. In section 2, we describe the parametric model of the dose-toxicity relationship and the adaptive design to conduct the trial for continuous dose levels. In section 3, we study the performance of the proposed methodology in terms of safety of the trial and efficiency of the estimate of the MTD. Adapting the methodology to the setting of discrete dose combinations is illustrated in section 4. Section 5 contains some final remarks and discussion of practical implementations.

2 Method

2.1. Dose-Toxicity Model

Denote by $X_{\min,i}$, $X_{\max,i}$, $i=A, B$ the minimum and maximum doses available in the trial for cytotoxic agents A and B , respectively. Suppose that the doses of these agents are continuous and standardized to be in the interval $[0, 1]$, $X_{\min,i} = 0$, $X_{\max,i} = 1$, $i=A, B$. We consider the dose-toxicity model of the form

$$\text{Prob}(Z=1|x, y) = F(\mu + \beta x + \gamma y + \eta xy), \quad (2.1)$$

where Z is the indicator of dose limiting toxicity (DLT), $Z = 1$ if a patient given the dose combination (x, y) exhibits DLT within one cycle of therapy, and $Z = 0$ otherwise, x is the standardized dose level of agent A , y is the standardized dose level of agent B , and F is a known cumulative distribution function. 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. A necessary and sufficient condition for this to hold is to assume that $\beta > 0$, $\gamma > 0$, and $\eta \geq 0$. The MTD is defined as any dose combination (x^*, y^*) such that

$$\text{Prob}(Z=1|x^*, y^*) = \theta. \quad (2.2)$$

The value of the target probability θ is pre-specified by the clinicians and depends on the severity and clinical manageability of the DLT. It follows from (2.1) and (2.2) that the MTD is the set of dose combinations

$$C = \left\{ (x^*, y^*) : y^* = \frac{F^{-1}(\theta) - \mu - \beta x^*}{\gamma + \eta x^*} \right\}. \quad (2.3)$$

The 4 parameter model (2.1) is defined for dose combinations (x', y') within the range $[X_{\min,A}, X_{\max,A}] \times [X_{\min,B}, X_{\max,B}]$ with $X_{\min,A}, X_{\min,B} > 0$ and cannot be extrapolated outside this domain. Here, x' and y' are the unstandardized dose levels of agents A and B , respectively. In particular, $\pi_{A,x'} = \text{Prob}(Z = 1 | (x', 0))$ and $\pi_{B,y'} = \text{Prob}(Z = 1 | (0, y'))$ are not necessarily the marginal probabilities of DLT of single agents A and B . Gasparini [16] describes a large class of dose-toxicity models for drug combinations having $\pi_{A,x'}$ and $\pi_{B,y'}$ as the marginal probabilities of DLT of single agents. This important class of models is very convenient in eliciting informative prior distributions of the parameters associated with $\pi_{A,x'}$ and $\pi_{B,y'}$ when we have experience about the toxicity profile of each agent as in Thall et al. [1]. In our setting, knowledge of the toxicity profiles of single agent trials is not essential as demonstrated by the model reparameterization and vague prior distributions of the model parameters described below.

We reparameterize model (2.1) in terms of ρ_{10} , the probability of DLT when the levels of drugs A and B are 1 and 0, respectively, ρ_{01} , the probability of DLT when the levels of drugs A and B are 0 and 1, respectively, ρ_{00} , the probability of DLT when the levels of drugs A

and B are both 0, and the interaction parameter η . This reparameterization is convenient to clinicians since prior information on ρ_{00} , ρ_{10} , and ρ_{01} may be available from other trials. Furthermore, it extends the reparameterization in [10] by allowing the MTD of each agent when the other one is held at its minimum value to lie outside the range of available doses in the trial. It can be shown that

$$\begin{cases} \mu = F^{-1}(\rho_{00}) \\ \beta = (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00})) \\ \gamma = (F^{-1}(\rho_{01}) - F^{-1}(\rho_{00})) \end{cases} \quad (2.4)$$

The MTD in (2.3) becomes

$$C = \left\{ (x^*, y^*) : y^* = \frac{(F^{-1}(\theta) - F^{-1}(\rho_{00})) - (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00})) x^*}{(F^{-1}(\rho_{01}) - F^{-1}(\rho_{00})) + \eta x^*} \right\}. \quad (2.5)$$

