

Research paper

Bayesian optimization for estimating the maximum tolerated dose in Phase I clinical trials

Ami Takahashi ^{a,b,*}, Taiji Suzuki ^{c,d}^a Department of Mathematical and Computing Science, School of Computing, Tokyo Institute of Technology, Tokyo, Japan^b Biometrics and Data Management, Clinical Statistics, Pfizer R&D Japan, Tokyo, Japan^c Department of Mathematical Informatics, Graduate School of Information Science and Technology, The University of Tokyo, Tokyo, Japan^d Center for Advanced Intelligence Project, RIKEN, Japan

ARTICLE INFO

Keywords:

Bayesian optimization
Dose-finding study
Maximum tolerated dose
Nonparametric method

ABSTRACT

We introduce a Bayesian optimization method for estimating the maximum tolerated dose in this article. A number of parametric model-based methods have been proposed to estimate the maximum tolerated dose; however, parametric model-based methods need an assumption that dose–toxicity relationships follow specific theoretical models. This assumption potentially leads to suboptimal dose selections if the dose–toxicity curve is misspecified. Our proposed method is based on a Bayesian optimization framework for finding a global optimizer of unknown functions that are expensive to evaluate while using very few function evaluations. It models dose–toxicity relationships with a nonparametric model; therefore, a more flexible estimation can be realized compared with existing parametric model-based methods. Also, most existing methods rely on point estimates of dose–toxicity curves in their dose selections. In contrast, our proposed method exploits a probabilistic model for an unknown function to determine the next dose candidate without ignoring the uncertainty of posterior while imposing some dose-escalation limitations. We investigate the operating characteristics of our proposed method by comparing them with those of the Bayesian-based continual reassessment method and two different nonparametric methods. Simulation results suggest that our proposed method works successfully in terms of selections of the maximum tolerated dose correctly and safe dose allocations.

1. Introduction

The primary goal of oncology Phase I clinical trials is to identify the maximum tolerated dose (MTD) defined as the highest dose that does not cause an unacceptable level of dose-limiting toxicity (DLT). A number of statistical methods have been proposed for identifying the MTD. The earliest rule-based method is the 3+3 method [1] and is still employed by researchers mainly because of its simplicity [2]; however, it is well-known that the 3+3 method has substantial limitations [3–6]. In addition, many authors have demonstrated that the continual reassessment method (CRM) [7] and its variations [8–11] provide superior performance to the 3+3 method in terms of accuracy of estimation, proper dose allocation, and flexibility [12–19]. Therefore, leading pharmaceutical companies have commonly applied variations of the CRM in oncology Phase I clinical trials [20].

The CRM needs to specify a theoretical model describing a dose–toxicity relationship before the beginning of a trial. It has been demonstrated that the CRM can achieve the primary goal with regards to the selection of the MTD even if the model is misspecified [21]; however,

sufficient conditions to converge to the true MTD established by the previous study [21] might be too restrictive [22]. It means that there are possible scenarios where the CRM cannot guarantee convergence. In addition, the choice of theoretical models in the CRM would affect the operating characteristics if sample sizes are small (e.g., 10 to 50 patients) [23].

As different approaches from parametric model-based methods, nonparametric Bayesian methods have been introduced to relax the effect of model selections. Toxicity probability interval methods might be major approaches of nonparametric Bayesian methods. The modified toxicity probability interval (mTPI) [24] is one of the most popular toxicity probability interval methods. The mTPI is designed to implement as simply as the 3+3 methods without any logistic burden, although it is a Bayesian adaptive design assisted by a beta-binomial model. Because of its simplicity and much better performance than the 3+3 method, the popularity of the mTPI has been growing in both research and industry entities during the relatively short period since it was proposed [25].

* Correspondence to: 2-12-1 Ookayama, Meguro-ku, Tokyo 152-8550, Japan.
E-mail address: takahashi.a.ax@m.titech.ac.jp (A. Takahashi).

The Bayesian optimal interval (BOIN) [26] uses the same type of decision rules as the mTPI except for escalation and de-escalation boundaries. An extended version of the mTPI named mTPI-2 [27] has been proposed to solve an undesirable issue, which may happen under specific situations, about dose-escalation and de-escalation rule based on the mTPI.

An existing study that compared parametric model-based methods (CRM, dose-escalation with overdose control and Bayesian logistic regression model) and toxicity probability interval methods (mTPI, BOIN, and keyboard design which operates in the same way as the mTPI-2) reported that the CRM outperformed the other parametric model-based methods in terms of accuracy of identifying the MTD, and the BOIN outperformed the mTPI and provided comparable performance with the CRM [28]. Similarly, another study that compared the operating characteristics of the 3+3 method, the CRM, the BOIN, and the keyboard design [29] suggested that both the BOIN and the keyboard design provided comparable performance with the CRM. It was also reported that the CRM tended to outperform the BOIN and mTPI as the number of dose levels increases [30]. There are other studies that compared performance between the CRM and toxicity probability interval methods. Conclusions differ slightly among studies; however, the common point is that the CRM would consistently perform well in various scenarios and toxicity probability interval methods are attractive in their ease of use.

Other than toxicity probability interval methods, a nonparametric Bayesian method based on a product-of-beta-prior (PBP) [31,32] has been cited as a curve-free method in many articles. It models a toxicity probability at each dose level directly without assuming a specific dose-toxicity curve. It assumes a prior distribution of a toxicity probability as a product-of-beta prior owing to the reparametrization of a toxicity probability at each dose level by another parameter with a beta-prior distribution. Although the PBP might cause rigidity in situations where a low toxicity rate is targeted due to its vague priors, the solutions have been discussed and given by the existing study [33]. Also, the PBP has sometimes been pointed for its numerical problem on exact posterior toxicity distributions, while this is avoidable if Markov chain Monte Carlo (MCMC) is used instead of the exact computation. A different type of Bayesian curve-free methods [34] has been introduced to consider probabilities falling into pre-specified toxicity risk categories on each dose level (hereafter, we abbreviate this method WTW based on the initial of the authors). There is also a curve-free method based on Dirichlet process prior [35]. An existing study indicated that curve-free methods would prefer to the CRM if there is little evidence of high enough quality about dose-toxicity relationships [18]. There are already several curve-free methods. Nevertheless, it would have been yet a great challenge to develop a sophisticated curve-free method because little is known about dose-toxicity relationships in general.

The purpose of this article is to introduce a Bayesian nonparametric approach named Bayesian optimization method that provides a different approach from existing curve-free methods and realizes sophisticated dose selection procedures by utilizing a Bayesian optimization framework. The Bayesian optimization [36,37] has emerged as an efficient optimization strategy of unknown functions that are expensive to evaluate and has recently impacted a wide range of areas such as machine learning, sensor networks, environmental monitoring [38]. Our proposed method nonparametrically models dose-toxicity relationships; therefore, it realizes flexible modeling. Also, it allows identifying the MTD in relatively few evaluations while utilizing all available information that includes uncertainties of estimates from observations without simply relying on local optimal points in accordance with a Bayesian optimization framework. To the best of our knowledge, no literature exists on investigating a Bayesian optimization approach for estimating the MTD in Phase I clinical trials.

Once patient outcomes are observed, our proposed method updates a distribution over an unknown dose-toxicity function via Bayes' rule. The next dose is then guided by an acquisition function derived from

the updated distribution. Similar to other dose-finding methods, we repeat the steps of observing patient outcomes, updating the distribution, and selecting the next dose until either pre-specified stopping criteria is met. As for safety monitoring during a trial, our proposed method imposes some dose-escalation restrictions through an admissible dose set defined to ensure patient safety. Based on the final distribution for the dose-toxicity relationship at the end of a trial, the MTD is selected from the final MTD candidate set defined from the view of safety perspective which aims to reduce overdose determination. Our simulation study suggests that our proposed method provides comparable or better performance on the MTD estimation and safer dose allocations compared with the Bayesian-based CRM and also provides more stable results than the PBP and the WTW.

We organize this article as follows: In Section 2, we describe statistical modelings and dose-finding strategies of our proposed method. We also introduce a brief example with some illustrations when our proposed method applies. In Section 3, we describe simulation frameworks to compare our proposed method, the CRM, the PBP, and the WTW. We briefly explain each competitor as well. The simulation results are shown after the explanations of the simulation framework. The last section presents our conclusion as well as some discussions for our future works.

2. A Bayesian optimization design

Sections 2.1 and 2.2 describe how our proposed method models dose-toxicity relationships and selects the MTD candidate. Section 2.3 mentions overdose control imposed in our proposed method to ensure patient safety. Section 2.4 describes the implementation steps of our proposed method. Section 2.5 illustrates a brief example with some drawings that we can obtain in our proposed method.

