Published in final edited form as:

J R Stat Soc Ser C Appl Stat. 2019 January ; 68(1): 235–250. doi:10.1111/rssc.12294.

Two-stage design for phase I–II cancer clinical trials using continuous dose combinations of cytotoxic agents

Mourad Tighiouart

Cedars-Sinai Medical Center, Los Angeles, USA

Summary

We present a two-stage phase I/II design of a combination of two drugs in cancer clinical trials. The goal is to estimate safe dose combination regions with a desired level of efficacy. In stage I, conditional escalation with overdose control is used to allocate dose combinations to successive cohorts of patients and the maximum tolerated dose curve is estimated as a function of Bayes estimates of the model parameters. In stage II, we propose a Bayesian adaptive design for conducting the phase II trial to determine dose combination regions along the MTD curve with a desired level of efficacy. The methodology is evaluated by extensive simulations and application to a real trial.

Keywords

Cancer Phase I/II trials; Maximum tolerated dose; Escalation with overdose control; Drug combination; Bayesian adaptive design; Cubic Splines

1. Introduction

The primary objective of cancer phase I/II clinical trials is to determine a tolerable dose level that maximizes treatment efficacy. For treatments where efficacy is ascertained in a relatively short period of time such as one or two cycles of therapy, sequential designs for updating the joint probability of toxicity and efficacy and estimating the optimal dose have been studied extensively in the literature for single agent trials, see e.g. Murtaugh and Fisher (1990); Thall and Russell (1998); Braun (2002); Ivanova (2003); Thall and Cook (2004); Chen et al. (2015); Sato et al. (2016). These methods are designed to identify a safe dose that maximizes the probability of treatment response or a pre-defined utility function. More recently, these designs have been extended to dose combination trials where the dose levels of two drugs are allowed to vary during the trial and explored for toxicity and efficacy (Yuan and Yin (2011); Wages and Conaway (2014); Cai et al. (2014), Riviere et al. (2015)) or to determine the optimal treatment regimen Wages et al. (2016). These methods accommodate cytotoxic and biologic drugs and are appropriate for short term efficacy endpoints to allow sequential updating of the probabilities of toxicity and efficacy.

For treatments where response evaluation takes few cycles of therapy, it is standard practice to perform a two-stage design where a maximum tolerable dose (MTD) of a new drug or combinations of drugs is first determined, then this recommended phase II dose is studied in stage II and evaluated for treatment efficacy, possibly using a different population of cancer

patients from stage I, see Rogatko et al. (2008); Le Tourneau et al. (2009); Chen et al. (2012) for a review of such a paradigm. The two-stage design is routinely used in early phase cancer trials investigating a single agent or combinations of two or more drugs except that in stage II, a single dose level is investigated. For drug combination phase I trials where the dose levels of two or more drugs are allowed to vary during the trial, more than one MTD can be recommended at the conclusion of the trial, e.g. Thall et al. (2003); Wang and Ivanova (2005); Braun and Wang (2010); Tighiouart et al. (2014); Mander and Sweeting (2015); Wages (2016); Tighiouart et al. (2017). Choosing a single MTD combination for efficacy study may result in a failed phase II trial since other MTDs may present higher treatment efficacy. Hence, adaptive or parallel phase II trials may be more suitable for searching an optimal dose combination that is well tolerable with desired level of efficacy as mentioned by Thall et al. (2003).

In this article, we present a two-stage design for drug combination trials with continuous dose levels when treatment efficacy is evaluated after three or more cycles of therapy. Examples of early phase trials that use continuous dose levels can be found in Cheng et al. (2004); Borghaei et al. (2009). In the first stage, the method of Tighiouart et al. (2017) is used to estimate the MTD curve. In stage II, a Bayesian adaptive design is proposed to allocate patients to drug combinations along the MTD curve obtained from stage I with the goal of determining a dose with maximum probability of treatment efficacy. To accommodate various functional forms of the efficacy curve, a flexible family of cubic splines is used to model the dose-response relationship. Adaptive randomization is used sequentially after the efficacy status of each cohort of patients is resolved and the efficacy curve is updated in order to minimize the number of patients treated at sub-therapeutic dose combinations. At the end of the trial, a dose combination with maximum probability of response *a posteriori* is selected for further phase III studies. The method is illustrated by an application to a phase I/II trial of the combination of cisplatin and cabazitaxel (ciscab) in patients with advanced prostate cancer. A previous phase I trial with dose expansion cohort explored three dose combinations of these agents and declared the combination cabazitaxel/cisplatin of 15/75 mg/m² to be the MTD although no DLT was observed at this dose in the phase I part of the study. Based on safety and efficacy data from the phase I study and expansion cohort, we hypothesize that a series of tolerable and active dose combinations exist.

The manuscript is organized as follows. In Section 2, we review the stage I design described in Tighiouart et al. (2017) for estimating the MTD curve. Modeling the efficacy curve along the MTD curve and the adaptive design in stage II of the trial are described in Section 3 along with operating characteristics to evaluate the performance of the phase II trial. In Section 4, we apply this methodology to the ciscab trial described above with the goal of finding a tolerable dose combination with maximum benefit rate. Some concluding remarks and a discussion are found in Section 5.

2. Stage I

2.1. Model

Consider the dose-toxicity model of the form

$$P(T = 1 \mid x, y) = F(\eta_0 + \eta_1 x + \eta_2 y + \eta_3 xy), \quad (1)$$

where T is the indicator of DLT, $T = 1$ if a patient given the dose combination (x, y) exhibits DLT within one cycle of therapy, and $T = 0$ otherwise, $x \in [X_{\min}, X_{\max}]$ is the dose level of agent A_1 , $y \in [Y_{\min}, Y_{\max}]$ is the dose level of agent A_2 , and F is a known cumulative distribution function. Here, $X_{\min}, X_{\max}$ and $Y_{\min}, Y_{\max}$ are the lower and upper bounds of the continuous dose levels of agents A_1 and A_2 , respectively. Suppose that the doses of agents A_1 and A_2 are continuous and standardized to be in the interval $[0, 1]$ using the transformations $h_1(x) = (x - X_{\min})/(X_{\max} - X_{\min})$, $h_2(y) = (y - Y_{\min})/(Y_{\max} - Y_{\min})$ and the interaction parameter $\eta_3 > 0$.

