Published in final edited form as:

J Biopharm Stat. 2020 November 01; 30(6): 964–978. doi:10.1080/10543406.2020.1818251.

Continual Reassessment Method with Regularization in Phase I Clinical Trials

Xiang Li^a, Anastasia Ivanova^b, Hong Tian^a, Pilar Lim^c, Kevin Liu^d

^aStatistics and Decision Sciences, Janssen Research & Development, LLC, Raritan, NJ, USA

^bDepartment of Biostatistics, University of North Carolina at Chapel Hill, NC, USA

^cStatistics and Decision Sciences, Janssen Research & Development, LLC, Titusville, NJ, USA

^dGenmab, Princeton, NJ, USA

Abstract

Many Phase I trial designs have been developed to improve upon the standard 3 + 3 design. These designs can be classified as long-memory designs, for example, the continual reassessment method (CRM), and short-memory designs such as the modified toxicity probability interval (mTPI) design. Long-term memory designs use all data but their performance can be negatively affected by the model misspecification. Short-term memory designs only use data at the current dose and might lose efficiency as a result. To overcome these issues, we propose a regularized CRM (rCRM). The rCRM offers a trade-off between long-term memory and short-term memory methods. The rCRM gives more weight to data obtained at the doses with the estimated probability of toxicity closer to the target toxicity rate. The addition of a regularization term has an effect of shrinking the dimension of the model and leads to improved performance of the 2-parameter CRM. The rCRM is a good design choice to guide assignments in an expansion cohort phase of a dose-finding trial since dose assignments do not seem to change as often as in corresponding CRMs.

Keywords

continual reassessment method; dose-expansion; dose-finding; regularization

1. Introduction

The primary goal of Phase I clinical trials is to find the maximum-tolerated dose (MTD), the dose with a certain probability of dose-limiting toxicity (DLT). The standard 3+3 design is the most frequently used design in Phase I oncology studies due to its simplicity. It has been widely criticized for its inflexibility due to a built-in stopping rule that often results in insufficient number of patients to estimate the MTD with desired precision (Iasonos, Wilton, Riedel, Seshan, & Spriggs, 2008; Nie et al., 2016). More flexible designs have been proposed allowing for including as many patients as needed to estimate the MTD with

desired precision. Up-and-down designs, such as cumulative cohort design (Ivanova, Flournoy, & Chung, 2007), t-statistic design (Ivanova & Kim, 2009), mTPI (Ji & Wang, 2013) and mTPI-2 (Guo, Wang, Yang, Lynn, & Ji, 2017), are flexible alternatives to the 3+3 design, as they can be used with any desired sample size. The continual reassessment method (CRM) (O'Quigley & Shen, 1996) is a model based design that estimate a dose-toxicity curve using 1- or 2-parameter model. It can be more efficient compared to up-and-down designs because it uses all data to make the next assignment, not only data from the current dose.

In Phase I clinical trials, subjects are assigned sequentially starting from a very low dose. In the CRM, the dose-toxicity curve is estimated based on all data. Data collected at lower doses with no or few DLTs can bias the estimation of the dose-toxicity curve since the CRM works the best when assignments are close to the true MTD. Estimates in the CRM can be also influenced by DLTs observed at doses above the true MTD. This boundary effect is also related to the “winner-takes-all” paradigm (Oron & Hoff, 2013), and may be worse due to small sample sizes in Phase I trials. Most CRM models employ Bayesian approach with a prior distribution chosen to yield good operational characteristics of the design. Oron and Hoff (2013) discussed a substantial impact of prior distribution on the dose-finding performance. Up-and-down designs are short-memory or memoryless (O'Quigley & Zohar, 2006) and make the recommendation of next dose based on the data at the current dose only. Not using all the information, however, may result in imprecise estimation of the MTD.