2.1. Dose-toxicity relationship

In a Bayesian optimization method, dose-toxicity relationships are modeled through a nonparametric approach. Let us define dose-toxicity relationships as follows:

$$f(x_j) = \text{logit}(\pi_j) = \log\left(\frac{\pi_j}{1 - \pi_j}\right), \quad (1)$$

where π_j denotes a toxicity probability at a conceptual dose x_j corresponding to a dose level $j \in \{1, \dots, J\}$. The logit transformation for π_j plays a role to guarantee that π_j bounds within the range 0 to 1 on a finite dose range. We note that a conceptual dose x_j is not necessary to be an actual dose because dose-toxicity relationships in the estimation process only rely on distances between conceptual doses but not actual doses; however, conceptual doses should be as equally spaced with regards to toxicity probabilities as possible so that we could avoid skewed estimation of toxicity.

Based on typical Bayesian optimization frameworks, a Gaussian process prior is put over the unknown function f to estimate it in the Bayesian manner:

$$f \sim \text{GP}(m, k). \quad (2)$$

The Gaussian process prior is specified by a mean function $m(x)$ and a covariance function $k(x, x')$, where $x \in \{x_1, \dots, x_J\}$. A Gaussian process makes a model easy to treat because it leaves setting a prior distribution to designing a kernel function [39]. (As a side note, there is a method that utilizes a Gaussian process for dose-finding studies, although its main purpose is modeling population pharmacokinetics based on data observed in dose-finding studies [40].)

The prior mean function for $m(x)$ derives from pre-specified initial guesses for toxicity probabilities. Because little is known about dose-toxicity relationships in general, one option for setting initial guesses might be utilizing an indifference interval determined by a systematic approach proposed for the CRM [41]. The indifference interval for a

given dose level is defined as an interval of toxicity rates associated with the neighboring doses, such that these neighboring doses may be selected instead of the true MTD. The systematic approach provides an indifference interval to maximize the average percentage of correct selection in the selected model of the CRM across a set of scenarios of true toxicity probabilities. We will describe detailed steps about how to set initial guesses based on the indifference interval approach at Step 1 for the model phase in Section 2.4. The covariance function $k(x, x')$, which is a kernel function, determines the smoothness properties of samples drawn from it. We apply the squared exponential kernel given by

$$k(x, x') = \sigma_f^2 \exp\left(-\frac{1}{2\rho^2} |x - x'|^2\right), \quad (3)$$

where σ_f is called a signal variance and ρ is a scale parameter. A signal variance σ_f determines the variation of function values from their mean, and a scale parameter ρ controls the width of the kernel. The squared exponential kernel is a very popular choice for the Bayesian optimization and gives a suitable solution for problems with a smooth unknown function. Dose–toxicity curves are sufficiently smooth functions in general; therefore, we adopt the squared exponential kernel. In addition, a small value of ξ , which is similar to noise in a regression model, is added on the diagonal elements of the covariance function in order to provide computational stability [42]. Each element in the covariance matrix K for an arbitrary conceptual dose x and x' is expressed as $K_{x,x'} = k(x, x') + \xi 1[x = x']$.

The number of patients experienced DLT follows a binomial distribution; therefore, the likelihood function of the observed values up to the t th cohort is given by

$$L(D_{1:t} | f) = \prod_{i=1}^t \pi_{(i)}^{y_{(i)}} (1 - \pi_{(i)})^{n_{(i)} - y_{(i)}}, \quad (4)$$

where $D_{1:t} = \{(n_{(1)}, x_{(1)}, y_{(1)}), \dots, (n_{(t)}, x_{(t)}, y_{(t)})\}$; $y_{(i)}$ denotes the number of patients experienced DLT out of $n_{(i)}$ patients treated with $x_{(i)}$ corresponding to the dose level (i) at the i th test; $\pi_{(i)}$ denotes a toxicity probability corresponding to the dose level (i).

Once patient outcomes at the t th test are observed, a posterior distribution for f is updated based on the Bayes' rule. In practice, we obtain posterior samples of f generated by MCMC.

2.2. Dose selection strategy

Suppose that we are interested in finding a dose that produces the closest toxicity to a target toxicity rate θ . Because the Bayesian optimization needs to define an objective function to be minimized (or maximized) for its optimization, we set an objective function as follows:

$$g(x_j) = |\pi_j - \theta|, \quad (5)$$

where π_j derives from $f(x_j)$. Soon after posterior samples for f are obtained based on Section 2.1, posterior samples for g are calculated. The exact form of g is still unavailable; however, we can leverage those probabilistic beliefs in order to reach $x_{(t+1)}$ through designing an acquisition function \hat{g} that is an alternative of the true objective function g . In this article, we utilize the expected improvement (EI) [37] as \hat{g} , which is one of the most popular strategies in the Bayesian optimization and has been shown to be efficient in the number of function evaluations required to find the global optimum [43].

In the EI algorithm, what we calculate is how much a value of the objective function g can be expected to improve over our current best point while taking the uncertainty derived from the posterior distribution of the objective function g into account. A prior distribution we set gradually shrinks to the true values as the data is obtained; however, there is a large width of uncertainties when the distribution has not been sufficiently converged. Because the EI algorithm considers the width of uncertainties, a dose with a large variance is not selected

even if it is an optimum in the sense of average. The following are the specifics for the EI algorithm. Firstly, an improvement function is given by

$$I(x) = \max\{0, g^+ - g(x)\}, \quad (6)$$

where $g^+ = \min_x [E_g\{g(x)\} | D_{1:t}]$ that means the current best point providing the minimum value on $g(x)$ among all available doses [44]. Accordingly, $I(x)$ provides a positive value if $g(x)$ turns out to be less than g^+ . Otherwise, $I(x)$ is set to be zero. Secondly, EI(x) is calculated as the expectation of $I(x)$ as follows:

$$EI(x) = E\{I(x) | D_{1:t}\} = \int_0^1 I(x)p\{g(x) | D_{1:t}\}dg, \quad (7)$$

where $p\{g(x) | D_{1:t}\}$ is a probability density function for g on an arbitrary dose x after $D_{1:t}$ is obtained. Finally, the next dose is found by maximizing the expected improvement function:

$$x_{(t+1)} = \arg \max_{x \in \mathcal{A}_t} \{EI(x)\}, \quad (8)$$

where \mathcal{A}_t is an admissible dose set defined at Section 2.3. It imposes some dose-escalation restrictions for overdose control to ensure patient safety.

2.3. Overdose control for patient safety

Patients in the next cohort are treated with the selected dose according to Eq. (8). In order to ensure patient safety, we impose the following overdose control during a trial. The admissible dose set \mathcal{A}_t is refreshed at each test and includes doses that satisfy all conditions as follows:

1. If the lowest dose does not satisfy $P(\pi_1 > \theta | D_{1:t}) \leq \tau_1$, \mathcal{A}_t includes only x_1 .
2. All doses in \mathcal{A}_t require to satisfy $P(\pi_j > \theta | D_{1:t}) < \tau_2$.
3. No dose skip is allowed in this article; therefore, the highest dose level in \mathcal{A}_t is up to one dose level higher than $x_{(t)}$. If $x_{(t)} = x_j$, the highest dose is up to x_j .
4. If two or more patients experience DLT at $x_{(t)}$, the highest dose level is up to one dose level lower than $x_{(t)}$. If $x_{(t)} = x_1$, \mathcal{A}_t includes only x_1 .

For the first and second conditions, we assume $\tau_1 \leq \tau_2$ in order not to miss unsafe situations under the assumption that toxicity increases monotonically with increasing dose levels. The third condition might be a typical setting when the CRM is implemented in actual clinical trials.

2.4. Implementation steps

We employ a start-up phase before implementing model estimation procedures because the information available at the beginning of a trial may be too limited to rely entirely on the model estimation part based on a Bayesian optimization method when little is known about the dose–toxicity relationship.

Start-up phase

The start-up phase is described as follows:

1. Patients at the first cohort are treated with the lowest dose x_1 .
2. If no DLT is observed, patients in the next cohort are treated with one level higher dose than the current dose level.
3. If one patient experiences DLT for the first time in the trial, patients at the next cohort are treated at the same dose level.
4. The start-up phase is stopped when the trial meets one of the following conditions:
 - (a) Two or more patients experience DLT.
 - (b) The test dose reaches the highest dose x_j .

Model phase

After the start-up phase, the trial proceeds to the model phase that is implemented as follows.

