The Annals of Applied Statistics
2016, Vol. 10, No. 4, 2349–2376
DOI: 10.1214/16-AOAS982

BOOTSTRAP AGGREGATING CONTINUAL REASSESSMENT METHOD FOR DOSE FINDING IN DRUG-COMBINATION TRIALS[1]

BY RUITAO LIN AND GUOSHENG YIN

University of Hong Kong

Phase I drug-combination trials are becoming commonplace in oncology. Most of the current dose-finding designs aim to quantify the toxicity probability space using certain prespecified yet complicated models. These models need to characterize not only each individual drug's toxicity profile, but also their interaction effects, which often leads to multi-parameter models. We propose a novel Bayesian adaptive design for drug-combination trials based on a robust dimension-reduction method. We continuously update the order of dose combinations and reduce the two-dimensional searching space to a one-dimensional line based on the estimated order. As a result, the common approaches to single-agent dose finding, such as the continual reassessment method (CRM), can be applied to drug-combination trials. We further utilize the ensemble technique in machine learning, the so-called bootstrap aggregating (bagging) in conjunction with Bayesian model averaging, to enhance the efficiency and reduce the variability of the proposed method. We conduct extensive simulation studies to examine the operating characteristics of the proposed method under various scenarios. Compared with existing competitive designs, the bagging CRM demonstrates its precision and robustness in terms of pinning down the correct dose combination. We apply the proposed bagging CRM to two recent cancer clinical trials with combined drugs for dose finding.

1. Introduction. In oncology, effective treatments through combining multiple drugs are becoming a routine practice. A combination of drugs can induce treatment synergy, target cancer cells with different drug susceptibilities, and achieve higher dose intensity and thus improve cure rates. A phase I drug-combination trial in cancer research often focuses on identifying the maximum tolerated dose (MTD) combination, which is defined as the one that has a dose-limiting toxicity (DLT) probability closest to the target toxicity rate. Unlike single-agent trials, dose finding for drug combinations faces many challenges due to complicated drug–drug interactions and a multiplicatively increasing number of dose pairs. More importantly, the toxicity order of the combined doses is only partially known, that is, the toxicity probability increases with the dose level of one drug when fixing that of the other drug.

Received April 2016; revised August 2016.

[1]Supported in part by a Grant (17125814) from the Research Grants Council of Hong Kong.

Key words and phrases. Bagging, Bayesian model averaging, dimension reduction, drug combination, maximum tolerated dose.

These new challenges in drug-combination trials have stimulated extensive development of statistical methods, most of which are built upon the continual reassessment method (CRM) [O'Quigley, Pepe and Fisher (1990)]. Existing model-based designs often rely on complicated and heavily parameterized statistical models. Korn and Simon (1991) described a mathematical model for combination therapies, which can account for the differing toxicity profiles of multiple drugs. Thall et al. (2003) proposed a six-parameter regression model to characterize a two-dimensional dose space. Wang and Ivanova (2005) introduced a relatively parsimonious working model for the dose–toxicity relationship. Yin and Yuan (2009a) linked the individual toxicity rates of the two agents using a Clayton-type copula function. Houede et al. (2010) formulated a generalized Aranda-Ordaz model to facilitate dose finding in two-agent phase I–II trials. Hirakawa et al. (2013) developed a shrinkage logistic regression model to estimate the joint toxicity probabilities, and the logistic model is also considered in Riviere et al. (2014). Tighiouart et al. (2014) reparameterized the logistic regression model to simplify the model interpretation. In contrast to the model-based methods, algorithm-based methods can locate the MTD combination in a "nonparametric" way without imposing any parametric assumption on the underlying dose–toxicity relationship. Conaway, Dunbar and Peddada (2004) proposed to use the pool-adjacent-violators algorithm (PAVA) to determine the dose allocation in a drug-combination trial. Huang et al. (2007) applied the traditional $3 + 3$ method to drug-combination trials by partitioning the dose space into zones along the diagonal direction from the lowest to the highest dose level. Fan et al. (2009) proposed a searching strategy to locate the MTDs in a two-agent toxicity space. Mander and Sweeting (2015) assigned a product of independent beta priors to the toxicity probabilities, and then developed a curve-free method for a dual-agent dose-finding trial. Lin and Yin (2016) proposed an algorithm-based two-agent Bayesian optimal interval design. For comprehensive reviews and comparisons on the existing designs for drug-combination trials, see Thall (2010), Harrington et al. (2013), Mandrekar (2014), Riviere et al. (2015a) and Hirakawa et al. (2015).

Iasonos and O'Quigley (2014) recently conducted a comprehensive review of adaptive phase I clinical trials in oncology, and they promoted the wide usage of model-based designs in practice. Not only is this true for single-agent trials, but it is also sensible with drug-combination trials due to limited sample size yet a larger dose searching space. Riviere et al. (2015b) reviewed 162 published articles, and concluded that the designs of current phase I drug-combination trials in oncology can be improved by utilizing more model-based methods. Most of the model-based methods directly quantify the toxicity profile based on the partial order information. However, the sample size in a phase I dose-finding trial is rather small relative to the number of unknown parameters, which usually leads to unstable estimation for the model parameters, especially at the beginning of a trial. To increase the stability as well as to reduce the number of unknown parameters, a common approach is to reduce the dimensionality of the dose searching space to

a one-dimensional searching line. Nevertheless, very limited work has been carried out along the line of dimension reduction, which is in fact a major theme of high-dimensional modeling. To overcome the challenge of multi-dimension, Korn and Simon (1993) constructed a tolerable-dose diagram to target specific MTD combinations. Kramar et al. (1999) studied a selected subset of drug combinations that are monotonically ordered. Yuan and Yin (2008) proposed a simple sequential design based on the partial orders of the joint toxicities. Wages et al. (2011) introduced a partial ordering continual reassessment method (POCRM) by laying out several possible orders for the joint toxicity rates and then selecting the best order based on the data. Recognizing the limitation of eliciting a fixed number of initial orders, we propose to estimate the toxicity order based on isotonic regression and then reduce the two-dimensional space to one dimension so that the standard CRM can be applied directly. By doing so, the toxicity order is not prefixed but can be adjusted dynamically during the trial. To enhance the robustness of the proposed method, we further incorporate novel ensemble methods in machine learning to our dose-finding procedure. Specifically, we estimate the toxicity order of the reduced one-dimensional space through the bootstrap aggregating (bagging) approach [Breiman (1996), Hastie, Tibshirani and Friedman (2009)] and the Bayesian model averaging procedure [Raftery et al. (1997), Hoeting et al. (1999)]. Our proposed dimension reduction technique and estimation procedure are novel and fundamentally different from the POCRM which prespecifies a fixed number of orders and uses model selection for dose finding. Not only can the proposed design pin down the MTD more accurately, but it is also much safer in terms of preventing patients from unnecessary toxicities.

The rest of the paper is organized as follows. In Section 2, we present two recent cancer clinical trials that have motivated the proposed method, and the trend of more drug-combination trials is also discussed. In Section 3, we develop the bagging CRM by incorporating both bagging and Bayesian model averaging techniques. In addition, we formulate the decision rules and the dose-finding algorithm. In Section 4, we examine the operating characteristics of the new designs based on simulation studies and sensitivity analysis. We illustrate the proposed method with two clinical trials in Section 5, and Section 6 provides some concluding remarks.

2. Motivating trials.

2.1. *Neratinib and temsirolimus combination trial.* Gandhi et al. (2014) conducted a phase I dose-finding trial for the combination of neratinib and temsirolimus in patients with human epidermal growth factor receptor 2 (HER2), which plays an important role in the development and progression of breast and lung cancers. HER2 activation via overexpression or kinase domain mutation is oncogenic *in vitro* and *in vivo*. Neratinib is a small-molecule irreversible pan-HER tyrosine kinase inhibitor. However, preclinical studies showed that neratinib alone is insufficient for complete extinguishing of mammalian target of rapamycin

(mTOR) activity and may cause both primary and acquired resistances. The addition of an mTOR inhibitor to a HER2 inhibitor has been demonstrated via *in vivo* models to result in synergistic tumor growth inhibition and regression. The primary goal of the trial is to study the synergistic benefits of the combined drugs of neratinib and temsirolimus (an mTOR inhibitor) in HER2-driven tumors as well as to identify the MTD combination with a target DLT rate of 33%. The inability to maintain the prescribed doses during the first 28 days is considered as a DLT. In the trial, 60 patients were treated by possible dosing combinations of four dose levels of neratinib and four dose levels of temsirolimus.

The dose assignment of this trial is determined by a modified up-and-down design [Ivanova and Wang (2004)], and at the end of the trial two MTD combinations were identified. However, the up-and-down design for drug-combination trials suffers from several limitations. First, the cohort size cannot be easily adapted due to the one-to-one correspondent relationship between the cohort size and the target toxicity rate. Second, the allocation rule is ad hoc and is strictly bundled with the cohort size. Last but not least, only the cumulative information of the current dose level is utilized in determining the next assigned dose level, which indicates the inefficiency of the design as it does not fully utilize the information across different dose levels. As a result, the up-and-down design may not be able to identify the true MTD combinations accurately, and it has been demonstrated by simulations to be inferior to existing model-based methods [Riviere et al. (2015a)].

2.2. *Capecitabine and bosutinib combination trial*. Capecitabine, an oral 5-fluorouracil (5-FU) prodrug, has been shown to be effective in treating metastatic colorectal cancer and metastatic breast cancer. Unfortunately, resistance to capecitabine has been observed in several solid tumor models. Preclinical studies showed that the combined therapy of a Src inhibitor and a 5-FU inhibitor would result in synergistic tumor growth inhibition and regression. Isakoff et al. (2014) conducted a phase I trial to study the safety and efficacy of the dosing combinations of capecitabine and bosutinib (a Src inhibitor). The primary objective of the trial is to investigate the MTD dose levels among nine dose combinations consisting of three dose levels of capecitabine and three dose levels of bosutinib. Tumor assessments were performed for all patients at the screening stage, and response was evaluated every six weeks until disease progression or treatment discontinuation. The target toxicity rate was 33%, and each cohort contained two patients. A total of thirty-two patients were enrolled with twenty-four patients assessed in the dose escalation stage, which was guided by an up-and-down design, to determine the MTD.