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 property to hold is to assume $\eta_1, \eta_2 > 0$. The MTD is defined as any dose combination (x^*, y^*) such that

$$\text{Prob}(T = 1 \mid x^*, y^*) = \theta. \quad (2)$$

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. We reparameterize model (1) in terms of parameters clinicians can easily interpret. One way is to use ρ_{10} , the probability of DLT when the levels of drugs A_1 and A_2 are 1 and 0, respectively, ρ_{01} , the probability of DLT when the levels of drugs A_1 and A_2 are 0 and 1, respectively, and ρ_{00} , the probability of DLT when the levels of drugs A_1 and A_2 are both 0. It can then be shown that the MTD is

$$C = \left\{ (x^*, y^*): y^* = \frac{(F^{-1}(\theta) - F^{-1}(\rho_{00})) - (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00}))x^*}{(F^{-1}(\rho_{01}) - F^{-1}(\rho_{00})) + \eta_3 x^*} \right\}. \quad (3)$$

Following the work of Tighiouart et al. (2017), we assume that $\rho_{01}, \rho_{10}, \eta_3$ 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}) , $\rho_{00}/\min(\rho_{01}, \rho_{10}) \sim \text{beta}(a_3, b_3)$. The prior distribution on the interaction parameter η_3 is a gamma with mean a/b and variance a/b^2 . If $D_k = \{(x_i, y_i, T_i)\}$ is the data after enrolling k patients to the trial, the posterior distribution of the model parameters is

$$\begin{aligned} \pi(\rho_{00}, \rho_{01}, \rho_{10}, \eta_3) &\propto \prod_{i=1}^k (G(\rho_{00}, \rho_{01}, \rho_{10}, \eta_3; x_i, y_i))^{T_i} (1 - G(\rho_{00}, \rho_{01}, \rho_{10}, \eta_3; x_i, y_i))^{1-T_i} \\ &\times \pi(\rho_{01})\pi(\rho_{10})\pi(\rho_{00} \mid \rho_{01}, \rho_{10})\pi(\eta_3), \end{aligned}$$

where

$$G(\rho_{00}, \rho_{01}, \rho_{10}, \eta_3; x_i, y_i) = F(F^{-1}(\rho_{00}) + (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00}))x_i + (F^{-1}(\rho_{01}) - F^{-1}(\rho_{00}))y_i + \eta_3 x_i y_i).$$

(4)

2.2. Trial Design

Dose escalation/de-escalation proceeds 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 , see e.g. Babb et al. (1998); Tighiouart et al. (2005); Tighiouart and Rogatko (2010); Tighiouart and Rogatko (2012). For a given cohort, one subject receives a new dose of agent A_1 for a given dose of agent A_2 that was previously assigned and the other patient receives a new dose of agent A_2 for a given dose of agent A_1 that was previously assigned. The algorithm is described in Tighiouart et al. (2017) and is reviewed here for convenience.

- i. The first two patients receive the same dose combination $(x_1, y_1) = (x_2, y_2) = (0, 0)$ and let $D_2 = \{(x_1, y_1, T_1), (x_2, y_2, T_2)\}$.
- ii. In the second cohort, patients 3 and 4 receive doses (x_3, y_3) and (x_4, y_4) , respectively, where $y_3 = y_1$, $x_4 = x_2$, x_3 is the α -th percentile of $\pi(\Gamma_{A_1|A_2=y_1}|D_2)$, and y_4 is the α -th percentile of $\pi(\Gamma_{A_2|A_1=x_2}|D_2)$. Here, $\pi(\Gamma_{A_1|A_2=y_1}|D_2)$ is the posterior distribution of the MTD of drug A_1 given that the level of drug A_2 is y_1 , given the data D_2 .
- iii. In the i -th cohort of two patients, if i is even, then patient $(2i-1)$ receives dose (x_{2i-1}, y_{2i-3}) , patient $2i$ receives dose (x_{2i-2}, y_{2i}) , where

$$x_{2i-1} = \prod_{\Gamma_{A_1|A_2=y_{2i-3}}}^{-1}(\alpha | D_{2i-2}), y_{2i} = \prod_{\Gamma_{A_2|A_1=x_{2i-2}}}^{-1}(\alpha | D_{2i-2}).$$

Similar dose allocations can be derived when i is odd. Here, $\prod_{\Gamma_{A_1|A_2=y}}^{-1}(\alpha | D)$

denotes the inverse cdf of the posterior distribution $\pi(\Gamma_{A_1|A_2=y}|D)$.

Repeat step (iii) until N patients are enrolled to the trial subject subject to the following stopping rule.

Stopping Rule: We stop enrollment to the trial if $P(P(DLT)|(x, y) = (0, 0)) > \theta + \xi_1 | \text{data} > \xi_2$, 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. The design parameters ξ_1 and ξ_2 are chosen to achieve desirable model operating characteristics.