1. Design parameters for a Gaussian process prior are determined as follows:
 - (a) If there is little information about dose–toxicity relationships, an indifference interval approach can be utilized for initial guesses that are associated with a prior mean function for $m(x)$.
 - i. According to the systematic approach proposed for the CRM [41], we calculate an indifference interval named an optimal δ while using a power model and the initial MTD at the center of the dose range (i.e., $J/2$), which we could maximize the average percentage of correct MTD selection when applying the CRM with a power model. The main purpose of utilizing an optimal δ is to obtain one of the good guesses of the slope of initial guesses pertaining to a prior mean function for $m(x)$ in our proposed method. Smaller δ makes the slope more gentle. As a result, dose-escalation becomes more aggressive than larger δ in our proposed method. In contrast, initial guesses based on a large δ tends to offer more conservative dose-escalation. Empirically, our proposed method would perform well when using values close to the optimal δ (e.g., a range of $\delta \pm 0.02$).
 - ii. The initial MTD location ν could be determined based on the start-up phase. Otherwise, we recommend putting the initial MTD location on the center of the dose range (i.e., $\nu = J/2$) to ensure enough space within the range both below and above the dose.
 - iii. We generate the initial guesses using `getprior` function in the R package `dfcrm` with the above δ and ν under a power model. We calculate a prior mean function for $m(x)$ by logit transformation of the initial guesses.

If there is an informative belief on dose–toxicity relationships, initial guesses should reflect it.
 - (b) A signal variance σ_f is set as one because it is a typical setting in the Bayesian optimization and works well in most cases. A scale parameter ρ indicates a typical distance between turning points and depends on the conceptual dose range. Because dose–toxicity curves do not have many turning points, an appropriate ρ provides up to two turning points in the range. In our proposed method, a good value is often comparable to the length of the dose range.
2. With Eqs. (1) through (4), posterior samples for the function f are calculated based on the Bayes' rule through MCMC. Simultaneously, they are transformed with the inverse logit function to obtain posterior samples of toxicity probabilities π to assess dose–toxicity relationships (e.g., by drawing the posterior distribution for π). In addition, posterior samples for the function g are calculated with the posterior samples for π according to Eq. (5).
3. Patients in the next cohort are treated with the selected dose according to Eq. (8).
4. Steps 2 and 3 continue until either pre-specified stopping rules is met.

5. At the end of a trial, the MTD is determined based on the final posterior distribution for toxicity probabilities as follows:

$$\text{MTD} = \arg \max_{x_j \in \{x_j | \hat{\pi}_j < (\theta + \epsilon_2)\}} P(\theta - \epsilon_1 < \pi_j < \theta + \epsilon_1 \mid D_{1:t}), \quad (9)$$

where $\hat{\pi}_j$ is a posterior mean estimate of a toxicity probability at a dose level j ; ϵ_1 and ϵ_2 are pre-specified small values ($\epsilon_1 \leq \epsilon_2$). The MTD is selected based on the acceptable range for θ provided by ϵ_1 from the final MTD candidate set that does not retain doses with excessive toxicities based on ϵ_2 .

2.5. An example with illustrations

We illustrate a specific example of the model phase under $\theta = 0.3$ after the end of the start-up phase that evaluated five cohorts with a cohort size of three in Fig. 1. The upper section provides posterior distributions on dose–toxicity relationships at each test as its posterior mean (dashed line) and 10 to 90 percentiles area (filled area) along with observed toxicity probabilities at each tested dose level (filled circle). It also provides curves of initial toxicity guesses (dotted line) and true toxicity probabilities (solid line). The acquisition function EI corresponding to each upper figure is shown in the lower section. The lower section presents doses included in \mathcal{A}_t after each test as ranges with horizontal lines and points the next dose level selected by Eq. (8) with arrow marks as well as the dotted lines.

In the start-up phase with five tests, tested dose levels were 1, 1, 2, 3, and 4 with the number of patients who experienced DLT of 1, 0, 0, 0, and 1, respectively. The upper-left section illustrates the updated toxicity probability distribution at the end of the start-up phase. The posterior mean function drew still similar curve to the initial toxicity probabilities. Based on the corresponding EI shown in the lower-left section and \mathcal{A}_5 , the next dose became x_4 . There was no patient with DLT at the 6th test. As shown in the figures in the middle section, the posterior distribution slightly changed and the EI selected x_5 . In the 7th test, there was one patient who experienced DLT. The upper-right section shows the posterior mean function approached the true dose–toxicity curve. After the 7th test, the EI selected x_5 again as the next dose.

3. Simulation studies and results

We conducted a simulation study to examine the performance of a Bayesian optimization method (BO). As benchmarks for the performance evaluation, we set the Bayesian-based CRM, the product-of-beta-prior method (PBP), and another curve free method (WTW). In Section 3.1, we explain the simulation frameworks for each method and a brief explanation for the competitors. Section 3.2 presents the simulation results.

3.1. Simulation settings

Suppose that we aimed to find an MTD that has a toxicity probability closest to the target toxicity θ among eight dose levels ($J = 8$) under the maximum sample size of 36 and a cohort size of three. The sample size of 36 was based on the average number of patients who enrolled in model-guided Phase I clinical trials reported by a review article [45]. As shown in Table 1, fifteen scenarios were used for the simulation study, which included six scenarios excerpted from the existing article (scenario 1 to 6) [35]. We evaluated the operating characteristics of each method under two different target toxicity rate settings. For $\theta = 0.3$ that might be a typical setting in dose-finding studies, we used scenarios 1 to 10. We might encounter a lower target toxicity rate when investigational agents could present severe symptoms as DLT; therefore, we set $\theta = 0.1$ and evaluated scenarios 3, 10 to 15.

All methods started to test from the lowest dose. In addition, all methods terminated trials when the maximum sample size was reached unless otherwise specified. The number of trials to evaluate the operating characteristics of each method was 1000 in the simulation study.

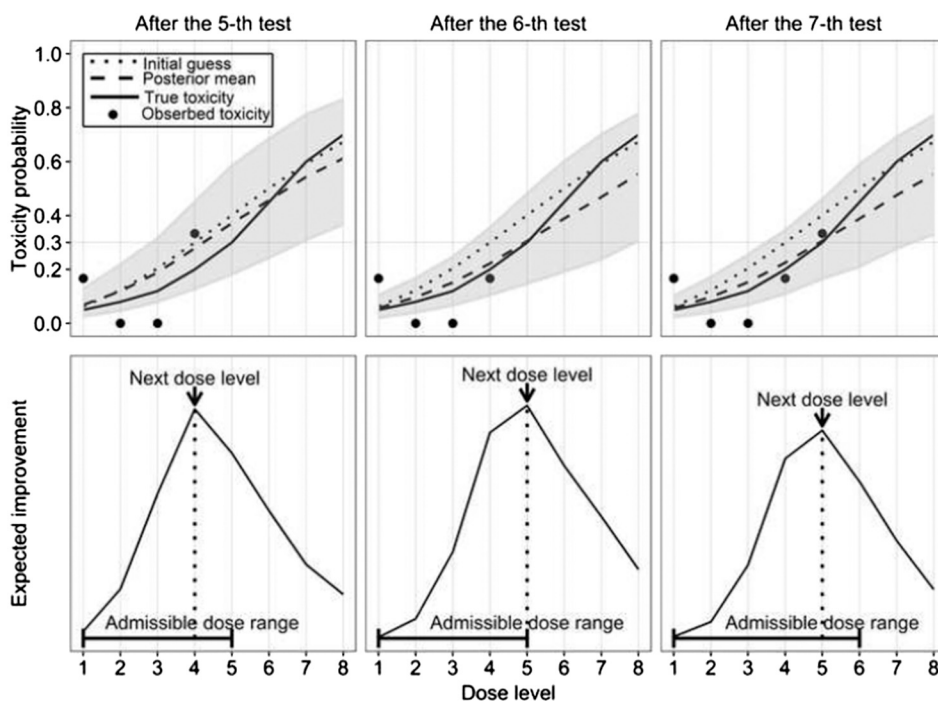


Fig. 1. Posterior distributions and the EI in a Bayesian optimization method.

Table 1
True toxicity scenarios for the simulation study.