2.3. *Trend of drug-combination trials*. In the era of precision medicine, combination therapies are becoming more and more popular since the benefit with single-agent treatment is generally modest [Papadatos-Pastos et al. (2015)]. Most of the real trials in practice used inefficient designs, such as the conventional

$3 + 3$ design, to determine the dose escalation strategy. For example, Wilky et al. (2015) applied a $3 + 3$ design to study the combined treatment of cixutumumab and selumetinib. Saura et al. (2014) enrolled 33 patients in a multinational, open-label, phase I/II trial to determine the MTD among the dosing combinations with neratinib (three levels) and capecitabine (two levels). A modified $3 + 3$ design was utilized in dose assignment, which led to approximately one-third of the patients being treated at the over-toxic dose levels. Ullenhag et al. (2015) studied the lenalidomide in combination with gemcitabine in patients with advanced pancreatic cancer based on an ad hoc dose-escalation rule. Siegel et al. (2014) used the $3+3$ design to determine the MTD of vorinostat in combination with lenalidomide and dexamethasone in patients with relapsed or refractory multiple myeloma, and Bendell et al. (2015) also applied the $3+3$ design to identify the MTD of linsitinib in combination with everolimus in patients with refractory metastatic colorectal cancer. Although the $3 + 3$ design dominates practical applications, its limitations have been widely recognized and its performance has been criticized; for instance, see Korn et al. (1994) and Ahn (1998). The drawbacks of the $3+3$ design would be more prominent in drug-combination trials due to small sample sizes yet a larger searching space. In addition to its poor operating characteristics and tendency for an inconsistent estimate of the MTD in single-agent trials, the application of the $3 + 3$ design to drug-combination trials has a higher risk of excluding the true MTDs because often only a subset of dose combinations is selected for investigation. These recent trial examples and limitations of the current practice demonstrate the urgent need for a novel dose-finding design which not only is easy to implement but also has desirable properties.

3. Methodology.

3.1. *Dynamic ordering.* In a two-dimensional dose-finding study, we consider combining J dose levels of drug A and K levels of drug B, which leads to a total of $J \times K$ dose levels in combination. Let p_{jk} denote the true toxicity probability of dose combination (A_j, B_k), $j = 1, \ldots, J$; $k = 1, \ldots, K$. The aim of dose finding in phase I drug-combination trials is to identify the dose combination that has the toxicity probability closest to the target toxicity rate ϕ. For the ith patient, $i = 1, \ldots, n$, let t_i denote the toxicity outcome; that is, $t_i = 1$ if the patient has experienced DLT, and $t_i = 0$ otherwise. Let $d_i = (A_j, B_k)$ if patient i is treated at the dose combination (A_j, B_k). The data collected up to the nth patient can be represented as

$$D = \begin{bmatrix} t_1 & \cdots & t_n \\ d_1 & \cdots & d_n \end{bmatrix}.$$

To reduce the number of unknown parameters in a model-based approach, we propose a dimension-reduction technique that reduces the two-dimensional space to one dimension. Our procedure can be split into two steps: First, we estimate

the toxicity orders of all the dose combinations and sort these dose levels into a one-dimensional vector. Second, we apply a commonly used single-agent model to quantify the ordered toxicity probabilities. The model for a single-agent trial usually involves fewer parameters, which can increase the stability of the estimation procedure.

Based on the partial ordering information, we can specify an initial toxicity order for the dose combinations from the lowest to the highest,

$$\mathcal{O}_0 : (A_1, B_1) \prec (A_1, B_2) \prec (A_2, B_1) \prec \cdots \prec (A_J, B_K),$$

where $(A_j, B_k) \prec (A_{j'}, B_{k'})$ means that dose combination $(A_{j'}, B_{k'})$ is more toxic than (A_j, B_k). Let $\mathcal{O}_0(k)$ denote the dose level that ranks the kth place in $\mathcal{O}_0$. For example, $\mathcal{O}_0(1) = (A_1, B_1)$ is the lowest dose level in the drug-combination space, and $\mathcal{O}_0(3) = (A_2, B_1)$. The initial order can be arbitrary but must follow the partial order requirement; that is, for $j = 1, \ldots, J$; $k = 1, \ldots, K$,

$$(A_j, B_k) \prec (A_{j+1}, B_k) \quad \text{and} \quad (A_j, B_k) \prec (A_j, B_{k+1}).$$

As demonstrated by the sensitivity analysis in Section 4.2, the specified initial order has a minimal effect on the performance of the proposed method. As the trial proceeds, the toxicity order of the two-dimensional dose space is dynamically adjusted via a nonparametric method.

Under the Bayesian paradigm, we assign a beta prior distribution to p_{jk}, that is, $p_{jk} \sim \text{Beta}(a, b)$, where a and b are hyperparameters taking small values to ensure noninformativeness. Thus, the posterior mean of p_{jk} is

$$\bar{p}_{jk} = \frac{y_{jk} + a}{n_{jk} + a + b},$$

where $n_{jk} = \sum_{i=1}^{n} I\{d_i = (A_j, B_k)\}$ denotes the number of patients assigned to the dose combination (A_j, B_k), $y_{jk} = \sum_{i=1}^{n} t_i I\{d_i = (A_j, B_k)\}$ is the number of observed DLTs, and $I\{\cdot\}$ represents the indicator function. As an alternative, one can also assign different beta priors to p_{jk}'s based on the partial order constraint. Let a_{jk} and b_{jk} denote the hyperparameters of the beta prior distribution for the dose combination (A_j, B_k). We vary the value of a_{jk} according to the partial orders, while fixing the sum $a_{jk} + b_{jk} = n_0$,

$$a_{jk} = n_0 p_{\min} + \frac{n_0(j + k - 2)(p_{\max} - p_{\min})}{J + K - 2}, \qquad j = 1, \ldots, J; k = 1, \ldots, K,$$

where a small number n_0 indicates a small effective sample size and thus weak prior information, and $p_{\min}$ and $p_{\max}$ are the prespecified toxicity probabilities for the lowest dose combination (A_1, B_1) and the highest dose combination (A_J, B_K), respectively. For example, considering $n_0 = 0.1$, $p_{\min} = 0.15$ and $p_{\max} = 0.55$ for a 3×3 drug-combination trial, the values of (a_{jk}, b_{jk}) are given by

$$\{(a_{jk}, b_{jk})\} = \begin{bmatrix} (0.035, 0.065) & (0.045, 0.055) & (0.055, 0.045) \\ (0.025, 0.075) & (0.035, 0.065) & (0.045, 0.055) \\ (0.015, 0.085) & (0.025, 0.075) & (0.035, 0.065) \end{bmatrix},$$

which leads to the prior means of p_{jk},

$$\{\bar{p}_{jk}\} = \begin{bmatrix} 0.35 & 0.45 & 0.55 \\ 0.25 & 0.35 & 0.45 \\ 0.15 & 0.25 & 0.35 \end{bmatrix}.$$

Based on a simulation study (not reported here due to the space limit), there is no gain in the performance by considering the partial orders in the prior specification, as long as the prior distributions on p_{jk}'s are noninformative, for example, setting $n_0 = 0.1$ corresponds to only $1/10$ subject's information.

To ensure the estimated toxicity rates satisfying the partial ordering constraint when fixing one drug at a certain dose level, we perform the isotonic regression on $\{\bar{p}_{jk}\}$. The isotonic regression is a model-free method to fit the data with monotone or partial ordering constraints, which ensures the monotonic dose–response relationship [Yuan and Chappell (2004)]. Specifically, we apply the two-dimensional PAVA algorithm [Bril et al. (1984)] to $\{\bar{p}_{jk}\}$, and denote the isotonically transformed values by $\{\tilde{p}_{jk}\}$. However, there are often ties for these $\tilde{p}_{jk}$'s due to the sparsity of the data. In a sequential way, we can utilize the information in the previous round of ordering to break the ties through the adjustment as follows:

$$\tilde{p}^{\dagger}_{jk} = \tilde{p}_{jk} + r_{jk}\varepsilon,$$

where r_{jk} denotes the rank of dose level (A_j, B_k) in the estimated order based on the data up to the previous cohort ($\mathcal{O}_0$ is utilized for the first cohort of patients), and ε is a very small positive number, for example, $\varepsilon = 0.001$. Thus, the dynamically estimated order is given by

$$\mathcal{O} = \text{order}(\{\tilde{p}^{\dagger}_{jk}\}),$$

where order$(\cdot)$ is a function that returns a permutation coordinate vector that rearranges a two-dimensional matrix into an ascending order. We break the ties using the previous ordering after the PAVA algorithm, which is a sequential updating process as newly enrolled patients may lead to substantial change in the toxicity ordering.

As an illustration, we consider an example with 2 levels of drug A and 3 levels of drug B. The initial order is specified as

$$\mathcal{O}_0 : (A_1, B_1) \prec (A_1, B_2) \prec (A_2, B_1) \prec (A_1, B_3) \prec (A_2, B_2) \prec (A_2, B_3).$$

Suppose that after treating the first fifteen patients, the collected data are given by

$$\{y_{jk}\} = \begin{bmatrix} 1 & - & - \\ 0 & 2 & 3 \end{bmatrix}, \qquad \{n_{jk}\} = \begin{bmatrix} 6 & - & - \\ 3 & 3 & 3 \end{bmatrix},$$

where "$-$" means no patient assigned to the corresponding dose combination. If we assign Beta$(0.05, 0.05)$ as the prior distribution for each p_{jk}, $j = 1, 2$, $k =$

1, 2, 3, then the isotonically transformed estimates of the toxicity rates are given by

$$\{\tilde{p}_{jk}\} = \begin{bmatrix} 0.169 & 0.644 & 0.956 \\ 0.010 & 0.644 & 0.956 \end{bmatrix},$$

and the adjusted estimates after breaking the ties are

$$\{\tilde{p}^{\dagger}_{jk}\} = \begin{bmatrix} 0.172 & \searrow & 0.649 & \searrow & 0.962 \\ \uparrow & & \uparrow & & \uparrow \\ 0.011 & & 0.646 & & 0.960 \end{bmatrix}.$$

As a result, the dynamically adjusted order $\mathcal{O}$ is

$$\begin{aligned} \mathcal{O} &= \text{order}(\{\tilde{p}^{\dagger}_{jk}\}) \\ &= \{(A_1, B_1) \prec (A_2, B_1) \prec (A_1, B_2) \prec (A_2, B_2) \prec (A_1, B_3) \prec (A_2, B_3)\}. \end{aligned}$$

The dynamic adjustment is an adaptive procedure: If the initial guess satisfies the order implied by the observed data, the initial order retains; otherwise, the initial order is restructured according to the observed information.

