

論文 / 著書情報
Article / Book Information

題目(和文)	用量探索試験におけるベイズ最適化デザイン
Title(English)	Bayesian Optimization Design for Dose-Finding Studies
著者(和文)	高橋亜実
Author(English)	Ami Takahashi
出典(和文)	学位:博士(理学), 学位授与機関:東京工業大学, 報告番号:甲第12060号, 授与年月日:2021年9月24日, 学位の種類:課程博士, 審査員:金森 敬文,高安 美佐子,三好 直人,渡邊 澄夫,中野 張,鈴木 大慈
Citation(English)	Degree:Doctor (Science), Conferring organization: Tokyo Institute of Technology, Report number:甲第12060号, Conferred date:2021/9/24, Degree Type:Course doctor, Examiner:,,,,,
学位種別(和文)	博士論文
Type(English)	Doctoral Thesis

Doctoral Thesis

Bayesian Optimization Design for Dose-Finding Studies

(用量探索試験におけるベイズ最適化デザイン)

Ami Takahashi

Supervisor Taiji Suzuki
Takafumi Kanamori

July, 2021

Department of Mathematical and Computing Science
School of Computing
Tokyo Institute of Technology

Abstract

Dose-finding studies primarily aim to identify a maximum tolerated dose or an optimal dose that determines a therapeutic dose for subsequent studies. For cancer treatment, chemotherapies with single cytotoxic agents still remain fundamental depending on cancer types. On the other hand, drug development for combination and targeted therapies has become increasingly commonplace than that for chemotherapies because they are potentially more effective and less toxic. Considering room for improvement on even statistical designs dealing with chemotherapies and no standard solutions for dose findings of combination and targeted therapies, we propose a Bayesian optimization design for identifying target doses in three types of dose-finding studies for oncology: (1) monotherapies with cytotoxic agents, (2) combination drug therapies, and (3) targeted therapies with biological agents. Our simulation-based evaluations compared with some popular statistical designs suggest that our design works successfully on the selections of target doses and safe dose allocations.

Contents

Chapter 1	Introduction	1
1.1	Drug development and statistical methodologies	1
1.1.1	Drug development process	1
1.1.2	Expectation of exploiting Bayesian approaches in clinical trials	2
1.2	Motivation and contribution	3
1.3	Roadmap	5
Chapter 2	Transition and problems of statistical designs for dose-finding studies	9
2.1	Common procedures in dose-finding studies	9
2.2	The earliest statistical design	10
2.3	Parametric model-based designs	12
2.4	Nonparametric designs	14
2.4.1	Toxicity probability interval designs	14
2.4.2	Curve-free designs	16
2.5	Preparation for leveraging Bayesian optimization	18
2.5.1	Typical Bayesian optimization	19
Chapter 3	Bayesian optimization design of MTD estimation for mono-therapies	21
3.1	Basic structure of Bayesian optimization design	21
3.1.1	Statistical model for dose–toxicity relationships	22
3.1.2	Dose selection strategy	23
3.1.3	Overdose control for patient safety	24
3.1.4	Implementation steps	25
3.1.5	Example with illustrations	27
3.2	Performance evaluation	28
3.2.1	Simulation settings	28
3.2.2	Simulation results	32
3.3	Summary and consideration	40
A	Appendices for Chapter 3	43
A.1	Toxicity risk approach	43
A.2	Initial idea for expanding BOD to address continuous toxicity variables	44
A.3	BOD simulation code for MTD estimation	46
Chapter 4	Bayesian optimization design of MTDC estimation for combination drug therapies	53
4.1	Overview of drug development and statistical designs for dose-finding studies of combination drug therapies	53
4.1.1	Dose-finding studies for combination drug therapies	53
4.1.2	MTDC definition	54
4.1.3	Bayesian optimal interval design	54
4.1.4	Partial ordering continual reassessment method	57