2.2. Prior and Posterior Distributions

Equation (2.4) implies that $0 < \rho_{00} < \min(\rho_{01}, \rho_{10})$ since β and γ are positive. We assume that ρ_{01} , ρ_{10} , η are independent *a priori* with $\rho_{01} \sim \text{beta}(a_1, b_1)$, $\rho_{10} \sim \text{beta}(a_2, b_2)$, and conditional on (ρ_{01}, ρ_{10}) , $\rho_{00} / \min(\rho_{01}, \rho_{10}) \sim \text{beta}(a_3, b_3)$. A gamma distribution with mean $E(\eta) = a / b$ and variance $\text{Var}(\eta) = a / b^2$ is placed on the interaction term η . Let $D_n = \{(x_i, y_i, z_i), i = 1, \dots, n\}$ be the data after enrolling n patients in the trial. Using Bayes rule, the posterior distribution of the model parameters is proportional to the product of the likelihood and prior distribution

$$\begin{aligned} \pi(\rho_{00}, \rho_{01}, \rho_{10}, \eta | D_n) &\propto \prod_{i=1}^n (G(\rho_{00}, \rho_{01}, \rho_{10}, \eta; x_i, y_i))^{z_i} (1 - G(\rho_{00}, \rho_{01}, \rho_{10}, \eta; x_i, y_i))^{1-z_i} \\ &\times \pi(\rho_{01}) \pi(\rho_{10}) \pi(\rho_{00} | \rho_{01}, \rho_{10}) \pi(\eta), \end{aligned} \quad (2.6)$$

where

$$G(\rho_{00}, \rho_{01}, \rho_{10}, \eta; x_i, y_i) = F\left(F^{-1}(\rho_{00}) + (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00}))x_i + (F^{-1}(\rho_{01}) - F^{-1}(\rho_{00}))y_i + \eta x_i y_i\right).$$

We used WinBUGS [17] and JAGS [18] to estimate features of the posterior distribution of these parameters and estimate design operating characteristics of cancer phase I trials described below.

2.3. Trial Design

We describe a dose escalation/de-escalation algorithm by treating cohorts of two patients simultaneously. It is based on the escalation with overdose control (EWOC) principle where at each stage of the trial, the posterior probability of overdosing a future patient is bounded by a feasibility bound α [13-15, 19]. For a given cohort of two patients, one subject receives a new dose of agent A for a given dose of agent B that was previously assigned and the other patient receives a new dose of agent B for a given dose of agent A that was previously assigned. Specifically, the design proceeds as follows:

1. Each patient in the first cohort of two patients receives the same dose combination $(x_1, y_1) = (x_2, y_2) = (X_{\min,A}, X_{\min,B}) = (0, 0)$. Let $D_2 = \{(x_1, y_1, z_1), (x_2, y_2, z_2)\}$.
2. In the second cohort of two patients, patient 3 receives dose (x_3, y_3) and patient 4 receives dose (x_4, y_4) , where $y_3 = y_1$, $x_4 = x_2$, x_3 is the α -th percentile of $\pi(\Gamma_{A|B=y_1} | D_2)$, and y_4 is the α -th percentile of $\pi(\Gamma_{B|A=x_2} | D_2)$. Here, $\pi(\Gamma_{A|B=y_1} | D_2)$ is the posterior distribution of the MTD of drug A given that the level of drug B is y_1 , given the data D_2 . $\pi(\Gamma_{B|A=x_2} | D_2)$ is defined similarly. These posterior distributions are easily obtained from the MCMC output since $\Gamma_{A|B=y}$ and $\Gamma_{B|A=x}$ can be expressed in terms of the model parameters $\rho_{00}, \rho_{01}, \rho_{10}, \eta$.
3. In the i -th cohort of two patients, if i is even, then patient $(2i-1)$ receives dose (x_{2i-1}, y_{2i-1}) , patient $2i$ receives dose (x_{2i}, y_{2i}) , where $x_{2i-1} = \Pi_{\Gamma_{A|B=y_{2i-3}}}^{-1}(\alpha | D_{2i-2})$, $y_{2i} = \Pi_{\Gamma_{B|A=x_{2i-2}}}^{-1}(\alpha | D_{2i-2})$. If i is odd, then patient $(2i-1)$ receives dose (x_{2i-1}, y_{2i-1}) , patient $2i$ receives dose (x_{2i}, y_{2i}) , where $y_{2i-1} = \Pi_{\Gamma_{B|A=x_{2i-3}}}^{-1}(\alpha | D_{2i-2})$, $x_{2i} = \Pi_{\Gamma_{A|B=y_{2i-2}}}^{-1}(\alpha | D_{2i-2})$. Here, $\Pi_{\Gamma_{A|B=y}}^{-1}(\alpha | D)$ is the inverse cdf of the posterior distribution $\pi(\Gamma_{A|B=y} | D)$.
4. Repeat step 3 until n patients are enrolled to the trial subject to the following stopping rule.