To overcome these issues, we propose a regularized CRM (rCRM) that adds a constraint on the model parameters. This technique has been widely used for feature selection in the “small n large p ” (Friedman, Hastie, & Tibshirani, 2010; Li, Xie, Zeng, & Wang, 2018). The use of adaptive weight in the regularization provides differential regularization degrees for different dose levels. The dose levels with the probability of DLT closer to the target rate are given more weight in the model estimation, therefore diminishing the impact of the “winner-takes-all” paradigm. The regularization provides a general framework and could be easily incorporated to other dose-finding models to improve the performance.

The rest of this article is organized as follows. In Section 2, we present the details of the proposed rCRM. We provide the estimation algorithm for 1-parameter and 2-parameter rCRM models. In Section 3, we compare CRM and rCRM in a simulation study. We present two hypothetical examples with expansion cohort in Section 4 and concluding remarks in Section 5.

2. Methodology and Computational Algorithm

We consider the dose selection within n_D pre-specified dose levels, $\mathcal{D} = \{D_j: j = 1, \dots, n_D\}$. Denoted by y_i the binary DLT status of subject i . The probability of DLT at the dose $d_i \in \mathcal{D}$ for subject i , is given as a function of parameter θ (θ could be a scalar or vector):

$$y_i \sim \text{Bernoulli}(\pi(d_i | \theta)), \quad i = 1, \dots, n,$$

and the likelihood function is given by

$$l_n(\theta) = n^{-1} \sum_{i=1}^n (y_i \log(\pi(d_i | \theta)) + (1 - y_i) \log(1 - \pi(d_i | \theta))), \quad (1)$$

where the function $\pi(d_i | \theta)$ characterizes the dose-toxicity curve. The long-term and short-term memory methods use different approaches to estimate this function: the long-term memory design models the $\pi(d_i | \theta)$ using all the collected data, while the short-term memory design estimates it based on the data at the dose d_i only, e.g. the percentage of DLT at the dose d_i .

In dose escalation studies, subjects are enrolled in sequential cohorts with size of m , which makes the model estimation complicated due to very limited data available for the first few cohorts. The convergence of model estimation could be an issue. Bayesian CRM is commonly used to resolve this issue by incorporating prior information. The prior is carefully constructed based on previous studies, relevant external data, or expertise knowledge. However, the question on the use of prior remains with an untestable assumption of exchangeability, especially for Phase I trials, and the prior distribution could have a substantial impact on the dose-finding performance (Oron & Hoff, 2013).

An alternative approach is to incorporate regularization to handle the small sample size. We denote γ (e.g., $\gamma = 0.3$) as the target toxicity probability. Our proposed regularized CRM (rCRM) model is given by

$$\min -l_n(\theta) + \lambda \sum_{j=1}^{n_D} P_j(|g(\pi(D_j | \theta)) - g(\gamma)|), \quad (2)$$

where λ is a non-negative tuning parameter and $P_j(x)$ is a convex function satisfying $P_j(0) = 0$ and $P_j(x) > 0$. The penalty term $P_j(x)$ enforces some similarity between the probability of DLT at dose D_j and the target rate γ , and the tuning parameter λ reflects the degree of similarity. A large value of λ gives more weight on the penalty term and the estimation of $\pi(D_j | \theta)$ would be closer to γ . The likelihood-based CRM model is a special case when λ is set at zero. The $g(\cdot)$ is a monotonic function and could be selected for ease of calculation. For example, we use the logarithm function for the 1-parameter CRM (see Section 2.1) and logit function for the 2-parameter CRM (see Section 2.2). The selection of λ will be discussed in Section 2.3.