At the end of the trial, we estimate the MTD curve using (3) as

$$C_{\text{est}} = \left\{ (x^*, y^*) : y^* = \frac{(F^{-1}(\theta) - F^{-1}(\hat{\rho}_{00})) - (F^{-1}(\hat{\rho}_{10}) - F^{-1}(\hat{\rho}_{00}))x^*}{(F^{-1}(\hat{\rho}_{01}) - F^{-1}(\hat{\rho}_{00})) + \hat{\eta}_3 x^*} \right\}, \quad (5)$$

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

Uncertainty about the estimated MTD curve is evaluated by the pointwise average bias and percent selection described in Tighiouart et al. (2017) when assessing the operating characteristics of the design. Specifically, for $i = 1, \dots, m$, let C_i be the estimated MTD curve and C_{true} be the true MTD curve, where m is the number of simulated trials. For every point $(x, y) \in C_{\text{true}}$, let

$$d_{(x,y)}^{(i)} = \text{sign}(y' - y) \times \min_{\{(x^*, y^*) : (x^*, y^*) \in C_i\}} ((x - x^*)^2 + (y - y^*)^2)^{1/2}, \quad (6)$$

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 . Let

$$d_{(x,y)} = \frac{1}{m} \sum_{i=1}^m d_{(x,y)}^{(i)}. \quad (7)$$

Equation (7) can be interpreted as the pointwise average bias in estimating the MTD.

Let (x, y) be the Euclidean distance between the minimum dose combination $(0, 0)$ and the point (x, y) on the true MTD curve and $0 < p < 1$. Let

$$P_{(x,y)} = \frac{1}{m} \sum_{i=1}^m I(|d_{(x,y)}^{(i)}| \leq p \Delta(x, y)). \quad (8)$$

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 . This is interpreted as the pointwise percent selection for a given tolerance p . In this manuscript, we present the operating characteristics of stage I in the context of the ciscab trial in Section 4.

3. Stage II

Let C_{est} be the estimated MTD curve obtained from stage I and suppose it is defined for $x \in [X_1, X_2]$ and $y \in [Y_1, Y_2]$. Here, $[X_1, X_2] \subset [X_{\min}, X_{\max}]$ and $[Y_1, Y_2] \subset [Y_{\min}, Y_{\max}]$. Let E be the indicator of treatment response such as tumor shrinkage, $E = 1$ if we have a positive response after a pre-defined number of treatment cycles, and $E = 0$ otherwise. Let p_0 be the probability of efficacy of the standard of care treatment. We propose to carry out a phase II study to identify dose combinations $(x, y) \in C_{\text{est}}$ such that $P(E = 1 | (x, y)) > p_0$.

3.1. Model

For every dose combination $(x, y) \in C_{\text{est}}$, let x be the unique vertical projection of (x, y) on the interval $[X_1, X_2]$. Next, denote by $z \in [0, 1]$ the standardized dose of $x \in [X_1, X_2]$ using the transformation $z = h_3(x) = (x - X_1)/(X_2 - X_1)$. In the sequel, we will refer to z as dose combination since there is a one-to-one transformation mapping $z \in [0, 1]$ to $(x, y) \in C_{\text{est}}$, $x \in [X_1, X_2]$, $y \in [Y_1, Y_2]$. This process is further illustrated in the supplement using the CisCab trial example described in Section 4. We model the probability of treatment response given dose combination z in C_{est} as

$$P(E = 1 | z, \boldsymbol{\psi}) = F(f(z; \boldsymbol{\psi})), \quad (9)$$

where F is a known link function, $f(z; \boldsymbol{\psi})$ is an unknown function and $\boldsymbol{\psi}$ is an unknown parameter. A flexible way to model the probability of efficacy along the MTD curve is the cubic spline function

$$f(z; \boldsymbol{\psi}) = \beta_0 + \beta_1 z + \beta_2 z^2 + \sum_{j=3}^k \beta_j (z - \kappa_j)_+^3, \quad (10)$$

where $\boldsymbol{\psi} = (\boldsymbol{\beta}, \boldsymbol{\kappa})$, $\boldsymbol{\beta} = (\beta_0, \dots, \beta_k)$, $\boldsymbol{\kappa} = (\kappa_3, \dots, \kappa_k)$ with $\kappa_3 = 0$. Let $D_m = \{(z_i, E_i), i = 1, \dots, m\}$ be the data after enrolling m patients in the trial, where E_i is the response of the i -th patient treated with dose combination z_i and let $\pi(\boldsymbol{\psi})$ be a prior density on the parameter $\boldsymbol{\psi}$. The posterior distribution is

$$\pi(\boldsymbol{\psi} | D_m) \propto \prod_{i=1}^m [F(f(z_i; \boldsymbol{\psi}))]^{E_i} [1 - F(f(z_i; \boldsymbol{\psi}))]^{1-E_i} \pi(\boldsymbol{\psi}). \quad (11)$$

Let p_z be the probability of treatment efficacy at dose combination z and denote by p_0 the probability of efficacy of a poor treatment or treatment not worthy of further investigation. We propose an adaptive design to conduct a phase II trial in order to test the hypothesis

$$H_0: p_z = p_0 \text{ for all } z \text{ versus } H_1: p_z > p_0 \text{ for some } z.$$

3.2. Trial Design

- i. Randomly assign n_1 patients to dose combinations $z_1, \dots, z_{n_1}$ equally spaced along the MTD curve C_{est} so that each combination is assigned to one and only one patient.
- ii. Obtain a Bayes estimate $\hat{\psi}$ of ψ given the data D_{n_1} using (11).
- iii. Generate n_2 dose combinations from the standardized density $R(\ell(z; \hat{\psi}))$ and assign them to the next cohort of n_2 patients.
- iv. Repeat steps (ii) and (iii) until a total of n patients have been enrolled to the trial subject to pre-specified stopping rules.

This algorithm can be viewed as an extension of a Bayesian adaptive design to select a superior arm among a finite number of arms Berry et al. (2011) to selecting a superior arm from an infinite number of arms.

Decision Rule: At the end of the trial, we accept the alternative hypothesis if

$$\text{Max}_z[P(F(f(z; \psi)) > p_0 \mid D_n)] > \delta_u, \quad (12)$$

where δ_u is a design parameter.