If the α -th percentile of $\pi(\Gamma_{A|B=y} | D)$ or $\pi(\Gamma_{B|A=x} | D)$ is less than 0 or greater than 1, the recommended dose for the next patient is 0 or 1, respectively. In steps 2 and 3 of the algorithm, any dose escalation is further restricted to be no more than a pre-specified fraction of the dose range of the corresponding agent.

Stopping rule—We stop enrollment to the trial if $P(\text{DLT} | (x, y) = (0, 0)) - \theta + \zeta_1 | \text{data}) > \zeta_2$, 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. ζ_1 and ζ_2 are design parameters chosen to achieve desirable model operating characteristics.

At the end of the trial, we estimate the MTD using (2.5) as

$$\hat{C} = \left\{ (x^*, y^*) : 0 \leq x^*, y^* \leq 1, y^* = \frac{(F^{-1}(\theta) - F^{-1}(\hat{\rho}_{00})) - (F^{-1}(\hat{\rho}_{10}) - F^{-1}(\hat{\rho}_{00}))x^*}{(F^{-1}(\hat{\rho}_{01}) - F^{-1}(\hat{\rho}_{00})) + \hat{\eta}x^*} \right\},$$

(2.9)

where $\hat{\rho}_{00}, \hat{\rho}_{01}, \hat{\rho}_{10}, \hat{\eta}$ are the posterior medians given the data D_n .

3. Simulation Studies

3.1. Simulation set up and Scenarios

In all simulation studies, a logistic link function $F(u) = (1 + \exp(-u))^{-1}$ is taken for the working model. In order to assess the performance of the method under model misspecification, DLT responses are generated from the working mode, three link functions (1) the probit link $F(u) = \Phi(u)$, where $\Phi(\cdot)$ is the cdf of the standard normal distribution, (2) the normal link $F(u) = \Phi(u/\sigma)$, and (3) the complementary log-log link $F(x) = 1 - e^{-e^x}$, and the six-parameter model used in [1], see the supplement. In all simulations, the target probability of DLT is fixed at $\theta = 0.33$, the trial sample size is $n = 40$ patients, $\zeta_1 = 0.05$ and $\zeta_2 = 0.8$. We present results based on eight scenarios for the true MTD curve (a)–(h). Due to space limitation, the last four scenarios (e)–(h) are included in the supplement. Scenario (a) is a case where the two drugs are very safe within the range of available doses in the trial but the true MTD curve lies near the upper-right corner of the x - y plane, see Figure 1(a). The corresponding true parameters are $\rho_{00} = 10^{-7}$, $\rho_{01} = \rho_{10} = 3 \times 10^{-6}$, $\eta = 10$ and are included in Table 1. In scenario (b), the MTD of agent A when agent B is at its minimum dose level is within the range of doses of drug A but the MTD of agent B when drug A is at its minimum dose level is above the maximum dose level of agent B , see Figure 1(b). Scenario (c) is a case where drug A is safe but the MTD of agent B when drug A is at its minimum dose level is just above 0.8 and scenario (d) is a case where the two drugs are highly interactive with $\eta = 100$. In all scenarios, vague priors for $\rho_{00}, \rho_{01}, \rho_{10}$ were selected by taking $a_i = b_i = 1$, $i = 1, 2, 3$. A similar argument used in [10] lead to the choice of a vague prior for η with mean 21 and variance 540. For each scenario, $m = 2000$ trials were simulated using the logistic link function for the working model and logistic, probit, normal with $\sigma = 2$, and complementary log-log link functions for the true model. The parameter values μ, β, γ, η of these models were selected in such a way that they all have the same true MTD curve. The extents of departure of the true model from the working model are shown in Figure 1 of the supplement under scenarios (f) and (h). We also include operating characteristics under model misspecification when the true model uses the normal link function with $\sigma = 2.5, 3$ to explore greater departure of the true model from the working model in the supplement. We used a variable feasibility bound α , starting with $\alpha = 0.25$ and increase this value in increments of 0.05 each time we compute the dose for the next patient until $\alpha = 0.5$, see [15, 20] for the rationale for selecting a variable feasibility bound. Any dose escalation is further restricted to be no more than 0.2, that is, 20% of the dose range of the corresponding agent.