Scenario	Dose level							
	1	2	3	4	5	6	7	8
1	0.05	0.08	0.12	0.20	0.30	0.45	0.60	0.70
2	0.05	0.08	0.12	0.20	0.30	0.60	0.80	0.90
3	0.01	0.05	0.10	0.14	0.18	0.22	0.25	0.30
4	0.01	0.05	0.08	0.12	0.16	0.20	0.24	0.26
5	0.30	0.40	0.50	0.60	0.70	0.80	0.90	0.95
6	0.40	0.45	0.50	0.55	0.60	0.65	0.70	0.80
7	0.01	0.05	0.10	0.14	0.17	0.20	0.30	0.40
8	0.01	0.05	0.30	0.45	0.55	0.70	0.80	0.90
9	0.15	0.30	0.45	0.50	0.55	0.60	0.65	0.70
10	0.05	0.10	0.15	0.30	0.40	0.55	0.65	0.70
11	0.10	0.20	0.30	0.40	0.50	0.60	0.70	0.80
12	0.01	0.02	0.05	0.10	0.15	0.20	0.30	0.40
13	0.01	0.01	0.02	0.03	0.10	0.20	0.35	0.50
14	0.01	0.02	0.03	0.05	0.07	0.10	0.15	0.20
15	0.01	0.01	0.01	0.02	0.04	0.06	0.10	0.15

3.1.1. Bayesian optimization method (BO)

We generated initial toxicity guesses using the indifference interval approach as described at Step 1 for the model phase in Section 2.4. We derived an optimal δ under each θ by evaluating possible indifference intervals from 0.01 to 0.15 by 0.01 running 2000 simulations, where these settings (i.e., the evaluation range and the number of iterations) were based on the original article of the systematic approach [41]. In addition, we performed preliminary evaluations on values around the optimal δ from the view of correct MTD selection probabilities and dose allocations, because the optimal δ for the CRM might not be optimal for our proposed method. As a result, 0.05 which was equal to the optimal value, and 0.02 which was 0.01 lower than the optimal value seemed suitable for scenarios under $\theta = 0.3$ and $\theta = 0.1$, respectively.

For the initial MTD location ν in the initial guesses under $\theta = 0.3$, as a similar approach in the 3+3 method, we set ν as either the last-tested dose level at the start-up phase or one dose level lower than the last-tested dose level if more than two patients experienced DLT at the last tested dose. Although the 3+3 method algorithm could apply to scenarios for $\theta = 0.3$ to set our initial thought of an MTD location,

it is not usually applicable to cases with different toxicity rates from 0.3; therefore, we used the center of the dose range (i.e., $\nu = 4$) when $\theta = 0.1$.

A covariance function was calculated with $\sigma_f = 1$, $\rho = 1.4$ and $\xi = 0.08$ under $(x_1, \dots, x_8) = (0.0, 0.2, 0.4, 0.6, 0.8, 1.0, 1.2, 1.4)$ regardless of θ . The value of ρ was the same length as the conceptual dose range. This indicates that the first turning point of the dose-toxicity function was placed on the end of the dose range. The value of ξ was decided from the view of computational speed, while the smaller value is generally better in terms of less impact on the operating characteristics.

For overdose control, $\tau_1 = 0.5$ and $\tau_2 = 0.9$ were used under $\theta = 0.3$. The assumption of $\tau_1 = 0.5$ means to stay at the lowest dose level and not to escalate to higher doses when the lowest dose level exceeds θ with a probability of more than a half. Because we set a relatively small value for δ under $\theta = 0.1$ in addition to the lower target toxicity rate itself, overdose allocations tended to be easier occurred than the settings for $\theta = 0.3$; therefore, both τ_1 and τ_2 had to be smaller values than those sets for $\theta = 0.3$. As a result, we used $\tau_1 = \tau_2 = 0.4$ under $\theta = 0.1$ to avoid overdose allocations. We determined the values of τ_2 under each θ based on a balance between MTD selections and overdose allocations.

For MTD determination described in Eq. (9), $\epsilon_1 = 0.05$ and $\epsilon_2 = 0.1$ were used under $\theta = 0.3$, where the proper range with $\epsilon_1 = 0.05$ for MTD determination is often assumed as a typical setting in dose-finding studies for $\theta = 0.3$. The final MTD candidate set composed of doses with toxicity probabilities that were equal to or less than 0.4 could be acceptable when $\theta = 0.3$ considering the estimation accuracy of the point estimate. Also, $\epsilon_1 = \epsilon_2 = 0.015$ was used for $\theta = 0.1$, where $\theta = 0.1$ and $\epsilon_1 = 0.015$ provided almost as small a ratio as $\theta = 0.3$ and $\epsilon_1 = 0.05$. When the target toxicity is lower than a typical setting, investigators might consider controlling overdose selections as much as possible from a clinical perspective; therefore, we conservatively set ϵ_2 as the same value of ϵ_1 , which also aimed to reduce the possible effect on overdose selections due to a small δ .

As exploratory analyses under $\theta = 0.3$, we evaluated a Bayesian optimization method that provided only monotonically increasing dose-toxicity functions (BO-mono) by assuming the following equation instead of Eq. (2):

$$f' \sim \text{GP}(m, k) \times \mathbb{1}[f' : \text{monotonically increasing function}]. \tag{10}$$

In practical implementation, posterior samples of toxicity probabilities were composed of only monotonically increasing functions (i.e., all posterior samples of toxicity probability functions met a condition of $\pi_1 \leq \dots \leq \pi_8$) so as to achieve the monotonically increasing restriction described in Eq. (10).

As another exploratory analysis for a Bayesian optimization method (without monotonically increasing restriction) under $\theta = 0.3$, we evaluated the effect of the slope on initial guesses by applying different values of δ ($\delta = 0.03$ and 0.07).

Supplementary material of this article provides the R code used to implement BO in this simulation study.

3.1.2. Continual reassessment method (CRM)

The CRM models dose-toxicity relationships with simple one-parameter monotonically increasing functions. For example, a power model describes a toxicity probability at a dose level j as follows:

$$\pi_j(a) = (\pi_j^0)^{\exp(a)}, \tag{11}$$

where a is an unknown model parameter, and $\pi_1^0 < \dots < \pi_J^0$ denote initial toxicity guesses at each dose level. The initial toxicity guesses are pre-specified constants and called skeletons for the model. In our simulation study, we evaluated a power model and a logistic model with an intercept of three that are commonly employed for the CRM. Based on the systematic approach [41], an optimal δ was 0.05 for $\theta = 0.3$ and 0.03 for $\theta = 0.1$ when assuming an initial MTD location $v = J/2$ under each model. In addition, assuming a normal formulation to the prior distribution of a model parameter a (i.e., $a \sim N(0, \sigma_a^2)$), σ_a was calibrated by another systematic approach [46] to obtain a least informative prior variance σ_a^{LI} . A least informative prior variance provides that the probabilities with each dose being MTD follow a uniform distribution as a prior distribution in order to minimize the effect of the prior in the estimation process because little is known about the dose-toxicity relationship at the start of a trial. As a result, the power model set $\sigma_a^{LI} = 0.76$ and 0.72 , and the logistic model set 0.34 and 0.36 under $\theta = 0.3$ and 0.1 , respectively.

The MTD was determined as a dose with the closest toxicity to θ based on the final posterior mean of toxicity probabilities. Dose selections did not allow either to skip doses in dose-escalation or to escalate doses immediately after a toxic outcome to ensure patient safety.

The above key calculations can be implemented by R package `dfrm`.

3.1.3. A product-of-beta prior method (PBP)

In the PBP, a toxicity probability π_j at a dose level j is reparameterized with b_j as follows:

$$b_1 = 1 - \pi_1 \text{ and } b_j = \frac{1 - \pi_j}{1 - \pi_{j-1}} \text{ for } j = 2, \dots, J. \tag{12}$$

A beta prior distribution is independently assumed to each b_j ; that is, $b_j \sim \text{Beta}(\gamma_j, \eta_j)$. Eq. (12) can be converted to $\pi_s = 1 - \prod_{j=1}^s b_j$, where $s \in \{1, \dots, J\}$; therefore, the prior distribution for π_j is called a product-of-beta-prior distribution.

The PBP treats π_j as if it has a beta distribution; that is, $\pi_j \sim \text{Beta}(A_j, B_j)$, although the product-of-betas are not betas themselves. This is because a product of an independent beta distribution is determined by its moments, and a beta approximation for the first and second moments is known to provide good fitting results for a product of an independent beta distribution.

In our simulation study, hyperparameters for π_j were calculated by equations in the original articles [31,32]. Because the hyperparameter calculations needed initial toxicity guesses, they were also generated by an indifference interval approach. Assuming the initial MTD location of the center of the dose range, we generated initial guesses with $\delta = 0.05$ for $\theta = 0.3$ and with $\delta = 0.02$ and 0.03 for $\theta = 0.1$, where δ was the same settings in the CRM and our proposed method.

The MTD was determined as a dose with the closest toxicity to θ based on the final posterior mean of toxicity probabilities.

