



Published in final edited form as:

Clin Trials. 2011 August ; 8(4): 380–389. doi:10.1177/1740774511408748.

Dose-finding design for multi-drug combinations

Nolan A Wages^a, Mark R Conaway^b, and John O'Quigley^c

^aDepartment of Mathematics and Computer Science, Hampden-Sydney College, Hampden-Sydney, Virginia, USA

^bDivision of Translational Research and Applied Statistics, Department of Public Health Sciences, University of Virginia, Charlottesville, Virginia, USA

^cInserm, Université Paris VI, Place Jussieu, Paris, France

Abstract

Background—Most of the current designs used for Phase I dose finding trials in oncology will either involve only a single cytotoxic agent or will impose some implicit ordering among the doses. The goal of the studies is to estimate the maximum tolerated dose (MTD), the highest dose that can be administered with an acceptable level of toxicity. A key working assumption of these methods is the monotonicity of the dose–toxicity curve.

Purpose—Here we consider situations in which the monotonicity assumption may fail. These studies are becoming increasingly common in practice, most notably, in phase I trials that involve combinations of agents. Our focus is on studies where there exist pairs of treatment combinations for which the ordering of the probabilities of a dose-limiting toxicity cannot be known a priori.

Methods—We describe a new dose-finding design which can be used for multiple-drug trials and can be applied to this kind of problem. Our methods proceed by laying out all possible orderings of toxicity probabilities that are consistent with the known orderings among treatment combinations and allowing the continual reassessment method (CRM) to provide efficient estimates of the MTD within these orders. The design can be seen to simplify to the CRM when the full ordering is known.

Results—We study the properties of the design via simulations that provide comparisons to the Bayesian approach to partial orders (POCRM) of Wages, Conaway, and O'Quigley. The POCRM was shown to perform well when compared to other suggested methods for partial orders. Therefore, we compare our approach to it in order to assess the performance of the new design.

Limitations—A limitation concerns the number of possible orders. There are dose-finding studies with combinations of agents that can lead to a large number of possible orders. In this case, it may not be feasible to work with all possible orders.

Conclusions—The proposed design demonstrates the ability to effectively estimate MTD combinations in partially ordered dose-finding studies. Because it relaxes the monotonicity assumption, it can be considered a multivariate generalization of the CRM. Hence, it can serve as a link between single and multiple-agent dose-finding trials.

Introduction

The primary objective of a Phase 1 clinical trial in oncology is the estimation of the maximum tolerated dose (MTD). The MTD is defined as the highest dose that can be administered with a ‘tolerable’ level of toxicity. This ‘tolerable’ level is based upon the probability that a patient in the trial experiences a dose-limiting toxicity (DLT), which is typically defined by side-effects that are considered severe, and in certain cases, potentially life threatening. The majority of the statistical methods underlying the experimental design of dose-finding studies will entail the assumption that the probability of toxicity increases monotonically with the dose. This assumption is generally a reasonable one, in particular in studies involving a single cytotoxic agent. Here, the administration of higher doses can be expected to result in DLTs in a higher percentage of patients. In the language of Robertson, Wright and Dykstra [1], the probabilities of DLTs follow a ‘simple order.’ When studying multiple agents, the goal of the trial is to locate a dose combination or combinations with an acceptable toxicity rate. The set of dose combinations with acceptable toxicity forms a contour in two dimensions. The monotonicity assumption may not hold since the ordering of the toxicity probabilities could possibly be unknown for several of the available drug combinations. This is clearly seen in the simplest case where we have two drugs at two levels. Suppose that the second level corresponds to an increase in one of the drugs but a decrease in the other. Then, the ordering of the toxic probabilities is not known. Returning to multiple combinations and drug levels, some of the orderings between doses are known while others are not. We describe the whole situation as that of a partial ordering. Various possible partial order scenarios are described in [1], whereas parameter estimation subject to order restrictions is discussed in Dunbar *et al.* [2], as well as in Hwang and Peddada [3]. In this article we investigate a method for trials in which the toxicity order isn't fully known. The proposed method builds off of the continual reassessment method (CRM), introduced by O'Quigley *et al.* [4], which has demonstrated near-optimal properties in trials when the order is known.

In its original form the CRM is a Bayesian method based on the use of a simple working model and sequential updating of the dose–toxicity relationship to estimate the dose level at which to treat the next available patient. O'Quigley and Shen [5] suggested a two-stage design for MTD estimation that employs the same likelihood as that for the original CRM but uses a distinct initial dose escalation stage until we have at least one toxicity and one nontoxicity. The first stage of escalation is decided by the investigators and makes no appeal to any model. The design we propose, in part, resembles that of the two-stage CRM of [5] in that we have an initial escalation stage involving a scheme resembling that of the classic Up-and-Down, followed by a second stage that leans upon maximum likelihood estimation. As a consequence the approach has the advantage of building off the CRM. Suppose that we have a discrete set of k preset dose levels, d_1, \dots, d_k . The usual CRM begins by assuming a simplified working dose–toxicity curve, $\psi(d_j, a)$, that is monotonic in both dose levels, d_j and the parameter, a , for instance, the logistic curve or power model. After having included j subjects, we obtain an estimate, \hat{a}_j based on the likelihood function of a . The dose given to; the $(j + 1)$ th patient is the level, d_j that minimizes $\Delta(\psi(d_j, \hat{a}_j), \theta)$, where θ is the target probability of a DLT and $\Delta(v, w)$ is some measure of distance, such as Euclidean distance. As in other areas of statistics the notion ‘distance’ can be taken in a looser sense than the given mathematical one and, for example, we may allow it to be asymmetric [6]. The procedure continues until some predetermined sample size of patients is exhausted or a stopping rule takes effect.

