Dose Finding for Drug Combination in Early Cancer Phase I Trials Using Conditional Continual Reassessment Method

Márcio Augusto Diniz, Quanlin-Li and Mourad Tighiouart*

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

Abstract

We describe a dose escalation algorithm for drug combinations in cancer phase I clinical trials. Parametric models for describing the association between the doses and the probability of dose limiting toxicity are used assuming univariate monotonicity of the dose-toxicity relationship. Trial design proceeds using the continual reassessment method, where at each stage of the trial, we seek the dose of one agent with estimated probability of toxicity closest to a target probability of toxicity given the current dose of the other agent. A Bayes estimate of the maximum tolerated dose (MTD) curve is proposed at the conclusion of the trial for continuous doses or a set of MTDs is determined in the case of discrete dose levels. 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 under various model generated scenarios and misspecification. The method is further assessed for varying algorithms enrolling cohorts of two and three patients receiving different doses and compared to previous approaches such as escalation with overdose control and two-dimensional design.

Introduction

Cancer management and treatment has seen major advances over the last two decades with the widespread use of targeted therapy and more recently immunotherapy. These treatments often combine cytotoxic agents with biologic and immunotherapy drugs, and possibly radiation because it is known that combining several drugs can help reduce tumor resistance by targeting different signaling pathways simultaneously. Using the optimal combination of these drugs in large phase II and III trials is a challenging problem and require safe, efficient, and robust designs in early phase cancer trials. Unlike single agent phase I trials where the ordering of the doses with respect to the probability of dose limiting toxicity (DLT) is completely specified [1-3], dose combination trials where the doses of two or more agents are allowed to vary during the trial imply a partial ordering of the doses. Hence, many strategies for exploring the space of dose combinations safely are possible and not one algorithm described in the literature seems to perform uniformly better than the others in estimating the maximum tolerated dose. In general, parametric model based designs that link the dose combination-toxicity relationship described [4-13], the partial ordering approach [14,15], and nonparameteric method [16] all proceed by treating successive cohorts of patients with dose escalation starting from the lowest dose combination and the model parameters and estimated probabilities of toxicities sequentially updated. Dose allocation to the next cohort proceeds by using variations of the continual reassessment method (CRM) [17-21] or escalation with overdose control (EWOC) [22-30].

In this manuscript, we examine the performance of the EWOC based designs described in Tighiouart et al. [11,13] when the dose combinations for the next cohort of two patients are determined according to the CRM scheme. In addition, we propose a new dose escalation algorithm that treats cohorts of three patients receiving different dose combinations. Briefly, cohorts of two patients that receive doses along the dose levels of each drug as described in Tighiouart et al. [13] are supplemented by a third patient who will be treated with a dose combination along the diagonal defined by the two drugs. The severity of such an escalation scheme depends on the nature

of the agents under consideration. For instance, this approach may not be appropriate for two cytotoxic drugs but can be an option when studying two biologic agents since trial duration may be shortened considerably. The performance of these extensions will be evaluated by extensive simulations and compared to the approach of Tighiouart et al. [13] and Wang and Ivanova [5] with respect to safety of the trial and efficiency of the estimated MTD under a large number of practical scenarios.

This paper is organized as follows. In Section 4, we review the class of dose-toxicity models from the previous work and introduce the dose escalation/de-escalation algorithm in the case of continuous dose levels. In Section 5, we derive the operating characteristics of the designs under various scenarios for the location of the true MTD curve and under model misspecification. The algorithm is adapted to the case of discrete dose combinations in Section 4 and compared to previous approaches. Section 6 contains some discussion and plans for future work.

Model

Dose-toxicity model

The class of linear logit models described [9-13] for synergistic drugs will be adopted to describe the dose-toxicity relationship:

$$Prob(Z = 1 \mid x, y) = F(\mu + \beta x + \gamma y + \eta xy). \tag{1}$$

Here, the interaction coefficient η is non-negative, the link function $F(\cdot)$ is known, Z is the binary indicator of DLT, and the continuous dose levels (x, y) of two drugs A and B are bounded in the space $[X_{min}, X_{max}] \times$

$[Y_{min}, Y_{max}]$. In the remaining of this paper, dose levels of these two drugs are standardized to be in the unit interval [0,1]. Following the usual assumption of monotonicity of the probability of DLT as a function of dose for cytotoxic or biologic agents, we presume that for all $y \in [0,1]$, $Prob(Z=1|x,y)$ is increasing in x and for all $x \in [0,1]$, $Prob(Z=1|x,y)$ is increasing in y. These conditions are satisfied if and only if $\beta > 0$ and $\gamma > 0$. By definition, the MTD is the set C of dose combinations (x^*, y^*) that satisfy:

$$Prob(Z=1 | x^*, y^*) = \theta. \qquad (2)$$

The parameter θ is known as the target probability of DLT and must be pre-specified by the clinician. Typically, θ is set relatively high when the nature of DLTs anticipated in the trial are reversible and low when these are hazardous or life-threatening with common values selected in [0.1,0.4]. It follows from (1) and the definition of the MTD in (2) that

$$C = \left\{ (x^*, y^*) \in [0,1]^2 : y^* = \frac{F^{-1}(\theta) - \mu - \beta x^*}{\gamma + \eta x^*} \right\}. \qquad (3)$$

In order to facilitate prior elicitation on model parameters that have a practical interpretability, model (1) is reparametrized in terms of ρ_{kl}, the probability of DLT when the levels of drugs A and B are k and l, respectively, for $(k,l) \in \{(0,0),(0,1),(1,0)\}$, and the interaction parameter η. This is a one-to-one transformation with inverse function given by

$$\begin{cases} \mu = F^{-1}(\rho_{00}) \\ \beta = F^{-1}(\rho_{10}) - F^{-1}(\rho_{00}) \\ \gamma = F^{-1}(\rho_{01}) - F^{-1}(\rho_{00}). \end{cases} \qquad (4)$$

The MTD set can be re-expressed in terms of these new parameters:

$$C = \left\{ (x^*, y^*) : y^* = \frac{(F^{-1}(\theta) - F^{-1}(\rho_{00})) - (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00}))x^*}{(F^{-1}(\rho_{01}) - F^{-1}(\rho_{00})) + \eta x^*} \right\} \qquad (5)$$