Remark: For computation purposes, we recommend using the squared difference between the probability of DLT and the target rate as the penalty term, given by

$$P_j = w_j (g(\pi(D_j | \theta)) - g(\gamma))^2, \quad (3)$$

where

$$w_j = 1 / (g(\pi_0(D_j)) - g(\gamma))^2$$

and $\pi_0(D_j)$ is some good estimate. In the simulation studies, we are using the maximum likelihood estimate (MLE) from (1) based on the cumulative subjects, i.e. 1-parameter or 2-parameter CRM model without regularization. The $\pi_0(D_j)$ could also be obtained from simple probability calculation at each dose level, some prior information/expertise knowledge or other CRM models. The weight w_j is calculated based on the difference with the target rate and provides differential regularization. More weight is given to the dose levels that are closer to the target rate, therefore addressing the “winner-takes-all” issue in the long-memory method.

2.1. Algorithm for 1-Parameter Power Model

A skeleton needs to be pre-specified for each dose level in order to use the 1-parameter power model and could be obtained using the method proposed by Lee and Cheung (2009). Denote by p_i the skeleton value for dose d_i for subject i . The probability of DLT at the dose d_i for subject i is given by

$$\pi_1(d_i) = p_i^a, \quad \text{subject to } a > 0,$$

and the log-likelihood function is written as

$$l_{n1}(a) = n^{-1} \sum_{i=1}^n y_i \log p_i^a + (1 - y_i) \log(1 - p_i^a).$$

For the non-linear structure of the log-likelihood function, we form a quadratic approximation by second order Taylor expansion at the current estimates, $\tilde{a}$, which is

$$l_{Q1}(a) = -\frac{1}{2n} \sum_{i=1}^n c_i (a - \tilde{a} - b_i)^2$$

up to a constant, where

$$c_i = \frac{q_i^2 \tilde{\pi}_1(d_i)(1 - y_i)}{(1 - \tilde{\pi}_1(d_i))^2}, \quad b_i = \frac{(1 - \tilde{\pi}_1(d_i))(y_i - \tilde{\pi}_1(d_i))}{q_i \tilde{\pi}_1(d_i)(1 - y_i)},$$

$q_i = \log p_i$ and $\tilde{\pi}_1(d_i)$ is evaluated at the current estimates. We recommend using the logarithm function for function $g(\cdot)$. By incorporating the regularization term, the objective function is given by

$$\min \frac{1}{2n} \sum_{i=1}^n c_i (a - \tilde{a} - b_i)^2 + \lambda \sum_{j=1}^{nD} w_j (a \log P_j - \log \gamma)^2, \quad (4)$$

where P_j is the skeleton value for dose level D_j . For each given value of λ , the objective function (4) could be solved iteratively until convergence.

When there is no heterogeneity of responses, the estimate of α is bounded and set to a fixed value (e.g. 10 when no DLT).

2.2. Algorithm for 2-Parameter Logistic Model

Based on the 2-parameter logistic model, the probability of DLT at the dose d_i for subject i is given by

$$\pi(d_i) = \frac{\alpha d_i^\beta}{1 + \alpha d_i^\beta}, \quad \text{subject to } \alpha > 0, \beta > 0,$$

and the log-likelihood function is written as

$$l_{n2}(\alpha, \beta) = n^{-1} \sum_{i=1}^n y_i (\log \alpha + \beta \log d_i) - \log(1 + \alpha d_i^\beta) \quad (5)$$

For the non-linear structure of the log-likelihood function, we form a quadratic approximation by second order Taylor expansion at the current estimates, $\tilde{\gamma}$, $\tilde{\beta}$, which is

$$l_{Q2}(\eta, \beta) = -\frac{1}{2n} \sum_{i=1}^n \tau_i (z_i - \eta - \beta \log d_i)^2,$$

up to a constant, where

$$\tau_i = \tilde{\pi}(d_i)(1 - \tilde{\pi}(d_i)), \quad z_i = \frac{y_i - \tilde{\pi}(d_i)}{\tilde{\pi}(d_i)(1 - \tilde{\pi}(d_i))} + \tilde{\gamma} + \tilde{\beta} \log d_i$$

and $\tilde{\pi}(d_i)$ is evaluated at the current estimates. We recommend using the logit function for function $g(\cdot)$. By incorporating the regularization term, the objective function to be minimized is given by

$$\min \frac{1}{2n} \sum_{i=1}^n \tau_i (z_i - \eta - \beta \log d_i)^2 + \lambda \sum_{j=1}^{nD} w_j (\eta + \beta \log D_j - \text{logit}(\gamma))^2, \quad (6)$$

For each given value of λ , we use the coordinate descent algorithm to solve the objective function (6).