Although the number of Phase 1 trials that use multiple-agent combinations is increasing, there are relatively few statistical methods for designing these trials. Several themes emerge from the published methods. These include the need to use as much of the ordering

information as possible and the need to reduce the dimension of the problem in terms of the number of combinations to be considered. In many trials involving combinations of agents, investigators assume an ordering for a specific set of combinations at the start of the trial and do standard dose escalation methods for the chosen ordering [7]. The disadvantage of this method is that it limits the number of combinations that can be considered and could produce misleading results if the initial ordering is incorrect. A second method fixes the second agent at its lowest dose and does standard dose escalation for the first agent. When a 'MTD' combination is reached for the first agent by the usual '3+3' rules, the first agent is fixed at this point, and a standard '3+3' scheme is performed for the second agent. The final dose of this second phase is taken as the 'MTD' combination. It has been established that the CRM has superior statistical properties to those of traditional escalation schemes [8]. Given these properties, this approach is likely to have poor statistical properties for combinations of agents and be wasteful of resources.

Wang and Ivanova [9] propose a Bayesian method for two agent combinations where the prior information is based on the single agent toxicity profiles. The allocation of patients to doses is done in a way that is similar to the CRM of O'Quigley *et al.* [4]. They lay out an initial grid of points designed to move quickly towards a 'solution', specifically, a dose combination with acceptable toxicity. These initial allocations are essentially a series of individual 'standard' Phase 1 trials where the dose-toxicity relationship can be assumed to be monotone. Yin and Yuan [10] propose a Bayesian dose-finding design that incorporates a copula-type model for estimating the toxicity probabilities of drug combinations. Another design proposed by Yin and Yuan [11] is a Bayesian design for dose-finding based on latent contingency tables. The approach of Wages *et al.* [12], as well as the one described in this current work, leans to some degree upon the framework of Conaway *et al.* [13], which identifies all possible simple orders for the toxicity probabilities that are consistent with the known orderings among the treatment combinations. Each of these simple orders consistent with a partial order can be thought of as a model. The idea here is to focus estimation of toxicity probabilities within a small number of simple orders, and allow the properties of the CRM [14] to provide efficient estimates of MTD combinations within these orders. When the toxicity order is fully known, the method we propose reduces to the CRM. The remainder of the article is organized as follows: Second section gives an example of a partially ordered trial. Third section presents the probability model and inference associated with the new design, while fourth section provides simulation results comparing the new method to the Bayesian approach to partial orders of Wages *et al.* [12]. We conclude this article with some discussion on the implications of the new design together with some additional topics of interest for further research.

An example of a partially ordered trial

Consider the dose finding study on combinations of topotecan and irinotecan described by Lokich [15]. The treatment combinations chosen by the investigators are displayed in Table 1. The toxicity ordering between some of the treatment combinations is not known. Specifically, the trial consists of eight drug combinations (doses), d_1, \dots, d_8 . The toxicity ordering between doses d_1 and d_2 is known due to the fact that the dose of Irinotecan remains the same while the dose of Topotecan increases. This is also the case for the ordering between doses d_3 and d_4 . However, the order relationships between doses d_2 and d_3 and between d_4 and d_5 are not known because the dose of Topotecan decreases while the dose of Irinotecan increases. If we continue to assess the known and unknown toxicity order relationships in this way, we can determine that the following order relationships hold among those that are known: (1) $d_1 \rightarrow d_2$, (2) $d_3 \rightarrow d_4$, (3) $d_5 \rightarrow d_6$, and (4) $d_7 \rightarrow d_8$. In these diagrams, combinations whose orderings are known are connected by arrows, with the treatment to the right being more toxic, that is, it is known that combination d_8 is more toxic

than d_7 , combination d_6 is more toxic than d_5 , and so on. Escalation to a previously untried dose depends on a specification prior to the trial of ‘possible escalation combinations’ associated with each dose. For example, the possible escalation combinations for d_1 are d_2 and d_3 , meaning that if d_1 was tried and found to be well tolerated, then escalation could proceed to the previously untried levels d_2 or d_3 . Taking into account the subset of drug combinations for which we know the toxicity order, we aim to formulate the possible orderings of the toxicity profile. Denoting the probability of a DLT at combination d_i by $R(d_i)$, $i = 1, \dots, 8$, the simple order:

$$R(d_1) \leq R(d_3) \leq R(d_2) \leq R(d_4) \leq R(d_5) \leq R(d_6) \leq R(d_7) \leq R(d_8)$$

is consistent with the partial order. In general, we suppose that the dose combinations follow a partial order for which there are M possible simple orders consistent with the partial order.

Therefore, we have a class of M models of interest indexed by $m = 1, \dots, M$. To gain further generality, we can take account any prior information concerning the plausibility of each model and so introduce $p(m) = \{p(1), \dots, p(M)\}$, where $p(m) \geq 0$ and $\sum_m p(m) = 1$. This would then assign a set of prior model weights so that if any initial knowledge allows us to consider that some orders are more likely than others, then this knowledge can be used to sharpen our inference. Even when there is no prior information available, we can formally proceed in the same way by placing equal probability on each model and hence use $p(m) = 1/M$. In the context of the current example, the trial requires the investigation of the eight simple orders in Table 2. In the table, we suppress the ‘ $R(\cdot)$ ’ notation and just display the labels for the combinations. Our preference is to allow the toxicity probabilities, given an ordering, to be modeled by a parametric model from the CRM class of models. Given the properties of the CRM, it is reasonable to believe that a method that quickly reduces to the CRM on the correct ordering would have excellent properties in terms of identifying dose combinations with acceptable toxicity.