Stopping Rules: For ethical considerations and to avoid exposing patients to sub-therapeutic doses, we stop the trial for futility after j patients are evaluable for efficacy if there is strong evidence that none of the dose combinations are promising, i.e. $\text{Max}_z[P(R(\ell(z; \psi)) > p_0 \mid D_j)] < \delta_0$ where δ_0 is a small pre-specified threshold. In cases where the investigator is interested in stopping the trial early for superiority, the trial can be terminated after j patients are evaluable for efficacy if $\text{Max}_z[P(R(\ell(z; \psi)) > p_0 \mid D_j)] > \delta_1$ where $\delta_1 > \delta_u$ is a pre-specified threshold and the corresponding dose combination $z^* = \text{argmax}_u\{P(R(\ell(u; \psi)) > p_0 \mid D_j)\}$ is selected for future randomized phase II or III studies.

4. Application to the CisCab Trial

The methodology described in Sections 2 and 3 was used to design a phase I/II trial of the combination cisplatin and cabazitaxel in patients with prostate cancer with visceral metastasis. A recently published phase I trial of this combination by Lockhart et al. (2014) identified the MTD of cabazitaxel/cisplatin as 15/75 mg/ m^2 . This trial used a “3+3” design exploring 3 pre-specified dose levels 15/75, 20/75, 25/75. In part 1 of the trial, 9 patients were evaluated for safety and no DLT was observed at 15/75 mg/ m^2 . In part 2 of the study, 15 patients were treated at 15/75 mg/ m^2 and 2 DLTs were observed. Based on these results and other preliminary efficacy data, we hypothesize that there exists a series of active dose combinations which are tolerable and active in prostate cancer. Cabazitaxel dose levels will be selected in the interval [10, 25] and cisplatin dose levels were selected in the interval [50, 100] administered intravenously. For stage I, we plan to enroll $N = 30$ patients and estimate the MTD curve and in stage II, $n = 30$ patients will be enrolled to identify dose combinations

along the MTD curve with maximum clinical benefit rate. Clinical benefit is defined as either a complete response, partial response, or stable disease within three cycles of treatment. The probability of a poor clinical benefit is $p_0 = 0.15$ and we expect that a tolerable dose combination achieves a clinical benefit rate of $p = 0.4$.

4.1. Stage I

For stage I, trial design proceeds as described in Section 2.2 with a target probability of DLT $\theta = 0.33$ and a logistic link function for $F(\cdot)$ in (1). DLT is resolved within one cycle (3 weeks) of treatment. The starting dose combination for the first cohort of two patients is 15 mg/ m^2 cabazitaxel and 75 mg/ m^2 cisplatin. The prior distributions were calibrated so that the prior mean probability of DLT at the dose combination 15/75 mg/ m^2 equals the target probability of DLT. Specifically, informative priors were used for the model parameters ρ_{01} , $\rho_{10} \sim \text{beta}(1.4, 5.6)$, and conditional on ρ_{01} , ρ_{10} , $\rho_{00}/\min(\rho_{01}, \rho_{10}) \sim \text{beta}(0.8, 7.2)$ and a vague prior for η_3 was used as in Tighiouart et al. (2017) so that $E(P(DLT|(15; 75))) \approx 0.33$ *a priori*. We derived the operating characteristics of stage I by simulating 2000 trial replicates as in Tighiouart et al. (2017) under four scenarios for the true MTD curve. Due to space consideration, we only present the estimated MTD curves from two scenarios A and B that are expected by the principal investigator of the trial. Figure 1 shows the true and estimated MTD curve obtained using (3) with the parameters ρ_{00} , ρ_{01} , ρ_{10} , η_3 replaced by their posterior median averaged across all 2000 simulated trials. Scenario A shown on the left panel of Figure 1 is a case where the true MTD curve passes through a point very close to the dose combination (15, 75) identified as the MTD from the previous trial. Scenario B shown on the right panel is a case where the MTD curve is way above this dose combination. In each case, the estimated MTD curves are very close to the true MTD curves. This is also evidenced by the pointwise bias and percent selection (graphs included in the supplement). The trial was also safe since the percent of trials with DLT rate above $\theta + 0.1$ were 3.5% for the scenario on the left and 5.0% for the scenario on the right.

4.2. Stage II

4.2.1. Simulation Set-up and Scenarios—For stage II, we present simulations based on six scenarios that include three situations favoring the alternative hypothesis and three instances supporting the null hypothesis. A logistic link function $F(u) = (1 + \exp(u))^{-1}$ is used in (9) and $f(z, \boldsymbol{\psi})$ is modeled as a cubic spline function with two knots in (0,1). This is a very flexible class of efficacy curves and accommodates cases of constant probability of efficacy along the MTD curve, high probability of efficacy around the middle of the MTD curve and high probability of efficacy at one or both edges of the MTD curve. Vague priors are placed on the model parameters by assuming that $\boldsymbol{\beta} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_6)$ with $\sigma^2 = 10^4$ and $(\kappa_4, \kappa_5) \sim \text{Unif}\{(u, v): 0 \leq u < v \leq 1\}$. It can be shown that the induced prior mean and variance of the probability of treatment response are $E_{\text{prior}}(F(f(z, \boldsymbol{\psi}))) \approx 0.5$ and $\text{Var}_{\text{prior}}(F(f(z, \boldsymbol{\psi}))) \approx 0.25$ for all dose combinations $z \in [0, 1]$. The trial sample size is $n = 30$, the initial number of patients enrolled to the trial was set to $n_1 = 10$, and $n_2 = 5$ was used in the adaptive randomization phase of the design. The design parameter for the decision rule in (12) was taken as $\delta_u = 0.8$. In each scenario, we simulated $M = 2000$ trial replicates. The true probability of response curves under scenarios (a,b,c) are shown in blue in Figure 2. The black horizontal lines correspond to the probability of a poor treatment response $p_0 = 0.15$

