Biostatistics (2017) **18**, 1, *pp.* 180–194
doi:10.1093/biostatistics/kxw038
Advance Access publication on August 22, 2016

Nonparametric overdose control with late-onset toxicity in phase I clinical trials

RUITAO LIN, GUOSHENG YIN*

Department of Statistics and Actuarial Science, The University of Hong Kong, Pokfulam Road, Hong Kong
gyin@hku.hk

Summary

Under the framework of Bayesian model selection, we propose a nonparametric overdose control (NOC) design for dose finding in phase I clinical trials. Each dose assignment is guided via a feasibility bound, which thereby can control the number of patients allocated to excessively toxic dose levels. Several aspects of the NOC design are explored, including the coherence property in dose assignment, calibration of design parameters, and selection of the maximum tolerated dose (MTD). We further propose a fractional NOC (fNOC) design in conjunction with a so-called fractional imputation approach, to account for late-onset toxicity outcomes. Extensive simulation studies have been conducted to show that both the NOC and fNOC designs have robust and satisfactory finite-sample performance compared with the existing dose-finding designs. The proposed methods also possess several desirable properties: treating patients more safely and also neutralizing the aggressive escalation to overly toxic doses when the toxicity outcomes are late-onset. The fNOC design is exemplified with a real cancer phase I trial.

Keywords: Bayesian posterior probability; Dose finding; Late-onset toxicity; Maximum tolerated dose; Overdose control.

1. Introduction

The American Society of Clinical Oncology (ASCO) recently published an update of the ASCO policy statement to address the critical roles of phase I clinical trials in cancer research and treatment (Weber *and others*, 2015). The update calls for innovative adaptive phase I trial designs, which may pin down the correct dose quickly, decrease patient risk, expose more patients to therapeutic dose levels, and so on, to cope with the changing landscape in cancer research. The primary objective of phase I oncology trials is to identify the maximum tolerated dose (MTD) whose induced dose-limiting toxicity (DLT) probability is closest to the target toxicity rate. Numerous phase I dose-finding designs have been proposed in the literature, which can be generally classified as algorithm- and model-based methods (Yin, 2012). The algorithm-based methods, which do not impose any parametric assumption on the dose–toxicity curve, are often considered as nonparametric or curve-free (Sverdlov *and others*, 2014), for example, the $3 + 3$ design (Storer, 1989), the accelerated titration design (Simon *and others*, 1997), the biased coin design (Durham *and others*, 1997), the group up-and-down design (Gezmu and Flournoy, 2006), and the class

*To whom correspondence should be addressed.

of interval designs (Ivanova *and others*, 2007; Liu and Yuan, 2015). By contrast, model-based methods explicitly assume some parametric dose–toxicity relationship, such as the continual reassessment method (CRM) by O'Quigley *and others* (1990) and the escalation with overdose control (EWOC) method by Babb *and others* (1998), which can fully incorporate the observed data across all dose levels and thus are usually more efficient than the algorithmic counterparts. The CRM tends to assign the dose based on the mean or the mode of the MTD posterior distribution (Heyd and Carlin, 1999; Leung and Wang, 2002; Yuan *and others*, 2007), while EWOC selects the αth quantile, with α specified as the feasibility bound for the overdosing control (Zacks *and others*, 1998; Tighiouart *and others*, 2005; Tighiouart and Rogatko, 2010; Chen *and others*, 2012). Model-based designs are in general fully parametric and thus vulnerable to model misspecifications.

In the aforementioned dose-finding methods, toxicity is assumed to be quickly observable after treatment relative to the speed of patient arrival, which, however, is not true for many types of DLTs. Postel-Vinay *and others* (2011) reported that 57% of the grade 3 and 4 toxicities among a total of 445 patients involved in 36 trials were late-onset. In a phase I concurrent radiation therapy trial (Desai *and others*, 2007), 44 patients with pancreatic cancer were enrolled. Patients were examined at least once weekly for a maximum of nine weeks, while the enrollment rate was approximately one patient every two weeks. Consequently, toxicity outcomes for some previously treated patients might not have been yet available upon the arrival of a new patient. Liu *and others* (2013) demonstrated that such missing data are nonignorable due to the fact that at each decision-making time the probability of nonresponse or missingness is larger for the patients who would not experience toxicity during the assessment period. If each new patient is required to wait for dose assignment until all the toxicity outcomes of the previously treated patients are available, it would cause a treatment delay for each newly arrived patient and result in a lengthy trial. On the other hand, if the dose assignment is made by simply discarding the missing data, it would lead to biased estimates and deteriorate the trial performance. Cheung and Chappell (2000) proposed the time-to-event CRM (TITE-CRM) to incorporate information of follow-up times into toxicity evaluation via a weighting scheme, and similarly Ivanova *and others* (2007) introduced the TITE cumulative cohort design (TITE-CCD) method. Bekele *and others* (2008) proposed to monitor late-onset toxicities using predicted risks. Yuan and Yin (2011) utilized the expectation–maximization algorithm to estimate the dose toxicity probabilities based on the incomplete data to direct dose assignment. Mauguen *and others* (2011) explored the performance of time-to-event EWOC method (TITE-EWOC). Yin *and others* (2013) modeled the underlying time-to-toxicity data through the well-known Kaplan–Meier estimator (Kaplan and Meier, 1958), and imputed the missing data by their fractional DLT information.

Our research is motivated by a phase I, open-label, first-in-human, dose-finding clinical trial, which involves an oral smoothened inhibitor sonidegib therapy. The primary objective of the phase I trial was to determine the MTD of sonidegib in adult patients with advanced solid tumors. A total of thirty patients in cohorts of size three were enrolled and treated by five dose levels of sonidegib. Adverse events with grade ≥ 3 in severity were considered as DLTs with the target DLT rate set at 33%. Table 1 summarizes the accumulated data for the first four cohorts. The enrollment rate was approximately three patients per month, but it generally took three months to fully assess DLT for each patient. Therefore, there would be some missing DLT data by the time of dose assignment for a newly enrolled patient, because some of the previous patients might not have completed the monitoring period for toxicity. The EWOC method was used under a two-parameter Bayesian logistic regression model with a feasibility bound $\alpha = 0.25$.