Dose escalation in partially ordered trials

Toxicity probability model

Our main focus is on Phase 1 trials with a small number of possible orders as illustrated in the example in second section. In general, we assume that there are k drug combinations, d_1, \dots, d_k , to be studied, and assume that the combinations follow a partial order for which there are M distinct simple orders consistent with the partial order. A key idea to this design is that each possible simple order can be thought of as a model. This ties in with the ideas of Bayesian model choice. For a particular ordering, $m(m = 1, \dots, M)$, we model $R(x_j)$, the true but unknown probability of dose-limiting toxic response, Y_j , at drug combination $X_j = x_j$, $x_j \in \{d_1, \dots, d_k\}$, via:

$$R(x_j) = Pr(Y_j = 1 | X_j = x_j) = E(Y_j | x_j) \cong \psi_m(x_j, a) \quad (1)$$

for some one parameter model $\psi_m(x_j, a)$ and a defined on the set A . For each dose d_i , there exists some $a_i \in A$ such that $R(d_i) = \psi_m(d_i, a_i)$. Specifically, the model is rich enough to exactly reproduce the true probability of toxicity at each dose. There is a wide variety of choices for potential working models and the simple power model:

$$\psi_m(d_i, a) = \alpha_{mi}^a, \quad i = 1, \dots, k \quad (2)$$

has been used by several investigators and appears to work well in practice. Here, $0 < a < \infty$ and $0 < \alpha_{m1} < \dots < \alpha_{mk} < 1$ represents the skeleton of the model [16,17]. Once a model has been chosen and we have data in the form of $\Omega_j = \{x_1, y_1, \dots, x_j, y_j\}$, we obtain an estimate for the parameter a and subsequently generate estimates $\widehat{R}(d_i)$, $i = 1, \dots, k$ of the true unknown toxicity probabilities at each of the k drug combinations. The target dose level is that having a corresponding toxicity probability as close as possible to the target rate θ . The drug combination $x_j = d_i$ assigned to the j th patient enrolled in the trial is such that:

$$\Delta(\widehat{R}(d_i), \theta) = |\widehat{R}(d_i) - \theta|, \quad i = 1, \dots, k \quad (3)$$

is minimized. For partial orders, there may exist more than one dose with toxicity probability closest to the target. In this case, we would randomly choose among the set of candidate treatments. Equation (3) translates the idea that the overall objective of the study is also the objective for each included patient. The design is an iterative sequential design whereby the treatment combination chosen for the hypothetical $(n + 1)$ th patient is also the current best estimate for an MTD combination.

Inference

The estimates for the toxicity probabilities at each of the available drug combinations can either be likelihood based or Bayesian. In this article, we lean upon the likelihood approach to the CRM of O'Quigley and Shen [5]. After inclusion of the first j patients into the study, the logarithm of the likelihood under dose–toxicity ordering m is given by:

$$L_m(a|\Omega_j) = \sum_{\ell=1}^j y_\ell \log \psi_m(x_\ell, a) + \sum_{\ell=1}^j (1 - y_\ell) \log(1 - \psi_m(x_\ell, a)) \quad (4)$$

where any terms not involving the parameter a have been equated to zero. For each of the M distinct orderings, the above expression can be maximized in order to generate an estimate, \hat{a}_m , for a . When the method can effectively estimate a single ‘correct’ ordering, then the approach reduces to a standard phase I trial. We weight each of the M candidate models as we make progress. A plausible choice for model weights is then given by:

$$\pi(m) = \frac{\exp\{L_m(\hat{a}_m|\Omega_j)\}}{\sum_{m=1}^M \exp\{L_m(\hat{a}_m|\Omega_j)\}},$$

where a particular $\pi(m)$ is considered as the weight of evidence in favor of model m . The expression $L_m(\hat{a}_m|\Omega_j)$ is the value of the log-likelihood evaluated by substituting the maximum likelihood estimate of the parameter. This maximum point on the log-likelihood function corresponds to the value of the maximum likelihood estimate for each given value of m .

Prior weights can be specified and Bayesian methods could be used to down-weight models that are less plausible. Let $p(m)$ be the prior probability that model m is the ‘best’ model. Specifically, $p(m)$ is our prior state of belief that model m , fitted to the data, provides the ‘best’ model for the design. For a set of prior probabilities, $p(m)$, $m = 1, \dots, M$, generalized weights are then given by:

$$\pi(m) = \frac{\exp\{L_m(\hat{a}_m | \Omega_j)\} p(m)}{\sum_{m=1}^M \exp\{L_m(\hat{a}_m | \Omega_j)\} p(m)}. \quad (5)$$

The inclusion of prior probabilities, $p(m)$, in the $\pi(m)$ is not a true Bayesian approach. A true Bayesian approach to model selection requires both the prior $p(m)$ on the model and a prior probability distribution on the parameter a in model m for each model. Then the derivation of posterior results requires integration. When a new patient is to be enrolled, we will choose a single ordering h such that:

$$h = \underset{m}{\operatorname{argmax}} \pi(m), \quad m = 1, \dots, M$$

among the M contending models. Given h , we take the working model associated with this ordering and apply the likelihood approach of the CRM to obtain estimates of the toxicity probabilities at each of the k available treatment combinations. Using the dose-toxicity model (1), the estimated probability of toxicity at each dose combination is given by:

$$\widehat{R}(d_i) = \psi_h(d_i, \widehat{a}_h), \quad i = 1, \dots, k. \quad (6)$$

On the basis of this formula, the dose combination to be given to the next patient enrolled in the study is determined based on the minimization of (3). For partial orders, there may be more than one treatment combination with toxicity probability closest to the target. If there is a 'tie' between two or more dose combinations, the patient will be randomized to one of the dose combinations with DLT probability closest to the target. A requirement to be able to maximize the log-likelihood on the interior of the parameter space is that we have heterogeneity (at least one toxic and one nontoxic) among the responses. Otherwise the likelihood is maximized on the boundary of the parameter space and our estimates of $R(d_i)$ are trivially either zero or one, or, depending on the model we are working with, may be undefined.