and the green horizontal lines represent the target probability of response $p = 0.4$. Scenario (a) is a case where the probability of efficacy is maximized near the middle of the estimated MTD curve with dose combinations in the interval (0.03, 0.76) having probability of efficacy greater than $p_0 = 0.15$. The target probability of response is achieved at a single dose combination $z = 0.42$. Scenario (b) is a case where higher doses of cisplatin and lower doses of cabazitaxel achieve higher efficacy. Specifically, standardized dose combinations in the interval (0.00, 0.49) have probability of efficacy greater or equal to $p_0 = 0.15$. Scenario (c) is an unusual situation where the probability of efficacy is maximized at the edges of the MTD curve. In this case, dose combinations in the interval (0.00, 0.41) $\cup$ (0.90, 1.00) have probability of efficacy greater or equal to $p_0 = 0.15$. Corresponding to these scenarios are situations (d–f) favoring the null hypothesis shown in Figure 2 (d,e,f). The true probability of response curves shown in blue have been shifted downward so that the probability of response equals to p_0 at one dose combination only for scenarios (d), (e), and (f).

4.2.2. Operating Characteristics—For each scenario favoring the alternative hypothesis, we estimate the Bayesian power as

$$\text{Power} \approx \frac{1}{M} \sum_{i=1}^M I[\text{Max}_z \{P(F(f(z; \boldsymbol{\psi}_i)) > p_0 \mid D_{n,i})\} > \delta_u], \quad (13)$$

where $P(F(f(z; \boldsymbol{\psi}_i)) > p_0 \mid D_{n,i})$ is estimated using an MCMC sample of $\boldsymbol{\psi}_i$,

$$P(F(f(z; \boldsymbol{\psi}_i)) > p_0 \mid D_{n,i}) \approx \frac{1}{L} \sum_{j=1}^L I[F(f(z; \boldsymbol{\psi}_{i,j})) > p_0], \quad (14)$$

where $\boldsymbol{\psi}_{i,j}$, $j = 1, \dots, L$ is an MCMC sample from the i -th trial. For scenarios favoring the null hypothesis, (13) is the estimated Bayesian type I error probability. The optimal or target dose from the i -th trial is

$$z_i^* = \text{argmax}_v \{P(F(f(v; \boldsymbol{\psi}_i)) > p_0 \mid D_{n,i})\}. \quad (15)$$

We also report the estimated efficacy curve by replacing $\boldsymbol{\psi}$ in (9) by the average posterior medians across all simulated trials

$$F(f(z; \bar{\boldsymbol{\psi}})), \quad (16)$$

where $\bar{\psi} = (\bar{\beta}, \bar{\kappa})$, $\bar{\beta}_l = M^{-1} \sum_{i=1}^M \hat{\beta}_{i,l}$, $l = 0, \dots, 5$, $\bar{\kappa}_k = M^{-1} \sum_{i=1}^M \hat{\kappa}_{i,k}$, $l = 4, 5$ and $\hat{\beta}_{i,l}$, $\hat{\kappa}_{i,k}$ are the posterior medians from the i -th trial. Finally, we also report the mean posterior probability of declaring the treatment as efficacious for all dose combination z as

$$\frac{1}{M} \sum_{i=1}^M P(F(f(z; \psi_i)) > p_0 \mid D_{n,i}). \quad (17)$$

4.2.3. Results—The estimated efficacy curves shown in black dashed-lines in Figure 2 computed using equation 16 are fairly close to the true probability of efficacy curve in all scenarios except for scenario (c) near the lower edge of the MTD curve. The mean posterior probability of efficacy curve shown in red dashed-line computed using equation 17 is 80% or more at dose combinations where the true probability of efficacy is maximized for scenarios (a,b) and close to 80% for scenario (c). Similar conclusions can be drawn for scenarios favoring the null hypothesis where the maximum of the mean posterior probability of efficacy is less than 50%. Figure 3 is the estimated density of the target dose z^* defined in equation 15 under scenarios favoring the alternative hypothesis (ac) and the shaded region corresponds to dose combinations with corresponding true probability of efficacy greater than $p_0 = 0.15$. The mode of these densities are close to the target doses. Moreover, the estimated probabilities of selecting a dose with true probability of efficacy greater than $p_0 = 0.15$ varies between 0.90 and 0.96 across the three scenarios. The Bayesian power for scenarios (a–c) and type I error probability for scenarios (d–f) estimated using equation 13 using a threshold $\delta_u = 0.8$ are reported in Table 1. Power varies between 0.81 and 0.92 and the type I error probability varies between 0.10 and 0.19. The coverage probability in the last column of Table 1 is the estimated probabilities of selecting a dose with true probability of efficacy greater than $p_0 = 0.15$.

We carried out a sensitivity analysis with respect to the initial number of patients enrolled to the trial n_1 and the cohort size used in the adaptive randomization phase of the design n_2 . Table 2 gives the power, type I error and coverage probabilities as n_1 varies between 5 and 20 and n_2 varies between 3 and 7. For scenario (a), power is about 2% higher when using $n_2 = 3, 7$ relative to the other three cases but the type I error (scenario (d)) increases significantly. Under scenario (b), there is a 3% drop in power when $n_1 = 20$ relative to the case $n_1 = 10$, $n_2 = 5$. The type I error under scenario (e) is fairly constant. For scenario (c), an even higher drop in power of 6% is observed when $n_1 = 20$ relative to the case $n_1 = 10$, $n_2 = 5$ with no significant changes in the type I error probabilities under scenario (f). In all cases, the coverage probabilities are not affected by n_1 and n_2 . These findings show that the effects of n_1 and n_2 on the power depend on the scenarios under study, at least for a final sample size $n = 30$. The drop in power of 6% under scenario (c) suggests that in practice, the initial number of patients n_1 to be randomized to dose combinations equally spaced along the estimated MTD curve be selected between 1/4 and 1/3 of the total sample size in stage II. For binary efficacy outcomes, it is well known that outcome adaptive randomization is not suitable for fast accruing trials or for trials with a longer time τ to resolve efficacy outcome, e.g., response at 6 months or one year disease free survival. Otherwise, we recommend that

the number of patients n_2 to use for the adaptive phase of the trial to be around the number of patients expected to be accrued within the interval $[0, \tau]$. For this trial, the PI expects to enroll about 2 patients per month. Since the length of a cycle is 3 weeks and clinical benefit is resolved withing 3 cycles of therapy, n_2 should be selected between 4 and 5 patients. It is important to note that these guidelines do depend on p_0 , the effect size, and total trial sample size and the investigators have to evaluate the corresponding operating characteristics taking into account the accrual rate and length of the trial.