3.2. Design Operating Characteristics

For each scenario, we present an estimate of the MTD curve

$$\bar{C} = \left\{ (x^*, y^*) : 0 \leq x^*, y^* \leq 1, y^* = \frac{(F^{-1}(\theta) - F^{-1}(\bar{\rho}_{00})) - (F^{-1}(\bar{\rho}_{10}) - F^{-1}(\bar{\rho}_{00}))x^*}{(F^{-1}(\bar{\rho}_{01}) - F^{-1}(\bar{\rho}_{00})) + \bar{\eta}x^*} \right\}, \quad (3.1)$$

where $F(\cdot)$ is the logistic function and $\bar{\rho}_{00}, \bar{\rho}_{01}, \bar{\rho}_{10}, \bar{\eta}$ are the average posterior medians of the parameters $\rho_{00}, \rho_{01}, \rho_{10}, \eta$ from all $m = 2000$ trials. For trial safety, we present the percent of DLTs across all patients and all trials, the percent of trials with DLT rate greater than $\theta + 0.05$ and $\theta + 0.1$. For trial efficiency, we present the pointwise average relative minimum distance from the true MTD curve to the estimated MTD curve $d_{(x,y)}$ described in [10]. For $i = 1, \dots, m$, let C_i be the estimated MTD curve and let C_{true} be the true MTD curve. For every point $(x, y) \in C_{true}$, let

$$d_{(x,y)}^{(i)} = \text{sign}(y' - y) \times \min_{\{(x^*, y^*) : (x^*, y^*) \in C_i\}} \left((x - x^*)^2 + (y - y^*)^2 \right)^{1/2} \quad (3.2)$$

where y' is such that $(x, y') \in C_i$. Let

$$d_{(x,y)} = m^{-1} \sum_{i=1}^m d_{(x,y)}^{(i)}. \quad (3.3)$$

This is the pointwise average relative minimum distance from the true MTD curve to the estimated MTD curve which can be interpreted as the pointwise average bias in estimating the MTD. The last measure of efficiency we consider is

$$m^{-1} \sum_{i=1}^m I \left(|d_{(x,y)}^{(i)}| \leq p \Delta(x, y) \right). \quad (3.4)$$

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 can be interpreted as the pointwise percent of MTD recommendation for a given tolerance p , see [10].

3.3. Results

Table 1 shows that the average percent of DLTs varies between 16% and 34% across scenarios (a)–(d). These rates are similar when using the true and misspecified models. Moreover, the percent of trials with an excessive DLT rate is small. Based on these results and from scenarios (e)–(h) in the supplement, we found that the trial is safe and that there is

a high probability of stopping the trial when the true MTD curve is near or below the minimum dose combination (results not shown).

Figure 1 shows the plots of the true and estimated MTD curves under scenarios (a)–(d). The estimated MTD curve shown by a dashed line was obtained using (3.1) and DLT responses were simulated using the true logistic model. In general, the estimated MTD curve is fairly close to the true MTD except perhaps at the edges of the curve. This is probably due to the fact that we are using uniform priors for ρ_{01} and ρ_{10} . The gray scatter plot represents the last two dose combinations from each of the $m = 2000$ simulated trials along with the 90% confidence regions. These statistics are useful if the dose combinations given to the last cohort of two patients are used in a phase II trial [6]. Figure 2 displays the pointwise average bias under scenarios (a)–(d) as defined by (3.3). Under scenario (a), Figure 1(a) shows that the average bias varies between -0.03 and 0.00 when the true and misspecified models are used. Under scenario (b), Figure 2(b) shows that the average bias is still negligible with the highest absolute value achieved at the edges of the MTD curve. Similar conclusions can be made for scenarios (c) and (d). Figure 3 shows the plots of the pointwise percent of MTD recommendation for tolerances $p = 0.1$ and $p = 0.2$. Under scenario (a), Figure 3(a) shows that 95% or more of the trials will recommend an MTD estimate that is within 10% of the true value of the MTD when the DLT responses are generated from the true and misspecified model. The corresponding percent of MTD recommendation under scenario (b) is reasonably high except near the true MTD of drug A when the level of drug B is at its minimum value. With a tolerance of $p = 0.2$, the percent of trials with correct MTD recommendation is very high under all four scenarios. Based on these results and the scenarios shown in the supplement, we conclude that the design is practically efficient in general in recommending the MTD curve estimate and is fairly robust to model misspecification under small deviations of the true model from the working model using the family of dose-toxicity models of the form $\text{Prob}(\text{DLT} \mid x, y) = R\mu + \beta x + \gamma y + \eta xy$ and the six-parameter model used in [1].