Prior and posterior distributions

Denote by $Dn = \{(x_i, y_i, z_i), i=1,\ldots,n\}$ the data after the DLT status of n patients has been resolved. The likelihood function is Bernoulli

$$L(\rho_{00}, \rho_{01}, \rho_{10}, \eta | D_n) = \prod_{i=1}^{n} (H(\rho_{00}, \rho_{01}, \rho_{10}, \eta; x_i, y_i))^{z_i} \qquad (6)$$
$$\times (1 - H(\rho_{00}, \rho_{01}, \rho_{10}, \eta; x_i, y_i))^{1-z_i}$$

with probability of success

$$H(\rho_{00}, \rho_{01}, \rho_{10}, \eta; x_i, y_i)$$
$$= F\left(F^{-1}(\rho_{00}) + (F^{-1}(\rho_{10}) - F^{-1}(\rho_{00}))x_i + (F^{-1}(\rho_{01}) - F^{-1}(\rho_{00}))y_i + \eta x_i y_i \right). \qquad (7)$$

Since $\beta, \gamma > 0$, the inverse transformation in (4) implies that $0 < \rho_{00} < \min(\rho_{01}, \rho_{10})$. Assuming that $\rho_{01}, \rho_{10}, \eta$ are independent *a priori*, $\rho_{01} \sim beta(a_2, b_2), \rho_{10} \sim beta(a_2, b_2)$, we model the ratio $\rho_{00}/\min(\rho_{01}, \rho_{10}) \sim beta(a_3, b_3)$, given ρ_{01}, ρ_{10}. A gamma prior with mean $E(\eta) = a/b$ and variance $Var(\eta) = a/b^2$ is placed on the interaction coefficient η. The posterior distribution of the four-model parameters is

$$\pi(\rho_{00}, \rho_{01}, \rho_{10}, \eta | D_n) \propto \prod_{i=1}^{n} (H(\rho_{00}, \rho_{01}, \rho_{10}, \eta; x_i, y_i))^{\delta_i} (1 - H(\rho_{00}, \rho_{01}, \rho_{10}, \eta; x_i, y_i))^{1-\delta_i} \qquad (8)$$
$$\times \pi(\rho_{01})\pi(\rho_{10})\pi(\rho_{00} | \rho_{01}, \rho_{10})\pi(\eta).$$

We use JAGS [31] to estimate features of this posterior distribution.

Trial design

The dose escalation/de-escalation algorithm is similar to the one described [13]. The trial enrolls consecutive cohorts of two patients receiving different dose combinations determined using univariate CRM. Specifically,

1. The two patients in the first cohort receive the minimum dose combination available in the trial $(x_i, y_i) = (0,0)$.

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

(a) If i is even, patient $(2i-1)$ receives dose (x_{2i-1}, y_{2i-3}) and patient $2i$ receives dose (x_{2i-2}, y_{2i}), where

$$x_{2i-1} = \underset{u}{argmin} | H(\hat{\rho}_{00}, \hat{\rho}_{01}, \hat{\rho}_{10}, \hat{\eta}; u, y_{2i-3}) - \theta |$$
$$y_{2i} = \underset{v}{argmin} | H(\hat{\rho}_{00}, \hat{\rho}_{01}, \hat{\rho}_{10}, \hat{\eta}; x_{2i-2}, v) - \theta |.$$

(b) If i is odd, patient $(2i-1)$ receives dose (x_{2i-3}, y_{2i-1}) and patient $2i$ receives dose (x_{2i}, y_{2i-2}), where

$$x_{2i} = \underset{u}{argmin} | H(\hat{\rho}_{00}, \hat{\rho}_{01}, \hat{\rho}_{10}, \hat{\eta}; u, y_{2i-2}) - \theta |$$
$$y_{2i-1} = \underset{v}{argmin} | H(\hat{\rho}_{00}, \hat{\rho}_{01}, \hat{\rho}_{10}, \hat{\eta}; x_{2i-3}, v) - \theta |.$$

In (a) and (b) above, the Bayes estimates $(\hat{\rho}_{00}, \hat{\rho}_{01}, \hat{\rho}_{10}, \hat{\eta})$ are the medians of the posterior distribution $\pi((\rho_{00}, \rho_{01}, \rho_{10}, \eta | D_{2i-2})$ and the recommended doses are obtained by minimizing the distance between the plug in Bayes estimate of the probability of DLT and the target probability of DLT θ. This algorithm implies that a patient in the current cohort can be treated at a dose (x,y) if and only if a patient in the previous cohort was treated at a dose on the same horizontal (along drub A) or vertical (along drug B) line within the dose range.

When treating consecutive cohorts of three patients, the first two patients in any given cohort receive dose combinations as in (a) or (b) above and the third patient is given a dose combination (x_{3i}, y_{3i}) along the diagonal defined by the doses of the two agents according to the CRM principle:

$$x_{3i} = \underset{u}{argmin} | H(\hat{\rho}_{00}, \hat{\rho}_{01}, \hat{\rho}_{10}, \hat{\eta}; u, u) - \theta |$$
$$y_{3i} = x_{3i}.$$

3. Repeat step 2 and terminate the trial when a total of a pre-specified number of patients n are enrolled or the following stopping rule holds.

Stopping rule: Due to the monotonicity assumption of the dose-toxicity relationship with respect to each drug, a stopping rule for safety can be tested at the minimum dose combination (0,0). Specifically, enrollment to the trial is stopped if there is statistical evidence that the minimum dose combination available in the trial is too toxic, i.e, $P(P(Z=1|(x,y)=(0,0) > \theta + \delta_1 | data) > \delta_2$. The design parameters δ_1 and δ_2 are chosen to achieve good operating characteristics for a given set of scenarios.