Two-stage design

O'Quigley and Shen [5] recommended including an initial escalation stage utilizing traditional or non-traditional Up-and-Down schemes: that is, starting at the lowest available dose, three patients are treated and only if all three fail to experience a DLT do we escalate to a higher dose level. As soon as the first DLT is observed, the first stage is closed and the second stage is subsequently opened based on CRM modeling, using the data accumulated thus far in the trial. Such a design could be varied in several ways, for example, including a cohort of one patient at the lowest level, a cohort of two patients at the second lowest and then proceeding as above. Even though the first stage is closed, the toxicity response information accrued by the initial scheme is retained and used in the second stage.

When the trial is being conducted subject to a partial order, it is necessary to use a variant of the traditional Up-and-Down scheme in the initial stage. After establishing the M simple orders that the toxicity probabilities can assume, the drug combinations can be partitioned into sets of possible escalation treatments consistent with the partial order. For instance, suppose a cohort of patients is observed at one of the k discrete dose combinations, d_j , and none of the three experience a DLT. A standard Up-and-Down scheme, in the context of a simple monotonic order, would require that the next cohort of patients be enrolled at the next highest dose, d_{j+1} . However, when two drugs are being combined, the toxicity order of

doses d_i and d_{i+1} may not be clear. As a result, the most ‘appropriate’ dose to which the trial should escalate could consist of more than one treatment combination.

Allocation algorithm in the first stage—Returning to the example described in second section, there are eight possible simple orders of the dose-toxicity relationship. Accordingly, we can use Table 2 to partition the available doses into sets A through E displayed in Table 3. For instance, we can see from Table 2 that d_1 is the lowest dose in each possible ordering. Therefore, the trial could begin in set $A = \{d_1\}$. At the first observation of a toxicity in one of the patients, the first stage is closed and the second stage is opened. As long as no toxicities occur, cohorts of patients are examined at each dose within the currently occupied set, before escalating to the next highest set. We can see in Table 2, if d_1 was tried and deemed ‘safe’, the trial would escalate to d_2 in four of the possible orders and d_3 in four of the orders. In other words, the second ‘level’ of doses consists of d_2 and d_3 . Therefore, if no toxicity is observed in a cohort of patients at d_1 , the trial proceeds to set $B = \{d_2, d_3\}$. If more than one dose is contained within a particular set, we can sample without replacement from the doses available within the set. Therefore, the next cohort is enrolled on a dose that is chosen randomly from d_2 and d_3 . The trial is not allowed to advance to set $C = \{d_4, d_5\}$ in the first stage until a cohort of patients has been observed at both doses d_2 and d_3 . This procedure continues until a toxicity is observed or all available sets have been exhausted. Subsequent to a DLT being observed, the second stage of the trial begins.

Allocation algorithm in the second stage—Based on heterogeneous data in form of the set Ω_j , the expression for the likelihood function (4) can be maximized with respect to a . Given the ordering h , we make use of the maximum likelihood estimate, \hat{a}_h that was generated through the maximization of (4) for model h . The resulting parameter estimate for a can be used to generate an estimate of the toxicity probability at each of the k treatment combinations by calculating $R^*(d_j) = \psi_h(d_j, \hat{a}_h) = 1, \dots, k$, from which the drug combination, x_{j+1} , given to the $(j+1)$ th patient is determined according to some loss function such as (3). Once the toxicity data, Ω_n , has accumulated from the predetermined sample size of n patients entered into the trial, the recommended dose for the hypothetical $(n+1)$ th patient will be the estimate for an MTD combination. That is, an MTD is the treatment combination, x_{n+1} , that satisfies (3) based on Ω_n .

Simulations

Illustration

In order to illustrate the proposed method, consider the example given in second section involving $k = 8$ discrete treatment combinations, d_1, \dots, d_8 . The partial order for this trial has eight possible simple orders. For this simulation, we assume that the true order is $d_1 \rightarrow d_2 \rightarrow d_3 \rightarrow d_4 \rightarrow d_5 \rightarrow d_6 \rightarrow d_7 \rightarrow d_8$ and that the true toxicity probabilities are $R(d_1) = 0:05$, $R(d_2) = 0:10$, $R(d_3) = 0:20$, $R(d_4) = 0:30$, $R(d_5) = 0:45$, $R(d_6) = 0:58$, $R(d_7) = 0:70$ and $R(d_8) = 0:81$. The targeted toxicity probability is $\theta = 0:20$, indicating that the ‘correct’ treatment combination for an MTD is given by d_3 . Before getting the trial underway, the skeletons, α_{mi} , $m = 1, \dots, 8$, $i = 1, \dots, 8$, for the toxicity probabilities at each treatment combination need to be specified for each ordering. We implemented the systematic approach of Lee and Cheung [16] in order to establish the skeleton for the monotonic order $m = 1$. These values were adjusted to correspond to each of the eight simple orders consistent with the partial. Implementing the algorithm yielded the results given in Table 4. For the eight drug combinations under investigation, the simulations use the simple power parameter working model (2):

$$\psi_m(d_i, a) = \alpha_{mi}^a$$

where the α_{mi} have the values provided in Table 4.