Based on the order $\mathcal{O}$, the two-dimensional space in a drug-combination trial can be reduced to a one-dimensional line with a monotonically increasing toxicity order, and thus the conventional CRM can be utilized to characterize the dose–toxicity curve. Owing to its desirable properties and superior performance in single-agent dose-finding trials [Heyd and Carlin (1999), Yin and Yuan (2009b), O'Quigley and Conaway (2010), Cheung (2011)], we expect the CRM to preserve its advantages in drug-combination trials as long as the toxicity ordering can be pinned down correctly. Specifically, in the power model [Shen and O'Quigley (1996)], let $q_1 < \cdots < q_{JK}$ be the prespecified toxicity probabilities of a set of $J \times K$ dose levels, which is typically known as the skeleton of the CRM. Under the estimated order $\mathcal{O}$, the toxicity probability at dose level $\mathcal{O}(l)$ $(l = 1, \ldots, JK)$ is

$$\Pr\big(\text{toxicity at dose level } \mathcal{O}(l)\big) = q_l^{\exp(\alpha)}, \tag{3.1}$$

where α is the only unknown parameter that needs to be estimated. Such a power model is simple yet flexible enough to characterize the underlying dose–toxicity curve in dose-finding studies [Yin (2012)]. The likelihood function based on the observed data D is given by

$$L(\alpha \mid \mathcal{O}, D) \propto \prod_{l=1}^{JK} \{q_l^{\exp(\alpha)}\}^{y_{\mathcal{O}(l)}} \{1 - q_l^{\exp(\alpha)}\}^{n_{\mathcal{O}(l)} - y_{\mathcal{O}(l)}},$$

where $y_{\mathcal{O}(l)}$ and $n_{\mathcal{O}(l)}$ represent the number of observed DLTs and the number of treated patients at dose level $\mathcal{O}(l)$, respectively. Let $f(\alpha)$ be a proper prior

distribution for α, and then the posterior mean of the toxicity probability at dose level (A_j, B_k) under the aforementioned dynamic ordering scheme is given by

$$\hat{p}_{jk} = \int q_{r_{jk}}^{\exp(\alpha)} \frac{L(\alpha \mid \mathcal{O}, D) f(\alpha)}{\int L(\alpha \mid \mathcal{O}, D) f(\alpha)\, d\alpha}\, d\alpha, \tag{3.2}$$

where r_{jk} is the rank of dose combination (A_j, B_k) in $\mathcal{O}$. In addition to the estimate of the toxicity probability, the posterior probability that dose level (A_j, B_k) is overly toxic is given by

$$\Pr(p_{jk} > \phi \mid \mathcal{O}, D) = \int_{-\infty}^{\log\{\log(\phi)/\log(q_{r_{jk}})\}} \frac{L(\alpha \mid \mathcal{O}, D) f(\alpha)}{\int L(\alpha \mid \mathcal{O}, D) f(\alpha)\, d\alpha}\, d\alpha,$$

which can be used for decision-making on dose escalation or de-escalation.

3.2. *Bagging CRM*. Due to the sparsity of the data, the toxicity order of the drug-combination space may vary dramatically for consecutive cohorts, which means the estimated toxicity probabilities $\hat{p}_{jk}$ in (3.2) tend to suffer from high variability. On the other hand, Bootstrap aggregating (bagging), a powerful machine learning ensemble technique in statistical classification and regression, averages the inference results based on multiple bootstrap samples [Breiman (1996), Hastie, Tibshirani and Friedman (2009)]. As a result, the combined classifier or predictor possesses low variance and high accuracy, and also helps to avoid overfitting. To enhance the stability of the estimation, we consider the bagging procedure to quantify the toxicity order.

Based on the cumulative data D, we construct B bootstrap samples of size n,

$$D^b = \begin{bmatrix} t_1^b & \cdots & t_n^b \\ d_1^b & \cdots & d_n^b \end{bmatrix}, \qquad b = 1, \ldots, B,$$

where each pair of (t_i^b, d_i^b) is uniformly drawn from the data D with replacement. For each bootstrap sample D^b, we apply the dynamic ordering approach described in Section 3.1 to obtain an estimated order $\mathcal{O}_b$, $b = 1, \ldots, B$. After removing the duplicate orders, we collect B_U ($B_U \le B$) sets of unique orders, $\mathcal{O}_1, \ldots, \mathcal{O}_{B_U}$. For each of the unique orders $\mathcal{O}_b$, $b = 1, \ldots, B_U$, we obtain the posterior estimates of the toxicity probabilities based on $\mathcal{O}_b$ and D (i.e., we use the bagging order and the original data to fit the CRM model), which are denoted by $\{\hat{p}_{jk}^b\}$, $b = 1, \ldots, B_U$. Finally, we apply a Bayesian model averaging (BMA) procedure [Raftery et al. (1997), Hoeting et al. (1999)] to combine these bagging estimators. Specifically, the bagging estimate of p_{jk} is a weighted average of $\hat{p}_{jk}^b$,

$$\hat{p}_{jk}^{\text{Bagging}} = \sum_{b=1}^{B_U} \Pr(\mathcal{O}_b \mid D)\, \hat{p}_{jk}^b,$$

where the weight $\Pr(\mathcal{O}_b \mid D)$ is the posterior probability of order $\mathcal{O}_b$,

$$\Pr(\mathcal{O}_b \mid D) = \frac{\Pr(\mathcal{O}_b) \int L(\alpha \mid \mathcal{O}_b, D) f(\alpha)\, d\alpha}{\sum_{b'=1}^{B} \Pr(\mathcal{O}_{b'}) \int L(\alpha \mid \mathcal{O}_{b'}, D) f(\alpha)\, d\alpha},$$

and $\Pr(\mathcal{O}_b)$ is the prior order probability. For simplicity, we set $\Pr(\mathcal{O}_1) = \cdots = \Pr(\mathcal{O}_{B_U}) = 1/B_U$ for equal prior preference on all the orders. In addition, the bagging estimate of the posterior probability of $p_{jk} > \phi$ is given by

$$\Pr\bigl(p_{jk}^{\text{Bagging}} > \phi \mid D\bigr) = \sum_{b=1}^{B_U} \Pr\bigl(p_{jk}^{b} > \phi \mid \mathcal{O}_b, D\bigr) \Pr(\mathcal{O}_b \mid D),$$

where $\Pr(p_{jk}^{b} > \phi \mid \mathcal{O}_b, D)$ is obtained based on the bth bootstrapped order. We can also consider the Bayesian bagging approach [Clyde and Lee (2001)] under the Bayesian bootstrap framework [Rubin (1981)]. For ease of implementation, the standard bootstrap approach is more straightforward without compromising the performance.

3.3. *Dose-finding algorithm.* At the beginning of a trial, the posterior estimates are often unstable as the collected data are very limited. Usually, a start-up phase is required for the model-based dose-finding methods to enhance the reliability of the posterior estimates. For the proposed method, we initiate a start-up phase by treating the first cohort of patients at the lowest dose combination. If no toxicity is observed for that cohort, then we consider a diagonal escalation rule, that is, escalating drug A and drug B by one dose level simultaneously. If one drug reaches its maximum dose level, then we increase the dose level of the other drug until both agents arrive at their individual maximum dose levels. Such an escalation rule may be aggressive at the very beginning of the trial, but it can be more efficient in locating the target toxicity level [Riviere et al. (2014)]. As soon as the first toxicity outcome is observed, the prephase stage is completed, and the proposed dose-finding algorithm kicks in seamlessly for the rest of the trial as follows.

We define $c_e > 0$ and $c_d > 0$ ($c_e + c_d > 1$) as the probability cutoffs for dose escalation and de-escalation, respectively, and their values can be calibrated through simulation studies. Suppose the current dose level is (A_j, B_k), and the target toxicity rate is ϕ.

(1) If $\Pr(p_{jk}^{\text{Bagging}} > \phi \mid D) > c_d$, then we define an admissible dose de-escalation set as

$$\mathcal{A}_D = \bigl\{(A_j - 1, B_k), (A_j, B_k - 1), (A_j - 1, B_k + 1), (A_j + 1, B_k - 1)\bigr\},$$

which is constituted by the dose levels adjacent to (A_j, B_k). The next cohort of patients is assigned to dose combination (A_{j^*}, B_{k^*}), which has a toxicity probability lower than (A_j, B_k) but closest to ϕ,

$$(j^*, k^*) = \underset{(A_{j'}, B_{k'}) \in \mathcal{A}_D}{\arg\min} \bigl\{\bigl|\hat{p}_{j'k'}^{\text{Bagging}} - \phi\bigr| / I\bigl(\hat{p}_{j'k'}^{\text{Bagging}} \le \hat{p}_{jk}^{\text{Bagging}}\bigr)\bigr\}.$$

(2) If $\Pr(p_{jk}^{\text{Bagging}} < \phi \mid D) > c_e$, then we define an admissible dose escalation set as

$$\mathcal{A}_E = \{(A_j + 1, B_k), (A_j, B_k + 1), (A_j - 1, B_k + 1), (A_j + 1, B_k - 1)\}.$$

The next cohort of patients is assigned to dose combination (A_{j^*}, B_{k^*}), which has a toxicity probability higher than (A_j, B_k) but closest to ϕ,

$$(j^*, k^*) = \underset{(A_{j'}, B_{k'}) \in \mathcal{A}_E}{\arg\min} \{|\hat{p}_{j'k'}^{\text{Bagging}} - \phi| / I(\hat{p}_{j'k'}^{\text{Bagging}} \geq \hat{p}_{jk}^{\text{Bagging}})\}.$$

(3) Otherwise, we retain the dose at the same level.

If an escalation decision is made at the highest dose level (A_J, B_K) or a de-escalation decision is made at the lowest dose level (A_1, B_1), then we retain the same dose level for the next cohort. To prevent the dose assignment from being trapped into some local suboptimal levels, we impose an additional rule: If dose level $(A_{j'}, B_{k'})$ belongs to the admissible dose escalation or de-escalation set while that level has never been tested with patients, then we update the difference between the estimated toxicity probability at $(A_{j'}, B_{k'})$ and the target ϕ to be one-fourth of the original value. This adaptation rule can facilitate the proposed method to visit untried dose combinations more aggressively, such that the entire dose searching space can be explored.