Finally, we assessed the uncertainty about the estimated MTD curve in stage II in terms of safety of the phase II trial as determined by first cycle DLT under scenario A shown in Figure 1 (A). To estimate how likely a dose used in the phase II trial is toxic taking into account the pointwise percent selection given in (8), for each dose $x_{i,j} \in (0, 1)$ from the i -th simulated trial allocated to the j -th patient using the algorithm in Section 3.2, $i = 1, \dots, M, j = 1, \dots, n$, let $(x_{i,j}, y_{i,j})$ be the dose combination on the true MTD curve shown in Figure 1 (A). Next, generate a dose combination $(x_{i,j}^*, y_{i,j}^*)$ uniformly inside the circle with center $(x_{i,j}, y_{i,j})$ and radius $p(x_{i,j}, y_{i,j})$, where the tolerance probability p and distance (x, y) are defined in Section 2.2. A first cycle DTL $T_{i,j}$ is then generated from model (1) with $(\eta_0, \eta_1, \eta_2, \eta_3)$ replaced by the true values defining scenario A and $(x, y) = (x_{i,j}^*, y_{i,j}^*)$. Safety of the phase II trial is summarized by reporting the average DLT rate $(Mn)^{-1} \sum_{i=1}^M \sum_{j=1}^n T_{i,j}$ and the percent of trials with an excessive DLT rate, i.e., a DLT rate exceeding $\theta + 0.1$. These statistics are reported in Table 3 under the six scenarios and various values of n_1 and n_2 when the tolerance probability is $p = 0.05$. As expected, the mean first cycle DLT rate is essentially equal to the target probability of DLT $\theta = 0.33$. The percent of trials with excessive DLT rate varies between 16% and 19% when $n_1 = 10, n_2 = 5$. These rates are not constant due to the dependence of the radius of the circle on the location of the dose combination on the MTD curve. We conclude that the phase II trial is safe with respect to first cycle DLT rate. Since treatment in the ciscab consists of three cycles of therapy, a separate stopping rule for safety was included in the clinical protocol. Specifically, let θ^* be the true probability that a DLT occurs within three cycles of therapy. The trial will be stopped for safety if $P(\theta^* > 0.43 | \text{data}) > 0.9$. A Bayesian continuous monitoring was used by evaluating this rule after every 6-th patient is evaluable for DLT with a uniform prior on θ^* . Other simulations testing the same hypotheses showed similar results. We conclude that the design has good operating characteristic in identifying tolerable dose combinations with maximum benefit rate. The trial design has recently been approved by the scientific review committee of Cedars-Sinai Medical Center and is scheduled to open enrollment. Operating characteristics of the phase II trial design under a smaller effect size and different values for the probability of a poor treatment response p_0 can be found in the supplement.

5. Discussion

We described a two-stage Bayesian adaptive design for cancer phase I clinical trials using two drugs with continuous dose levels. The goal is to (1) estimate the MTD curve in the two-dimensional Cartesian plane and (2) search for dose combination regions along the MTD curve that yield a desired probability of treatment response. Design of the phase I trial and

estimation of the MTD curve in stage I was carried out using EWOC (Tighiouart et al. (2014, 2017)). In the context of the ciscab trial, we showed that good operating characteristics are obtained using informative prior distributions and sample size of 30 patients. In stage II, we modeled treatment efficacy as a binary indicator of treatment response using a cubic spline form of the dose combination-treatment response relationship. This is a very flexible form for the efficacy curve since it accommodates cases of constant probability of efficacy along the MTD curve, high probability of efficacy around the middle of the MTD curve, high probability of efficacy at the edges of the MTD curve. In this stage, a Bayesian adaptive design is proposed to conduct a phase II trial with the goal of identifying dose combination regions that yield a desired probability of treatment response. Initially, a number of patients are treated with dose combinations equally spaced along the estimated MTD curve from stage I and after resolving their treatment response status, the estimated probability of efficacy curve is updated. A small number of patients are then allocated to dose combinations generated from this updated efficacy curve. The trial continues until we reach the final sample size. This design can be viewed as an extension of the Bayesian adaptive design comparing a finite number of arms (Berry et al. (2011)) to comparing an infinite number of arms. In particular, if the dose levels of the two agents are discrete, then methods such as the ones described in Thall et al. (2003); Wang and Ivanova (2005); Wages (2016) can be used to identify a set of MTDs in stage I and the trial in stage II can be done using adaptive randomization to select the most efficacious dose. Unlike phase I/II designs that use toxicity and efficacy data simultaneously and require a short period of time to resolve efficacy status, the use of a two-stage design is sometimes necessary in practice if it takes few cycles of therapy to resolve treatment efficacy or if the populations of patients in phase I and II are different. In fact, for the ciscab trial described in Section 4, efficacy is resolved after three cycles (9 weeks) of treatment and patients in stage I must have metastatic, castration-resistant prostate cancer whereas patients in stage II must have visceral metastasis.