In the Bayesian model selection paradigm, we propose a nonparametric overdose control (NOC) method for dose finding. By continuously estimating the posterior model probabilities, the proposed method can control the posterior probability that each successive dose assignment exceeds the MTD without imposing any parametric assumption on the dose–toxicity relationship. In addition, we propose a fractional NOC (fNOC) design, which combines the NOC design with the so-called fractional imputation method (Yin *and*

Table 1. *Arrival time (in days), dose level, time to toxicity (in days), and imputed fractional toxicity outcome for the first 12 patients in the sonidegib trial*

Cohort No.	Patient No.	Day of arrival	Dose level	Time to DLT	Fractional DLT
1	1	4	1	0	0
	2	7	1	0	0
	3	19	1	0	0
2	4	29	2	0	0
	5	31	2	0	0
	6	50	2	0	0.000
3	7	58	3	1 (123)	1
	8	67	3	0	0.143
	9	78	3	0	0.143
4	10	91	3	1 (120)	1
	11	100	3	0	0.143
	12	118	3	0	0.221

Note: If a patient has experienced the DLT, then his/her "Time to DLT" outcome is 1 with the event time in parentheses, and 0 otherwise. The "Fractional DLT" is evaluated upon the arrival of the 13th patient.

others, 2013), to resolve the late-onset toxicity problem. Both the NOC and fNOC are fully nonparametric, and have desirable properties as well as competitive operating characteristics.

The remainder of this article is organized as follows. In Section 2, we formulate the NOC design and present its theoretical properties, and further we extend NOC to fNOC by incorporating late-onset toxicities. In Section 3, we examine the performance of the new designs based on simulation studies and make extensive comparisons with existing methods. As an illustration, we apply fNOC to a phase I dose-finding trial with sonidegib therapy in Section 4, and Section 5 concludes with some remarks. The supplementary material available at *Biostatistics* online contains details on computation, simulations, and the proof of Theorem 2.1.

2. METHODOLOGY

2.1. *Nonparametric overdose control*

Let p_j denote the toxicity probability of dose level $j, j = 1, \ldots, J$, and ϕ denote the target toxicity rate. The toxicity probability is assumed to increase monotonically with the dose level, i.e., $0 < p_1 < \cdots < p_J < 1$. Let $d_i \in \{1, \ldots, J\}$ represent the dose level at which the ith patient is treated, and the toxicity outcome $y_i = 1$ if the patient has experienced DLT, and $y_i = 0$ otherwise. We cast dose finding in a Bayesian model selection problem. Specifically, we consider J models,

$$M_k : \begin{cases} |p_j - \phi| \le \epsilon, & j = k, \\ |p_j - \phi| > \epsilon, & j \ne k, \end{cases} \quad k = 1, \ldots, J,$$

where $\epsilon \ge 0$ is a prespecified small positive number, indicating that dose level k is the MTD when M_k is the true model. Under each M_k, we specify a joint uniform prior distribution for $p_1, \ldots, p_J$ to maintain

the monotone increasing relationship of the toxicity probabilities,

$$\begin{cases} p_j|M_k \sim \text{Unif}(\phi-\epsilon, \phi+\epsilon), & j=k, \\ p_j|p_{j-1}, M_k \sim \text{Unif}(\max(\phi+\epsilon, p_{j-1}), p_{J+1}), & j>k, \\ p_j|p_{j+1}, M_k \sim \text{Unif}(p_0, \min(\phi-\epsilon, p_{j+1})), & j<k, \end{cases} \tag{2.1}$$

where p_0 and p_{J+1} ($0 < p_0 < p_{J+1} < 1$) are the lower and upper boundaries of the prior distribution, respectively. Typically, we take $\epsilon = 0.05$, while $\epsilon = 0$ indicates a point mass at ϕ, and $p_0 = 0$ and $p_{J+1} = 0.8$ would yield desired operating characteristics for most cases. Let $D_n = \{(d_1, y_1), \ldots, (d_n, y_n)\}$ denote the observed data prior to the arrival of the $(n+1)$th patient, then $x_j = \sum_{i=1}^{n} y_i I\{d_i = j\}$ and $m_j = \sum_{i=1}^{n} I\{d_i = j\}$ represent the number of responses and number of patients at dose level j, respectively. The marginal likelihood under model M_k is

$$P(D_n \mid M_k) \propto \int f(p_1, \ldots, p_J \mid M_k) \prod_{j=1}^{J} \left\{ p_j^{x_j} (1-p_j)^{m_j - x_j} \right\} \mathrm{d}p_1 \cdots \mathrm{d}p_J,$$

where $f(p_1, \ldots, p_J \mid M_k)$ is the joint prior distribution of $p_1, \ldots, p_J$ under M_k given in (2.1). The posterior probability that model M_k is true is given by

$$P(M_k \mid D_n) = \frac{P(D_n \mid M_k) P(M_k)}{\sum_{j=1}^{J} P(D_n \mid M_j) P(M_j)},$$

where $P(M_k)$ is the prior probability of M_k. We specify a discrete uniform distribution for the prior model probability; that is, $P(M_k) = 1/J$, $k = 1, \ldots, J$.

Based on the posterior probabilities, the next dose assignment $d_{n+1} = j^*$ is determined according to a prespecified feasibility bound α ($0 < \alpha < 1$),

$$j^* = \arg\min_j \Big| \sum_{k=1}^{j} P(M_k \mid D_n) - \alpha \Big|,$$

where $\sum_{k=1}^{j} P(M_k \mid D_n)$ is the posterior probability that the MTD is not above dose level j. If we denote γ as the dose level of the MTD, then the cumulative probability mass function in the discrete probability space is $P(\gamma \le j \mid D_n) = \sum_{k=1}^{j} P(M_k \mid D_n)$. Therefore, the optimal dose level j^* can be equivalently written as

$$j^* = \arg\min_j |P(\gamma \le j \mid D_n) - \alpha|, \tag{2.2}$$

which characterizes the overdose control rule. Such an allocation scheme selects approximately the dose level that minimizes the risk with respect to an asymmetric loss function,

$$\mathcal{L}(j, \gamma) = \begin{cases} \alpha(\gamma - j), & \text{for } j \le \gamma, \ (j \text{ is an underdose level}), \\ (1-\alpha)(j - \gamma), & \text{for } j > \gamma, \ (j \text{ is an overdose level}). \end{cases}$$

If we choose the feasibility bound $\alpha < 0.5$, it corresponds to placing a lower penalty on underdosing than on overdosing, so that the posterior probability of overdosing is small.