When there is no heterogeneity in response, we restrict the estimate of β to a very small positive value and set α to a fixed value (e.g. small value when there was no DLT).

2.3. Selection of Tuning Parameter

The regularization provides a solution path along with the λ value, and the optimal solution is determined based on minimizing the prediction error using cross-validation. For the

tuning parameter, we suggest searching for a grid of values. The grid points need to be able to cover a wide range of possible values. For example, we construct a sequence of K values of λ decreasing from 10^3 to 10^{-2} on the log scale. Due to the small sample size, we suggest using the leave-one-out method. Given a fixed λ , we use the data excluding subject i to estimate the model parameter $\hat{\theta}_{-i}$ and then evaluate the likelihood of the data from subject i under the estimated model, $\hat{\pi}_{-i}(d_i | \lambda) = \pi(d_i | \hat{\theta}_{-i}, \lambda)$, which is

$$y_i \log(\hat{\pi}_{-i}(d_i | \lambda)) + (1 - y_i) \log(1 - \hat{\pi}_{-i}(d_i | \lambda)). \quad (7)$$

We select the optimal tuning parameter $\hat{\lambda}$ with the largest prediction deviance,

$$\hat{\lambda} = \underset{\lambda}{\operatorname{argmax}} \sum_{i=1}^n y_i \log(\hat{\pi}_{-i}(d_i | \lambda)) + (1 - y_i) \log(1 - \hat{\pi}_{-i}(d_i | \lambda)), \quad (8)$$

and report the corresponding estimates $\hat{\theta}$. Note the estimate based on the likelihood-based CRM also lies on the solution path (when $\lambda = 0$).

Remark: In the Bayesian point of view, the penalty term looks like a prior distribution on the $\pi(D_j | \theta)$, which is proportional to $\exp(-\lambda p(|\logit(\pi(D_j | \theta)) - \logit(\gamma)|))$ with a precision parameter λ . However, a major difference between the proposed rCRM and Bayesian CRM is that the penalty term changes with each new data point in rCRM as it's a function of unknown parameters, while most Bayesian models assume a prior distribution on the model parameter θ (Neuenschwander, Branson, & Gsponer, 2008), which does not change with data. Imposing a constraint (3) on the $\pi(D_j | \theta)$ allows us to focus on the dose levels around the MTD and benefit from the advantages of the short-memory method. Another difference is that the evaluation of rCRM is based on the likelihood function alone while the Bayesian CRM is based on the posterior distribution. Our proposed model framework could be readily extended to incorporate any Bayesian prior on the model parameter. Assuming θ has a prior Gaussian distribution with mean μ_θ and covariance matrix Σ , we can write the model as

$$\min -l_n(\theta) + \lambda \sum_{j=1}^{nD} P_j(|g(\pi(D_j | \theta)) - g(\gamma)|) + \rho(\theta - \mu_\theta)^T \Sigma^{-1}(\theta - \mu_\theta)/2.$$

where ρ is another tuning parameter that reflects how much prior information the data likelihood can borrow from. The Bayesian CRM is a special case when $\rho = 1$ and $\lambda = 0$. In the simulation studies, the prior is not incorporated by specifying $\rho = 0$ in the proposed rCRM as the focus is on the regularization term.