3.1.4. Toxicity risk approach (WTW)

The WTW deals with a toxicity risk r_j that is a probability that a patient experiences DLT at a dose level j . The r_j is modeled directly and the model assumes r_j is equal to one of a grid of h values ($c_1 < \dots < c_h$). Following the setting that the original article exemplified, we employed $h = 5$ regardless of θ , and then set $\{c_1, \dots, c_5\} = \{0.1, 0.2, 0.3, 0.4, 0.6\}$ for $\theta = 0.3$ and $\{c_1, \dots, c_5\} = \{0.01, 0.05, 0.1, 0.2, 0.3\}$ for $\theta = 0.1$. The setting of $\{c_1, \dots, c_5\} = \{0.1, 0.2, 0.3, 0.4, 0.6\}$ under $\theta = 0.3$ means that, for example, a dose level j is interpreted as very safe, safe, ideal, risky, and toxic, respectively, if r_j is equal to each c_h .

The distributions of a risk r_j are linked by a monotonicity constraint; that is, higher doses have a risk of toxicity greater than or equal to that of lower doses. For a prior distribution of the joint distribution of r_j , a uniform joint prior π^0 respecting the monotonicity constraint was assumed in the simulation (i.e., $\pi^0(d_1, \dots, d_J) = P(r_1 = d_1, \dots, r_J = d_J) = e_0$, where $d_1, \dots, d_J \in \{c_1, \dots, c_h\}$ and $d_1 \leq \dots \leq d_J$, with $\pi^0(d_1, \dots, d_J) = 0$ otherwise). The value of e_0 was calculated so that the joint probability π had to sum to 1 for all combinations of d_j .

A dose maximizing the marginal posterior probability that the toxicity risk was equal to c_3 (i.e., the ‘‘ideal’’ risk) was selected as the next dose during a trial and as the MTD at the end of the trial. As overdose control, the probability of the ‘‘toxic’’ risk was taken into account in the dose selection to judge whether the dose was admissible or not. Only doses that satisfied a condition of $P(r_j = c_5 \mid D_{1:j}) < 0.2$ were included in an admissible dose set for the next dose selection in our simulation study, where 0.2 was the same setting in the original article.

3.2. Simulation results

3.2.1. BO and the other methods

Table 2 shows operating characteristics under $\theta = 0.3$. For observed toxicity, BO shows the lowest toxicity percentages in all the methods under all scenarios due to lower overdose allocations. BO treats less patients with overdose levels than the other methods in most scenarios. In particular, BO successfully controls overdose allocations in scenarios 5 and 6 where the true MTD is at the lowest dose level compared with the other methods due to the effect of τ_1 . BO tends to select safer doses than the other methods owing to overdose control such as an admissible dose set. Also, BO identifies the MTD correctly approximately 10% more than the CRMs (CRM-p (the CRM with a power model) and CRM-l (the CRM with a logistic model)) even if the MTD locates at the end of the dose range (i.e., scenarios 3 and 4).

As shown in Table 2, BO shows higher correct MTD selection probabilities than the CRM-p in all scenarios. Compared between BO and the CRM-l, BO shows higher correct MTD selection probabilities in most scenarios, while scenarios 2 and 7 are comparable results between the two methods, and a lower correct selection probability is shown in BO under scenario 1. Given the difference between the results of the CRMs in scenario 1, a logistic model might more fit this scenario. Also, the fixed MTD location strategy at the center of the dose range is more efficient to address scenarios where the MTD locates on near the center of the dose range than the changeable strategy BO employed. In addition, in scenarios 1 and 2 where the advantages of BO are minimal, initial guesses for the CRMs at around MTD including adjacent doses when parallel-shifted in the vertical direction are overlapped with or close to the true dose-toxicity curves. In such a case, it is highly likely that an estimated curve around the MTD would successfully approach the true dose-toxicity curve by updating the model parameter. In scenarios 3, 4, and 7 where the MTD is close to the highest dose level, the CRMs especially the CRM-p seem to be harder to reach the MTD than BO. The CRMs did not allow dose-escalation at the next cohort when patients in the current cohort experienced DLT. Due to this overdose control, the CRMs could not reach the MTD that locates on higher doses in the dose range as quickly as BO.

A comparison of BO and the PBP regarding correct MTD selection probabilities shows that BO provides lower MTD correct probabilities

Table 2

Operating characteristics under a typical target toxicity rate ($\theta = 0.3$) by each method and scenario (Selection probabilities of MTD determination (correct and overdose selections), average percentages of dose allocations at the MTD and overdoses, and average percentages of observed patients with DLT).

Method	Scenario	MTD determination		Dose allocation (%)		Toxicity (%)	Scenario	MTD determination		Dose allocation (%)		Toxicity (%)
		Correct	Overdose	MTD	Overdose			Correct	Overdose	MTD	Overdose	
CRM-p	1	0.519	0.163	27.6	10.5	21.7	6	0.886	0.114	78.5	21.5	41.3
CRM-l		0.535	0.173	28.3	11.2	22.0		0.884	0.116	78.2	21.9	41.3
PBP		0.512	0.212	29.5	17.2	24.1		0.887	0.113	81.5	18.6	41.6
WTW		0.614	0.173	36.5	13.5	24.1		0.545	0.455	32.8	67.2	47.0
BO		0.520	0.189	24.8	8.9	20.2		0.940	0.060	89.2	10.8	40.5
CRM-p	2	0.590	0.082	30.1	7.8	22.2	7	0.300	0.072	11.0	2.9	16.0
CRM-l		0.603	0.092	31.1	8.2	22.6		0.332	0.082	12.1	3.4	16.3
PBP		0.638	0.073	34.1	11.5	24.4		0.331	0.226	14.7	9.4	18.3
WTW		0.701	0.063	41.7	8.6	24.6		0.385	0.000	13.1	0.0	16.0
BO		0.595	0.097	26.4	6.7	20.6		0.329	0.147	11.3	4.2	15.8
CRM-p	3	0.120	0.000	3.9	0.0	15.7	8	0.609	0.334	44.2	33.6	29.8
CRM-l		0.138	0.000	4.5	0.0	15.9		0.608	0.335	43.8	33.9	29.8
PBP		0.376	0.000	11.8	0.0	17.2		0.627	0.290	43.6	33.1	29.5
WTW		0.000	0.000	0.0	0.0	16.1		0.572	0.387	37.8	43.2	32.0
BO		0.237	0.000	5.8	0.0	15.2		0.695	0.285	48.0	23.2	26.3
CRM-p	4	0.181	0.000	5.6	0.0	14.5	9	0.631	0.258	47.8	29.9	31.3
CRM-l		0.208	0.000	6.3	0.0	14.7		0.628	0.261	47.6	30.3	31.5
PBP		0.487	0.000	14.3	0.0	15.8		0.568	0.260	40.5	33.3	31.3
WTW		0.000	0.000	0.0	0.0	14.9		0.522	0.411	36.4	51.8	37.5
BO		0.312	0.000	7.6	0.0	14.2		0.658	0.232	46.8	19.1	27.4
CRM-p	5	0.662	0.338	59.6	40.4	35.1	10	0.570	0.296	36.3	21.2	24.8
CRM-l		0.660	0.340	59.0	41.0	35.2		0.553	0.319	36.0	22.3	25.0
PBP		0.679	0.321	61.1	38.9	35.1		0.479	0.352	32.7	28.1	26.1
WTW		0.315	0.685	21.9	78.1	42.2		0.537	0.368	37.0	30.5	27.6
BO		0.759	0.241	76.2	23.8	33.4		0.592	0.279	32.5	16.0	22.4

in scenarios 2, 3, and 4, but better or comparable results in the other scenarios. While the correct MTD selection under scenario 7 is comparable between the two methods, the PBP shows a higher probability of overdose selection and higher observed toxicity than BO in this scenario.

The WTW performs very well in one scenario but then provides much poorer performance than the other methods. The WTW seems to be not good at dealing with situations where the highest dose level is the MTD because of the admissible dose criterion (i.e., 0.2) in the current settings. This low value limits dose-escalation under such scenarios 3 and 4, while it seems not to control overdose allocations in the opposite scenarios 5 and 6 where the lowest dose is close to θ . This result implies that the cutoff value on an admissible dose set should be carefully decided.