We studied the properties of this design under various scenarios for the true probability of efficacy as a function of dose combinations and we found that the method yields high power and small type I error probability using a sample size of 40 patients with a target effect size of 0.2. For the ciscab trial, a sample size of 30 patients with effect size of 0.25 yield reasonable operating characteristics. Based on these scenarios and proposed models, we conclude that this two-stage design is feasible with sample size ranging between 60 and 80 patients. The uncertainty of the estimated MTD curve in stage I is not taken into account in stage II of the design in the sense that the MTD curve is not updated as a result of observing DLTs in Stage 2. This is a limitation of this approach since patients in stage II may come from a different population and may have different treatment susceptibility relative to patients in stage I. This problem is also inherent to single agent two-stage designs where the MTD from the phase I trial is used in phase II studies and safety is monitored continuously during this phase. An alternative design would account for first, second, and third cycle DLT in addition to efficacy outcome at each cycle. In addition, the nature of DLT (reversible vs. non-reversible) should be taken into account since patients with a reversible DLT are usually treated for that side effect and kept in the trial with dose reduction in subsequent cycles. For single agent trial and a discrete set of doses, Lee et al. (2015) introduced such designs that

account for cycles one and two toxicities and efficacy. The authors noted the increase in model complexity if the number of cycles is increased to three and this would ultimately require a more parsimonious model. For drug combinations with continuous dose levels and three cycles of therapy, another layer of model complexity will be introduced and such designs are beyond the scope of this manuscript and are subjects of future research. For the ciscab trial, a separate stopping rule using Bayesian continuous monitoring for excessive toxicity is included in the clinical protocol. We further emphasize the fact that phase I/II designs such as the ones described in Wages and Conaway (2014); Cai et al. (2014) or Wages et al. (2016) require a short period of time to resolve efficacy endpoint and may not be appropriate in the setting of the ciscab trial. Finally, we note that our proposed method is applicable to any combination of cytotoxic and biologic agents since the models used in the phase I part of the study still assume monotonicity of the probability of DLT of either agent for fixed level of the other agent. We plan to extend this work to time to event endpoints such as progression free survival and to accommodate patients baseline characteristics such as biomarker expression that is known to be a target of one or both drugs.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

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 R01 CA188480-01A1, the National Center for Research Resources, Grant UL1RR033176, and is now at the National Center for Advancing Translational Sciences, Grant UL1TR000124, and P01 CA098912.

References

- Babb J, Rogatko A, Zacks S. 1998; Cancer phase I clinical trials: efficient dose escalation with overdose control. *Statistics in Medicine*. 17:1103–1120. [PubMed: 9618772]
- Berry, S, Carlin, B, Lee, J, Muller, P. Bayesian Adaptive Methods for Clinical Trials. Boca Raton, FL: Chapman & Hall; 2011.
- Borghaei H, Alpaugh K, Hedlund G, Forsberg G, Langer C, Rogatko A, Hawkins R, Dueland S, Lassen U, Cohen R. 2009; Phase I dose escalation, pharmacokinetic and pharmacodynamic study of naptumomab estafenatox alone in patients with advanced cancer and with docetaxel in patients with advanced non-small-cell lung cancer. *Journal of Clinical Oncology*. 27:4116–4123. [PubMed: 19636016]
- Braun T. 2002; The bivariate continual reassessment method: extending the CRM to phase I trials of two competing outcomes. *Controlled Clinical Trials*. 23:240–256. [PubMed: 12057877]
- Braun TM, Wang SF. 2010; A hierarchical Bayesian design for phase I trials of novel combinations of cancer therapeutic agents. *Biometrics*. 66:805–812. [PubMed: 19995354]
- Cai C, Yuan Y, Ji Y. 2014; A Bayesian dose finding design for oncology clinical trials of combinational biological agents. *Applied Statistics*. 63:159–173. [PubMed: 24511160]
- Chen Z, Yuan Y, Li Z, Kutner M, Owonikoko T, Curran W, Khuri F, Kowalski J. 2015; Dose escalation with over-dose and under-dose controls in phase I/II clinical trials. *Contemporary Clinical Trials*. 43:133–141. [PubMed: 26012358]
- Chen Z, Zhao Y, Cui Y, Kowalski J. 2012; Methodology and application of adaptive and sequential approaches in contemporary clinical trials. *Journal of Probability and Statistics*. 2012:20.