4 Discrete Dose Combinations

In this section, we show how the methodology can be adapted to the case of discrete dose levels of agents A and B .

4.1 Approach

Let $(x_1, \dots, x_r)$ and $(y_1, \dots, y_s)$ be the doses of agents A and B , respectively. Following the notation of Section 2.1, $X_{\min, A} = x_1$, $X_{\min, B} = y_1$, $X_{\max, A} = x_r$, and $X_{\max, B} = y_s$ and the doses are standardized to be in the interval $[0, 1]$. Trial design proceeds using the algorithm describes in Section 2.3 where the continuous doses recommended in steps 2 and 3 of the algorithm are rounded to the nearest discrete dose level. At the end of the trial, a discrete set Γ of dose combinations satisfying (i) and (ii) below is selected as MTDs. Let C_j be the estimated MTD curve at the end of the trial and denote by $d((x_j, y_k), C_j)$ the Euclidian distance between the dose combination (x_j, y_k) and C_j as in (3.2).

- (i) Let $\Gamma_1 = \bigcup_{t=1}^r \left\{ (x_t, y) : y = \underset{y_j}{\operatorname{argmin}} d((x_t, y_j), C_i) \right\}$,
 $\Gamma_2 = \bigcup_{t=1}^s \left\{ (x, y_t) : x = \underset{x_j}{\operatorname{argmin}} d((x_j, y_t), C_i) \right\}$ and $\Gamma_0 = \Gamma_1 \cap \Gamma_2$.
- (ii) Let $\Gamma = \Gamma_0 \setminus \{ (x^*, y^*) : P(|P(DLT|(x^*, y^*) - \theta| > \delta_1 | D_n) > \delta_2 \}$.

In (i), dose combinations closest to the MTD curve are selected by first minimizing the distances across the levels of drug *B*, then across the levels of drug *A*. In (ii), we exclude MTDs from (i) that are either likely to be too toxic or sub therapeutic. This condition is needed because we are rounding the MTD curve to the nearest available dose combination in the trial and the probability of DLT of this combination may be too high or too low. The design parameter δ_1 is selected after consultation with a clinician and we typically use $\delta_1 = 0.1$. The parameter δ_2 is selected after exploring a large number of scenarios for a given prospective trial. In our experience with the sample sizes and scenarios used in Wang and Ivanova [2], we found that $\delta_2 = 0.3, 0.35$ result in good design operating characteristics.

4.2 Illustration

We study the performance of this approach using five scenarios presented in Wang and Ivanova [2] since their approach is designed to identify one or more MTDs. These are shown in Table 2 with $r = 6$ and $s = 3$ for scenarios 1—4 and $r = s = 4$ for scenario 5. The target probability of DLT is $\theta = 0.2$ for all scenarios. For each scenario, we simulated $m = 2000$ trials using the same vague priors for $\rho_{00}, \rho_{01}, \rho_{10}$ from Section 3.1 and a gamma(1,1) prior was selected for the interaction term since the model in [2] has two parameters with no interaction coefficient. The sample size for the first four scenarios is $n = 54$ and $n = 60$ for scenario 5. The performance of the method is evaluated by calculating the accuracy index introduced by Cheung [21]

$$A_n = 1 - K \times \frac{\sum_{k=1}^K \Delta_k \times p_{n,k}}{\sum_{k=1}^K \Delta_k}, \quad (4.1)$$