The NOC design imposes no parametric assumption on the dose–toxicity curve, which greatly enhances its robustness in comparison to EWOC. In addition, NOC is coherent in the sense that it does not de-escalate the dose if the most recently treated patient has not experienced the DLT, and it does not escalate the dose if a toxicity outcome has just been observed. Both the CRM and EWOC designs enjoy the coherence property (Cheung, 2005; Tighiouart and Rogatko, 2010). As shown in the next theorem, our NOC design also possesses such coherence principle.

THEOREM 2.1 Suppose that the dose level for the nth patient is $d_n = j$. The overdose control rule (2.2) is coherent in escalation and de-escalation, i.e., for the $(n+1)$th patient, $P(d_{n+1} = j-1 \mid y_n = 0, d_n = j, D_{n-1}) = 0$ and $P(d_{n+1} = j+1 \mid y_n = 1, d_n = j, D_{n-1}) = 0$.

The proof of Theorem 2.1 is given in Section C of the supplementary material available at *Biostatistics* online. Despite the coherence property, if the dose sequence follows this overdose control rule throughout the trial, there is a large probability that patients would be treated somewhere below the MTD with the accumulated information. For example, we consider six dose levels, and suppose that the posterior model probabilities are $(P(M_1 \mid D_n), \ldots, P(M_6 \mid D_n)) = (0.01, 0.02, 0.10, 0.70, 0.10, 0.07)$. If we set the feasibility bound $\alpha = 0.25$, then the next dose level is $j^* = 3$. Obviously, such dose assignment is suboptimal because the accumulated data indicate that dose level $j = 4$ is the best due to its largest posterior probability. To ensure that patients can receive the optimal dose level (i.e., the MTD), we impose a dose-switching rule,

$$j^* = \{j : P(M_j \mid D_n) > \eta\}, \tag{2.3}$$

where $0.5 \le \eta \le 1$ is a switching cutoff of the posterior model probability. If there is sufficient evidence to guarantee that dose level j is the MTD by (2.3), we assign the cohort to that dose; otherwise, we treat the next cohort at the dose level according to the overdose control rule in (2.2).

Each model M_k inherently assumes that only dose level k lies inside the ϵ-neighbourhood of ϕ, $k = 1, \ldots, J$. Nevertheless, there are situations where none of the dose levels has $p_j \in (\phi - \epsilon, \phi + \epsilon)$ or multiple dose levels are inside the interval. The proposed framework can still facilitate to identify the most six plausible dose level, as a model selection procedure is implemented to select the dose level that has a large posterior probability of $p_j \in (\phi - \epsilon, \phi + \epsilon)$. Furthermore, the proposed method can be generalized to allow multiple dose levels to satisfy $p_j \in (\phi - \epsilon, \phi + \epsilon)$, as shown in Section 3.2.

2.2. *Dose-finding algorithm*

In addition to the overdose control and dose-switching rules, we impose an additional safety rule to exclude dose level j and all the higher dose levels from the trial if

$$P(p_j > \phi \mid D_n) = \sum_{k=1}^{J} P(p_j > \phi \mid D_n, M_k) P(M_k \mid D_n) \ge \lambda,$$

where λ is a prespecified dose elimination cutoff probability, e.g., $\lambda = 0.85$. To simplify the practical implementation, those excluded dose levels are not reconsidered again for later patients. Although this safety rule may mistakenly eliminate some appropriate dose levels, we find by simulations the risk of erroneous elimination is negligible.

In summary, the proposed NOC design proceeds as follows:

(1) Treat the first cohort of patients at the lowest or the physician-specified dose level.

(2) Suppose the current dose level is j and the trial has treated a total of n patients, we determine the optimal dose level j^* according to (2.2) and (2.3),

(i) if there is a dose level that satisfies (2.3), then $j^* = \{j : P(M_j \mid D_n) > \eta\}$;

(ii) otherwise, the dose level j^* is selected using the overdose control rule (2.2) with a feasibility bound α, i.e., $j^* = \arg\min_j |P(\gamma \leq j \mid D_n) - \alpha|$.

Based on the optimal dose level j^*,

(i) if $j > j^*$, de-escalate to dose level $j - 1$;

(ii) if $j < j^*$ and dose level $j + 1$ has not been excluded, escalate to dose level $j + 1$;

(iii) otherwise, the dose stays at the same level j.

(3) The trial can be either stopped after exhaustion of the maximum sample size N, or be terminated early for safety if the lowest dose level is too toxic as noted by $P(p_1 > \phi \mid D_n) \geq \lambda$.

At the end of the trial, the dose level j^* with the maximum value of the posterior model probability is recommended as the MTD; that is,

$$j^* = \arg\max_{j \leq J^\dagger} P(M_j \mid D_N),$$

where $J^\dagger$ is the maximum dose level that has not been excluded. Section A of the supplementary material available at *Biostatistics* online provides the computational detail of the NOC design.

Most of the algorithm-based (nonparametric) phase I dose-finding designs select the MTD based on the isotonically transformed values of the estimated toxicity probabilities. However, there is no universal guideline on how to specify the weight utilized in the isotonic regression, and also it is often difficult to select the MTD from the tied estimates. By contrast, the proposed MTD selection procedure is automatic and straightforward from the dose escalation procedure, and thus avoids the aforementioned problems.

The NOC design has three tuning parameters: the feasibility bound α, the switching cutoff η, and the elimination cutoff λ, which make the overdose control framework very flexible. Section 3.2 examines the interplay between α and η in terms of the design performance. With a smaller value of α, more patients would be allocated to underdoses. The switching cutoff η controls the selection scheme of the optimal dose level j^*. The smaller the value of η, the more likely the patients are treated at the MTD, but a smaller value of η tends to lead to more aggressive dose escalation (see Figure S1 in the supplementary material available at *Biostatistics* online). As the default, we set $\alpha = 0.35$, $\eta = 0.5$ and $\lambda = 0.85$. Moreover, one can also use varying values of α to take into consideration the accumulated sample size. For example, a smaller value of α at the earlier stage of a trial would protect more patients from overdosing due to a lack of dose–response information, while a larger value of α at the later stage tends to keep more patients treated at the MTD.