2.4. Decision Rule

The estimation of the probability $\pi(d_j)$ is updated every time with new cohort data, and the dose level with the estimated probability of DLT closest to the target probability, e.g., $\gamma = 0.3$, is selected, which is given by

$$\hat{d} = \arg \min_{d \in \mathcal{D}} |\hat{\pi}(d) - \gamma|. \quad (9)$$

We follow the modified CRM and do not skip any dose in the recommendation of next dose. The MTD is also determined based on the rule (9) with all the data.

3. Simulation Studies

In the simulation studies, we compared the proposed rCRM models with 1-parameter and 2-parameter CRM models without regularization. For implementation of the proposed models, we have developed main routines in C++ with an R interface via R-package ‘‘Rcpp’’ (Eddelbuettel et al., 2011). The R-package ‘‘rCRM’’ will be available in CRAN soon. We followed the same default setting for the 1-parameter CRM as in a review paper (Horton, Wages, & Conaway, 2017). The skeleton was specified using the `getprior` (Lee & Cheung, 2009) function within the `dfcrm` R-package with a halfwidth of 0.06 and prior guess of MTD at dose 2. We used this setting for all the scenarios in the simulation. We also used this skeleton in the regularized approach and in the calculation of weight. In addition, we calculated the 2-parameter CRM model using the 2-parameter rCRM with $\lambda = 0$. We also considered the non-parametric optimal design (O’Quigley, Paoletti, & Maccario, 2002; Wages, Conaway, & O’Quigley, 2013) as a benchmark using the R codes available on the website http://www.faculty.virginia.edu/model-based_dose-finding/.

To assess the performance, we calculated the percentage of the final recommended dose (under target dose, on target dose or over target dose). To assess the model performance by incorporating the information at all dose levels, we calculated the accuracy index (Cheung, 2011) based on the selected dose level, which is given by

$$A = 1 - n_D \times \frac{\sum_{j=1}^{n_D} \rho_j * \text{Percentage of selecting dose level } D_j}{\sum_{j=1}^{n_D} \rho_j},$$

and the assignment index based on the percentage of subjects assigned, which is given by

$$B = 1 - n_D \times \frac{\sum_{j=1}^{n_D} \rho_j * \text{Percentage of subjects assigned to dose level } D_j}{\sum_{j=1}^{n_D} \rho_j},$$

where $\rho_j = |\pi_j - \gamma|$ and π_j is the true toxicity probability at dose level j . These two measures provide an assessment of model performance in terms of the accuracy of the MTD recommendation and accuracy of dose assignments. The assignment index is higher if there are many assignments at or around the true MTD. It penalizes assignments to doses further from the true MTD. Since the experimentation starts at the lowest dose, the assignment index also measures how fast the dose escalation reaches the true MTD.

We set the target probability at $\gamma = 0.3$ and the maximum sample sizes at 30 and 48 with cohort sizes of 1 and 3. In total, there are 4 combinations of cohort size and maximum sample size considered in the simulation, i.e. (1,30), (3,30), (1,48) and (3,48). Eighteen dose-toxicity scenarios (see Figure 1) were considered with six dose levels and with steep (scenarios 1 – 6), intermediate (scenarios 7 – 12) and flat (scenarios 13 – 18) curves. The MTD was set at various dose levels. We performed 10000 simulations for each scenario. For each scenario, subjects' responses were pre-generated. Simulations of all the methods, for each scenario, were run on the same set of pre-generated toxicity responses. Experimentation started at the lowest dose level and continued until the total sample size was reached for the study.

In the display of simulation results, the 1-parameter and 2-parameter CRM models are indicated as “1p CRM” and “2p CRM”, respectively; the 1-parameter rCRM based on weight calculated from 1-parameter and 2-parameter CRM models are indicated as “1p/1p rCRM” and “1p/2p rCRM”, respectively; the 2-parameter rCRM based on weight calculated from 1-parameter and 2-parameter CRM models are indicated as “2p/1p rCRM” and “2p/2p rCRM”, respectively; the optimal design is indicated as “Optimal”.