where n is the trial sample size, K is the number of discrete doses available in the trial, $p_{n,k}$ is the probability of selecting dose k in a trial with n patients, and Δ_k is a distance measure between the true probability of DLT p_k at dose k and the target probability of DLT θ . We consider three of the distances described in [21]; (i) square discrepancy (sq), $\Delta_k = (p_k - \theta)^2$, (ii) absolute discrepancy (abs), $\Delta_k = |p_k - \theta|$, and (iii) overdose error (od), $\Delta_k = \alpha^* (\theta - p_k)^+ + (1 - \alpha^*) (p_k - \theta)^+$. It can be shown that $A_N \rightarrow 1$ and higher values of A_N are desirable. However, the interpretation of nonpositive A_N for single agent trials given in [21] does not apply to dose combination trials since more than one MTD may exist. Furthermore, for designs aimed at identifying one or more MTDs, tables summarizing percent selection are useful but can be hard to interpret. For instance, the percent of MTD selection under scenario 1 in [2] for dose combinations (3, 3), (2, 4), and (1, 5) are 33%, 54%, and 27%, respectively. This means that there is a 0.33 probability that dose combination (3,3) was

selected correctly as an MTD but there are positive probabilities that this dose is selected along with other true MTDs, too toxic dose combinations, or dose combinations with low probability of DLT. Therefore, we propose to use the following measure of percent of MTD selection. For a given scenario, let $\Gamma_\delta = \{(x_i, y_j): |P(DLT|(x_i, y_j)) - \theta| < \delta\}$ be the set of true MTDs where the threshold parameter δ is fixed by the clinician. Typical values for δ are 0.05 or 0.1, see e.g. [6]. Let Γ_i be the set of MTDs at the end of the i -th trial as defined in (ii) of Section 4.1, $i=1, \dots, m$. The percent of MTDs selection is

$$\% \text{ Selection} = \frac{1}{m} \sum_{i=1}^m I(\Gamma_i \subset \Gamma_\delta). \quad (4.2)$$

The statistic in (4.2) is an estimate of the probability that for a given scenario, a prospective trial will recommend a set of dose combinations that are all MTDs. Table 3 gives the summary statistics of the accuracy index with $\alpha^* = 0.25$, percent selection with $\delta = 0.1$, and safety of the trial for the proposed approach (conditional EWOC) and the two-dimensional design of Wang and Ivanova [2]. Our approach seems to perform well relative to the two-dimensional design under scenarios 1, 3, and 5 according to the three measures of discrepancies. Scenario 2 is more complex due to the location of the true MTDs and this is reflected by the negative values of the accuracy index across the three measures of discrepancies for the two approaches. The two-dimensional design performs better than our approach under scenario 4 according to two of the 3 discrepancy measures. When the accuracy index A_N is averaged across the five scenarios, our approach performs better than the two-dimensional design for each discrepancy measure. The percent of true MTDs selection is higher using our approach under scenarios 1, 3, and 4 and higher on the average across all five scenarios relative to the two-dimensional design. The last 3 columns of Table 3 show that the trial is safe using both approaches under these 5 scenarios. Other simulation results using informative priors matching the priors used in [2] lead to similar conclusions and much higher percent MTDs selection under scenarios 1, 3, 4, and 5. Tables of percent MTD selection are also included in the supplement. Based on these results, we conclude that our approach has good operating characteristics when adapted to discrete dose combinations and performs better than the two-dimensional design according to the five entertained scenarios.

5 Conclusions

We proposed a Bayesian adaptive design for cancer phase I clinical trials using two drugs with continuous dose levels. The method extends the work of Tighiouart et al. [10] by allowing the MTD curve to lie anywhere in the Cartesian plane defined by the doses of the two agents. Furthermore, trial design proceeds by treating cohorts of two patients simultaneously and with different dose combinations to better explore the space of doses. We are currently developing a protocol using this approach for a phase I trial of dose combination using continuous dosing of cabazitaxel ranging between 10 mg/m² and 25 mg/m² and continuous dosing of cisplatin ranging between 45 mg/m² and 100 mg/m² in prostate cancer patients with visceral metastasis. We studied the operating characteristics of the design under many scenarios for the true location of the MTD curve and we found that

the trial is safe in general, and stops early with high probability when the true MTD curve is near or below the minimum dose combination. The method achieves good efficiency of the estimate of the MTD curve as assessed by the pointwise average bias and percent selection except perhaps at the edges of the curve under some scenarios. An alternative approach to quantify the percent selection globally is to estimate the probability that a large portion of the estimated MTD curve (say the middle 80% of the curve) is in a pre-specified two-dimensional neighborhood of the true MTD curve. We plan to study the properties of this measure in future work. We note that the operating characteristics were evaluated under vague prior distributions of the model parameters and no toxicity profiles of either agent were used *a priori*. In practice, if toxicity data from single agent trials are available, then informative priors can be placed on the parameters ρ_{00} , ρ_{10} , and ρ_{01} . We also showed how the method is adapted to the case of discrete dose combinations and proposed a statistic (percent selection) to estimate the probability that a prospective trial will recommend a set of dose combinations that are all MTDs for a given scenario. The percent selection can be used in conjunction with the accuracy index and safety measures when designing prospective trials and evaluating the performance of a design relative to other approaches. We also emphasize the fact that unlike previous approaches, increasing the number of discrete dose combinations in the trial does not require us to increase the trial sample size. We are currently extending this methodology to account for patients baseline characteristics as in [22] and late onset toxicity [23].