2.3. *Late-onset toxicity*

Operatively, the NOC design requires that the DLT outcome is ascertainable quickly after treatment. To allow for late-onset toxicity, we consider imputing the missing toxicity outcomes based on the Kaplan–Meier estimator. Suppose that patients enter a trial sequentially with an interarrival time T, and each patient is followed for a prespecified evaluation period τ to assess the drug's toxicity. During $[0, \tau]$, a

binary toxicity outcome is measured for each patient i,

$$y_i = \begin{cases} 1, & \text{if DLT has occurred in } [0, \tau], \\ 0, & \text{if no DLT observed in } [0, \tau]. \end{cases}$$

The larger the τ/T ratio, the more missing data would be induced due to the fact that more patients might have not responded yet upon a new patient's arrival. Let t_i denote the time to toxicity for subject i, and let $u_i \le \tau$ denote the actual follow-up time. If $u_i \ge t_i$, then the toxicity outcome of the ith patient has been fully observed, while the patient response is censored if $u_i < t_i$. For patients who do not experience DLT within $[0, \tau]$, we have $u_i = \tau < t_i$. If the toxicity outcome y_i is censored by the decision-making time u_i, we can calculate the conditional probability of experiencing DLT in (u_i, τ) given that the DLT has not occurred by time u_i,

$$P(t_i < \tau \mid t_i > u_i) = \frac{P(u_i < t_i < \tau)}{P(t_i > u_i)},$$

which can be viewed as the fractional contribution of patient i to the overall toxicity data. Based on the Kaplan–Meier estimator of the survival function, $\hat{S}(\cdot)$, the fractional contribution for a censored toxicity outcome can be estimated by

$$\hat{y}_i = \frac{\hat{S}(u_i) - \hat{S}(\tau)}{\hat{S}(u_i)}, \tag{2.4}$$

which is a fractional number between 0 (no toxicity) and 1 (toxicity). The Bayesian survival function estimator (Susarla and Van Ryzin, 1976) can also be implemented in (2.4), which, however, requires a complicated prior specification and the computation is much more intensive. For the imputation purpose, the conventional Kaplan–Meier estimator is much easier to use, while the performance is similar.

As a result, the fNOC design takes effect by imputing the missing toxicity outcome of the censored patient by his/her fractional DLT $\hat{y}_i$. For safety, a start-up phase, which completely follows the patients until the first toxicity outcome is observed, is needed for fNOC at the beginning of the trial to prevent aggressive escalation to overly toxic dose levels. In regard to the specification of the tuning parameters, we recommend a conservative configuration by setting $\eta = 0.6$, so that the selection scheme (2.2) plays a dominant role at the early stage of the trial when the data are sparse. Usually, the imputation of the fractional toxicity outcome accelerates the trial because the newly arrived patients do not need to wait for the toxicity outcomes of previous patients. On the other hand, more patients tend to be allocated to overdoses due to the incomplete data. The combination of overdose control and the fractional imputation would induce a neutralization effect, because the overdose control rule can bring more patients back to underdoses and thus lead to superior performance with late-onset toxicities.

3. SIMULATION STUDY

3.1. *Comparison with existing methods*

To assess the performance of the proposed NOC design, we make comparisons of its operating characteristics with four existing dose-finding designs including two model-based and two algorithm-based methods: the CRM, the EWOC, the cumulative cohort design (CCD), and the Bayesian optimal interval design (BOIN). We consider a total of $J = 6$ dose levels with the target toxicity probability $\phi = 0.3$. We

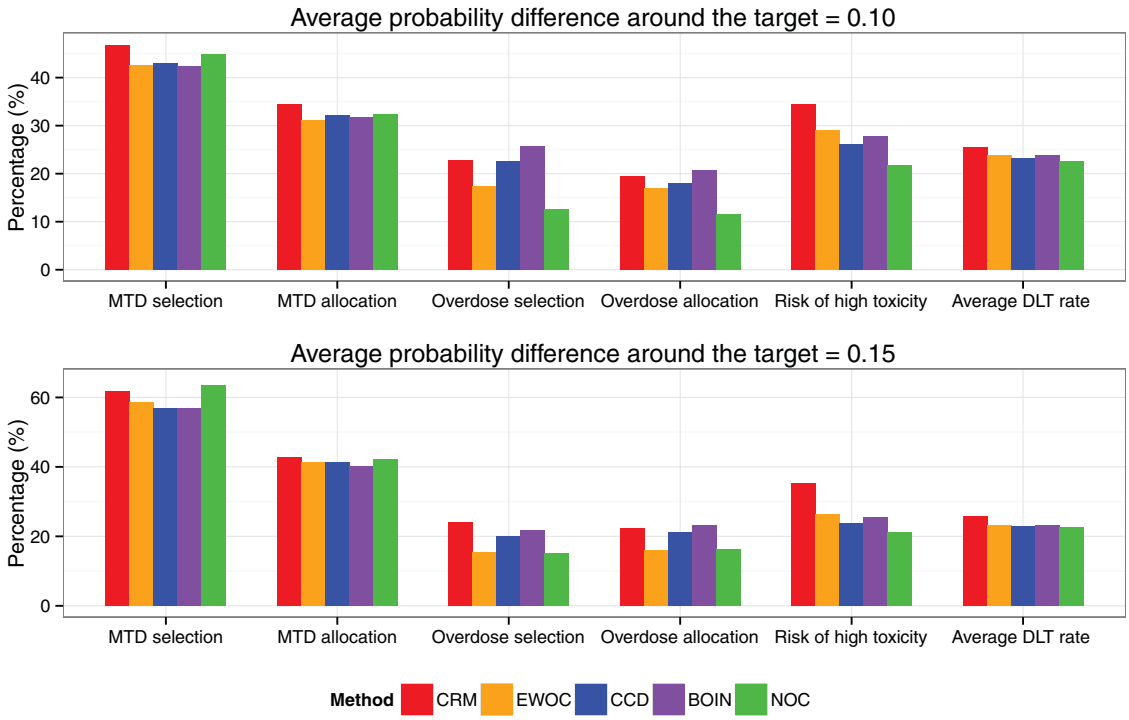

Fig. 1. Simulation results based on 10 000 randomly generated dose–toxicity scenarios with an average probability difference around the target of 0.10 and 0.15, respectively. The target toxicity probability is $\phi = 0.3$.

take the maximum sample size to be 36 and the cohort size to be three patients. More detailed configurations for this comparison study are provided in the Section B.1 of the supplementary material available at *Biostatistics* online.