The first entered patient will be treated on a treatment from among a set at the second lowest level, which in this particular example, consists of d_2 and d_3 . Since more than one treatment is in the set at the second lowest level, we randomly select from those doses in the set. For this example, d_2 is the second least toxic treatment in four of the eight possible orderings, so it is given probability 0.50 of being the dose administered to the first entered patient, with probability 0.50 given to d_3 . The trial begins at the second lowest level because we are observing responses in cohort sizes of one patient. Thus, if for the first entered patient, d_2 or d_3 is deemed too toxic, we have the ability to de-escalate to the lowest treatment set, which consists of d_1 , in order to continue the trial. The first patient is treated at d_2 and did not experience a dose-limiting toxic response. Before the method escalates to the next highest set of treatment combinations, it must exhaust the treatments in the current set as long as non-DLT's are observed. That is, the next patient is treated at d_3 before escalating to the next highest set of treatments, which consists of d_4 and d_5 . The second patient is treated on d_3 and does not experience a DLT. The trial then proceeds to the next highest set of treatments and the next enrolled patient is treated on a dose that is randomly selected from d_4 and d_5 . The third and fourth patients are treated on d_5 and d_4 , respectively, and neither experiences a DLT. Next, the fifth patient is administered d_7 and experiences a DLT. At this point in the trial, heterogeneity in the responses exists because we have at least one toxic and one nontoxic response. Consequently, the first stage of the trial is now closed and the second stage is subsequently opened.

The model selection techniques described in third section were implemented and generated estimated model weights $\pi(1) = 0:15$, $\pi(2) = 0:15$, $\pi(3) = 0:15$, $\pi(4) = 0:10$, $\pi(5) = 0:15$, $\pi(6) = 0:10$, $\pi(7) = 0:100$ and $\pi(8) = 0:10$. Therefore, based on the model weights, the ordering is randomly chosen from those with $\pi(m) = 0:15$ and determined to be $h = 1$. Given $h = 1$, we take the working model associated with this ordering from Table 4 and apply the likelihood based CRM. Heterogeneity in the responses means that the maximum likelihood estimate for $\hat{\alpha}_1$ now exists and can be seen to be equal to 1.5. We then have that $\hat{R}(d_i) = \alpha_{1i}^{\hat{\alpha}_1}$, $i = 1, \dots, 8$ is given by $R^*(d_1) = 0:001$, $R^*(d_2) = 0:005$, $R^*(d_3) = 0:032$, $R^*(d_4) = 0:089$, $R^*(d_5) = 0:190$, $R^*(d_6) = 0:322$, $R^*(d_7) = 0:465$ and $R^*(d_8) = 0:586$. The 6th entered patient is then administered d_5 for which $R^*(d_5) = 0:19$ since, from the available estimates, this is the closest to the target $\theta = 0:20$. The 6th included patient suffers toxic side-effects and the new model weights become $\pi(1) = 0:16$, $\pi(2) = 0:16$, $\pi(3) = 0:10$, $\pi(4) = 0:13$, $\pi(5) = 0:10$, $\pi(6) = 0:13$, $\pi(7) = 0:13$ and $\pi(8) = 0:09$ from which h is chosen to be 2. The new maximum likelihood estimate, $\hat{\alpha}_2$, becomes 0.90. Therefore, d_4 becomes the treatment combination with an estimated toxicity probability closest to the target. The dose that is recommended as the treatment that would have been administered to the 22nd patient, had one been included, is d_3 . Therefore it is identified (estimated) to be an MTD combination with an estimated probability of toxicity of $R^*(d_3) = 0:23$. Table 5 shows the results of the simulation for which d_3 is chosen as an MTD combination on the basis of all available knowledge. In terms of the example provided in Lokich [15], d_3 would correspond to a dose combination of 75 mg/m²/week of Irinotecan and 1.0mg/m²/week of Topotecan.

Simulation studies

The results in this section compare the new methodology for partial ordering, which henceforth will be referred to as POCRML, with the design of Wages *et al.* [12] for the trial

example of Lokich [15] described in second section, which has the eight possible orders given in Table 2. We assume that the true order is $d_1 \rightarrow d_2 \rightarrow d_3 \rightarrow d_4 \rightarrow d_5 \rightarrow d_6 \rightarrow d_7 \rightarrow d_8$. The approach of Wages *et al.* [12] (POCRM) makes use of prior information on the parameter a and begins modeling with the first entered patient, as opposed to the two-stage design presented here. POCRM was shown to perform well against other suggested methods for partial orders, such as Conaway *et al.* [13], Yin and Yuan [10], and Yin and Yuan [11], so we compare our approach to it in order to assess the performance of POCRML.