The trial can be stopped after the exhaustion of the maximum sample size N or be terminated early for safety if the lowest dose combination is still overly toxic as indicated by $\Pr(p_{11}^{\text{Bagging}} > \phi \mid D) > \lambda$, where λ is a threshold value close to 1. At the end of the trial, we select the MTD as the dose combination (A_{j^*}, B_{k^*}) that has been tested in the trial and attains the largest posterior probability of falling inside the ϵ-neighborhood of the target ϕ,

$$(A_{j^*}, B_{k^*}) = \underset{(A_j, B_k) \in \mathcal{N}}{\arg\max} \Pr\big(\phi - \epsilon < p_{jk}^{\text{Bagging}} < \phi + \epsilon \mid D\big),$$

where $\mathcal{N}$ is the set that contains all the tried dose combinations, and ϵ is a small positive number, for example, $\epsilon = 0.1$.

4. Simulation study.

4.1. *Performance with finite samples.* We conduct extensive simulation studies to investigate the operating characteristics of the bagging CRM in comparison with four existing methods: the partial ordering CRM (POCRM) [Wages et al. (2011)], the Clayton copula-type regression [Yin and Yuan (2009a)], the logistic method [Riviere et al. (2014)] and the two-dimensional Bayesian optimal interval design (2d-BOIN) [Lin and Yin (2016)]. The POCRM reduces the two-dimensional searching space to one dimension based on several prespecified toxicity orders; the copula design links the individual toxicity probabilities of the two

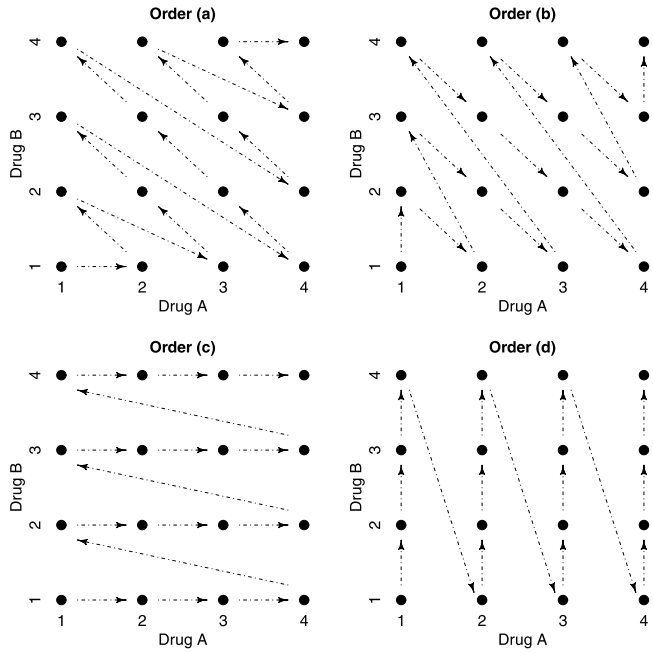

FIG. 1. *Four different specifications of the initial order.*

drugs via a copula function; the logistic design quantifies the joint toxicity probabilities of the combined drugs using a four-parameter logistic regression model; and the 2d-BOIN design utilizes a predetermined toxicity tolerance interval to guide the dose-finding procedure. Following the guidance of the POCRM, we select six possible orders with an equal prior probability, and we use the R package `pocrm` with a slight modification to allow for a cohort size of 3. The skeleton of the CRM is chosen using the `getprior` function in the R package `dfcrm` by specifying $\eta = 0.03$ and an initial guess of the MTD position at 8. The simulation configuration of the copula design follows the recommendation in Yin and Lin (2015). The results of logistic and 2d-BOIN designs are obtained using the R packages `dfcomb` and `boin` under their default settings, respectively. For the proposed bagging CRM, we set the skeletons to be the same as those in the POCRM and assign a noninformative prior distribution $\alpha \sim \mathrm{N}(0, 2)$ in the power model (3.1). We assign a noninformative prior distribution $\mathrm{Beta}(0.05, 0.05)$ to p_{jk}, and the initial order is specified as in Figure 1(a). To implement the bagging CRM, we take the cutoff probabilities $c_e = 0.7$ and $c_d = 0.5$, and draw 50 bootstrap samples, that is, $B = 50$. We also examine a reduced form of the bagging CRM with $B = 0$ (denoted as the bagging-0 CRM), which only estimates the toxicity order once based on the observed data according to the dynamic ordering procedure described in Section 3.1. We investigate whether there is any gain of the bagging-50

CRM (with 50 bootstrap samples) over the bagging-0 CRM. The diagonal escalation rule described in Section 3.3 is implemented as the start-up phase for all the model-based methods. For a fair comparison, we do not impose any early termination for all the considered designs.

The target toxicity probability ϕ is set at 0.3, and the trial starts from the lowest dose level (A_1, B_1). The total sample size is 60 with a cohort size of 3. For a patient treated at the dose combination (A_j, B_k), the toxicity outcome is generated as a binary random variable that takes a value of 1 with probability p_{jk} and 0 otherwise. We consider 16 scenarios involving eight 4×4 and eight 3×5 drug-combination spaces, where the true toxicity probabilities, p_{jk}'s, are given in Table 1. These 16 toxicity scenarios cover a wide range of dose–response relationships commonly encountered in real trials. In particular, scenario 1 has four MTD combinations along the diagonal line. There exist three MTD combinations in scenarios 2 and 3, and scenario 3 is typically considered more difficult as the MTDs in this scenario are far away from the starting dose and (A_1, B_4) is isolated from the other two MTDs. In each of scenarios 4 to 7, there are two MTD combinations which are located at different positions. The goal of scenario 4 is to examine the design's performance when the MTD combinations are located at the boundaries of the dose searching space and the MTD contour is not complete, while scenario 5 mainly aims to investigate the safety aspect of the design as the MTDs are more concentrated at the lower dose combinations. Only one MTD exists in scenario 8, and the model-based approaches might yield a large incorrect selection percentage if the model is misspecified. Scenarios 9–16 investigate the performance of our design under an asymmetric dose-combination space with five dose levels of drug A and three dose levels of drug B.

To quantify the operating characteristics of the dose-finding methods, we consider six performance statistics, including the percentage of correct MTD selection, the percentage of patients treated at the true MTDs, an accuracy index that represents the entire distribution of selected doses, the percentage of trials that select over-toxic dose combinations as the MTDs, the percentage of patients allocated to the over-toxic dose combinations, and the percentage of patients experiencing toxicities. The definition of the accuracy index is given by

$$\mathcal{AI} = 1 - JK \frac{\sum_{j=1}^{J} \sum_{k=1}^{K} w_{jk} |p_{jk} - \phi|}{\sum_{j=1}^{J} \sum_{k=1}^{K} |p_{jk} - \phi|},$$

where w_{jk} is the probability of selecting dose (A_j, B_k) as the MTD [Cheung (2011)]. The first three statistics, for which the larger the better, reflect the accuracy and efficiency of a design, while the remaining three are related to the safety aspects of a trial, and thus a design with smaller values of these three statistics is considered more ethical and desirable.

The simulation study is focused on the comparisons among the bagging-50 CRM, bagging-0 CRM and POCRM, as these three methods implement dimension reduction techniques for dose finding. For each scenario, we replicate 1000

TABLE 1
Sixteen toxicity scenarios for two-drug combinations, with a target toxicity probability of 30% in shaded cells

		Agent A									
	Dose level	**1**	**2**	**3**	**4**	**5**	**1**	**2**	**3**	**4**	**5**
		Scenario 1					Scenario 2				
Agent B	4	0.30	0.45	0.60	0.70		0.50	0.55	0.60	0.70	
	3	0.15	0.30	0.48	0.60		0.30	0.50	0.55	0.60	
	2	0.10	0.20	0.30	0.50		0.12	0.30	0.50	0.55	
	1	0.08	0.14	0.19	0.30		0.10	0.15	0.30	0.45	
		Scenario 3					Scenario 4				
	4	0.30	0.52	0.60	0.70		0.20	0.30	0.50	0.60	
	3	0.19	0.40	0.45	0.60		0.15	0.20	0.45	0.55	
	2	0.10	0.21	0.30	0.40		0.10	0.15	0.20	0.30	
	1	0.02	0.11	0.20	0.30		0.05	0.10	0.15	0.20	
		Scenario 5					Scenario 6				
	4	0.48	0.52	0.55	0.58		0.50	0.55	0.60	0.70	
	3	0.42	0.45	0.50	0.52		0.15	0.30	0.45	0.60	
	2	0.30	0.40	0.48	0.50		0.10	0.18	0.30	0.45	
	1	0.15	0.30	0.40	0.45		0.06	0.08	0.10	0.15	
		Scenario 7					Scenario 8				
	4	0.42	0.48	0.60	0.65		0.50	0.60	0.67	0.73	
	3	0.20	0.30	0.45	0.61		0.40	0.54	0.62	0.68	
	2	0.10	0.20	0.40	0.46		0.20	0.30	0.50	0.60	
	1	0.05	0.15	0.30	0.45		0.15	0.20	0.40	0.56	
		Scenario 9					Scenario 10				
	3	0.15	0.30	0.45	0.50	0.60	0.30	0.40	0.45	0.55	0.60
	2	0.10	0.15	0.30	0.45	0.55	0.07	0.10	0.20	0.30	0.55
	1	0.05	0.10	0.15	0.30	0.45	0.02	0.07	0.10	0.15	0.30
		Scenario 11					Scenario 12				
	3	0.30	0.45	0.50	0.65	0.75	0.30	0.50	0.60	0.70	0.80
	2	0.15	0.30	0.45	0.52	0.60	0.20	0.45	0.50	0.60	0.75
	1	0.07	0.10	0.12	0.30	0.50	0.15	0.20	0.30	0.50	0.60
		Scenario 13					Scenario 14				
	3	0.15	0.30	0.45	0.50	0.60	0.50	0.60	0.70	0.80	0.90
	2	0.09	0.12	0.15	0.30	0.45	0.45	0.55	0.65	0.75	0.85
	1	0.05	0.08	0.10	0.13	0.15	0.30	0.45	0.60	0.70	0.80
		Scenario 15					Scenario 16				
	3	0.07	0.09	0.12	0.15	0.30	0.08	0.20	0.40	0.55	0.65
	2	0.03	0.05	0.10	0.13	0.15	0.05	0.08	0.15	0.30	0.50
	1	0.01	0.02	0.08	0.10	0.11	0.02	0.05	0.10	0.19	0.45