Table 3 that shows the operating characteristics under a lower target toxicity ($\theta = 0.1$) supports Table 2. BO provides better or comparable results than the CRMs in all scenarios except for scenario 13. The same would be applied for the results in scenario 13 as we explained for scenarios 1 and 2 under $\theta = 0.3$. Also, the same trend shown in scenarios 3, 4, and 7 under $\theta = 0.3$ is provided in scenario 15. Regarding why BO arrives at the true MTD more than the CRMs when the MTD locates near the highest dose, there might be another reason other than the effect of overdose control. BO considers uncertainties of the posterior toxicity probabilities to seek the MTD without concentrating on the current best point. In contrast, the CRMs rely on the exploitation of the estimated dose-toxicity curve. As a result, possible searching space for the next dose candidates might be wider in BO than that in the CRMs. This exploration property seems to perform well in particular under such a scenario in terms of correct MTD selections. Compared with the PBPs that are displayed as PBP (δ) in Table 3, BO performs better than the PBPs in most scenarios. Although concern has been raised that the PBP causes undesirable rigidity when a target toxicity rate is low due to its vague priors [33,47], the PBPs did not get stuck at suboptimal dose level under our simulation settings; however, BO and the CRMs provides better results than the PBPs in most scenarios. In only scenario 11 where the lowest dose level is the true MTD, the PBPs outperforms the other methods. Considering that a toxicity probability at j is composed of the product of the toxicity probabilities of all dose

levels up to $j-1$, it would be reasonable to be able to estimate a toxicity probability at the lowest dose more accurately than the other doses. The WTW provides the best performance in scenarios 12, 13, and 14; however, the correct MTD selection probability for scenario 11 is half as high as the other methods and the overdose selections are occurred with higher percentages during and at the end of the trial.

3.2.2. BO with different settings

It might be reasonable to assume that a toxicity probability increases monotonically with increasing dose levels; however, Table 4 shows the monotonicity restriction on a GP prior (i.e., BO-mono) does not improve the performance compared with BO in terms of correct MTD selections and dose allocations (i.e., Correct selection probabilities and dose allocations of MTD decrease while overdose selection probabilities and overdose allocations increase). The posterior distributions by BO based on Eq. (2) include partially non-monotonically increasing functions; however, posterior distributions, as well as their posterior mean functions, show a monotonically increasing shape owing to follow the trend that true dose-toxicity curves and initial guesses draw. On the other hand, toxicity probabilities in the posterior distribution provided by BO-mono tend to go down in the lower dose range and go up in the higher dose range than BO. That is because the posterior distribution is composed of only functions with a positive slope. As a result, the effect of the monotonicity restriction only skews the posterior distribution but does not improve the performance of BO.

Table 5 evaluates impact of initial toxicity guesses with a different δ on the operating characteristics for BO that uses equation (2). When the slope becomes gentle by using $\delta = 0.03$, initial toxicities at the higher dose range approach θ . The correct selection probability is higher in scenarios where the true MTD locates at the end of the dose range (i.e., scenarios 3 and 4). However, the allocation percentage to overdose levels also increases because of the effect of a gentle slope. In contrast, increasing the slope by using $\delta = 0.07$ lowers correct selection probability in scenarios 3 and 4. On the other hand, BO with $\delta = 0.07$ performs well when the true MTD locates at lower than the middle of the dose range (e.g., scenarios 8 and 9), because the steeper slope makes dose-escalation restrict stronger. For scenarios 5 and 6 where the toxicity at the lowest dose level is equal to or higher than θ , the

Table 3

Operating characteristics under a lower target toxicity rate ($\theta = 0.1$) by each method and scenario (Selection probabilities of MTD determination (correct and overdose selections), average percentages of dose allocations at the MTD and overdoses, and average percentages of observed patients with DLT).

Method	Scenario	MTD determination		Dose allocation (%)		Toxicity (%)	Scenario	MTD determination		Dose allocation (%)		Toxicity (%)
		Correct	Overdose	MTD	Overdose			Correct	Overdose	MTD	Overdose	
CRM-p	3	0.369	0.420	30.5	35.6	10.3	13	0.539	0.145	28.3	16.5	8.0
CRM-l		0.369	0.436	27.3	39.5	10.6		0.547	0.163	28.7	17.1	8.1
PBP(0,02)		0.322	0.307	22.9	27.2	8.6		0.437	0.129	23.3	14.7	7.4
PBP(0,03)		0.331	0.261	22.3	28.6	8.4		0.370	0.092	22.7	14.1	6.8
WTW		0.337	0.577	27.6	50.1	11.3		0.606	0.195	38.5	13.4	7.6
BO		0.409	0.362	25.6	32.7	9.7		0.462	0.054	13.1	17.1	7.7
CRM-p	10	0.464	0.324	35.0	38.7	12.6	14	0.249	0.180	14.0	12.0	6.5
CRM-l		0.456	0.343	35.2	39.2	12.9		0.281	0.184	14.6	12.0	6.6
PBP(0,02)		0.327	0.248	28.3	25.5	10.4		0.212	0.181	11.5	12.2	6.1
PBP(0,03)		0.377	0.204	30.0	25.2	10.4		0.216	0.123	15.3	7.6	5.5
WTW		0.378	0.578	28.4	61.0	15.5		0.383	0.077	17.7	1.9	5.7
BO		0.495	0.272	38.2	27.7	11.8		0.249	0.144	13.4	13.5	6.5
CRM-p	11	0.783	0.217	61.1	38.9	15.5	15	0.276	0.156	15.8	6.5	5.0
CRM-l		0.786	0.214	60.5	39.5	15.7		0.289	0.148	15.3	7.0	5.1
PBP(0,02)		0.863	0.137	75.0	25.1	13.5		0.229	0.162	11.7	10.7	4.9
PBP(0,03)		0.874	0.126	74.1	25.9	13.6		0.223	0.134	13.2	4.1	4.2
WTW		0.430	0.570	27.7	72.3	20.2		0.240	0.000	4.6	0.0	3.8
BO		0.780	0.220	68.5	31.5	14.8		0.320	0.089	14.1	6.4	4.6
CRM-p	12	0.400	0.325	26.9	26.7	9.2						
CRM-l		0.404	0.349	28.9	27.7	9.4						
PBP(0,02)		0.307	0.259	20.2	21.7	7.9						
PBP(0,03)		0.251	0.223	19.4	21.8	7.4						
WTW		0.459	0.378	34.5	29.1	9.4						
BO		0.421	0.212	21.7	23.3	8.7						

Table 4

Differences of the operating characteristics between Bayesian optimization methods with the monotonicity restriction (BO-mono) and without the restriction (BO) under $\theta = 0.3$ (For BO, the results of BO in Table 2 are re-displayed here.)

Scenario	Method	MTD determination		Dose allocation (%)		Toxicity (%)
		Correct	Overdose	MTD	Overdose	
1	BO	0.520	0.189	24.8	8.9	20.2
	BO-mono	0.495	0.190	23.4	9.4	20.2
2	BO	0.595	0.097	26.4	6.7	20.6
	BO-mono	0.589	0.087	25.9	6.0	20.7
3	BO	0.237	0.000	5.8	0.0	15.2
	BO-mono	0.208	0.000	4.9	0.0	15.0
4	BO	0.312	0.000	7.6	0.0	14.2
	BO-mono	0.287	0.000	7.2	0.0	14.0
5	BO	0.759	0.241	76.2	23.8	33.4
	BO-mono	0.698	0.302	72.6	27.4	33.6
6	BO	0.940	0.060	89.2	10.8	40.5
	BO-mono	0.920	0.080	87.0	13.0	41.1
7	BO	0.329	0.147	11.3	4.2	15.8
	BO-mono	0.331	0.104	10.3	3.1	15.7
8	BO	0.695	0.285	48.0	23.2	26.3
	BO-mono	0.674	0.301	47.4	23.1	26.5
9	BO	0.658	0.232	46.8	19.1	27.4
	BO-mono	0.681	0.219	48.0	18.2	28.2
10	BO	0.592	0.279	32.5	16.0	22.4
	BO-mono	0.592	0.277	33.7	15.3	22.6

correct selection probabilities and the average toxicity percentages are not affected by δ . Sensitivity analyses for δ , for example around the optimal $\delta \pm 0.02$, might be needed considering the effect of δ .

4. Discussion and conclusion

In this article, we have evaluated a Bayesian optimization method that applied a Bayesian optimization framework to dose-finding studies. If a theoretical model fits a true dose-toxicity relationship at around the MTD through updating a model parameter, the CRM has good performance. Even in that case, a Bayesian optimization method could have almost comparable performance on correct MTD selections with

lower toxicity percentages than the CRM. On the other hand, if a theoretical model is far or different from a true dose-toxicity curve at around the MTD, the CRM tends to decrease performance. The simulation results show that a Bayesian optimization method provides better performance than the CRM, especially in the latter cases. In general, little is known about dose-toxicity relationships; therefore, a Bayesian optimization method has a potential to provide better results than or at least comparable results to the CRM regardless of the shapes of true dose-toxicity curves. Compared with the other two curve-free methods, a Bayesian optimization method provides more stable results in terms of correct MTD selections. Also, overdose control works successfully

Table 5
Impact of the slope of initial guesses (i.e., $\delta = 0.3, 0.7$) on the operating characteristics of a Bayesian optimization method under $\theta = 0.3$ at each scenario.