At the conclusion of the trial, we propose the following estimate of the MTD curve

$$\hat{C} = \left\{ (x^*, y^*) : y^* = \frac{(F^{-1}(\theta) - F^{-1}(\hat{\rho}_{00})) - (F^{-1}(\hat{\rho}_{10}) - F^{-1}(\hat{\rho}_{00}))x^*}{(F^{-1}(\hat{\rho}_{01}) - F^{-1}(\hat{\rho}_{00})) + \hat{\eta} x^*} \right\}, \qquad (9)$$

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

Simulation Studies

Simulation set-up and scenarios

To facilitate comparison with the work of Tighiouart et al. [13], we derive the operating characteristics of this design under the same scenarios [13] using the logistic link function $F(u) = (1+e^{-u})^{-1}$ for the working model. The corresponding true parameters $(\rho_{00}, \rho_{01}, \rho_{10}, \eta)$ are shown in Table 1 and the true MTD curves are shown by a solid line in Figure 1. Model misspecification was investigated under the

Scenario (ρ_{00}, ρ_{01}, ρ_{10}, η)	True model	CRM 2p	CRM 3p	EWOC 2p
(a) (10^{-7}, 3×10^{-6}, 3×10^{-6}, 10)	Logistic	14.35 (0.0)	12.07 (0.0)	15.83 (0.0)
	Probit	12.29 (0.0)	11.42 (0.0)	15.72 (0.0)
	Normal	15.01 (0.0)	12.12 (0.0)	15.87 (0.0)
	CLog-log	13.54 (0.0)	11.84 (0.0)	15.85 (0.0)
(b) (0.01, 0.2, 0.9, 20)	Logistic	32.58 (0.1)	31.83 (0.3)	31.53 (0.7)
	Probit	31.99 (0.1)	31.19 (0.0)	31.03 (0.0)
	Normal	32.82 (0.8)	32.00 (0.4)	31.25 (0.5)
	CLog-log	32.23 (0.0)	31.17 (0.0)	30.92 (0.2)
(c) (0.001, 0.6, 0.01, 10)	Logistic	27.68 (0.0)	26.26 (0.0)	26.01 (0.0)
	Probit	27.56 (0.0)	25.71 (0.0)	25.93 (0.0)
	Normal	27.57 (0.0)	25.95 (0.0)	26.14 (0.0)
	CLog-log	27.43 (0.0)	25.78 (0.0)	25.84 (0.0)
(d) (0.01, 0.2, 0.9, 100)	Logistic	34.44 (0.4)	33.79 (0.4)	33.57 (0.4)
	Probit	33.77 (0.0)	32.00 (0.1)	32.66 (0.0)
	Normal	34.62 (0.8)	34.51 (0.6)	34.04 (0.4)
	CLog-log	33.61 (0.2)	32.27 (0.2)	32.89 (0.0)

Table 1: Operating characteristics summarizing average DLT rate (% trials: DLT rate>θ+0.10) for three designs and scenarios (a)-(d) under true models (i)-(iv).
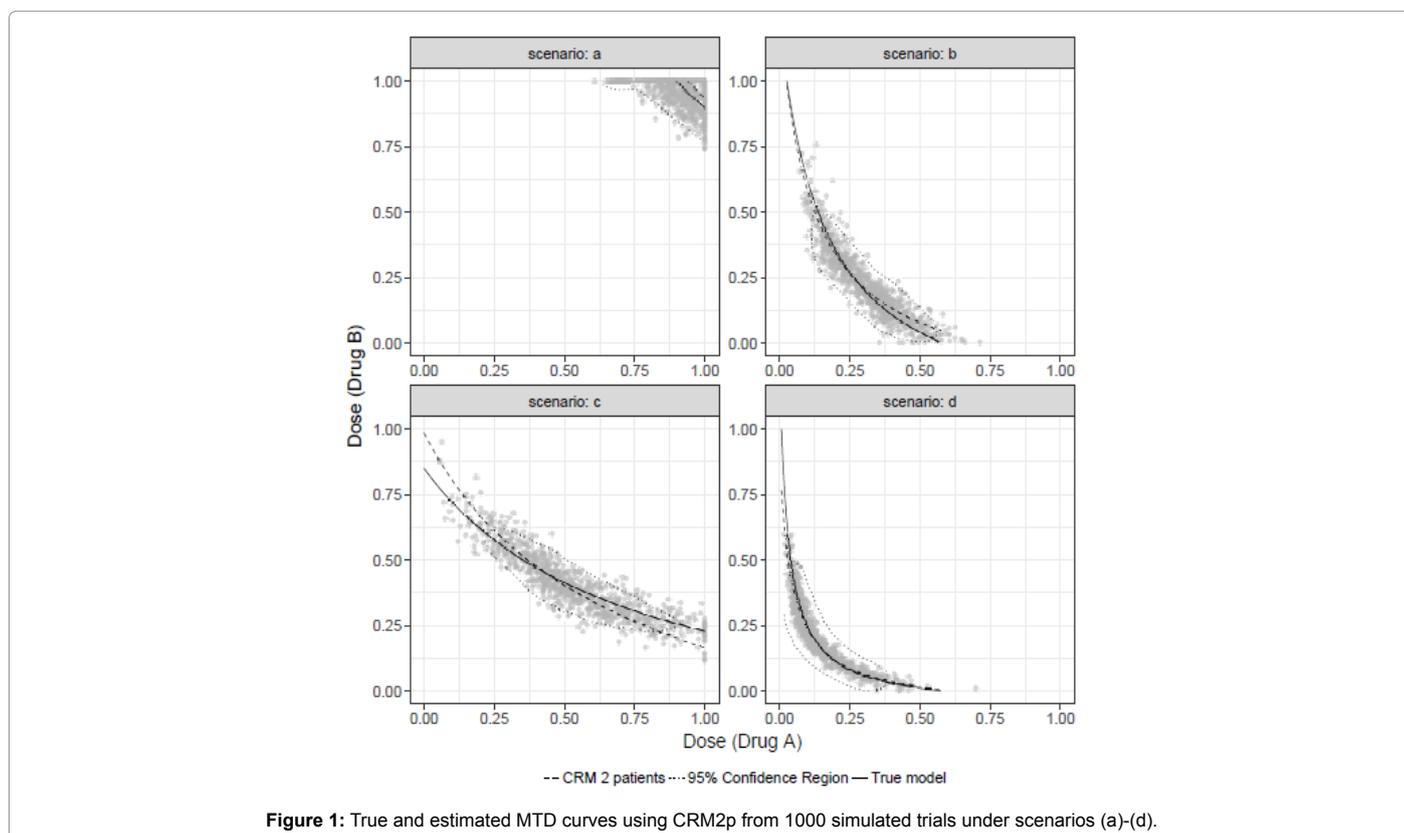
Figure 1: True and estimated MTD curves using CRM2p from 1000 simulated trials under scenarios (a)-(d).

probit link $F(u)=\Phi(u)$, where $\Phi(.)$ is the c.d.f. of the standard normal distribution, normal link $F(u)=\Phi(u/\sigma)$, , and the complementary log-log link $F(u)=1-e^{-e^{u}}$. The target probability of DLT was set to θ=0.33 and the number of trial replicates is m=1000. When treating cohorts of size 2, the trial sample size is n=40 and we choose n=42 when using the algorithm that treats consecutive cohorts of three patients. Vague priors for ρ_{00}, ρ_{01}, ρ_{10} were selected by setting $a_i=b_i=1$, i=1,…,3 and a diffuse prior for η was selected by taking $E(\eta)$=21 and $Var(\eta)$=542, see Tighiouart et al. [13] for the rationale behind this choice.