trials and summarize the respective averages of the six performance statistics in Table 2. Overall, the bagging-50 CRM and bagging-0 CRM show improvement of 5% to 20% in terms of the MTD selection percentage over the POCRM in most of the scenarios, and the superiority of the bagging CRMs is most striking in scenario 16. In reality, it may happen that all the six prespecified orders in the POCRM deviate far from the truth and, as a result, such a fixed number of initial orders cannot cover all the possible ordering profiles. By contrast, the bagging CRM does not limit the possibilities of ordering, and it continually re-estimates the toxicity order based on the cumulative data, even though the initial order might be incorrectly specified. Averaged over all the sixteen scenarios, the bagging-50 CRM has 4.3% more chance to identify the true MTDs than the POCRM, which is a real gain that should not be undervalued given how critical it is to pin down the right dose. Comparing the bagging-0 CRM and bagging-50 CRM, as expected, the latter dominates the former due to reduction of the variation. The performances of the other two statistics related to accuracy, including the percentage of patients treated at the true MTDs and the accuracy index, are in line with the MTD selection criterion, that is, both demonstrate improvement over the POCRM. In addition, by comparing with other existing methods in Table 2, the bagging-50 CRM also has superior performances in terms of accuracy and efficiency: The bagging-50 CRM is more stable than the logistic method, especially in scenarios 3, 4, 8, 10, 12 and 16; and it is more efficient in identifying the MTD combinations than the copula and 2d-BOIN designs. Furthermore, Table 3 shows that the standard deviations for the percentages of patients assigned to the MTDs under the bagging-50 CRM are generally lower than others. As the variability quantifies the reliability of a design [Oron and Hoff (2013)], the highest overall MTD selection percentage in conjunction with low variability demonstrates the accuracy and reliability of our bagging CRM.

The three safety statistics are of great practical importance, and a trial with poor safety control should be deemed to be of high risk. Among all the six designs, the logistic approach appears to be the safest, as it allocates fewer patients to over-toxic dose combinations and also tends to avoid selecting high-risk dose levels. The bagging-50 CRM ranks the second, under which on average 4.6 fewer patients are allocated to over-toxic dose combinations than the POCRM. The safety advantage is especially prominent in scenarios 5, 7, 8 and 10, where the bagging-50 CRM reduces the percentage of patients allocated to over-toxic dose combinations about 10% in comparison with the POCRM. Due to the higher variability of the bagging-0 CRM, the posterior estimates of the toxicity probabilities are less stable, and thus the bagging-0 CRM tends to select a higher dose combination for each treatment assignment than the bagging-50 CRM. All the six considered designs exhibit similar patterns in terms of the percentage of patients experiencing the DLT. In summary, the simulation results demonstrate that both the proposed bagging-0 and bagging-50 CRMs outperform POCRM substantially, and the bagging-50

TABLE 2

Operating characteristics of the proposed bagging-50 CRM (with 50 bootstrap samples) in comparison with the bagging-0 CRM (without bootstrapping), POCRM, copula, logistic and 2d-BOIN designs under sixteen scenarios

	Scenarios																
Designs	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**Ave**
	% Correct MTD selection																
Bagging-50	68.6	82.1	44.1	30.6	61.2	67.0	45.2	34.2	76.1	55.5	74.0	49.4	56.2	86.8	72.1	44.0	59.2
Bagging-0	71.7	74.6	49.1	31.8	56.6	60.7	48.3	33.8	73.0	58.3	70.6	45.2	56.0	79.4	72.3	34.3	57.2
POCRM	72.8	75.5	50.4	38.4	51.8	59.4	37.7	23.6	70.6	48.1	68.4	52.4	55.6	75.4	75.4	23.0	54.9
Copula	65.1	77.1	31.2	26.1	60.8	45.2	42.4	34.1	62.6	48.1	58.1	30.2	45.7	83.3	92.0	24.4	51.2
Logistic	68.9	83.5	36.1	15.1	71.6	68.7	44.6	18.9	75.7	23.5	68.3	35.3	65.8	86.6	79.3	26.0	54.2
2d-BOIN	67.9	79.0	45.8	48.7	65.3	54.4	44.5	37.4	66.1	56.5	70.1	48.6	57.2	79.5	74.8	23.2	57.4
	% Patients treated at true MTDs																
Bagging-50	42.5	53.8	28.6	15.4	42.6	35.3	28.1	21.8	46.1	33.1	44.8	28.1	29.1	73.2	37.5	22.1	36.4
Bagging-0	43.8	52.0	30.5	16.5	40.9	34.6	28.9	22.7	45.8	34.6	45.7	24.1	29.6	71.2	37.6	19.4	36.0
POCRM	50.0	51.1	35.1	22.1	32.9	36.0	25.9	19.7	45.7	33.7	49.9	31.5	32.3	65.2	52.7	14.8	37.4
Copula	41.0	50.4	23.5	16.0	41.4	31.2	26.4	25.1	36.6	31.8	40.5	20.2	27.6	82.5	53.1	17.9	35.3
Logistic	41.6	53.8	23.1	10.0	43.4	37.1	25.8	17.8	44.8	14.0	41.9	22.3	33.3	77.7	39.6	14.5	33.8
2d-BOIN	40.0	48.7	26.7	24.7	46.5	31.2	26.5	21.2	41.9	31.3	43.2	27.6	33.9	73.6	44.2	11.5	35.8
	Accuracy index																
Bagging-50	74.8	82.7	62.7	46.9	70.1	74.1	63.0	68.4	76.5	68.2	77.5	70.8	59.5	94.0	78.0	65.2	70.8
Bagging-0	75.1	75.5	65.5	46.6	66.8	67.6	63.3	67.0	73.3	68.2	74.3	67.7	59.1	90.5	78.3	58.5	68.6
POCRM	75.4	77.3	65.3	51.2	62.0	65.2	53.9	61.5	69.0	56.5	70.3	70.3	56.1	88.2	79.3	46.4	65.5
Copula	66.7	77.5	52.7	34.4	70.8	53.9	60.7	67.5	62.6	56.7	64.0	59.2	49.4	93.5	94.0	47.0	63.2
Logistic	75.1	84.2	59.2	31.6	78.2	75.6	63.9	61.7	76.0	48.5	72.7	65.8	68.2	94.1	83.9	56.8	68.5
2d-BOIN	70.3	80.3	63.3	57.4	72.5	63.4	60.0	68.0	66.3	65.3	73.4	70.3	59.8	93.5	80.7	52.8	68.6

TABLE 2
(*Continued*)

Designs	Scenarios																
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**Ave**
	% Trials selecting over-toxic doses																
Bagging-50	6.7	12.1	30.8	15.0	31.9	13.7	26.7	35.0	12.4	18.6	18.3	21.8	17.9	13.2	0.0	26.8	18.8
Bagging-0	9.9	19.7	28.7	17.5	39.4	20.0	34.0	42.6	16.7	22.6	22.7	26.5	20.3	20.6	0.0	35.1	23.5
POCRM	9.6	16.9	28.3	15.7	44.8	18.6	40.5	47.5	16.9	26.7	21.0	22.1	18.3	24.6	0.0	23.0	23.4
Copula	17.5	15.1	39.2	35.8	32.3	35.7	35.7	36.7	23.7	38.5	35.5	33.5	41.0	16.7	0.0	45.9	30.2
Logistic	4.9	8.2	27.6	28.0	23.3	12.4	25.1	43.2	9.8	40.8	22.1	20.2	16.5	13.4	0.0	33.0	20.5
2d-BOIN	14.2	13.1	27.2	14.9	27.3	20.6	31.7	33.4	18.0	25.0	18.3	20.2	20.1	14.2	0.0	34.4	20.8
	% Patients treated at over-toxic doses																
Bagging-50	11.3	21.9	27.0	15.0	38.0	19.5	30.5	35.6	18.5	20.9	24.0	26.6	17.2	26.8	0.0	25.7	22.4
Bagging-0	14.9	24.6	28.9	17.6	40.4	23.4	34.1	36.5	21.5	24.7	26.4	30.2	20.5	28.4	0.0	29.5	25.1
POCRM	16.5	30.3	32.9	17.4	52.5	27.0	42.6	46.8	25.9	32.3	29.1	33.0	22.6	34.8	0.0	37.8	30.1
Copula	14.6	17.5	27.9	20.1	31.0	23.6	27.4	35.2	23.0	23.4	20.9	23.6	22.1	17.5	0.0	31.2	22.4
Logistic	9.3	14.9	24.0	15.7	29.9	16.1	27.8	31.4	14.0	29.2	19.0	18.8	15.1	22.3	0.0	26.0	19.6
2d-BOIN	17.7	21.6	27.5	18.9	31.2	24.2	29.8	33.8	20.5	28.1	24.6	25.2	21.9	26.4	0.0	32.5	24.0
	% Observed toxicities																
Bagging-50	25.8	30.1	27.6	23.2	31.7	26.1	28.7	30.3	26.8	25.7	28.5	29.5	23.5	34.8	18.5	25.1	27.2
Bagging-0	26.8	30.8	28.7	24.4	32.1	27.2	29.6	30.9	27.6	27.0	29.3	30.2	24.2	35.2	18.5	25.8	28.0
POCRM	28.2	32.6	30.0	24.9	34.6	28.8	31.6	33.3	29.3	29.3	31.0	31.9	25.6	36.6	20.9	27.8	29.8
Copula	26.4	28.0	27.3	25.0	29.2	26.3	27.3	27.9	26.7	26.0	26.5	27.1	24.8	33.1	21.2	26.4	26.8
Logistic	25.3	27.8	23.1	23.2	30.0	25.9	27.8	28.8	25.4	25.5	25.8	26.5	23.6	34.1	18.9	24.6	26.0
2d-BOIN	27.1	29.1	27.1	25.1	30.3	26.8	27.7	29.8	26.5	26.3	28.2	29.0	25.3	34.5	18.7	25.4	27.3

Note: Ave represents the averaged value over 16 scenarios. Under each scenario, the best design for each performance statistic is highlighted by the darker shaded cell, and the second best by the lighter shaded cell.