The simulation results averaged over 18 scenarios are summarized in Table 1 with the total sample size of 30 and 48. The results for each scenario are provided in the appendix. The 2p/1p rCRM, 1p/1p rCRM and 1p CRM had similar performance in terms of the correct MTD selection and assignment, and better than other methods. In terms of the accuracy index, the 2p/1p rCRM was the closest to the optimal design, followed by 1p/1p rCRM and 1p CRM. Similar for the assignment index. The performance of the 2-parameter CRM was improved for all four measurements by using regularization.

4. Examples with Expansion Cohort

Phase I studies often have two parts, dose escalation followed by dose expansion. The MTD is first selected based on a small number of subjects in the dose escalation part. More subjects are then enrolled at the estimated MTD in the expansion part, to further confirm the safety profile of the MTD and gain early evidence of drug effectiveness (Food, Administration, et al., 2018; Iasonos & O'Quigley, 2016). Because the MTD estimate after the dose-escalation part might not be precise, one needs to monitor toxicity during the expansion part to decrease (or increase) the dose if necessary. Frequent changes of a dose during the expansion part are not desirable since the goal is to collect efficacy data at the estimated MTD. Below we illustrate the use of the CRM or rCRM in the expansion part.

Here, we monitor toxicity with information from dose expansion and update the MTD estimate use the CRM or rCRM. The true dose-toxicity curve is the same as in Scenario 3 with the true toxicity probabilities (0.01, 0.09, 0.30, 0.55, 0.73, 0.84) and the true MTD at dose 3. In the dose escalation, we used a stopping rule and the dose with at least 6 subjects treated was recommended as the MTD. In the dose expansion, a similar stopping rule was used and the expansion was stopped after 10 subjects were assigned to the same dose. The MTD was re-estimated using the CRM (or rCRM) after each subject based on all available data, that is, combined dose-escalation and dose-expansion data.

We illustrate the performance of the methods with two examples: (1) the dose escalation ended on the wrong dose and recommended dose 4, i.e. right above the MTD; (2) the correct MTD was selected the dose escalation portion.

Figure 2 summarizes the dose assignments based on the 1p CRM and 2p/1p rCRM models. Because no DLTs observed at the true MTD, dose 3, and 2 DLTs were observed at dose 4 out of 6 patients, both CRMs recommended the wrong MTD, dose 4. In the expansion, both models de-escalated to the correct dose 3 after observing the first DLT. Though both models made the correct recommendation of MTD at the end of expansion, the CRM jumped down to dose 2 four times while the rCRM only de-escalated to dose 2 one time, which may suggest less fluctuation using rCRM.

Figure 3 summarizes the dose assignments when both models recommended the correct MTD in the dose escalation. The CRM still changed the dose several times in the expansion and fluctuated around the correct MTD. However, the rCRM stayed at the MTD and did not change the dose. The result may suggest that the rCRM does not change the dose as often as the regular CRM model. This is advantageous if we use rCRM to guide assignments in the expansion cohort phase when we don't want to change the dose unless absolutely necessary.

5. Discussion

In this article, we proposed a regularized CRM (rCRM), the CRM with a penalty function proportional to the difference between the estimated and the target toxicity probabilities. The use of an adaptive penalty term gives more weight to dose levels that are closer to the MTD. The proposed rCRM adaptively seeks a trade-off between long-memory and short-memory methods in a data-driven fashion. The degree of regularization depends on the data and is selected by using cross-validation. Our simulation results show that the proposed 2-parameter rCRM with weight calculated based on 1-parameter CRM outperforms the other methods achieving comparable percentage of selecting and assigning the correct MTD but higher accuracy and assignment Indices based on the average across all the scenarios.

The use of regularization has an effect of shrinking the dimensions of the model and improves the performance of the 2-parameter CRM. The 2-parameter CRM has been shown to be rigid in some scenarios and it's challenging to set up a good prior to avoid it being stuck at a dose (Cheung, 2011). This leads to comparable or better performance of the 2-parameter rCRM on average compared to 1- and 2-parameter CRMs and other rCRMs.