Design operating characteristics

We evaluate the performance of the CRM based design treating cohorts of two patients (CRM2p) in terms of safety, efficiency of the estimate of the MTD, and under model misspecification. This design is compared to the EWOC based dose escalation algorithm described in Tighiouart et al. [13]. Finally, CRM2p is compared to the CRM based algorithm that treats consecutive cohorts of three patients (CRM3p).

Safety and efficiency: To evaluate trial safety, we report the oberved percent of DLTs averaged across all m=1000 trials and the percent of trials with an excessive rate of DLT, e.g., a DLT rate exceeding θ +, for δ=0.1. This provides an estimate of the probability that a prospective trial will result in an excessive rate of DLT. For efficiency of the estimate of the MTD curve, we use the same summary statistics proposed [13].

The first is an estimate of the MTD curve given by

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

where $\bar{\rho}_{00}, \bar{\rho}_{01}, \bar{\rho}_{10}, \bar{\eta}$ are the average posterior medians of the parameters $\rho_{00}, \rho_{01}, \rho_{10}, \eta$ from all $m=1000$ trials. The next measure is the pointwise average bias

$$d_{(x,y)} = \frac{1}{m} \sum_{i=1}^{m} d_{(x,y)}^{(i)}, \quad (11)$$

where

$$d_{(x,y)}^{(i)} = sign(y' - y) \times min_{\{(x^*,y^*):(x^*,y^*) \in C_i\}} ((x - x^*)^2 + (y - y^*)^2)^{1/2}. \quad (12)$$

In the above equation, C_i is the estimated MTD curve from the i-th trial, C_{true} is the true MTD curve, and y' is such that $(x,y') \in C_i$ for all $(x,y) \in C_{true}$. The last measure of efficiency is the pointwise percent selection for a selected tolerance probability P

$$P_{(x,y)} = \frac{1}{m} \sum_{i=1}^{m} I(|d_{(x,y)}^{(i)}| \le p\Delta(x,y)), \quad (13)$$

where $\Delta(x,y)$ is the Euclidean distance between the minimum dose combination $(0,0)$ and the point (x,y) on the true MTD curve. Detailed description of these summary statistics and their properties can be found [11,13].

Results

CRM2p design

Table 1 shows that the average DLT rates when the true and working models are logistic are close to the target probability of DLT $\theta=0.33$ except for scenario (a) where the true MTD curve is close to the maximum dose combination available in the trial. In this case, the DLT rate is around 14%, consistent with the results obtained [11,13]. These DLT rates remain unchanged under model misspecification and the percent of trials with an excessive DLT rate is very low under all scenarios. We conclude that the CRM2p design is safe for this class of practical scenarios. The estimated MTD curve (dashed line) in Figure 1 is very close to the true MTD curve (solid line) under all scenarios. The gray cloud of points consists of the last two doses of the two patients in the last cohort. The 95% confidence region is fairly tight around the true MTD curve and illustrates the uncertainty of the estimated MTD should a clinician choose to select the last dose given to the last patient in the trial for future phase II trials. Figure 2 shows the pointwise average bias using the true model (logistic) and model missspecification under the four scenarios. Given that the dose range for each agent is the unit interval [0,1], the average bias is negligible under scenarios (a–c). For scenario (d), the average bias tend to be higher near the top of the true MTD curve, consistent with previous results [13]. These results are also consistent with the pointwise percent selection shown in Figure 3. The pointwise percent selection with a tolerance $p=0.1$ are very high for phase I trials ranging from 55% to 100% across all scenarios. The extent of variability of these pointwise percent selections under model misspecification is small except when using the probit model. We conclude that the CRM2p design is safe and achieves good pointwise percent MTD recommendation even under model misspecification.