Scenario	Delta	MTD determination		Dose allocation (%)		Toxicity (%)
		Correct	Overdose	MTD	Overdose	
1	0.03	0.453	0.299	23.6	15.1	22.2
	0.07	0.438	0.139	21.1	7.4	19.2
2	0.03	0.610	0.110	27.3	9.8	22.9
	0.07	0.505	0.064	23.1	4.9	19.6
3	0.03	0.371	0.000	9.7	0.0	16.1
	0.07	0.141	0.000	4.2	0.0	14.9
4	0.03	0.480	0.000	12.3	0.0	15.1
	0.07	0.201	0.000	5.7	0.0	13.9
5	0.03	0.768	0.232	75.3	24.8	33.4
	0.07	0.762	0.238	76.1	23.9	33.2
6	0.03	0.935	0.065	88.6	11.4	40.7
	0.07	0.923	0.077	88.4	11.6	40.5
7	0.03	0.358	0.288	14.4	8.0	17.2
	0.07	0.260	0.090	8.6	2.8	15.0
8	0.03	0.637	0.343	45.1	26.4	27.1
	0.07	0.700	0.263	49.7	19.9	25.2
9	0.03	0.620	0.263	44.2	21.2	28.0
	0.07	0.688	0.158	47.9	14.6	27.2
10	0.03	0.488	0.376	30.6	22.1	24.2
	0.07	0.570	0.226	32.8	12.7	21.7

because a Bayesian optimization method provides lower or comparable observed toxicity percentages than the other methods in most scenarios.

One of the features in a Bayesian optimization method is to select a dose based on a posterior distribution of dose–toxicity relationships without neglecting its uncertainty. While the benefit of this feature is limited owing to an admissible dose set for overdose control, such restrictions are crucial in dose-finding studies from a safety perspective compared with other areas the Bayesian optimization applies to. A Bayesian optimization method achieves to allocate fewer patients to overdose levels than the CRMs in most scenarios under the overdose control while providing correct MTD selections that are better than or comparable to the CRMs.

Our model has five design parameters (i.e., $\delta, \nu, \sigma_f, \rho, \xi$). Although it is not mandatory to generate initial guesses by using an indifference interval approach, it might be familiar with statisticians who belong to pharmaceutical companies. In this case, δ that determines the slope of initial toxicity guesses can refer to an optimal δ derived by the systematic approach for the CRM. We recommend evaluating at least the range of the optimal $\delta \pm 0.02$ from the view of correct MTD selection probabilities and safe dose allocations. As exemplified in our simulation study, an initial MTD location ν can be selected based on the last-tested dose in the start-up phase or put at the center of the dose range. For the kernel parameters, σ_f is fixed as a value of 1 as mentioned in Section 2.4. The scale parameter ρ should be a value providing less than two turning points in the dose range considering common premises on dose–toxicity relationships. For general use, the same value as the conceptual dose range works well. The value of ξ relies on only computing speed, while a smaller value has less impact on the operating characteristics. In addition, an admissible dose set and MTD determination are defined with four design parameters (i.e., $\tau_1, \tau_2, \epsilon_1, \epsilon_2$). These values are adjusted based on a balance between correct MTD selections and safe dose allocations, and we should also consider clinical perspective as well as a typical setting (e.g., $\epsilon_1 = 0.05$ is often used when $\theta = 0.3$). The calibration approaches for the design parameters are still open discussion, and we should further evaluate them as our future works.

The observed values are generally assumed to follow a normal distribution in typical Bayesian optimization frameworks. Although the simulation results under the normal approximation assumption are

not provided in this article, they generally suggested higher doses than the true MTD with high probability and did not work well in various simulation patterns consistently. It implies that such a normal approximation for patient outcomes is inappropriate in dose-finding studies; hence, we applied the exact distribution for the observed values (i.e., a binomial distribution).

There are other options about acquisition functions instead of $EI(x)$ (e.g., the lower confidence bound criteria [48]). According to our exploratory simulation results that are not shown in this article, the effect of acquisition functions is minimal at least between the expected improvement and the lower confidence bound criteria. In the exploratory simulation, we also evaluated a Bayesian optimization method based on an unit probability mass that mTPI uses for its dose selection as an acquisition function. The strategy based on the unit probability mass was to search for a dose that maximized a probability falling within a small range around θ concerning toxicity distributions (e.g., $\theta \pm 0.05$). It had almost similar results to a Bayesian optimization method based on $EI(x)$; however, we note that the unit probability mass strategy is not a standard approach in the Bayesian optimization.

Although we put monotonically increasing constants on a prior mean function for $m(x)$ and estimate it through a nonparametric approach, there might be another option; for example, $m(x)$ could be modeled by a parametric model that provides monotonically increasing functions. Unlike a Bayesian optimization method, the model misspecification issue still leaves; however, it might be sometimes a reasonable approach when a theoretical model that is likely to fit true dose–toxicity relationships is known. We might extend a Bayesian optimization method by collaborating with such a parametric approach in our future works.

While further discussions and evaluations are required in particular on how to calibrate design parameters, the conclusion of our simulation study is that a Bayesian optimization method has a capability to perform well in terms of correct MTD selections and safe dose allocations even if little information is available about dose–toxicity relationships. In practice, dose selections might be rarely decided by only model information, and observed data is carefully reviewed by clinical experts at pre-specified timings. Illustrations of dose–toxicity relationships such as shown in Section 2.5 will be useful materials in the comprehensive review by clinical experts. While we conducted experimental comparisons, further numerical investigations would reveal more detailed properties for a Bayesian optimization method. The scenarios we addressed in this article assume a single-agent treatment with monotonically increasing dose–toxicity relationships. If unknown functions we would like to know have more complex situations such as two-dimensional inputs or outputs, it is expected that the advantages of nonparametric approaches appear more; therefore, we will address applications of Bayesian optimization frameworks to such complex situations for our future work.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors would like to thank the members of the Tokyo Institute of Technology for their valuable advice that has led to improvements in our study. TS was partially supported by JSPS, Japan Kakenhi (18K19793, 18H03201 and 20H00576), Japan Digital Design, and JST-CREST, Japan.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.conctc.2021.100753>.