To avoid subjectivity in choosing specific toxicity scenarios, we randomly generate 10 000 dose–toxicity scenarios following the approach of Paoletti *and others* (2004). In particular, the average probability difference around the target is controlled to be 0.1 or 0.15, and the MTD is set at one of the six dose levels with an equal probability. To adequately quantify the operating characteristics of the dose-finding methods, we consider six performance statistics: the percentage of MTD selection and the percentage of patients treated at the MTD, for which the larger the better, reflecting the accuracy and efficiency of a design. The remaining four statistics, including the percentage of trials that select overdoses as the MTD, the percentage of patients allocated to overdoses, the risk of high toxicity (which is defined as the percentage of trials that leads to the DLT rate greater than ϕ), and the percentage of patients experiencing DLT, are related to the safety aspects of a trial, and therefore a design with smaller values of these four safety statistics should be considered more ethical and desirable.

Figure 1 shows that the CRM performs the best and NOC the second in terms of the percentages of MTD selection and patients treated at the MTD. Regarding to the safety aspects, the proposed NOC design substantially outperforms all the other competing designs. For example, under the NOC design, the percentage of patients allocated to overly toxic dose levels is approximately half of that based on BOIN when the average probability difference around the target is 0.1 (top panel). Although the CRM has the best accuracy in pinning down the MTD, on average it has a greater risk of high toxicity. Among the 10 000 replications, NOC has an approximately four percentage of chance to exclude the true MTD from

the trial based on the dose elimination rule. The differences of the six performance statistics among the five methods diminish to a certain extent when the average probability difference around the target is 0.15 (bottom panel). In this case, not only is the NOC design the safest due to its overdose control scheme, but it also has the highest MTD selection percentage.

For further investigation, we specify 12 scenarios which involve various numbers and locations of the MTD and different probability differences around the target. Specifically, only one MTD has $p_j \in (\phi-\epsilon, \phi+\epsilon)$ in the first eight scenarios, while the remaining four scenarios contradict the prespecified J models; that is, none of the dose levels or multiple dose levels satisfy $p_j \in (\phi-\epsilon, \phi+\epsilon)$. For each scenario, we replicate 5000 trials and summarize the simulation results in Table S1 of the supplementary material available at *Biostatistics* online. The conclusion remains the same that NOC ranks the best in terms of safety, while as to the MTD selection and patient allocation percentages, NOC tends to perform better when the MTD is near the lower dose levels. If one of the last two dose levels is the MTD, the NOC design is relatively conservative. When the toxicity probability of the dose right below the MTD is close to the target, both the EWOC and NOC designs would allocate relatively more patients to that dose level, as shown in scenarios 6 and 7. When the true dose–toxicity profiles contradict the model assumptions, NOC can still produce desirable operating characteristics. Typically, if none of the dose levels is inside the ϵ-neighbourhood of ϕ, as noted by scenarios 9–10, NOC tends to select the highest dose level that satisfies $p_j < \phi-\epsilon$; if multiple dose levels fall in the ϵ-neighbourhood, NOC is more likely to select the lowest one from the admissible dose levels, which is exhibited in scenarios 11–12.

3.2. *Sensitivity analysis*

We conduct a sensitivity analysis to examine the impact of the feasibility bound α and the switching cutoff η on the performance of NOC. Specifically, we consider three configurations of α and η: (i) $\alpha = 0.35$ and $\eta = 0.8$, (ii) $\alpha = 0.35$ and $\eta = 0.5$, and (iii) $\alpha = 0.5$ and $\eta = 0.5$, and summarize the simulation results in Figure S1 of the supplementary material available at *Biostatistics* online. There is a general trend that the percentage of overdosing decreases with a decreasing value of α and an increasing value of η. A larger value of η would prevent (2.3) from being the dose selection scheme, so that the overdose control rule (2.2) would dominate the decision making. On average, the activation (switching) rate of rule (2.3) when $\eta = 0.5$ is 34.3%, while that under $\eta = 0.8$ reduces to 4.5%, indicating that most of the dose assignment follows the overdose control rule (2.2). Typically, when the MTD is the lowest or the highest dose, the switching rate tends to be larger. On the other hand, the larger the value of α, the less penalty is assigned to overdosing, which leads to a relatively more aggressive dose escalation. Such a sensitivity analysis provides a practical guidance to choosing the tuning parameters of α and η.

To investigate the interplay between the model specification and the performance of NOC, we consider an alternative set of J models,

$$M_k : \begin{cases} |p_j - \phi| \le \epsilon, & j = k, \\ |p_j - \phi| > |p_k - \phi|, & j \ne k, \end{cases} \quad k = 1, \ldots, J, \tag{3.1}$$

which allows multiple dose levels in the $(\phi-\epsilon, \phi+\epsilon)$. Correspondingly, we specify a joint prior distribution for $p_1, \ldots, p_J$ under each M_k,

$$\begin{cases} p_j | M_k \sim \text{Unif}(\phi-\epsilon, \phi+\epsilon), & j = k, \\ p_j | p_{j-1}, M_k \sim \text{Unif}(p_{j-1} + 2(\phi - p_{j-1}) I\{p_{j-1} < \phi\}, p_{J+1}), & j > k, \\ p_j | p_{j+1}, M_k \sim \text{Unif}(p_0, p_{j+1} - 2(p_{j+1} - \phi) I\{p_{j+1} > \phi\}), & j < k, \end{cases} \tag{3.2}$$

which maintains the monotone increasing pattern of the dose–toxicity relationship. We conduct the simulation study with 5000 replications under each scenario in Tables S1 and S2 of the supplementary material available at *Biostatistics* online shows that model (3.1) with the prior specification of (3.2) produces similar operating characteristics of NOC.

3.3. *Late-onset toxicity*

To investigate the proposed fNOC design, we consider the toxicity assessment period $\tau = 3$ months and the interarrival time between two consecutive cohorts $T = 0.5$ months; that is, the accrual rate is six patients per month, and the τ/T ratio $= 6$. We extend the scenarios in Table S1 of the supplementary material available at *Biostatistics* online by simulating time-to-toxicity outcomes based on Weibull distributions, and choose the shape and scale parameters such that only 30% of the toxicities would occur in the first half of the assessment period. We compare the fNOC design with five methods: TITE-CRM (Cheung and Chappell, 2000), TITE-EWOC (Mauguen *and others*, 2011), TITE-CCD (Ivanova *and others*, 2007) and fBOIN, where fBOIN is the design that combines BOIN and the fractional imputation approach.