TABLE 3
Standard deviations of "% patients treated at true MTDs" in Table 2

	Scenarios															
Designs	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**
Bagging-50	19.5	15.6	15.4	13.2	21.2	15.6	12.9	15.8	17.1	16.4	13.6	15.7	13.6	24.8	28.0	14.0
Bagging-0	18.2	16.9	18.4	16.2	21.9	18.0	16.3	18.5	16.3	18.4	16.3	17.5	15.1	27.6	28.2	16.6
POCRM	22.7	20.4	21.4	19.5	23.5	22.1	19.2	19.8	21.5	21.2	20.5	18.6	18.3	30.1	26.8	17.9
Copula	17.1	13.9	16.3	13.6	17.4	15.4	12.7	13.9	14.0	13.6	12.0	13.5	11.3	19.3	20.1	11.6
Logistic	24.9	22.2	20.7	13.7	24.1	21.6	17.8	18.4	22.2	19.1	22.9	20.1	18.7	25.6	25.7	17.0
2d-BOIN	20.7	18.2	24.8	23.2	23.7	22.4	24.1	24.0	20.5	25.1	22.4	26.1	23.2	27.7	21.3	18.5

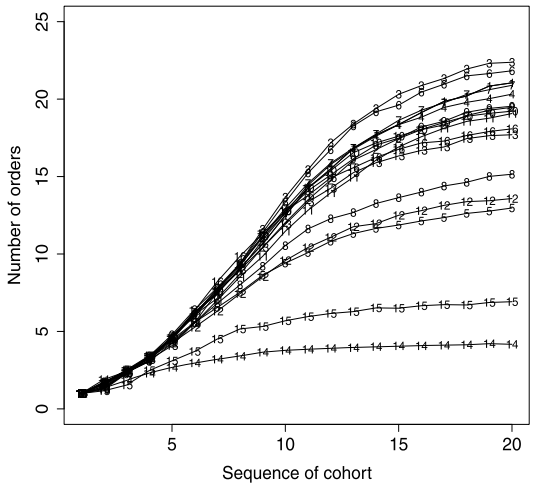

FIG. 2. *Average number of orders generated by the bagging-50 CRM along the sequence of cohorts.*

CRM possesses more stable and competitive operating characteristics in comparison with other existing methods. More importantly, the new approach starts with only one initial order and continually updates the toxicity order with every accrued cohort so that the standard CRM can be applied straightforwardly. The nonparametric estimates of the toxicity probabilities are used for adjusting the toxicity order only, and the parametric CRM model is implemented separately based on the updated toxicity order.

Figure 2 shows the average number of orders generated by the bagging-50 CRM after enrolling each cohort, which clearly increases with the number of accrued patients. However, if the MTDs are located in the lower- or upper-dose areas (i.e., the MTDs lie in the bottom-left or top-right corner of the two-dimensional space), the bagging-50 CRM generates a smaller number of orders, for example, scenarios 5, 8, 12, 14 and 15. Under these scenarios, patient allocation tends to be more concentrated, and the data are likely to be clustered in a small neighborhood of the MTDs. As a result, the regions far away from the MTDs would only experiment with very few patients, and thus the information from these isolated spots has a minimal effect on the bootstrap samples, which in turn leads to fewer orders being generated. For the other scenarios, more than 10 orders are generated after the 10th cohort of patients, which indicates to some extent that only six orders specified by the POCRM are not adequate for practical use.

4.2. *Sensitivity analysis*. To investigate the robustness of the bagging CRM to the specification of the initial order, we consider three other orders, as shown in Figures 1(b), (c) and (d) for the 4×4 drug-combination trials, while the initial order specifications for the 3×5 cases are similar. For each order, we simulate

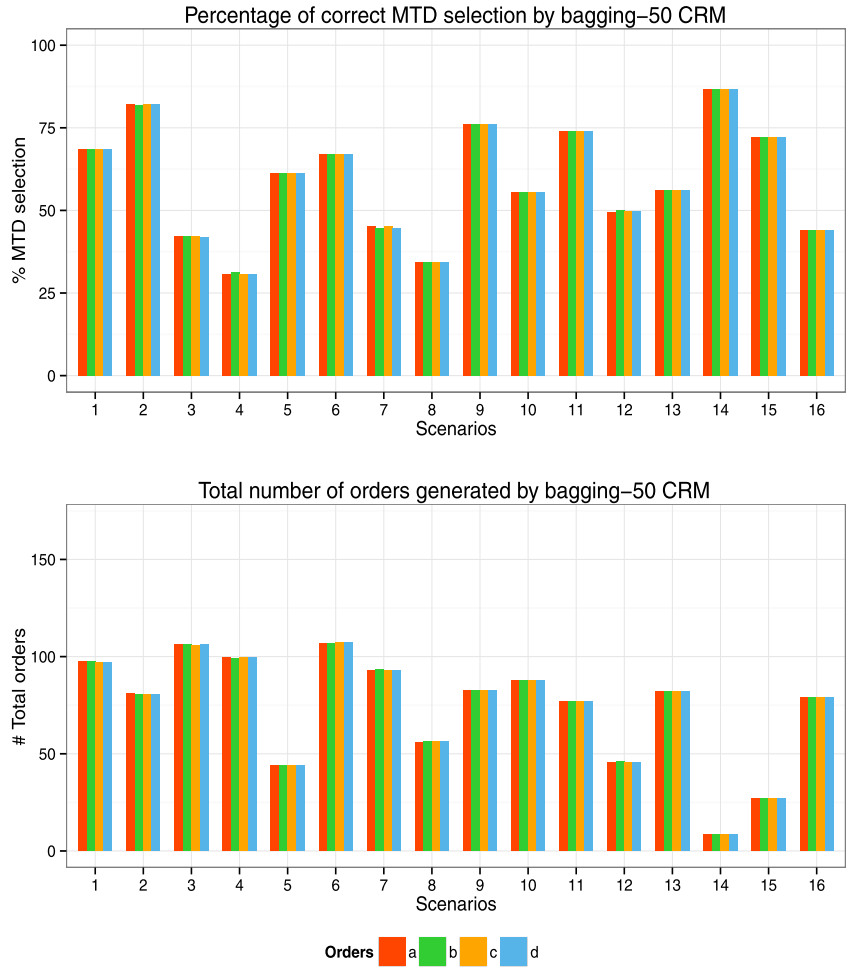

FIG. 3. *Sensitivity analysis of the bagging-50 CRM under four different initial orders in Figure 1.*

1000 trials based on the same settings as described in Section 4.1. Figure 3 displays the percentage of correct MTD selection and the total number of different orders generated by the bagging-50 CRM under four initial orders across 16 scenarios. We conclude that the initial order almost has no impact on both quantities.

To further study the influence of the bagging times, B, on the performance of the bagging CRM, we consider three values of B: 30, 50 and 70. Figure 4 shows the percentages of correct MTD selection with respect to B under the sixteen scenarios. There is a general trend that the percentage of correct MTD selection slowly increases with the number of bagging times, although some small variations are observed. Due to the small sample size in phase I trials, the default value of B is set to be 50, which is typically adequate for generating representative orders.

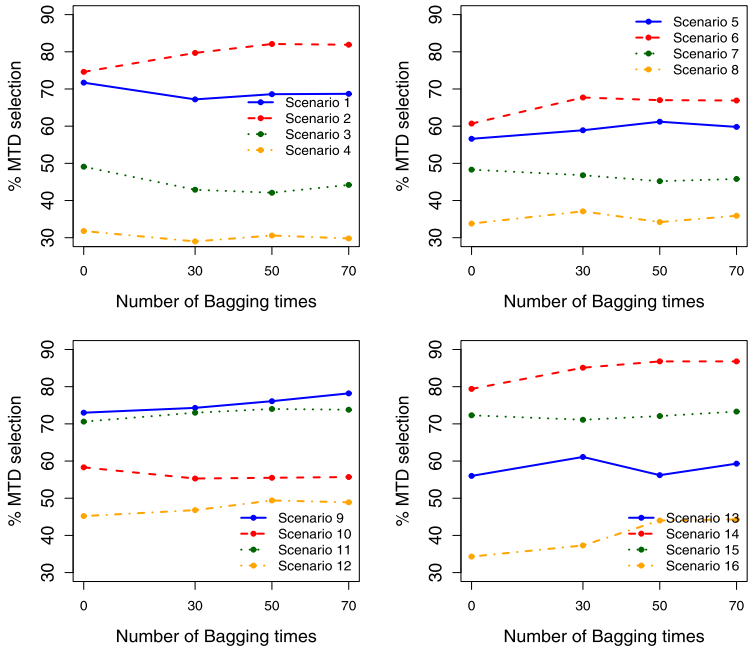

FIG. 4. *Percentages of correct MTD selection with respect to the number of bagging times.*

5. Trial application.

5.1. *Neratinib and temsirolimus combination trial.* For illustration, we apply our bagging CRM to the neratinib and temsirolimus combination trial [Gandhi et al. (2014)]. Four dose levels of neratinib $\{A_1, A_2, A_3, A_4\} = \{120, 160, 200, 240\}$ mg and four dose levels of temsirolimus $\{B_1, B_2, B_3, B_4\} = \{15, 25, 50, 75\}$ mg are investigated. The toxicity rate of each dose combination (A_j, B_k) is estimated by fitting a logistic regression model to the trial data,

$$\text{logit}(p_{jk}) = \beta_0 + \beta_1 x_j + \beta_2 y_k + \beta_3 x_j y_k, \qquad j, k = 1, \ldots, 4,$$

where x_j and y_k are the standardized dosages of neratinib and temsirolimus, respectively. The estimated toxicity rates of the combined doses are given by

		Temsirolimus			
		15	25	50	75
Neratinib	240	0.24	0.33	0.56	0.77
	200	0.14	0.19	0.33	0.55
	160	0.08	0.10	0.17	0.22
	120	0.04	0.05	0.07	0.10

which has two MTDs located in the higher dose region of the two-agent space. The observed DLTs can be generated from these estimated toxicity probabilities. The total sample size is 60 with a cohort size of 2, which is consistent with the up-and-down design that was used by the trial. We take the order in Figure 1(a) to be the initial order for the bagging-50 CRM (with 50 bootstrap samples), and the remaining parameters are specified according to Section 4.