Each of the tables in this section is based on 1000 simulated trials and the target toxicity probability is $\theta = 0.20$ for each of the twelve toxicity scenarios presented in Tables 6 and 7. The sample size for each simulated trial is $n = 21$. Our goal is to present a reasonably wide spectrum of toxicity situations in order to get a feel of how things might work in practice. For each toxicity scenario, the $R(d_i)$ denotes the probability of toxicity at treatment combination d_i , $i = 1, \dots, 8$. In this set of simulations, we do not present any results for the percentage of trials that each individual ordering was selected. The reason for this is that our ultimate goal is the identification (estimation) of an MTD, not to choose the ‘correct’ ordering in a finite number of patients. Consequently, we are concerned with the method’s ability to recommend an MTD at the conclusion of the trial as opposed to recommend an ordering. Ultimately, our goal with these simulations is to demonstrate the performance of the CRM when the dose–toxicity order is only partially known compared to the case where it is fully known. The first stage of the likelihood approach uses single patient cohorts during escalation and each trial enrolled a pre-specified sample size of 21 patients. We again used the skeletons provided in Table 4 that were generated using the algorithm of Lee and Cheung [16]. Since we assume there is no prior information available on the ‘correct’ ordering, we place equal probability on each model and use $p(m) = 1/8$, $m = 1, \dots, 8$ prior to the beginning of the trial.

Table 6 and 7 indicate that POCRML is performing quite well, correctly recommending MTD combinations in a large percentage of simulated trials (Table 6), as well as treating a large percentage of patients at and around MTD combinations (Table 7). Even in cases where POCRML performs less well, it is recommending MTD combinations and a neighboring dose in a large percentage of trials. For instance, in the fourth scenario, in 87% of the simulated trials, the POCRM identifies treatments as an MTD that have toxicity probabilities between 0.11 and 0.35. Similarly, in the fifth scenario, 89% of the time, the method chooses one of the treatments with toxicity probabilities between 0.10 and 0.35. Further, in terms of in-trial allocation, in scenario 4, 72% of patients are treated at doses with toxicity probabilities ranging from 0.08 to 0.22. Overall, the simulation results of POCRML are competitive with that of POCRM, a design that demonstrated a comparable performance to the methods of Conaway *et al.* [13], Yin and Yuan [10], and Yin and Yuan [11] in Wages *et al.* [12]. Therefore, POCRML can be considered a practical alternative as a multiple-agent trial design. Because it is an extension of the well known CRM, it is believed that our approach will be easily understood by clinicians and review boards. At the very least, the method gives the investigator an alternative to his or her design preference when presented with a multiple-agent Phase 1 trial.

Discussion

The new design which leans upon the CRM method shows itself to be effective in estimating MTD combinations in Phase 1 clinical trials when the toxicity order of the treatments is only partially known. The goal remains accurate estimation of MTD combinations, subject to the ethical constraints of treating as many patients as possible at and around MTD combinations. It is only of indirect interest to establish the correct order and, in as much as

the estimate of an MTD itself may remain unaffected by certain orders, we do not attempt to identify which order is the more correct. That is only an indirect goal.

Because the CRM for partial orders is a design that relaxes the monotonicity assumption of the CRM, it can be regarded as a multivariate generalization of CRM. When the order of toxicity probabilities is fully known as in single-agent trials, our approach reduces to the CRM. Our design is therefore suitable for dose-finding problems where toxicity orders are fully or partially known. Hence, it can serve as a link between single and multiple-agent dose-finding trials. Furthermore, in theoretical investigations or those involving simulations, we can use a CRM with known ordering to provide an upper bound as a gauge on how much information is lost as a result of the partial ordering. Operating characteristics appear to be good although further study, under a broader range of possible situations, may provide more insight into general behavior. Wages *et al.* [12] present a Bayesian CRM approach for partial orders, in which they provide simulation results comparing their approach to the designs of Conaway *et al.* [13] and Yin and Yuan [10,11].

There may be further alternatives to those already published and those outlined here. One possibility may be to continue the first stage allocation algorithm throughout the entire trial in order to estimate MTD combinations. This alternative warrants further study. Another possibility may be, instead of selecting the ordering with the greatest model weight, $\pi(m)$, and reducing the problem to a standard phase I trial, we could estimate the dose to recommend to the next entered patient for each distinct ordering and use the dose that is most agreed upon across all possible orders. That is, we could work with the cumulative probabilities for an MTD given the set of models. This would then correspond to using some sort of weighted average rather than the mode. Simulation results for this alternative approach are presented in the supplementary web appendix of Wages *et al.* [12]. A comparison between the POCRM and this alternative design show the two approaches to behave similarly.

Acknowledgments

The project was supported by the National Cancer Institute [grant Number 1R01CA142859].

References

1. Robertson, T.; Wright, F.; Dykstra, R. Order Restricted Statistical Inference. John Wiley & Sons; New York: 1988.
2. Dunbar S, Conaway M, Peddada S. On improved estimation of parameters subject to order restrictions. *Stat Appl.* 2001; 3:121–28.
3. Hwang J, Peddada S. Confidence interval estimation subject to order restrictions. *Ann Stat.* 1994; 22:67–93.
4. O'Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for Phase I clinical trials in cancer. *Biometrics.* 1990; 46:33–48. [PubMed: 2350571]
5. O'Quigley J, Shen LZ. Continual reassessment method: a likelihood approach. *Biometrics.* 1996; 52:163–74.
6. Babb J, Rogatko A, Zacks S. Cancer phase I clinical trials: efficient dose escalation with overdose control. *Stat Med.* 1998; 17:1103–20. [PubMed: 9618772]
7. Kramer A, Lebecqz A, Candalh E. Continual reassessment methods in phase I trials of the combination of two agents in oncology. *Stat Med.* 1999; 18:1849–64. [PubMed: 10407256]
8. Iasonos A, Wilton A, Riedel E, et al. A comprehensive comparison of the continual reassessment method to the standard 3+3 dose escalation scheme in Phase I dose-finding studies. *Clin Trials.* 2008; 5:465–77. [PubMed: 18827039]