Table S3 of the supplementary material available at *Biostatistics* online presents the simulation results based on 5000 replications under each scenario. All the five methods have similar trial durations, which are much less than the 36 months under the complete-data design. In scenarios 1–5, the fNOC design yields the highest MTD selection percentage. In terms of the number of patients treated at the MTD, fNOC on average yields a comparable performance with the other methods. For scenarios 9–12 which involve no or two MTDs, fNOC still performs satisfactorily. When evaluating trial designs with late-onset toxicities, patient safety is of great importance because ignoring late-onset toxicities would lead to overly aggressive dose escalation. By comparing Tables S1 and S3 in the supplementary material available at *Biostatistics* online, the average increment (with scenario 8 excluded due to no over-toxic doses) in the number of overdosing for the five methods (TITE-CRM, TITE-EWOC, TITE-CCD, fBOIN, and fNOC) under the late-onset toxicity are 2.9, 2.5, 7.1, 7.0, and 1.0 patients, respectively. The large increments of the two algorithm-based methods (TITE-CCD and fBOIN) are attributable to the fact that these two methods cannot borrow toxicity information from other dose levels for decision making: when the toxicity outcomes of the lower dose levels become available, it is impossible for the newly collected data to influence the dose escalation/de-escalation decision. Overall, the performance statistics with respect to the patient safety based on fNOC are notably better than those using the other methods. For example, fNOC assigns 16 fewer patients to overly toxic dose levels than TITE-EWOC in scenario 1; fNOC treats 10 fewer patients at overdoses than TITE-CRM; and the risk of high toxicity is about one third of that by TITE-CRM in scenario 3. Moreover, fNOC has comparable precision of the MTD selection with TITE-CRM, as the combination of the fractional imputation and the overdose control would induce a neutralization effect. When the toxicity data are updated, fNOC can borrow information from all the dose levels and quickly move toward underdoses if the current dose level is overly toxic.

4. Trial application

For illustration, we apply the proposed fNOC design to the sonidegib trial. Thirty patients in cohorts of size 3 were enrolled, and the trial started by treating the first cohort of patients at the lowest dose level. Table 1 shows that the first four cohorts were treated by dose levels 1, 2, 3, and 3, respectively. Under the fNOC design, we only update the posterior model probabilities upon the arrival of the first patient of each cohort. Based on the accumulated data in Table 1, two DLTs were observed at dose level 3 when patient 13 arrived on day 130. We estimate the survival function using the Kaplan–Meier estimator, as shown in Figure 2(a). At this moment, patients 6, 8, 9, 11, and 12 were still under the follow-up of evaluation without experiencing any DLT, which led to a total of five missing toxicity outcomes. We imputed the

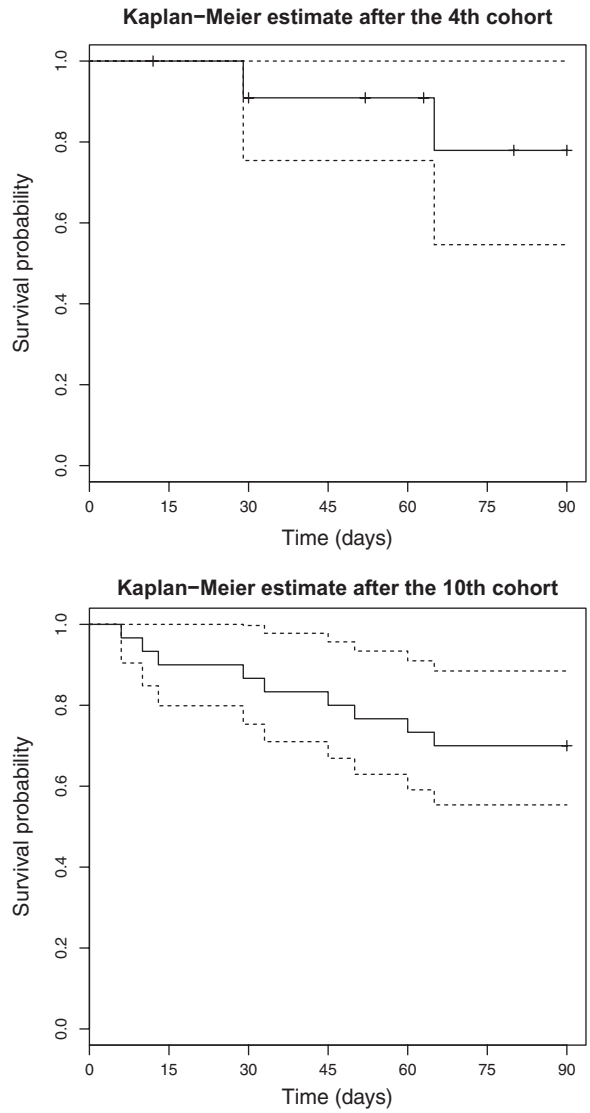

Fig. 2. Kaplan–Meier estimates of the survival function based on the observed data (a) after the 4th cohort and (b) after the 10th cohort, in the sonidegib trial.

missing data according to their fractional contribution (2.4), and as shown in Table 1, the corresponding fractional DLTs were 0.000, 0.143, 0.143, 0.143, and 0.221. The imputed dataset led to the posterior model probability $(P(M_1 \mid D_n), \ldots, P(M_5 \mid D_n)) = (0.02, 0.16, 0.55, 0.20, 0.07)$. Although dose level 3 had the largest posterior probability of being the MTD, our overdose control rule (2.2) determined the next dose to be dose level 2, and the additional dose-switching rule (2.3) was not activated. On the other hand, if we had treated the missing data as no DLTs, the posterior model probability would be $(0.01, 0.08, 0.49, 0.29, 0.13)$, i.e., the first three model probability would be underestimated, and then dose level 3 would be selected as the next treatment according to the overdose control rule, which coincides

Table 2. *Arrival time (days), dose level, and time to toxicity (days) for the remaining 18 patients in the sonidegib trial*