In the prephase, we started the trial by treating the first cohort of patients at the lowest dose combination (A_1, B_1), at which none of the two patients experienced the DLT. The subsequent dose assignment followed the diagonal escalation rule, that is, we simultaneously increased the dose levels of the two drugs until the first DLT was observed. During this escalation stage, dose combinations (A_2, B_2), (A_3, B_3) were visited once before one DLT was observed at (A_4, B_4). The trial then entered into the main phase with the bagging-50 CRM. By bootstrapping the prephase data 50 times, the posterior probability, $\Pr(p_{44}^{\text{Bagging}} > \phi \mid D)$, is 0.71, and thus the dose level for the next cohort should be de-escalated to (A_3, B_4), which has an estimated toxicity probability of 0.34. One of the two patients treated at (A_3, B_4) experienced the DLT, and then (A_2, B_4) was recommended for the next cohort. At each step of dose assignment with accrued trial information, the bagging CRM continually re-estimated the toxicity order and the toxicity probabilities of the dose combinations so that the new cohort can be assigned to the best combination thus far. After exhaustion of 30 cohorts, the bagging CRM stopped at the dose combination (A_3, B_3), where 8 out of 24 patients experienced DLTs. Given the collected data at the end of trial,

$$\{n_{ij}\} = \begin{bmatrix} 4 & 2 & 0 & 2 \\ 2 & 6 & 26 & 8 \\ 0 & 2 & 2 & 4 \\ 2 & 0 & 0 & 0 \end{bmatrix}, \qquad \{y_{ij}\} = \begin{bmatrix} 2 & 2 & 0 & 1 \\ 0 & 1 & 8 & 5 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

the dose combination (A_3, B_3) was recommended as the MTD, with an estimated toxicity probability of 0.31. Throughout the entire trial, 75 different toxicity orders are generated according to the bagging procedure. Since the toxicity profile of the 4×4 dose-combination space is presumably known, we can additionally conduct a simulation study to examine the potential gain using the proposed design. Based on 1000 replications, the up-and-down design utilized by the original trial only has a correct MTD selection percentage of 42.0%, and assigns about 11.6 patients to the two MTDs. By contrast, the bagging CRM has a 59.9% MTD selection percentage, and assigns 20.6 patients to the MTDs.

5.2. *Capecitabine and bosutinib combination trial*. As another application of the bagging-50 CRM, we redesign the phase I trial for advanced solid tumors [Isakoff et al. (2014)]. In the trial, the three doses of capecitabine were $\{A_1, A_2, A_3\} = \{625, 750, 1000\}$ mg/m^2 and those of bosutinib were $\{B_1, B_2,$

$B_3\} = \{200, 300, 400\}$ mg. Suppose the true toxicity probabilities of the combined doses are

		Bosutinib		
		200	300	400
Capecitabine	1000	0.20	0.33	0.55
	750	0.09	0.29	0.41
	625	0.05	0.20	0.33

which corresponds to scenario 2 in the supplemental simulation studies of Isakoff et al. (2014). To be consistent with the dose escalation stage of the original study, we set the total number of patients to be 24 with a cohort size of 2. As the total number of dose combinations is 9, in the specification of the CRM skeleton, we take the initial position of the MTD to be 5 and the half width of the indifference interval in the R function `getprior` $\eta = 0.05$ to allow for adequate spacing between toxicity probabilities of adjacent dose levels. The initial order is chosen in the same way as Figure 1(a):

$$\begin{aligned}\mathcal{O}_0 : (A_1, B_1) &\prec (A_1, B_2) \prec (A_2, B_1) \prec (A_1, B_3) \prec (A_2, B_2) \\ &\prec (A_3, B_1) \prec (A_2, B_3) \prec (A_3, B_2) \prec (A_3, B_3).\end{aligned}$$

All the other design parameters including the prior distributions are specified the same as those in Section 4. During the start-up phase of the trial, dose combinations of (A_1, B_1) and (A_2, B_2) were administrated along the diagonal escalation direction. Since one DLT occurred at the dose combination (A_2, B_2), the main phase with the bagging CRM was triggered immediately. Based on the cumulative data, we generated 50 bootstrap samples that resulted in three sets of unique orders:

$$\begin{aligned}\mathcal{O}_1 : (A_1, B_1) &\prec (A_1, B_2) \prec (A_2, B_1) \prec (A_1, B_3) \prec (A_2, B_2) \\ &\prec (A_3, B_1) \prec (A_2, B_3) \prec (A_3, B_2) \prec (A_3, B_3), \\ \mathcal{O}_2 : (A_1, B_1) &\prec (A_1, B_2) \prec (A_2, B_1) \prec (A_2, B_2) \prec (A_1, B_3) \\ &\prec (A_3, B_1) \prec (A_2, B_3) \prec (A_3, B_2) \prec (A_3, B_3), \\ \mathcal{O}_3 : (A_1, B_1) &\prec (A_1, B_2) \prec (A_2, B_1) \prec (A_1, B_3) \prec (A_3, B_1) \\ &\prec (A_2, B_2) \prec (A_2, B_3) \prec (A_3, B_2) \prec (A_3, B_3).\end{aligned}$$

We combined the estimated toxicity probabilities from the three orders using the BMA procedure, and found the next dose combination to be (A_1, B_3). The subsequent dose assignments are given in Table 4. Once the outcomes of all the 24 patients were observed and analyzed, the dose combination (A_1, B_3), that is,

TABLE 4
Application of the bagging-50 CRM to the bosutinib and capecitabine combination trial

Sequence of cohorts	Dose combination	# of DLTs	# of orders	Sequence of cohorts	Dose combination	# of DLTs	# of orders
1	(A_1, B_1)	0	1	7	(A_2, B_2)	1	4
2	(A_2, B_2)	1	3	8	(A_3, B_1)	1	8
3	(A_1, B_3)	1	8	9	(A_2, B_1)	0	18
4	(A_1, B_2)	0	8	10	(A_2, B_2)	1	15
5	(A_1, B_3)	0	5	11	(A_1, B_3)	2	20
6	(A_1, B_3)	0	4	12	(A_1, B_2)	0	17

capecitabine 625 mg/m^2 plus bosutinib 400 mg, was recommended as the MTD combination, with an estimated toxicity probability of 0.36.

6. Concluding remarks. To embrace the trend of drug-combination trials, we have proposed to dynamically estimate the toxicity order of two combined agents by two-dimensional isotonic regression. Based on the estimated toxicity order, we reduce the two-dimensional drug-combination searching space into a one-dimensional line and then apply the CRM to locate the MTD combination. Our design utilizes all the available data to continually reorder as well as to re-estimate the toxicity probabilities in the entire dose-combination space. To stabilize this dynamic procedure, we incorporate novel ensemble methods, bagging and Bayesian model averaging techniques, into the proposed method. Simulation studies show that our approach outperforms the competing methods and is rather robust to various prior ordering specifications.

The proposed method differs notably from the POCRM in many ways. First, our method requires only one initial order that satisfies the partial order constraint at the beginning of the trial, while the POCRM recommends to choose six to nine orders. If the initial orders in the POCRM cannot cover the true toxicity profile, the design may not perform well. In contrast, our method coherently updates the single toxicity "working" ordering with the cumulative data as more patients enter the trial, and thus it has a higher chance to estimate the order more accurately. Second, the POCRM utilizes the Bayesian model selection procedure to select only one order at each decision step, while the bagging CRM averages the estimates from all possible bagging-generated orders via Bayesian model averaging. Model selection based on a small sample may not be reliable, especially at the beginning of the trial when very few patients have been accrued. Third, our bagging CRM can be easily extended to clinical trial settings with more than two drugs in combination, as we only need to sort out one single ordering for all dose combinations. However, the POCRM may suffer from such a higher dimensional dose-finding space because the number of possible orders increases immensely

with three or more combined drugs. Many more initial orders are required to cover different scenarios, which may become infeasible. Last, the dose assignment rule of the bagging CRM is totally different from the POCRM: our method makes decisions based on the posterior distribution of the toxicity probability, while the POCRM solely relies upon the posterior mean estimate. We have focused on incorporation of bagging to the model-based dose-finding designs with combined drugs. It is of interest to investigate other ensemble methods in machine learning, such as the random forest [Breiman (2001)] and boosting [Freund and Schapire (1997), Friedman, Hastie and Tibshirani (2000)], in dose-finding trials.

Acknowledgments. We thank the Editor, Associate Editor and two referees for their many constructive and insightful comments that have led to significant improvements in the article.

REFERENCES

AHN, C. (1998). An evaluation of phase I cancer clinical trial designs. *Stat. Med.* **17** 1537–1549.

BENDELL, J. C., JONES, S. F., HART, L., SPIGEL, D. R., LANE, C. M., EARWOOD, C., INFANTE, J. R., BARTON, J. and BURRIS, H. A. (2015). A phase Ib study of linsitinib (OSI-906), a dual inhibitor of IGF-1R and IR tyrosine kinase, in combination with everolimus as treatment for patients with refractory metastatic colorectal cancer. *Invest. New Drugs* **33** 187–193.

BREIMAN, L. (1996). Bagging predictors. *Mach. Learn.* **24** 123–140.

BREIMAN, L. (2001). Random forests. *Mach. Learn.* **45** 5–32.

BRIL, G., DYKSTRA, R., PILLERS, C. and ROBERTSON, T. (1984). Algorithm AS 206: Isotonic regression in two independent variables. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **33** 352–357.

CHEUNG, Y. K. (2011). *Dose Finding by the Continual Reassessment Method*. Chapman & Hall/CRC, Boca Raton, FL.

CLYDE, M. A. and LEE, H. K. H. (2001). Bagging and the Bayesian bootstrap. In *Artificial Intelligence and Statistics* (T. Richardson and T. Jaakkola, eds.) 169–174. Elsevier, New York.