Comparison of CRM2p and EWOC2p designs

The average DLT rates for the CRM2p design shown in Table 1 are
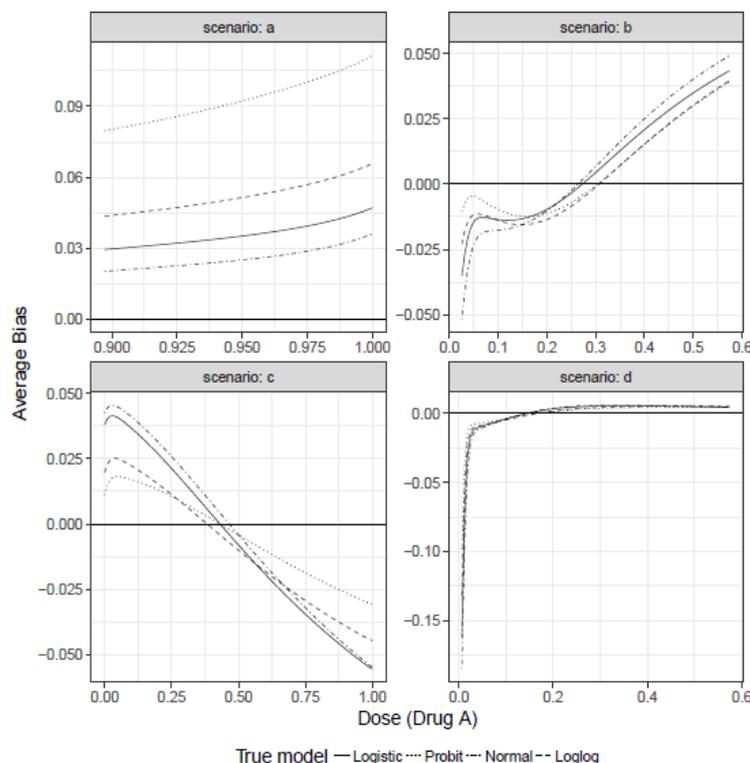
Figure 2: Pointwise average relative minimum distance from the true MTD curve to the estimated MTD curve using CRM2p under scenarios (a)-(d).
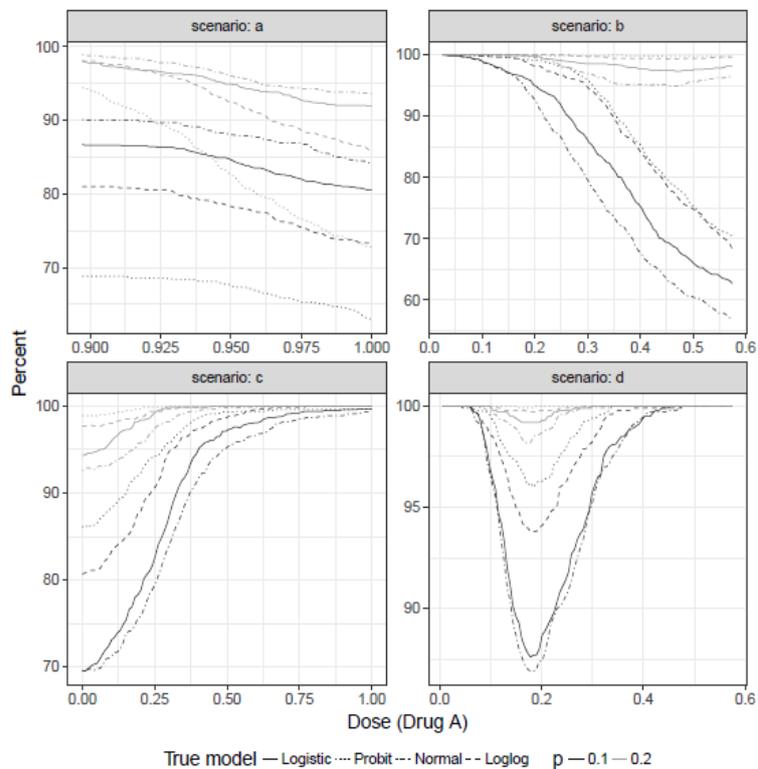
Figure 3: Pointwise percent of MTD recommendation using CRM2p for p=0.1, 0.2 under scenarios (a)-(d).
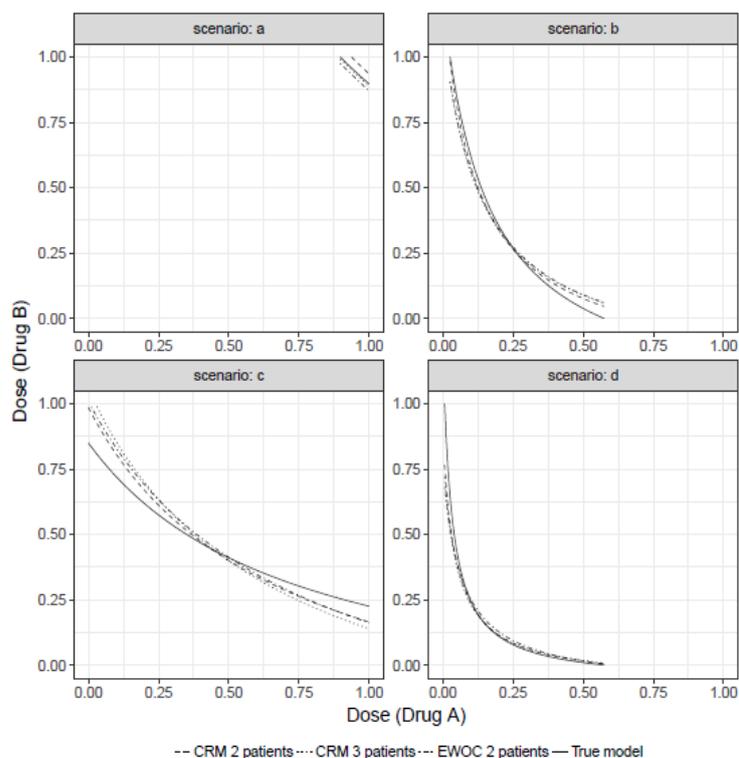
Figure 4: True and estimated MTD curves using CRM2p, CRM3p and EWOC2p from 1000 simulated trials under scenarios (a)-(d).
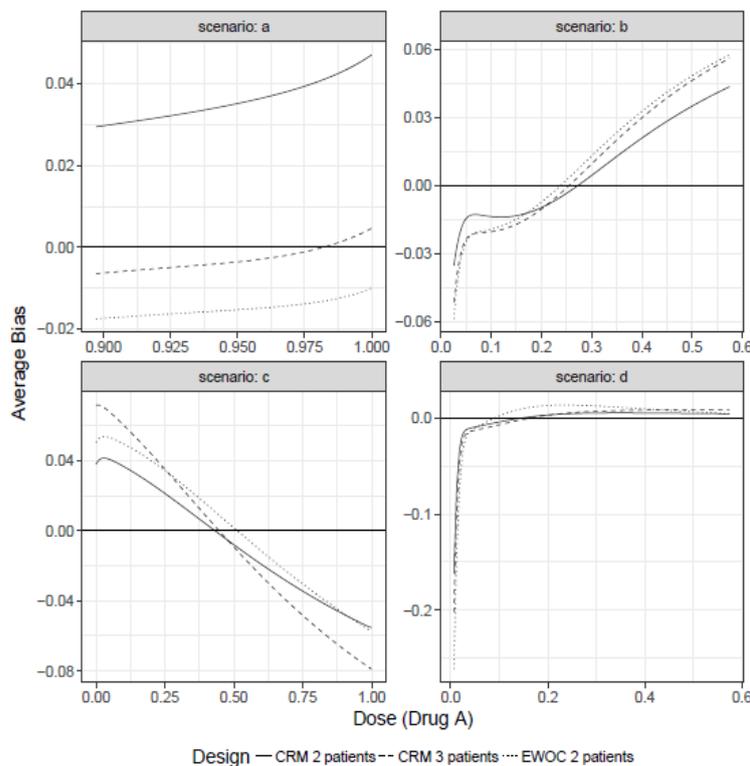
Figure 5: Pointwise average relative minimum distance from the true MTD curve to the estimated MTD curve using CRM2p, CRM3p and EWOC2p under scenarios (a)-(d).