Regularization gives more weight to dose levels that are closer to the estimated target dose, and therefore the proposed methods may not change the dose as often as an original CRM. This is advantageous if we use rCRM to guide assignments in the expansion cohort phase when we do not want to change the dose unless absolutely necessary.

The proposed model provides a general and flexible framework, that can be further customized based on project needs. The idea of using a regularization could be readily extended to other model-based approaches, and also be used for other endpoints (e.g. time-to-event) or a combination of safety and efficacy endpoints.

Acknowledgements

Ivanova's work was supported in part by the NIH grant P01 CA142538. The authors thank anonymous reviewers for their helpful comments.

Appendix A.: Simulation Results for Each of the 18 Scenarios

Figure A1.

Summary of the MTD recommendation (on target dose [top], under target dose [middle], and over target dose [bottom]) for cohorts of 3 and the total sample size of 48

Figure A2.

Summary of the MTD recommendation (on target dose [top], under target dose [middle], and over target dose [bottom]) for cohorts of 1 and the total sample size of 48

Figure A3.

Summary of the MTD recommendation (on target dose [top], under target dose [middle], and over target dose [bottom]) for cohorts of 3 and the total sample size of 30

Figure A4.

Summary of the MTD recommendation (on target dose [top], under target dose [middle], and over target dose [bottom]) for cohorts of 1 and the total sample size of 30

Figure A5.

The accuracy index for cohort sizes of 3 (top) and 1 (bottom) and the total sample size of 30

Figure A6.

The accuracy index for cohort sizes of 3 (top) and 1 (bottom) and the total sample size of 48

Figure A7.

The assignment index for cohort sizes of 3 (top) and 1 (bottom) and the total sample size of 30

Figure A8.

The assignment index for cohort sizes of 3 (top) and 1 (bottom) and the total sample size of 48

Figure A9.

Proportion of subjects assigned to the MTD for cohort sizes of 3 (top) and 1 (bottom) and the total sample size of 30

Figure A10.

Proportion of subjects assigned to the MTD for cohort sizes of 3 (top) and 1 (bottom) and the total sample size of 48

References

- Cheung YK (2011). Dose finding by the continual reassessment method. CRC Press.
- Eddelbuettel D, François R, Allaire J, Ushey K, Kou Q, Russel N, ... Bates D (2011). Rcpp: Seamless r and c++ integration. *Journal of Statistical Software*, 40(8), 1–18.
- Food U, Administration D, et al. (2018). Expansion cohorts: use in first-in-human clinical trials to expedite development of oncology drugs and biologics guidance for industry. FDA Maryland.
- Friedman J, Hastie T, & Tibshirani R (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1), 1. [PubMed: 20808728]
- Guo W, Wang S-J, Yang S, Lynn H, & Ji Y (2017). A bayesian interval dose-finding design addressing cockham's razor: mtpi-2. *Contemporary clinical trials*, 58, 23–33. [PubMed: 28458054]