Cohort No.	Patient No.	Day of arrival	Dose level	Time to DLT	Cohort No.	Patient No.	Day of arrival	Dose level	Time to DLT
5	13	130	2	0	8	22	205	2	1 (255)
	14	131	2	0		23	210	2	1 (223)
	15	155	2	0		24	222	2	0
6	16	158	2	0	9	25	239	2	1 (245)
	17	167	2	1 (200)		26	248	2	0
	18	178	2	0		27	265	2	0
7	19	185	3	0	10	28	280	2	0
	20	188	3	1 (233)		29	283	2	1 (343)
	21	189	3	1 (199)		30	285	2	0

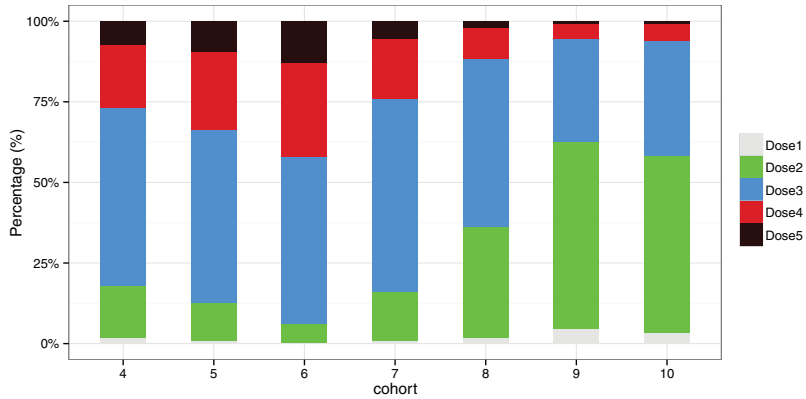

Fig. 3. Posterior probabilities of M_j based on the fNOC design in the sonidegib trial, $j = 1, \ldots, 5$. The result of the last cohort is obtained after all of the toxicity outcomes in the trial are completely observed.

with the decision made by TITE-EWOC. Typically, such dose assignment by treating missing data as no response would be too aggressive.

The subsequent dose assignment is provided in Table 2, and the posterior model probabilities by the arrival of the first patient of each new cohort are exhibited in Figure 3. The proposed fNOC design is able to update the posterior probabilities timely. On day 185 when patient 19 arrived, still two DLTs were observed, yet the toxicity outcome of patient 17 was censored at that moment. Given that no DLT had been observed at dose levels 1 and 2, the posterior probability of dose level 3 being the MTD was updated to be 0.60. Therefore, the dose-switching rule (2.3) was activated, and our algorithm recommended escalating to dose level 3. However, one DLT was observed prior to the arrival of patient 22, the proposed fNOC design took into account the new information and de-escalated back to dose level 2. For the last two sequences of enrolment, dose level 2 would remain, as its estimated cumulative posterior model probability is closest to the feasibility bound of 0.35. During the entire trial, the dose-switching rule (2.3) was not activated, as none of the posterior model probability was greater than 0.60. At the end of the trial, three dose levels had been explored with five observed DLTs out of 18 patients at dose level 2 and four DLTs out of nine patients at dose level 3. The observed toxicity probabilities at the five dose levels are $(0, 0.28, 0.44, -, -)$,

where "−" represents untried dose levels. Figure 2(b) presents the Kaplan–Meier estimator of the survival function based on the complete data after the trial is finished. The final posterior model probability is $(0.03, 0.55, 0.36, 0.05, 0.01)$, and therefore the second dose level would be selected as the MTD with the estimated toxicity probability of 0.28.

Figure 3 displays the posterior model probabilities throughout the trial. At the early stage of the trial, there was a small posterior probability of dose level 2 being the MTD since no DLT had been observed at that level, while the posterior probability of M_2 continuously increased by accumulating new information as the trial proceeded. It took the fNOC design approximately one year to finish the entire trial, while it would take over 2.5 years if each dose assignment is suspended until all of the toxicity outcomes in the trial are completely observed.

5. CONCLUDING REMARKS

We have proposed the NOC and fNOC designs to address the dose finding problems in phase I trials. In general, our methods have several prominent characteristics, which deviate from the existing dose-finding designs. First, NOC is built upon a nonparametric framework of Bayesian model selection, and each dose assignment is made to control overdosing by incorporating the complete information across all dose levels. As a result, NOC possesses both the robustness and safety properties. Second, by assembling NOC with the fractional imputation method under late-onset toxicities, the fNOC design can pull patients back to safe dose levels in case that late-onset toxicities lead to overly aggressive dose escalation. Such neutralization makes fNOC very appealing in dealing with delayed toxicity outcomes. Third, the proposed methods are flexible due to three tuning parameters: the feasibility bound α, the dose-switching cutoff η, and the dose elimination cutoff λ, and all of them have meaningful interpretations and can be easily calibrated. In addition to these unique features, simulation studies have shown that both the NOC and fNOC designs have competitive MTD selection and allocation percentages with the existing methods, and they can keep fewer patients from being treated at excessively toxic dose levels.

The proposed NOC method with prior (2.1) assumes that there exists only one dose level with $p_j \in (\phi-\epsilon, \phi+\epsilon)$ among the presepcified J dose levels, while simulations have shown that NOC still works well when this assumption is violated. On the other hand, model (3.1) with prior (3.2) can accommodate situations where multiple dose levels are assumed to lie in $(\phi-\epsilon, \phi+\epsilon)$ under M_k, $k = 1, \ldots, J$. The proposed methods have focused on the single-agent dose-finding trials, while the model selection framework of NOC can be generalized to drug-combination trials (Yin and Yuan, 2009; Lin and Yin, 2016) and phase I/II trials with (late-onset) efficacy outcomes (Lee *and others*, 2016). The R codes for implementation of the proposed NOC and fNOC designs are available on Github (https://github.com/ruitaolin/NOC).

SUPPLEMENTARY MATERIAL

Supplementary material is available at http://biostatistics.oxfordjournals.org.

ACKNOWLEDGMENTS

We thank the associate editor, the two referees, and the editor for their many constructive and insightful comments that have led to significant improvements in the article. *Conflict of Interest*: None declared.

FUNDING

The research was supported in part by a grant (17125814) from the Research Grants Council of Hong Kong.

References

Babb, J., Rogatko, A. and Zacks, S. (1998). Cancer phase I clinical trials: efficient dose escalation with overdose control. *Statistics in Medicine* **17**, 1103–1120.