very similar to the EWOC2p design with the largest difference of 3% obtained under scenario (a) using the probit link as the true model. Similarly, the percent of trials with an excessive rate of DLT are similar for the two algorithms. Figure 4 shows the plots of the estimated MTD curves for the CRM2p and EWOC2p designs. These curves are very close to one another across all scenarios. Although the pointwise average bias is still negligible for these two designs as shown in Figure 5, the bias for CRM2p is consistently lower (in absolute value) than the EWOC2p design except for scenarios (a) where the true MTD curve is near the maximum dose combination. This is also reflected in the pointwise percent selection shown in Figure 6 where the percent selection for CRM2p is higher relative to EWOC2p under scenarios (b–d) with the largest difference of 17% obtained near the middle of the MTD curve under scenario (d). Based on the results of these scenarios, CRM2p design outperforms EWOC2p in terms of the efficiency of the estimate of the MTD in general.

Comparison of CRM2p and CRM3p designs

Treating a third patient in a cohort along the diagonal does not seem to affect the safety of the trial. If anything, Table 1 shows that the average DLT rate for CRM3p design is slightly lower than that of CRM2p. Therefore, safety of the trial is not compromised by escalating both drugs along the diagonal according to these four scenarios. As in the previous section, the plots of the estimated MTD curves for the CRM2p and CRM3p designs in Figure 4 show no difference between these estimates. The pointwise average bias for CRM3p is negligible and is consistently higher than CRM2p design except under scenario (a). With respect to the pointwise percent selection shown in Figure 6, there is no clear distinction about the superiority of either design since

this depends on both the scenario under study and the value of the dose combination. For example, under scenario (a), CRM3p outperforms CRM2p across the dose range but in scenarios (b) and (c), CRM2p does better across half the dose range. Given the results for these scenarios, the main advantage of using CRM3p design over CRM2p is to shorten trial duration.

Discrete Dose Combinations

In this section, we review how the methodology can be adapted to a set of discrete pre-specified dose combinations and further compare the CRM2p design to the two-dimensional design [5] in addition to EWOC2p and CRM3p designs.

Approach

Denote by $(x_1,...,x_r)$ and $(y_1,...,y_s)$ the doses of agents A and B, respectively, standardized in the unit interval [0,1]. The CRM algorithm in Section 2.3 is applied by rounding the recommended continuous dose to the nearest discrete dose level sequentially for each patient in the trial. The estimated set of MTDs is obtained as described in Tighiouart et al. [13] and is briefly reviewed here. Let $d((x_j,y_k),\hat{C})$ be the Euclidean distance between dose combination (x_j,y_k) and the estimated MTD curve $\hat{C}$.

[label=()]

1. Let $\Gamma_A = \bigcup_{t=1}^{s}\left\{(x,y_t): x = \underset{x_j}{argmin}\ d((x_j,y_t),\hat{C})\right\}$,

$\Gamma_B = \bigcup_{t=1}^{r}\left\{(x_t,y): y = \underset{y_j}{argmin}\ d((x_t,y_j),\hat{C})\right\}$, and $\Gamma_0 = \Gamma_B \cap \Gamma_A$.
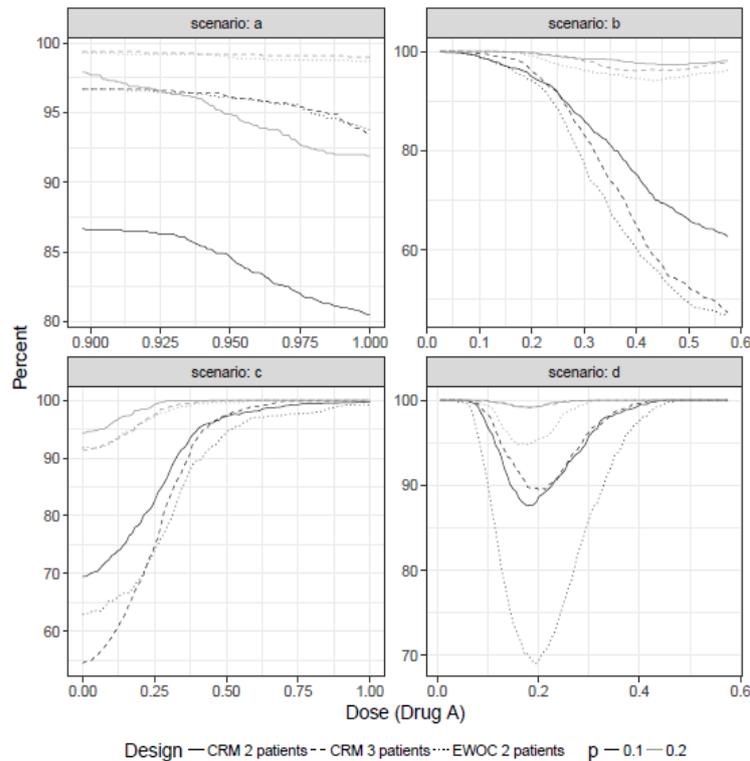
Figure 6: Pointwise percent of MTD recommendation using CRM2p, CRM3p and EWOC2p for p=0.1, 0.2 under scenarios (a)-(d).

2. Let $\Gamma = \Gamma_0 \setminus \{(x^*, y^*) : P(\mid P(DLT \mid (x^*, y^*)) - \theta \mid > \delta_1 \mid D_n) > \delta_2\}$.

The final set Γ consists of the recommended MTDs at the end of the trial with design parameters δ_1, δ_2 selected to achieve good operating characteristics for a pre-specified set of scenarios.

Operating characteristics

The performance of the method is evaluated by calculating the percent of MTDs selection introduced in Tighiouart et al. [13] estimating the probability that for a given scenario, a prospective trial will recommend a set of dose combinations that are all MTDS,

$$PS = \frac{1}{m}\sum_{i=1}^{m}I(\Gamma_i \subset \Gamma_\delta),\qquad(14)$$

where $\Gamma_\delta = \{(x_j, y_k) : \mid P(DLT \mid (x_j, y_k)) - \theta \mid < \delta\}$ is the set of true MTDs for a pre-specified threshold parameter δ set by a clinician independent of the scenarios under study and Γ_i is the estimated MTD set from the i-th trial obtained in (ii) above.