- Horton BJ, Wages NA, & Conaway MR (2017). Performance of toxicity probability interval based designs in contrast to the continual reassessment method. *Statistics in medicine*, 36(2), 291–300. [PubMed: 27435150]
- Iasonos A, & O’Quigley J (2016). Dose expansion cohorts in phase i trials. *Statistics in biopharmaceutical research*, 8(2), 161–170. [PubMed: 27516848]
- Iasonos A, Wilton AS, Riedel ER, Seshan VE, & Spriggs DR (2008). A comprehensive comparison of the continual reassessment method to the standard 3+ 3 dose escalation scheme in phase i dose-finding studies. *Clinical Trials*, 5(5), 465–477. [PubMed: 18827039]
- Ivanova A, Flournoy N, & Chung Y (2007). Cumulative cohort design for dose-finding. *Journal of Statistical Planning and Inference*, 137(7), 2316–2327.
- Ivanova A, & Kim SH (2009). Dose finding for continuous and ordinal outcomes with a monotone objective function: a unified approach. *Biometrics*, 65(1), 307–315. [PubMed: 18479486]
- Ji Y, & Wang S-J (2013). Modified toxicity probability interval design: a safer and more reliable method than the 3+ 3 design for practical phase i trials. *Journal of Clinical Oncology*, 31(14), 1785. [PubMed: 23569307]
- Lee SM, & Cheung YK (2009). Model calibration in the continual reassessment method. *Clinical Trials*, 6(3), 227–238. [PubMed: 19528132]
- Li X, Xie S, Zeng D, & Wang Y (2018). Efficient ℓ_0 -norm feature selection based on augmented and penalized minimization. *Statistics in medicine*, 37(3), 473–486. [PubMed: 29082539]
- Neuenschwander B, Branson M, & Gsponer T (2008). Critical aspects of the bayesian approach to phase i cancer trials. *Statistics in medicine*, 27(13), 2420–2439. [PubMed: 18344187]
- Nie L, Rubin E, Mehrotra N, Pinheiro J, Fernandes L, Roy A, ... de Alwis D (2016). Rendering the 3+ 3 design to rest: More efficient approaches to oncology dose-finding trials in the era of targeted therapy. *Clinical cancer research: an official journal of the American Association for Cancer Research*, 22(11), 2623. [PubMed: 27250933]
- O’Quigley J, Paoletti X, & Maccario J (2002). Non-parametric optimal design in dose finding studies. *Biostatistics*, 3(1), 51–56. [PubMed: 12933623]
- O’Quigley J, & Shen LZ (1996). Continual reassessment method: a likelihood approach. *Biometrics*, 673–684. [PubMed: 8672707]
- O’Quigley J, & Zohar S (2006). Experimental designs for phase i and phase i/ii dose-finding studies. *British Journal of Cancer*, 94(5), 609. [PubMed: 16434987]
- Oron AP, & Hoff PD (2013). Small-sample behavior of novel phase i cancer trial designs. *Clinical Trials*, 10(1), 63–80. [PubMed: 23345304]
- Wages NA, Conaway MR, & O’Quigley J (2013). Performance of two-stage continual reassessment method relative to an optimal benchmark. *Clinical Trials*, 10(6), 862–875. [PubMed: 24085776]

Figure 1.
Summary of simulation scenarios

Figure 2.
Example of dose expansion starting from a wrong MTD selected in the dose escalation

Figure 3.
Example of dose expansion starting from a correctly selected MTD in the dose escalation

Table 1.

Average over 18 scenarios of the proportion of subjects assigned to the true MTD (Correct MTD Selection), the correct selection measured by the accuracy index (Accuracy Index), the proportion of subjects assigned to the true MTD (Correct MTD Assignment), and the correct dose assignment measured by the assignment index (Assignment Index). Each average is given for a trial with the total sample size of 30 and 48. The best method (excluding the optimal design) is shown in bold.

Sample Size	1p CRM	1p/1p rCRM	1p/2p rCRM	2p CRM	2p/1p rCRM	2p/2p rCRM	Optimal
Correct MTD Selection							
30	0.45	0.45	0.42	0.32	0.46	0.39	0.50
48	0.50	0.50	0.47	0.34	0.51	0.43	0.56
Accuracy Index							
30	0.46	0.47	0.43	0.29	0.48	0.40	0.49
48	0.53	0.53	0.49	0.31	0.55	0.44	0.57
Correct MTD Assignment							
30	0.33	0.33	0.32	0.28	0.33	0.30	
48	0.39	0.39	0.37	0.30	0.39	0.35	
Assignment Index							
30	0.29	0.29	0.29	0.21	0.31	0.27	
48	0.36	0.37	0.35	0.24	0.39	0.33	