Bekele, B. N., Ji, Y., Shen, Y. and Thall, P. F. (2008). Monitoring late-onset toxicities in phase I trials using predicted risks. *Biostatistics* **9**, 442–457.

Chen, Z., Tighiouart, M. and Kowalski, J. (2012). Dose escalation with overdose control using a quasi-continuous toxicity score in cancer phase I clinical trials. *Contemporary Clinical Trials* **33**, 949–958.

Cheung, Y. K. (2005). Coherence principles in dose-finding studies. *Biometrika* **92**, 863–873.

Cheung, Y. K. and Chappell, R. (2000). Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* **56**, 1177–1182.

Desai, S. P., Ben-josef, E., Normolle, D. P., Francis, I. R., Greenson, J. K., Simeone D. M., Chang, A. E., Colletti, L. M., Lawrence, T. S. and Zalupski, M. M. (2007). Phase I study of oxaliplatin, full-dose gemcitabine, and concurrent radiation therapy in pancreatic cancer. *Journal of Clinical Oncology* **25**, 4587–4592.

Durham S. D., Flournoy N. and Rosenberger W. F. (1997). A random walk rule for phase I clinical trials. *Biometrics* **53**, 745–760.

Gezmu, M. and Flournoy, N. (2006). Group up-and-down designs for dose-finding. *Journal of Statistical Planning and Inference* **136**, 1749–1764.

Heyd, J. M. and Carlin, P. B. (1999). Adaptive design improvements in the continual reassessment method for phase I studies. *Statistics in Medicine* **18**, 1307–1321.

Ivanova, A., Flournoy, N. and Chung, Y. (2007). Cumulative cohort design for dose finding. *Journal of Statistical Planning and Inference* **137**, 2316–2327.

Kaplan, E. L. and Meier, P. (1958). Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association* **53**, 457–481.

Lee, J., Thall, P. F., Ji, Y. and Müller, P. (2016). A decision-theoretic phase I/II design for ordinal outcomes in two cycles. *Biostatistics* **17**, 304–319

Leung, D. H. Y. and Wang, Y. G. (2002). An extension of the continual reassessment method using decision theory. *Statistics in Medicine* **21**, 51–63.

Lin, R. and Yin, G. (2016). Bayesian optimal interval design for dose finding in drug-combination trials. *Statistical Methods in Medical Research*, doi: 10.1177/0962280215594494

Liu, S. and Yuan, Y. (2015). Bayesian optimal interval designs for phase I clinical trials. *Journal of Royal Statistical Society: Series C (Applied Statistics)* **64**, 507–523.

Liu, S., Yin, G. and Yuan, Y. (2013). Bayesian data augmentation dose finding with continual reassessment method and delayed toxicity. *The Annals of Applied Atatistics* **7**, 2138–2156.

Mauguen, A., Le Deley, M. C. and Zohar, S. (2011). Dose-finding approach for dose escalation with overdose control considering incomplete observations. *Statistics in Medicine* **30**, 1584–1594.

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.

Paoletti, X. , O'Quigley, J. and Maccario, J. (2004). Design efficiency in dose finding studies. *Computational Statistics & Data Analysis* **45**, 197–214.

Postel-Vinay, S., Gomez-Roca, C., Molife, L. R., Anghan, B., Levy, A., Judson, I., Bono, J. D., Soria, JC., Kaye, S. and Paoletti, X. (2011). Phase I trials of molecularly targeted agents: should we pay more attention to late toxicities? *Journal of Clinical Oncology* **29**, 1728–1735.

SIMON, R., RUBINSTEIN, L., ARBUCK, S. G., CHRISTIAN, M. C., FREIDLIN, B. AND COLLINS, J. (1997). Accelerated titration designs for phase I clinical trials in oncology. *Journal of the National Cancer Institute* **89**, 1138–1147.

STORER, B. E. (1989). Design and analysis of phase I clinical trials. *Biometrics* **45**, 925–937.

SUSARLA, V. AND VAN RYZIN, J. (1976). Nonparametric Bayesian estimation of survival curves from incomplete observations. *Journal of American Statistical Association* **71**, 897–902.

SVERDLOV, O., WONG, W. K. AND RYEZNIK, Y. (2014). Adaptive clinical trial designs for phase I cancer studies. *Statistics Surveys* **8**, 2–44.

TIGHIOUART, M. AND ROGATKO, A. (2010). Dose finding with escalation with overdose control (EWOC) in cancer clinical trials. *Statistical Science* **25**, 217–226.

TIGHIOUART, M., ROGATKO, A. AND BABB, J. S. (2005). Flexible Bayesian methods for cancer phase I clinical trials. Dose escalation with overdose control. *Statistics in Medicine* **24**, 2183–2196.

WEBER, J. S., LEVIT, L. A., ADAMSON, P. C., BRUINOOGE, S., BURRIS, H. A., CARDUCCI, M. A., DICKER, A.P., GÖNEN, M., KEEFE, S.M., POSTOW, M.A. *and others* (2015). American Society of Clinical Oncology policy statement update: the critical role of phase I trials in cancer research and treatment. *Journal of Clinical Oncology* **33**, 278–284.

YIN, G. (2012). *Clinical Trial Design: Bayesian and Frequentist Adaptive Methods.* New York: John Willey & Sons, Inc.

YIN, G. AND YUAN, Y. (2009). Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society Series C (Applied Statistics)* **61**, 211–224.

YIN, G., ZHENG, S. AND XU, J. (2013). Fractional dose-finding methods with late-onset toxicity in phase I clinical trials. *Journal of Biopharmaceutical Statistics* **23**, 856–870.

YUAN, Y. AND YIN, G. (2011). Robust EM continual reassessment method in oncology dose finding. *Journal of the American Statistical Association* **106**, 818–831.

YUAN, Z., CHAPPELL, R. AND BAILEY, H. (2007). The continual reassessment method for multiple toxicity grades: A Bayesian quasi-likelihood approach. *Biometrics* **63**, 173–179.

ZACKS, S., ROGATKO, A. AND BABB, J. (1998). Optimal Bayesian-feasible dose escalation for cancer phase I trials. *Statistics and Probability Letters* **38**, 215–220.

[*Received January 31, 2016; revised July 08, 2016; accepted for publication July 11, 2016*]