9. Wang K, Ivanova A. Two-dimensional dose finding in discrete dose space. *Biometrics*. 2005; 61:217–22. [PubMed: 15737096]
10. Yin G, Yuan Y. Bayesian dose finding for drug combinations by copula regression. *J Roy Stat Soc C Appl*. 2009; 581:211–24.
11. Yin G, Yuan Y. A latent contingency table approach to dose-finding for combinations of two agents. *Biometrics*. 2009; 65:866–75. [PubMed: 18759848]
12. Wages NA, Conaway MR, O'Quigley J. Continual reassessment method for partial ordering. *Biometrics*. 2011; 6710.1111/j.1541-0420.2011.01560.x
13. Conaway M, Dunbar S, Peddada S. Designs for single-or multiple-agent phase I trials. *Biometrics*. 2004; 60:661–69. [PubMed: 15339288]
14. O'Quigley J, Paoletti X, Maccario J. Non-parametric optimal design in dose-finding studies. *Biostatistics*. 2002; 3:51–56. [PubMed: 12933623]
15. Lokich J. Phase I clinical trial of weekly combined topotecan and irinotecan. *Am J Clin Oncol*. 2001; 24:336–40. [PubMed: 11474256]
16. Lee SM, Cheung YK. Model calibration in the continual reassessment method. *Clin Trials*. 2009; 6:227–38. [PubMed: 19528132]
17. Cheung YK, Chappell R. A simple technique to evaluate model sensitivity in the continual reassessment method. *Biometrics*. 2002; 58:671–74. [PubMed: 12230003]

\$watermark-text

\$watermark-text

\$watermark-text

Table 1

Drug combinations of Topotecan and Irinotecan in $mg/m^2/wk$

Agent	Drug combination							
	d_1	d_2	d_3	d_4	d_5	d_6	d_7	d_8
Irinotecan	50	50	75	75	100	100	100	125
Topotecan	1.0	1.5	1.0	1.5	1.0	1.5	1.0	1.5

Table 2
Eight possible simple dose–toxicity orders

Ordering (m)	Simple order
1	$d_1 \rightarrow d_2 \rightarrow d_3 \rightarrow d_4 \rightarrow d_5 \rightarrow d_6 \rightarrow d_7 \rightarrow d_8$
2	$d_1 \rightarrow d_3 \rightarrow d_2 \rightarrow d_4 \rightarrow d_5 \rightarrow d_6 \rightarrow d_7 \rightarrow d_8$
3	$d_1 \rightarrow d_2 \rightarrow d_3 \rightarrow d_5 \rightarrow d_4 \rightarrow d_6 \rightarrow d_7 \rightarrow d_8$
4	$d_1 \rightarrow d_2 \rightarrow d_3 \rightarrow d_4 \rightarrow d_5 \rightarrow d_7 \rightarrow d_6 \rightarrow d_8$
5	$d_1 \rightarrow d_3 \rightarrow d_2 \rightarrow d_5 \rightarrow d_4 \rightarrow d_6 \rightarrow d_7 \rightarrow d_8$
6	$d_1 \rightarrow d_3 \rightarrow d_2 \rightarrow d_4 \rightarrow d_5 \rightarrow d_7 \rightarrow d_6 \rightarrow d_8$
7	$d_1 \rightarrow d_2 \rightarrow d_3 \rightarrow d_5 \rightarrow d_4 \rightarrow d_7 \rightarrow d_6 \rightarrow d_8$
8	$d_1 \rightarrow d_3 \rightarrow d_2 \rightarrow d_5 \rightarrow d_4 \rightarrow d_7 \rightarrow d_6 \rightarrow d_8$

Watermark-text

Watermark-text

Watermark-text

Table 3

Treatment sets for allocation in initial stage

Set	A	B	C	D	E
		d_2	d_4	d_6	d_8
Dose	d_1		d_3	d_5	d_7

Table 4

Working model (α_{mi}) values for eight simple orders

m	Doses							
	d_1	d_2	d_3	d_4	d_5	d_6	d_7	d_8
1	0.01	0.03	0.10	0.20	0.33	0.47	0.60	0.70
2	0.01	0.10	0.03	0.20	0.33	0.47	0.60	0.70
3	0.01	0.03	0.10	0.33	0.20	0.47	0.60	0.70
4	0.01	0.03	0.10	0.20	0.33	0.60	0.47	0.70
5	0.01	0.10	0.03	0.33	0.20	0.47	0.60	0.70
6	0.01	0.10	0.03	0.20	0.33	0.60	0.47	0.70
7	0.01	0.03	0.10	0.33	0.20	0.60	0.47	0.70
8	0.01	0.10	0.03	0.33	0.20	0.60	0.47	0.70

Watermark-text

Watermark-text

Watermark-text

Table 5

Sequential trial of 21 patients

j	h	\hat{a}_h	x_j	y_j	j	h	\hat{a}_h	x_j	y_j
1	-	-	d_2	0	14	1	0.45	d_2	0
2	-	-	d_3	0	15	3	0.48	d_2	0
3	-	-	d_5	0	16	3	0.50	d_2	0
4	-	-	d_4	0	17	3	0.52	d_2	0
5	-	-	d_7	1	18	1	0.54	d_2	0
6	1	1.5	d_5	1	19	1	0.56	d_2	0
7	2	0.90	d_4	1	20	3	0.58	d_3	0
8	3	0.62	d_5	0	21	1	0.60	d_3	0
9	5	0.67	d_2	0	22	1	0.63	d_3	-
10	2	0.73	d_2	1					
11	5	0.55	d_3	1					
12	1	0.40	d_1	0					
13	1	0.43	d_2	0					