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

We thank the reviewers and associate editor for their thorough review and valuable comments which helped improve the quality of the manuscript. This work is supported in part by the National Institute of Health Grant Number 1 R01 CA188480-01A1 (M.T, A.R, and Q.L), the National Center for Research Resources, Grant UL1RR033176, and is now at the National Center for Advancing Translational Sciences, Grant UL1TR000124 (M.T and A.R), and 2 P01 CA098912 (M.T).

References

1. Thall PF, Millikan RE, Mueller P, Lee SJ. Dose-finding with two agents in phase I oncology trials. *Biometrics*. 2003; 59:487–496. [PubMed: 14601749]
2. Wang K, Ivanova A. Two-dimensional dose finding in discrete dose space. *Biometrics*. 2005; 61:217–222. [PubMed: 15737096]
3. Yuan Y, Yin G. Sequential continual reassessment method for two-dimensional dose finding. *Statistics in Medicine*. 2008; 27:5664–5678. [PubMed: 18618901]
4. Yin GS, Yuan Y. A Latent Contingency Table Approach to Dose Finding for Combinations of Two Agents. *Biometrics*. 2009; 65:866–875. [PubMed: 18759848]
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. [PubMed: 19995354]
7. Wages NA, Conaway MR, O'Quigley J. Continual Reassessment Method for Partial Ordering. *Biometrics*. 2011; 67:1555–1563. [PubMed: 21361888]

8. Wages NA, Conaway MR, O'Quigley J. Dose-finding design for multi-drug combinations. *Clinical Trials*. 2011; 8:380–389. [PubMed: 21652689]
9. Shi Y, Yin G. Escalation with overdose control for phase I drug-combination trials. *Statistics in Medicine*. 2013; 32:4400–4412. [PubMed: 23630103]
10. Tighiouart M, Piantadosi S, Rogatko A. Dose finding with drug combinations in cancer phase I clinical trials using conditional escalation with overdose control. *Statistics in Medicine*. 2014; 33:3815–3829. [PubMed: 24825779]
11. 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. [PubMed: 24828456]
12. Mander AP, Sweeting MJ. A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Statistics in Medicine*. 2015; 34:1261–1276. [PubMed: 25630638]
13. 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]
14. Tighiouart M, Rogatko A, Babb JS. Flexible Bayesian methods for cancer phase I clinical trials. Dose escalation with overdose control. *Statistics in Medicine*. 2005; 24:2183–2196. [PubMed: 15909291]
15. Tighiouart M, Rogatko A. Dose Finding with Escalation with Overdose Control (EWOC) in Cancer Clinical Trials. *Statistical Science*. 2010; 25:217–226.
16. Gasparini M. General classes of multiple binary regression models in dose finding problems for combination therapies. *Journal of the Royal Statistical Society Series C-Applied Statistics*. 2013; 62:115–133.
17. Lunn DJ, Thomas A, Best N, Spiegelhalter D. WinBUGS - A Bayesian modelling framework: Concepts, structure, and extensibility. *Statistics and Computing*. 2000; 10:325–337.
18. Plummer, M. JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. Technische Universit at Wien; City: 2003.
19. Tighiouart M, Rogatko A. Number of Patients per Cohort and Sample Size Considerations Using Dose Escalation with Overdose Control. *Journal of Probability and Statistics*. 2012; 2012:16.
20. Chu PL, Lin Y, Shih WJ. Unifying CRM and EWOC designs for phase I cancer clinical trials. *Journal of Statistical Planning and Inference*. 2009; 139:1146–1163.
21. Cheung, YK. Dose Finding by the Continual Reassessment Method. Chapman & Hall; Boca Raton: 2011.
22. 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.
23. Tighiouart M, Liu Y, Rogatko A. Escalation with overdose control using time to toxicity for cancer phase I clinical trials. *Plos One*. 2014; 9:e93070. [PubMed: 24663812]

Figure 01.