- Cheng JD, Babb JS, Langer C, Aamdal S, Robert F, Engelhardt LR, Fernberg O, Schiller J, Forsberg G, Alpaugh RK, Weiner LM, Rogatko A. 2004; Individualized patient dosing in phase I clinical trials: The role of EWOC in PNU-214936. *J Clin Oncol*. 22:602–609. [PubMed: 14966084]
- Ivanova A. 2003; A new dose-finding design for bivariate outcomes. *Biometrics*. 59:1001–1007. [PubMed: 14969479]
- Le Tourneau C, Lee JJ, Siu LL. 2009; Dose escalation methods in phase I cancer clinical trials. *Journal of the National Cancer Institute*. 101:708–720. [PubMed: 19436029]
- Lee J, Thall P, Ji Y, Müller P. 2015; Bayesian dose-finding in two treatment cycles based on the joint utility of efficacy and toxicity. *Journal of the American Statistical Association*. 110:711–722. [PubMed: 26366026]
- Lockhart A, Sundaram S, Sarantopoulos J, Mita M, Wang-Gillam A, Moseley J, Barber S, Lane A, Wack C, Kassalow L, JFD, Mita A. 2014; Phase I dose-escalation study of cabazitaxel administered in combination with cisplatin in patients with advanced solid tumors. *Investigational New Drugs*. 32:1236–1245. [PubMed: 25117475]
- Mander A, Sweeting M. 2015; A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Statistics in Medicine*. 34:1261–1276. [PubMed: 25630638]
- Murtaugh P, Fisher L. 1990; Bivariate binary models of efficacy and toxicity in dose-ranging trials. *Communications in Statistics Theory and Methods*. 19:2003–2020.
- Riviere M, Yuan Y, Dubois F, Zohar S. 2015; A Bayesian dose-finding design for clinical trials combining a cytotoxic agent with a molecularly targeted agent. *Journal of the Royal Statistical Society: Series C*. 64:215–229.
- Rogatko A, Gosh P, Vidakovic B, Tighiouart M. 2008; Patient-specific dose adjustment in the cancer clinical trial setting. *Pharmaceutical Medicine*. 22:345–350.
- Sato H, Hirakawa A, Hamada C. 2016; An adaptive dose-finding method using a change-point model for molecularly targeted agents in phase I trials. *Statistics in Medicine*. 35:4093–4109. [PubMed: 27221807]
- Thall P, Cook J. 2004; Dose-finding based on efficacy-toxicity trade-offs. *Biometrics*. 60:684–693. [PubMed: 15339291]
- Thall P, Russell K. 1998; A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics*. 54:251–264. [PubMed: 9544520]
- Thall PF, Millikan RE, Mueller P, Lee SJ. 2003; Dose-finding with two agents in phase I oncology trials. *Biometrics*. 59:487–496. [PubMed: 14601749]
- Tighiouart M, Li Q, Rogatko A. 2017; A Bayesian adaptive design for estimating the maximum tolerated dose curve using drug combinations in cancer phase I clinical trials. *Statistics in Medicine*. 36:280–290. [PubMed: 27060889]
- Tighiouart M, Piantadosi S, Rogatko A. 2014; Dose finding with drug combinations in cancer phase I clinical trials using conditional escalation with overdose control. *Statistics in Medicine*. 33:3815–3829. [PubMed: 24825779]
- Tighiouart M, Rogatko A. 2010; Dose finding with escalation with overdose control (EWOC) in cancer clinical trials. *Statistical Science*. 25:217–226.
- Tighiouart M, Rogatko A. 2012; Number of patients per cohort and sample size considerations using dose escalation with overdose control. *Journal of Probability and Statistics*. 2012:16.
- Tighiouart M, Rogatko A, Babb JS. 2005; Flexible Bayesian methods for cancer phase I clinical trials. Dose escalation with overdose control. *Statistics in Medicine*. 24:2183–2196. [PubMed: 15909291]
- Wages N. 2016; Identifying a maximum tolerated contour in two-dimensional dose finding. *Statistics in Medicine*. 36:242–253. [PubMed: 26910586]
- Wages N, Conaway M. 2014; Phase I/II adaptive design for drug combination oncology trials. *Statistics in Medicine*. 33:1990–2003. [PubMed: 24470329]
- Wages N, Slingsluff C, Petroni G. 2016; A Phase I/II adaptive design to determine the optimal treatment regimen from a set of combination immunotherapies in high-risk melanoma. *Contemporary Clinical Trials*. 41:172–179.
- Wang K, Ivanova A. 2005; Two-dimensional dose finding in discrete dose space. *Biometrics*. 61:217–222. [PubMed: 15737096]

Yuan Y, Yin G. 2011; Bayesian phase I/II adaptively randomized oncology trials with combined drugs. *Annals of Applied Statistics*. 5:924–942. [PubMed: 22375162]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Fig. 1.

True and estimated MTD curve under two different scenarios for the MTD curve. The grey diamonds represent the last dose combination from each simulated trial along with a 90% confidence region.

Fig. 2.
True and estimated efficacy curve under six scenarios favoring the null and alternative hypotheses.

Fig. 3.
Estimated density of the target dose combination under three scenarios favoring the alternative hypothesis.

Table 1

Bayesian power, type I error and coverage probabilities.

Scenarios	Power	Scenarios	Prob(Type I error)	Coverage Prob.
(a)	0.896	(d)	0.100	0.964
(b)	0.921	(e)	0.190	0.897
(c)	0.810	(f)	0.143	0.937

Table 2

Sensitivity analysis for n_1 and n_2 .

Scenario	n_1	n_2	Power	Scenario	Prob(Type I error)	Coverage Prob.
(a)	10	5	0.896	(d)	0.100	0.964
	9	3	0.916		0.170	0.946
	9	7	0.915		0.159	0.944
	5	5	0.898		0.091	0.954
	20	5	0.906		0.119	0.979
(b)	10	5	0.921	(e)	0.190	0.897
	9	3	0.931		0.183	0.899
	9	7	0.930		0.183	0.881
	5	5	0.94		0.183	0.892
	20	5	0.887		0.189	0.904
(c)	10	5	0.810	(f)	0.143	0.937
	9	3	0.810		0.150	0.934
	9	7	0.825		0.136	0.929
	5	5	0.861		0.135	0.959
	20	5	0.750		0.144	0.916

Table 3

Safety assessment on the uncertainty of the MTD curve.

Scenario	n_1	n_2	Ave. DLT rate	% Trials: DLT rate > $\theta + 0.1$
(a)	10	5	34.0	19.0
	9	3	34.0	19.0
	9	7	33.0	19.0
	5	5	34.0	18.0
	20	5	34.0	19.0
(b)	10	5	33.0	16.0
	9	3	33.0	17.0
	9	7	33.0	16.0
	5	5	33.0	17.0
	20	5	33.0	18.0
(c)	10	5	34.0	18.0
	9	3	33.0	17.0
	9	7	34.0	18.0
	5	5	34.0	18.0
	20	5	34.0	18.0
(d)	10	5	33.0	17.0
	9	3	33.0	16.0
	9	7	33.0	18.0
	5	5	33.0	18.0
	20	5	33.0	18.0
(e)	10	5	34.0	19.0
	9	3	33.0	19.0
	9	7	33.0	17.0
	5	5	33.0	16.0
	20	5	34.0	17.0
(f)	10	5	34.0	17.0
	9	3	34.0	16.0
	9	7	33.0	17.0
	5	5	33.0	19.0
	20	5	34.0	20.0