References

- [1] Barry E. Storer, Design and analysis of phase I clinical trials, *Biometrics* 45 (3) (1989) 925–937.
- [2] Christophe Le Tourneau, J. Jack Lee, Lillian L. Siu, Dose escalation methods in phase I cancer clinical trials, *J. Natl. Cancer Inst.* 101 (10) (2009) 708–720.
- [3] Mark J. Ratain, Rosemarie Mick, Richard L. Schilsky, Mark Siegler, Statistical and ethical issues in the design and conduct of phase I and II clinical trials of new anticancer agents, *J. Natl. Cancer Inst.* 85 (20) (1993) 1637–1643.
- [4] Ethan Reiner, Xavier Paoletti, John O'Quigley, Operating characteristics of the standard phase I clinical trial design, *Comput. Statist. Data Anal.* 30 (3) (1999) 303–315.
- [5] Anastasia Ivanova, Escalation, group and A + B designs for dose-finding trials, *Stat. Med.* 25 (21) (2006) 3668–3678.
- [6] Sarah Zohar, John O'Quigley, Re: Dose escalation methods in phase I cancer clinical trials, *J. Natl. Cancer Inst.* 101 (24) (2009) 1732–1733.
- [7] John O'Quigley, Margaret Pepe, Lloyd Fisher, Continual reassessment method: A practical design for phase I clinical trials in cancer, *Biometrics* 46 (1) (1990) 33–48.
- [8] Steven Goodman, Marianna L. Zahurak, Steven Piantadosi, Some practical improvements in the continual reassessment method for phase I studies, *Stat. Med.* 14 (11) (1995) 1149–1161.
- [9] Douglas Faries, Practical modifications of the continual reassessment method for phase I cancer clinical trials, *J. Biopharm. Statist.* 4 (2) (1994) 147–164.
- [10] Denis Heng-Yan Leung, You-Gan Wang, An extension of the continual reassessment method using decision theory, *Stat. Med.* 21 (1) (2002) 51–63.
- [11] Guosheng Yin, Ying Yuan, Bayesian model averaging continual reassessment method in phase I clinical trials, *J. Amer. Statist. Assoc.* 104 (487) (2009) 954–968.
- [12] Arzu Onar, Mehmet Kocak, James M. Boyett, Continual reassessment method vs. traditional empirically based design: Modifications motivated by phase I trials in pediatric oncology by the pediatric brain tumor consortium, *J. Biopharm. Statist.* 19 (3) (2009) 437–455.
- [13] Alexia Iasonos, Andrew S. Wilton, Elyn R. Riedel, Venkatraman E. Seshan, David R. Spriggs, A comprehensive comparison of the continual reassessment method to the standard 3 + 3 dose escalation scheme in Phase I dose-finding studies, *Clin. Trials* 5 (5) (2008) 465–477.
- [14] Arzu Onar-Thomas, Zang Xiong, A simulation-based comparison of the traditional method, Rolling-6 design and a frequentist version of the continual reassessment method with special attention to trial duration in pediatric phase I oncology trials, *Contemp. Clin. Trials* 31 (3) (2010) 259–270.
- [15] Revathi Ananthakrishnan, Stephanie Green, Mark Chang, Gheorghe Doros, Joseph Massaro, Michael LaValley, Systematic comparison of the statistical operating characteristics of various phase I oncology designs, *Contemp. Clin. Trials Commun.* 5 (2016) 34–48.
- [16] Philip S. Boonstra, Jincheng Shen, Jeremy M.G. Taylor, Thomas M. Braun, Kent A. Griffith, Stephanie Daignault, Gregory P. Kalemkerian, Theodore S. Lawrence, Matthew J. Schipper, A statistical evaluation of dose expansion cohorts in phase I clinical trials, *J. Natl. Cancer Inst.* 107 (3) (2015) dju429.
- [17] William F. Rosenberger, Linda M. Haines, Competing designs for phase I clinical trials: A review, *Stat. Med.* 21 (18) (2002) 2757–2770.
- [18] Thomas Jaki, Sally Clive, Christopher J. Weir, Principles of dose finding studies in cancer: A comparison of trial designs, *Cancer Chemother. Pharmacol.* 71 (5) (2013) 1107–1114.
- [19] Gareth D. James, Stefan N. Symeonides, Jayne Marshall, Julia Young, Glen Clack, Continual reassessment method for dose escalation clinical trials in oncology: A comparison of prior skeleton approaches using AZD3514 data, *BMC Cancer* 16 (1) (2016) 703.
- [20] Sharon B. Love, Sarah Brown, Christopher J. Weir, Chris Harbron, Christina Yap, Birgit Gaschler-Markefski, James Matcham, Louise Caffrey, Christopher Mckevitt, Sally Clive, Charlie Craddock, James Spicer, Victoria Cornelius, Embracing model-based designs for dose-finding trials, *Br. J. Cancer* 117 (3) (2017) 332–339.
- [21] Larry Z. Shen, John O'Quigley, Consistency of continual reassessment method under model misspecification, *Blometrika* 83 (2) (1996) 395–405.
- [22] Ying Kuen Cheung, Rick Chappell, A simple technique to evaluate model sensitivity in the continual reassessment method, *Biometrics* 58 (3) (2002) 671–674.
- [23] Xavier Paoletti, Andrew Kramar, A comparison of model choices for the continual reassessment method in phase I cancer trials, *Stat. Med.* 28 (24) (2009) 3012–3028.
- [24] Yuan Ji, Ping Liu, Yisheng Li, B. Nebiyou Bekele, A modified toxicity probability interval method for dose-finding trials, *Clin. Trials* 7 (6) (2010) 653–663.
- [25] Yuan Ji, Sue-Jane Wang, Modified toxicity probability interval design: A safer and more reliable method than the 3 + 3 design for practical phase I trials, *J. Clin. Oncol.* 31 (14) (2013) 1785–1791.
- [26] Suyu Liu, Ying Yuan, Bayesian optimal interval designs for phase I clinical trials, *J. R. Stat. Soc. Ser. C. Appl. Stat.* 64 (3) (2015) 507–523.
- [27] Wentian Guo, Sue-Jane Wang, Shengjie Yang, Henry Lynn, Yuan Ji, A Bayesian interval dose-finding design addressing Ockham's razor: MTPI-2, *Contemp. Clin. Trials* 58 (2017) 23–33.
- [28] Heng Zhou, Ying Yuan, Lei Nie, Accuracy, safety, and reliability of novel phase I trial designs, *Clin. Cancer Res.* 24 (18) (2018) 4357–4364.
- [29] Yaqian Zhu, Wei-Ting Hwang, Yimei Li, Evaluating the effects of design parameters on the performances of phase I trial designs, *Contemp. Clin. Trials Commun.* 15 (2019) 100379.
- [30] Bethany J. Horton, Nolan A. Wages, Mark R. Conaway, Performance of toxicity probability interval based designs in contrast to the continual reassessment method, *Stat. Med.* 36 (2) (2017) 291–300.
- [31] Mauro Gasparini, Jeffrey Eisele, A curve-free method for phase I clinical trials, *Biometrics* 56 (2) (2000) 609–615.
- [32] Mauro Gasparini, Jeffrey Eisele, Correction to “A curve-free method for phase I clinical trial” by M. Gasparini and J. Eisele; 56, 609–615, June 2000, *Biometrics* 57 (2) (2001) 659–660.
- [33] Ying Kuen Cheung, On the use of nonparametric curves in phase I trials with low toxicity tolerance, *Biometrics* 58 (1) (2002) 237–240.
- [34] John Whitehead, Helene Thygesen, Anne Whitehead, A Bayesian dose-finding procedure for phase I clinical trials based only on the assumptions of monotonicity, *Stat. Med.* 29 (17) (2010) 1808–1824.
- [35] Niansheng Tang, Songjian Wang, Gen Ye, A nonparametric Bayesian continual reassessment method in single-agent dose-finding studies, *BMC Med. Res. Methodol.* 18 (1) (2018) 172.
- [36] Jonas Mockus, Vytautas Tiesis, Antanas Zilinskas, The application of Bayesian methods for seeking the extremum, in: *Towards global optimisation 2*, Elsevier Science Ltd, North Holland, Amsterdam, 1978, pp. 117–129.
- [37] Jonas Mockus, On Bayesian methods for seeking the extremum, in: *Optimization techniques IFIP technical conference*, Springer, Berlin, Heidelberg, 1975, pp. 400–404.
- [38] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, Nando De Freitas, Taking the human out of the loop: A review of Bayesian optimization, in: *Proceedings of the IEEE*, vol. 104, (1) 2016, pp. 148–175.
- [39] Carl Edward Rasmussen, Christopher K.I. Williams, *Gaussian processes for machine learning*, The MIT Press, London, England, 2006.
- [40] Alberto Russu, Italo Poggesi, Roberto Gomeni, Giuseppe De Nicolao, Bayesian population modeling of phase I dose escalation studies: Gaussian process versus parametric approaches, *IEEE Trans. Biomed. Eng.* 58 (11) (2011) 3156–3164.
- [41] Shing M. Lee, Ying Kuen Cheung, Model calibration in the continual reassessment method, *Clin. Trials* 6 (3) (2009) 227–238.
- [42] Radford M. Neal, Monte carlo implementation of Gaussian process models for Bayesian regression and classification, *Tech. Rep.*, Department of Statistics, University of Toronto, 1997.
- [43] Adam D. Bull, Convergence rates of efficient global optimization algorithms, *J. Mach. Learn. Res.* 12 (88) (2011) 2879–2904.
- [44] Robert B. Gramacy, Herbert K.H. Lee, Optimization under unknown constraints, in: *Bayesian statistics 9*, Oxford University Press, NY, 2011, pp. 229–256.
- [45] Alexia Iasonos, John O'Quigley, Adaptive dose-finding studies: A review of model-guided phase I clinical trials, *J. Clin. Oncol.* 32 (23) (2014) 2505–2511.
- [46] Shing M. Lee, Ying Kuen Cheung, Calibration of prior variance in the Bayesian continual reassessment method, *Stat. Med.* 30 (17) (2011) 2081–2089.
- [47] Ying Kuen Cheung, *Dose finding by the continual reassessment method*, CRC Press, Boca Raton, FL, 2011.
- [48] Niranjan Srinivas, Andreas Krause, Sham M. Kakade, Matthias Seeger, Gaussian process optimization in the bandit setting: No regret and experimental design, in: *ICML'10: Proceedings of the 27th International conference on machine learning*, 2010, pp. 1015–1022.