True and estimated MTD curves from $m = 2000$ simulated trials under scenarios (a)–(d). The grey cloud of points represent the doses given to the last cohort of two patients from each simulated trial along with a 90% confidence region.

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 $p = 0.1, 0.2$ under scenarios (a)–(d).

Table 1

Operating characteristics summarizing trial safety for continuous doses.

Scenario ($\rho_{00}, \rho_{01}, \rho_{10}, \eta$)	True Model	Average % DLTs	% Trials: DLT rate > $\theta + 0.05$	% Trials: DLT rate > $\theta + 0.10$
(a) ($10^{-7}, 3 \times 10^{-6}, 3 \times 10^{-6}, 10$)	Logistic	15.83	0.00	0.00
	Probit	15.72	0.00	0.00
	Normal	15.87	0.00	0.00
	Log-Log	15.85	0.00	0.00
(b) (0.01, 0.2, 0.9, 20)	Logistic	31.53	6.90	0.65
	Probit	31.03	2.50	0.00
	Normal	31.25	7.85	0.50
	Log-Log	30.92	3.35	0.20
(c) (0.001, 0.6, 0.01, 10)	Logistic	26.01	0.00	0.00
	Probit	25.93	0.00	0.00
	Normal	26.14	0.35	0.00
	Log-Log	25.84	0.00	0.00
(d) (0.01, 0.2, 0.9, 100)	Logistic	33.57	8.10	0.35
	Probit	32.66	1.95	0.00
	Normal	34.04	12.35	0.40
	Log-Log	32.88	4.30	0.00

Table 2Dose limiting toxicity scenarios with $\theta = 0.2$.

Dose level	1	2	3	4	5	6
Scenario 1						
3	0.08	0.13	0.20	0.29	0.40	0.53
2	0.05	0.08	0.13	0.20	0.29	0.40
1	0.03	0.05	0.08	0.13	0.20	0.29
Scenario 2						
3	0.05	0.08	0.11	0.15	0.21	0.29
2	0.04	0.06	0.09	0.13	0.18	0.25
1	0.04	0.05	0.08	0.11	0.15	0.21
Scenario 3						
3	0.20	0.30	0.41	0.53	0.65	0.70
2	0.10	0.20	0.25	0.32	0.41	0.50
1	0.03	0.05	0.13	0.20	0.27	0.35
Scenario 4						
3	0.20	0.40	0.47	0.56	0.65	0.76
2	0.08	0.13	0.20	0.32	0.41	0.50
1	0.03	0.05	0.08	0.13	0.17	0.20
Scenario 5						
4	0.20	0.29	0.40	0.53		
3	0.13	0.20	0.29	0.40		
2	0.08	0.13	0.20	0.29		
1	0.05	0.08	0.13	0.20		

Table 3

Operating characteristics summarizing trial efficiency and safety for the two methods.

	Accuracy Index			% Selection	Mean % DLTs	% Trials: DLT rate > θ +0.05	% Trials: DLT rate > θ +0.10
	sq	abs	od				
Scenario 1							
Cond.ewoc	0.37	−0.13	0.07	75.2	15.05	0.30	0.00
Two−dimdesign	0.17	−0.45	−0.08	58.8	16.56	0.65	0.00
Scenario 2							
Cond.ewoc	−0.34	−0.63	−0.62	53.6	12.54	0.00	0.00
Two−dimdesign	−0.13	−0.72	−0.82	81.6	13.37	0.08	0.00
Scenario 3							
Cond.ewoc	0.80	0.40	0.46	92.65	19.08	6.30	0.00
Two−dimdesign	0.70	0.20	0.43	68.58	20.08	6.90	0.10
Scenario 4							
Cond.ewoc	0.57	0.10	0.20	45.95	17.10	2.30	0.00
Two−dimdesign	0.62	0.05	0.36	39.93	19.35	3.55	0.00
Scenario 5							
Cond.ewoc	0.54	−0.01	0.23	78.55	16.34	0.01	0.00
Two−dimdesign	0.27	−0.5	−0.12	87.40	17.72	0.88	0.00
Average							
Cond.ewoc	0.39	−0.05	0.07	69.19	16.02	1.78	0.00
Two-dim design	0.33	−0.28	−0.05	67.26	17.42	2.41	0.02