Another measure of percent selection is the probability of obtaining at least K MTDs,

$$PS_K = \frac{1}{m}\sum_{i=1}^{m}I(\mid \Gamma_i \cap \Gamma_\delta \mid \geq K).\qquad(15)$$

In addition, we also define the average proportion of the recommended set of dose combinations that are MTDs,

$$AV = \frac{1}{m}\sum_{i=1}^{m}\frac{\mid \Gamma_i \cap \Gamma_\delta \mid}{\mid \Gamma_i \mid}.\qquad(16)$$

However, such a measure is not well defined if the recommended set Γ_i is empty, thus we consider a weighted average proportion:

$$S_{\Gamma_\delta} = \sum_{i=1}^{m}\frac{\mid \Gamma_i \mid}{\sum_{j=1}^{m}\mid \Gamma_j \mid}\frac{\mid \Gamma_i \cap \Gamma_\delta \mid}{\mid \Gamma_i \mid}$$

$$= \frac{\sum_{=1}\mid \Gamma \cap \Gamma \mid}{\sum_{=1}\mid \quad \mid}.\qquad(17)$$

The statistic S_{Γ_δ} gives an estimate of the probability that any given dose in the set of recommend doses at the end of the trial is an MTD.

Simulation setup and scenarios

We consider five scenarios from Wang and Ivanova [5] presented 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 $\theta=0.2$ and the threshold parameter is fixed at $\delta=0.1$. For each scenario, we simulated $m=1000$ trials using the same vague priors for $\rho_{00}, \rho_{01}, \rho_{10}$ and η from Section 3. Trial sample size is $n=54$ for scenarios 1 - 4 and $n=60$ for scenario 5. The design parameters are $\delta_1=0.1, \delta_2=0.1$ as suggested [13].

Summary statistics for trial safety and efficiency of the estimated set of MTDs are shown for CRM2p, CRM3p, EWOC2p, and the two-dimensional design of Wang and Ivanova [5] (WI2p) in Table 3. The average percent of DLTs for all designs are all less than the target probability of DLT $\theta=0.2$ with WI2p achieving a DLT rate of 20.1% under scenario 3. The percent of trials with an excessive rate of DLT is also small with the highest rate of 7% achieved by CRM2p under scenario 2.

The percent selection when comparing CRM2p with EWOC2p designs are very similar for all five measures of efficiency except for

Dose level	1	2	3	4	5	6
Scenario 1						
3	0.08	0.13	*0.20*	0.29	0.4	0.53
2	0.05	0.08	0.13	*0.20*	0.29	0.4
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	0.2	0.3	0.41	0.53	0.65	0.7
2	0.1	0.2	0.25	0.32	0.41	0.5
1	0.03	0.05	0.13	*0.20*	0.27	0.35
Scenario 4						
3	*0.20*	0.4	0.47	0.56	0.65	0.76
2	0.08	0.13	*0.20*	0.32	0.41	0.5
1	0.03	0.05	0.08	0.13	0.17	0.2
Scenario 5						
4	*0.20*	0.29	0.4	0.53		
3	0.13	*0.20*	0.29	0.4		
2	0.08	0.13	*0.20*	0.29		
1	0.05	0.08	0.13	*0.20*		

Table 2: Dose limiting toxicity scenarios with θ=0.2. True MTD is shown in Italics.

Scenario	Algorithm	PS	PS_3	PS_2	PS_1	S_{18}	Average DLT rate (% Excessive DLT)
01	EWOC 2p	75.2	42.4	75.7	98.3	88.4	15.05 (0.30)
	CRM 2p	75.8	41.4	76	97.7	88.4	15.37 (0.30)
	CRM 3p	76.5	46.4	80.7	97.2	89.3	14.45 (0.10)
	WI 2p	58.8	58.8	93.4	98	83.4	16.6 (0.70)
02	EWOC 2p	53.6	13.7	45.4	97.6	74.7	12.54 (0.00)
	CRM 2p	57	15.6	49	97.6	76.1	12.49 (0.00)
	CRM 3p	71.3	19.2	60.7	96.7	82.6	11.37 (0.00)
	WI 2p	81.6	81.6	94.1	97.6	91.1	13.4 (0.08)
03	EWOC 2p	86.5	39.1	86.9	98.9	93.2	19.08 (6.30)
	CRM 2p	86.1	36.4	83.6	98.8	93.1	19.11 (7.00)
	CRM 3p	81	37.4	78.8	98.3	90.6	18.83 (5.20)
	WI 2p	68.6	68.6	97.6	100	88.7	20.1 (6.90)
04	EWOC 2p	47	27.1	68.8	92.8	73.9	17.10 (2.30)
	CRM 2p	56.2	28.4	69.7	94.6	78.7	17.18 (1.50)
	CRM 3p	58.9	31.9	72.9	94.6	80.7	17.51 (2.70)
	WI 2p	39.9	39.9	89.3	100	76.4	19.4 (3.55)
05	EWOC 2p	95.7	55.5	94.5	99.2	98.1	16.34 (0.50)
	CRM 2p	94.9	57	93.1	98.7	97.6	16.42 (0.30)
	CRM 3p	96.4	55.8	94.2	99.3	98.3	16.32 (0.30)
	WI 2p	87.4	97.6	100	100	96.3	17.7 (0.88)

Table 3: Operating characteristics summarizing trial efficiency and safety for CRM and EWOC using non-informative priors.

PS under scenario 4 where CRM2p exceeds EWOC2p by 9%. On the other hand, CRM3p seems to perform better than CRM2p on the average with the highest difference of 14% achieved under scenario 2 for the statistic *PS*. Similar to the results obtained in Tighiouart et al. [13], WI2p design performs uniformly better than EWOC2p, CRM2p, and CRM3p across all five statistics under scenario 2. For scenarios 1, 3–5, there is no clear advantage of the two-dimensional design WI2p and the results do depend on the statistic of interest. When using informative priors to match the priors [5] as was done [13], the three designs outperform the two-dimensional design for scenarios 1, 3-5 and for most of the efficiency statistics measures (data not shown). We note that here, the percent selections are exceptionally high due to the

fact that the number of true MTDs is very high for all scenarios since the threshold parameter δ=0.1. This means that any dose combination with probability of DLT within 0.1 of the target probability θ=0.2 is considered an MTD. Such an assumption is not uncommon in practice due to the heterogeneity of patients in cancer phase I trials and higher DLT rates are deemed acceptable in phase II trials relative to the target θ that is set in phase I.