Table 6
Percentage recommendation for the partial order continual reassessment method for eight possible orderings

Scenario	Dose	d_1	d_2	d_3	d_4	d_5	d_6	d_7	d_8	$\%_{tox}$
1	$R(d_i)$	0.09	0.20	0.26	0.44	0.58	0.74	0.83	0.90	-
	POCRM	0.23	0.38	0.34	0.03	0.02	0.00	0.00	0.00	23.8
	POCRM	0.16	0.40	0.37	0.05	0.02	0.00	0.00	0.00	27.9
2	$R(d_i)$	0.05	0.10	0.20	0.30	0.45	0.58	0.70	0.81	-
	POCRM	0.04	0.30	0.39	0.16	0.11	0.00	0.00	0.00	23.2
	POCRM	0.02	0.26	0.38	0.23	0.11	0.00	0.00	0.00	25.6
3	$R(d_i)$	0.06	0.12	0.22	0.30	0.45	0.55	0.70	0.84	-
	POCRM	0.06	0.29	0.39	0.15	0.11	0.00	0.00	0.00	23.1
	POCRM	0.03	0.27	0.36	0.23	0.11	0.00	0.00	0.00	25.6
4	$R(d_i)$	0.02	0.05	0.11	0.20	0.35	0.58	0.70	0.81	-
	POCRM	0.01	0.11	0.23	0.39	0.25	0.01	0.00	0.00	21.5
	POCRM	0.00	0.09	0.18	0.44	0.27	0.01	0.00	0.00	22.5
5	$R(d_i)$	0.01	0.04	0.07	0.10	0.20	0.35	0.50	0.64	-
	POCRM	0.00	0.02	0.05	0.34	0.38	0.16	0.06	0.00	19.7
	POCRM	0.00	0.02	0.03	0.34	0.41	0.14	0.05	0.00	19.1
6	$R(d_i)$	0.01	0.02	0.05	0.09	0.15	0.20	0.40	0.52	-
	POCRM	0.00	0.01	0.02	0.22	0.22	0.32	0.18	0.03	18.9
	POCRM	0.00	0.01	0.01	0.24	0.31	0.32	0.11	0.00	16.4
7	$R(d_i)$	0.01	0.03	0.05	0.08	0.10	0.22	0.41	0.55	-
	POCRM	0.00	0.01	0.02	0.18	0.24	0.37	0.17	0.01	19.2
	POCRM	0.00	0.01	0.01	0.23	0.29	0.34	0.13	0.00	16.1
8	$R(d_i)$	0.01	0.02	0.04	0.06	0.08	0.10	0.20	0.35	-
	POCRM	0.00	0.00	0.01	0.06	0.07	0.28	0.35	0.23	15.7
	POCRM	0.00	0.00	0.00	0.07	0.13	0.39	0.37	0.04	11.3

The target MTD combinations are in boldface.

Table 7
Percentage of in-trial allocation for the partial order continual reassessment method for eight possible orderings

Scenario	Dose	d_1	d_2	d_3	d_4	d_5	d_6	d_7	d_8
1	$R(d_i)$	0.09	0.20	0.26	0.44	0.58	0.74	0.83	0.90
	POCRM	0.25	0.31	0.28	0.07	0.06	0.01	0.01	0.00
	POCRM	0.15	0.30	0.30	0.14	0.11	0.00	0.00	0.00
2	$R(d_i)$	0.05	0.10	0.20	0.30	0.45	0.58	0.70	0.81
	POCRM	0.10	0.27	0.28	0.16	0.14	0.02	0.02	0.00
	POCRM	0.03	0.23	0.27	0.27	0.19	0.01	0.01	0.00
3	$R(d_i)$	0.06	0.12	0.22	0.30	0.45	0.55	0.70	0.84
	POCRM	0.11	0.27	0.27	0.16	0.14	0.02	0.02	0.00
	POCRM	0.05	0.23	0.25	0.26	0.19	0.01	0.01	0.00
4	$R(d_i)$	0.02	0.05	0.11	0.20	0.35	0.58	0.70	0.81
	POCRM	0.03	0.17	0.22	0.27	0.23	0.04	0.03	0.00
	POCRM	0.01	0.13	0.15	0.39	0.29	0.02	0.01	0.00
5	$R(d_i)$	0.01	0.04	0.07	0.10	0.20	0.35	0.50	0.64
	POCRM	0.02	0.10	0.12	0.25	0.26	0.14	0.10	0.02
	POCRM	0.00	0.04	0.05	0.35	0.38	0.11	0.06	0.00
6	$R(d_i)$	0.01	0.02	0.05	0.09	0.15	0.20	0.40	0.52
	POCRM	0.01	0.08	0.08	0.20	0.20	0.21	0.15	0.06
	POCRM	0.00	0.01	0.01	0.24	0.31	0.32	0.11	0.00
7	$R(d_i)$	0.01	0.03	0.05	0.08	0.10	0.22	0.41	0.55
	POCRM	0.01	0.07	0.08	0.19	0.19	0.24	0.16	0.06
	POCRM	0.00	0.02	0.03	0.29	0.32	0.22	0.13	0.00
8	$R(d_i)$	0.01	0.02	0.04	0.06	0.08	0.10	0.20	0.35
	POCRM	0.01	0.07	0.07	0.12	0.11	0.21	0.21	0.21
	POCRM	0.00	0.01	0.01	0.20	0.22	0.29	0.26	0.00

The target MTD combinations are in boldface.