CONAWAY, M. R., DUNBAR, S. and PEDDADA, S. D. (2004). Designs for single- or multiple-agent phase I trials. *Biometrics* **60** 661–669. MR2089441

FAN, S. K., VENOOK, A. P. and LU, Y. (2009). Design issues in dose-finding phase I trials for combinations of two agents. *J. Biopharm. Statist.* **19** 509–523. MR2668723

FREUND, Y. and SCHAPIRE, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. System Sci.* **55** 119–139. MR1473055

FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2000). Additive logistic regression: A statistical view of boosting. *Ann. Statist.* **28** 337–407. With discussion and a rejoinder by the authors. MR1790002

GANDHI, L., BAHLEDA, R., TOLANEY, S. M., KWAK, E. L., CLEARY, J. M., PANDYA, S. S., HOLLEBECQUE, A., ABBAS, R., ANANTHAKRISHNAN, R., BERKENBLIT, A., KRYGOWSKI, M., LIANG, Y., TURNBULL, K. W., SHAPIRO, G. I. and SORIA, J.-C. (2014). Phase I study of neratinib in combination with temsirolimus in patients with human epidermal growth factor receptor 2-dependent and other solid tumors. *J. Clin. Oncol.* **32** 68–75.

HARRINGTON, J. A., WHEELER, G. M., SWEETING, M. J., MANDER, A. P. and JODRELL, D. I. (2013). Adaptive designs for dual-agent phase I dose-escalation studies. *Nat. Rev. Clin. Oncol.* **10** 277–288.

HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2009). *The Elements of Statistical Learning*: *Data Mining*, *Inference*, *and Prediction*, 2nd ed. *Springer Series in Statistics*. Springer, New York. MR2722294

HEYD, J. M. and CARLIN, P. B. (1999). Adaptive design improvements in the continual reassessment method for phase I studies. *Stat. Med.* **18** 1307–1321.

HIRAKAWA, A., HAMADA, C. and MATSUI, S. (2013). A dose-finding approach based on shrunken predictive probability for combinations of two agents in phase I trials. *Stat. Med.* **32** 4515–4525. MR3118372

HIRAKAWA, A., WAGES, N. A., SATO, H. and MATSUI, S. (2015). A comparative study of adaptive dose-finding designs for phase I oncology trials of combination therapies. *Stat. Med.* **34** 3194–3213. MR3412626

HOETING, J. A., MADIGAN, D., RAFTERY, A. E. and VOLINSKY, C. T. (1999). Bayesian model averaging: A tutorial. *Statist. Sci.* **14** 382–417. With comments by M. Clyde, David Draper and E. I. George, and a rejoinder by the authors. MR1765176

HOUEDE, N., THALL, P. F., NGUYEN, H., PAOLETTI, X. and KRAMAR, A. (2010). Utility-based optimization of combination therapy using ordinal toxicity and efficacy in phase I/II trials. *Biometrics* **66** 532–540. MR2758833

HUANG, X., BISWAS, S., OKI, Y., ISSA, J.-P. and BERRY, D. A. (2007). A parallel phase I/II clinical trial design for combination therapies. *Biometrics* **63** 429–436. MR2370801

IASONOS, A. and O'QUIGLEY, J. (2014). Adaptive dose-finding studies: A review of model-guided phase I clinical trials. *J. Clin. Oncol.* **32** 2505–2511.

ISAKOFF, S. J., WANG, D., CAMPONE, M., CALLES, A., LEIP, E., TURNBULL, K., BARDY-BOUXIN, N., DUVILLIÉ, L. and CALVO, E. (2014). Bosutinib plus capecitabine for selected advanced solid tumours: Results of a phase 1 dose-escalation study. *Br. J. Cancer* **111** 2058–2066.

IVANOVA, A. and WANG, K. (2004). A non-parametric approach to the design and analysis of two-dimensional dose-finding trials. *Stat. Med.* **23** 1861–1870.

KORN, E. L. and SIMON, R. (1991). Selecting dose-intense drug combinations: Metastatic breast cancer. *Breast Cancer Res. Treat.* **20** 155–166.

KORN, E. L. and SIMON, R. (1993). Using the tolerable-dose diagram in the design of phase I combination chemotherapy trials. *J. Clin. Oncol.* **11** 794–801.

KORN, E. L., MIDTHUNE, D., CHEN, T. T., RUBINSTEIN, L. V., CHRISTIAN, M. C. and SIMON, R. M. (1994). A comparison of two phase I trial designs. *Stat. Med.* **13** 1799–1806.

KRAMAR, A., LEBECQ, A. and CANDALH, E. (1999). Continual reassessment methods in phase I trials of the combination of two drugs in oncology. *Stat. Med.* **18** 1849–1864.

LIN, R. and YIN, G. (2016). Bayesian optimal interval design for dose finding in drug-combination trials. *Stat. Methods Med. Res.*, DOI:10.1177/0962280215594494.

MANDER, A. P. and SWEETING, M. J. (2015). A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Stat. Med.* **34** 1261–1276. MR3322767

MANDREKAR, S. J. (2014). Dose-finding trial designs for combination therapies in oncology. *J. Clin. Oncol.* **32** 65–67.

O'QUIGLEY, J. and CONAWAY, M. (2010). Continual reassessment and related dose-finding designs. *Statist. Sci.* **25** 202–216. MR2789990

O'QUIGLEY, J., PEPE, M. and FISHER, L. (1990). Continual reassessment method: A practical design for phase 1 clinical trials in cancer. *Biometrics* **46** 33–48. MR1059105

ORON, A. P. and HOFF, P. D. (2013). Small-sample behavior of novel phase I cancer trial designs. *Clin. Trials* **10** 63–80.

PAPADATOS-PASTOS, D., LUKEN, M. D. M. and YAP, T. A. (2015). Combining targeted therapeutics in the era of precision medicine. *Br. J. Cancer* **112** 1–3.

RAFTERY, A. E., MADIGAN, D. and HOETING, J. A. (1997). Bayesian model averaging for linear regression models. *J. Amer. Statist. Assoc.* **92** 179–191. MR1436107

RIVIERE, M.-K., DUBOIS, F. and ZOHAR, S. (2015a). Competing designs for drug combination in phase I dose-finding clinical trials. *Stat. Med.* **34** 1–12. MR3286233

RIVIERE, M.-K., YUAN, Y., DUBOIS, F. and ZOHAR, S. (2014). A Bayesian dose-finding design for drug combination clinical trials based on the logistic model. *Pharm. Stat.* **13** 247–257.

RIVIERE, M. K., LE TOURNEAU, C., PAOLETTI, X., DUBOIS, F. and ZOHAR, S. (2015b). Designs of drug-combination phase I trials in oncology: A systematic review of the literature. *Ann. Oncol.* **26** 669–674.

RUBIN, D. B. (1981). The Bayesian bootstrap. *Ann. Statist.* **9** 130–134. MR0600538

SAURA, C., GARCIA-SAENZ, J. A., XU, B., HARB, W., MOROOSE, R., PLUARD, T., CORTES, J., KIGER, C., GERMA, C., WANG, K., MARTIN, M., BASELGA, J. and KIM, S. B. (2014). Safety and efficacy of neratinib in combination with capecitabine in patients with metastatic human epidermal growth factor receptor 2-positive breast cancer. *J. Clin. Oncol.* **32** 3626–3633.

SHEN, L. Z. and O'QUIGLEY, J. (1996). Consistency of continual reassessment method under model misspecification. *Biometrika* **83** 395–405. MR1439791

SIEGEL, D. S., RICHARDSON, P., DIMOPOULOS, M., MOREAU, P., MITSIADES, C., WEBER, D., HOUP, J., GAUSE, C., VUOCOLO, S., EID, J., GRAEF, T. and ANDERSON, K. C. (2014). Vorinostat in combination with lenalidomide and dexamethasone in patients with relapsed or refractory multiple myeloma. *Blood Cancer J.* **4** e182.

THALL, P. F. (2010). Bayesian models and decision algorithms for complex early phase clinical trials. *Statist. Sci.* **25** 227–244. MR2789992

THALL, P. F., MILLIKAN, R. E., MUELLER, P. and LEE, S.-J. (2003). Dose-finding with two agents in phase I oncology trials. *Biometrics* **59** 487–496. MR2004253

TIGHIOUART, M., PIANTADOSI, S. and 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. MR3260662

ULLENHAG, G. J., ROSSMANN, E. and LILJEFORS, M. (2015). A phase I dose-escalation study of lenalidomide in combination with gemcitabine in patients with advanced pancreatic cancer. *PLoS ONE* **10** e0121197.

WAGES, N. A., CONAWAY, M. R. and O'QUIGLEY, J. (2011). Continual reassessment method for partial ordering. *Biometrics* **67** 1555–1563. MR2872406

WANG, K. and IVANOVA, A. (2005). Two-dimensional dose finding in discrete dose space. *Biometrics* **61** 217–222. MR2135863

WILKY, B. A., RUDEK, M. A., AHMED, S., LAHERU, D. A., COSGROVE, D., DONEHOWER, R. C., NELKIN, B., BALL, D., DOYLE, L. A., CHEN, H., YE, X., BIGLEY, G., WOMACK, C. and AZAD, N. S. (2015). A phase I trial of vertical inhibition of IGF signalling using cixutumumab, an anti-IGF-1R antibody, and selumetinib, an MEK 1/2 inhibitor, in advanced solid tumours. *Br. J. Cancer* **112** 24–31.

YIN, G. (2012). *Clinical Trial Design*: *Bayesian and Frequentist Adaptive Methods*. John Wiley & Sons, Hoboken, New Jersey.

YIN, G. and LIN, R. (2015). Comments on 'Competing designs for drug combination in phase I dose-finding clinical trials' by M.-K. Riviere, F. Dubois, and S. Zohar [MR3286233]. *Stat. Med.* **34** 13–17. MR3286234

YIN, G. and YUAN, Y. (2009a). Bayesian dose finding in oncology for drug combinations by copula regression. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **58** 211–224. MR2649671

YIN, G. and YUAN, Y. (2009b). Bayesian model averaging continual reassessment method in phase I clinical trials. *J. Amer. Statist. Assoc.* **104** 954–968. MR2750228

YUAN, Z. and CHAPPELL, R. (2004). Isotonic designs for phase I cancer clinical trials with multiple risk groups. *Clin. Trials* **1** 499–508.

YUAN, Y. and YIN, G. (2008). Sequential continual reassessment method for two-dimensional dose finding. *Stat. Med.* **27** 5664–5678. MR2573775

DEPARTMENT OF STATISTICS AND
ACTUARIAL SCIENCE
THE UNIVERSITY OF HONG KONG
POKFULAM ROAD
HONG KONG
CHINA
E-MAIL: ruitaolin@gmail.com
gyin@hku.hk