Concluding Remarks

The main objectives of this manuscript are to evaluate the performance of the dose escalation algorithm described in Tighiouart et al. [13] for early phase drug combination cancer trials when the CRM scheme is used to estimate the recommended doses for consecutive cohorts of patients and the effect of treating an extra patient along the diagonal defined by the dose levels of the two drugs. We assessed the performance of these new designs under similar parametric models used in the previous approaches under vague prior distributions for the model parameters and settings of continuous and discrete dose levels.

We found that the trials are very safe under EWOC and CRM based designs, even when we treat cohorts of three patients with a third patient along the diagonal, i.e., when increasing the dose levels of both agents. With respect to trial efficiency for continuous dose levels, the CRM2p design outperforms EWOC2p in all scenarios except when the true MTD curve is near the maximum dose combination available in the trial. Given the extend of difference of the percent selection overall, we recommend the use of CRM2p design. Since none of CRM2p and CRM3p outperforms the other, CRM3p design should be used to shorten trial duration without compromising the safety of the trial. Similar conclusions can be drawn in the case of discrete dose combinations because CRM2p had better percent selection *PS* under scenario 4 relative to EWOC2p and CRM3p outperforms CRM2p on the average. We note that these recommendations hold for the class of scenarios used [5] and may not be generalized to settings with different number of dose levels, sample size, target probability of DLT, and prior information about the drugs when used as single agents. As in any early phase cancer trial, a set of all possible practical scenarios must be defined with the help of the clinician and operating characteristics using one or more algorithms derived to select the most appropriate design for the prospective trial.

We are currently working on extending this work to phase I/II trials where a binary efficacy endpoint is assessed relatively quickly and situations where baseline characteristics thought to be related to toxicity and efficacy outcomes are available.

Acknowledgments

References

1. Ting N (2006) Dose finding in drug development. Springer Science & Business Media.

2. Nembhard HB (2008) Statistical methods for dose-finding experiments. Technometrics 50: 409-409.

3. Le Tourneau C, Lee JJ, Siu LL (2009) Dose escalation methods in phase I cancer clinical trials. J Natl Cancer Inst 101: 708-720.

4. Thall PF, Millikan RE, Mueller P, Lee SJ (2003) Dose-finding with two agents in phase I oncology trials. Biometrics 59: 487-496.

5. Wang K, Ivanova A (2005) Two-dimensional dose finding in discrete dose space. Biometrics 61: 17-222.

6. Yin G, Yuan Y (2009) Bayesian dose finding in oncology for drug combinations by copula regression. Journal of the Royal Statistical Society Series C (Applied Statistics) 58: 211-224.

7. Yin G, Yuan Y (2009) A latent contingency table approach to dose finding for combinations of two agents. Biometrics 65: 866-875.

8. Braun TM, Wang S (2010) A hierarchical bayesian design for phase I trials of novel combinations of cancer therapeutic agents. Biometrics 66: 805-812.

9. Shi Y, Yin G (2013) Escalation with overdose control for phase I drug-combination trials. Stat Med 32: 4400-4412.

10. Riviere MK, Yuan Y, Dubois F, Zohar, S (2014) A bayesian dose-finding design for drug combination clinical trials based on the logistic model. Pharmaceutical statistics 13: 247-257.

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

12. Tighiouart M, Li Q, Piantadosi S, Rogatko A (2016) A bayesian adaptive design for combination of three drugs in cancer phase I clinical trials. Am J Biostat 9: 1-11.

13. 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. Stat Med 36: 280-290.

14. Wages NA, Conaway MR, O'Quigley J (2011) Continual reassessment method for partial ordering. Biometrics 67: 1555-1563.

15. Wages N (2016) Identifying a maximum tolerated contour in two-dimensional dose finding. Stat Med 36: 242-253.

16. Mander AP, Sweeting MJ (2015) A product of independent beta probabilities dose escalation design for dual-agent phase i trials. Stat Med 34: 1261-1276.

17. O'Quigley J, Pepe M, Fisher L (1990) Continual reassessment method: a practical design for phase I clinical trials in cancer. Biometrics 46: 33-48.

18. Faries D (1994) Practical modifications of the continual reassessment method for phase I cancer clinical trials. J Biopharm Stat 4: 147-164.

19. Goodman SN, Zahurak ML, Piantadosi S (1995) Some practical improvements in the continual reassessment method for phase I studies. Stat Med 14: 1149-1161.

20. O'Quigley J, Shen LZ (1996) Continual reassessment method: a likelihood approach. Biometrics 52: 673-684.

21. Piantadosi S, Fisher JD, Grossman S (1998) Practical implementation of a modified continual reassessment method for dose-finding trials. Cancer Chemother Pharmacol 41: 429-436.

22. Babb J, Rogatko A, Zacks S (1998) Cancer phase I clinical trials: efficient dose escalation with overdose control. Stat Med 17: 1103-1120.

23. Tighiouart M, Rogatko A, Babb JS (2005) Flexible bayesian methods for cancer phase I clinical trials. Dose escalation with overdose control. Stat Med 24: 2183-2196.

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

25. Chen ZJ, Tighiouart M, Kowalski J (2012) Dose escalation with overdose control using a quasi-continuous toxicity score in cancer phase i clinical trials. Contemp Clin Trials 33: 949-958.

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

27. Tighiouart M, Cook-Wiens G, Rogatko A (2012) Escalation with Overdose Control Using Ordinal Toxicity Grades for Cancer Phase I Clinical Trials. Journal of Probability and Statistics 2012: 18.

28. 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.

29. Tighiouart M, Liu Y, Rogatko A (2014) Escalation with Overdose Control using Time to Toxicity for Cancer Phase I Clinical Trials. PLoS One 9: e93070.

30. Tighiouart M, Cook-Wiens G, Rogatko A (2017) A bayesian adaptive design for cancer phase I trials using a flexible range of doses. J Biopharm Stat 31: 1-13.

31. Plummer M (2003) Jags: A program for analysis of bayesian graphical models using gibbs sampling.