4.2	Modeling frameworks and practical procedures of Bayesian optimization design for MTDC estimation	59
4.2.1	Statistical modeling of dose-toxicity relationships and dose selection strategy	60
4.2.2	BOD implementation steps for MTDC estimation	61
4.2.3	Illustration and example	64
4.3	Performance evaluation	65
4.3.1	General settings for the simulation study	66
4.3.2	Simulation settings of BOD for MTDC estimation	67
4.3.3	Simulation settings for the competitors	68
4.3.4	Simulation results	69
4.4	Summary and discussion	74
B	Appendices for Chapter 4	77
B.1	BOD simulation code for MTDC estimation	77
Chapter 5	Bayesian optimization design of OD estimation for targeted therapies	85
5.1	Overview of drug development and statistical designs for dose-finding studies of targeted therapies	85
5.1.1	Difference between mono-therapies with cytotoxic agents and targeted therapies with biologic agents in dose-finding studies	85
5.1.2	Parametric model-based designs	87
5.1.3	Optimal interval designs	89
5.1.4	Isotonic design	92
5.2	Modeling frameworks and practical procedures of Bayesian optimization design for OD estimation	93
5.2.1	Statistical modeling on dose-response relationships	93
5.2.2	Integrating toxicity and efficacy information	94
5.2.3	BOD implementation steps for OD estimation	95
5.3	Performance evaluation	97
5.3.1	General settings for the simulation study	98
5.3.2	Simulation settings of BOD for OD estimation	100
5.3.3	Simulation settings for the competitors	102
5.3.4	Simulation results	103
5.4	Summary and discussion	109
C	Appendices for Chapter 5	112
C.1	Supplemental simulation results	112
C.2	BOD simulation code for OD estimation	115
Chapter 6	Conclusion and future work	123
	Acknowledgment	126

Chapter 1

Introduction

1.1 Drug development and statistical methodologies

1.1.1 Drug development process

Drug development is the process of finding a new pharmaceutical drug in order to contribute to unmet medical needs. At the first step, researchers identify a promising compound for development from thousands of potential candidates through various laboratory testings. Those compounds proceed to preclinical tests such as *in vitro* and *in vivo* to obtain safety information. Based on the preclinical results, researchers determine whether the candidate compounds should be examined in humans. Preclinical research dispels basic concerns about the safety aspects of the candidate compounds; however, it cannot answer a question about how they interact with human bodies. Therefore, the compounds that have passed these various tests finally proceed to clinical trials performed in people as candidates for a new treatment.

Clinical trials to develop a new treatment are composed of a series of studies so-called Phase I, II, and III in general. The ultimate goal of Phase I clinical trials is to confirm if the new drug candidate is safe on human bodies. Researchers aim to learn the effect of the new treatment on human bodies, look at the ideal way of administrations, identify the highest dose that does not cause unacceptable toxicity for predicting a therapeutic dosage, and so on through several Phase I clinical trials. Phase I clinical trials typically involve a small number of patients or healthy volunteers.

When Phase I clinical trials show expected results, clinical trials move to the next phase; that is, Phase II clinical trials. Phase II clinical trials shift their focus on the efficacy aspect in addition to safety. They mainly aim to see if the new treatment works for patients who are under a specific condition that the new treatment targets while collecting and reviewing safety information. In Phase II clinical trials, researchers determine whether the new treatment has some benefit on targeted patients. While Phase II clinical trials involve more participants than Phase I clinical trials, it is still not large enough to statistically demonstrate the efficacy of the new drug.

If the new drug candidate shows some benefit to targeted patients and is safe enough from the view of a balance between safety and benefit, it goes to Phase III clinical trials to statistically demonstrate the effectiveness of the new treatment. In general, Phase III clinical trials set strict statistical criteria to determine the effectiveness of the new treatment. They often involve a large number of participants to compare the new treatment with the current standard treatment based on statistical rationale. The trial designs usually adopt double-blind for both participants and investigators in terms of the treatment for participants to eliminate any bias on assigning treatment and interpreting results. Phase III clinical trials could last for a long period of more than a year depending on their designs

according to clinical perspective and statistical rationale. Regulatory authorities in some countries or regions occasionally require at least two Phase III clinical trials that provide successful results on efficacy and safety in order to approve the new treatment.

When Phase III clinical trials are successfully done, researchers and a regulatory authority collaborate their works for new drug application and approval processes. Researchers still need to examine overall safety after the declaration of effectiveness of the new treatment based on a series of clinical trials; therefore, safety information will be continuously collected, monitored, and reviewed in a certain period as postmarketing surveillance after the new treatment is released on the market as a pharmacovigilance part. Fig. 1.1 shows a brief explanation of the main part of the drug development process.

There are a lot of articles summarizing and discussing drug development processes such as Robuck and Wurzelmann (2005) and Orloff et al. (2009). Also, the U.S Food & Drug Administration (FDA) summarizes the drug development process with some videos on their official website (<https://www.fda.gov/patients/learn-about-drug-and-device-approvals/drug-development-process>).




Fig. 1.1. Main part of drug development process

1.1.2 Expectation of exploiting Bayesian approaches in clinical trials

The author has been involved in drug development and been in charge of many projects as a biostatistician for more than 10 years in Pfizer Japan or Pfizer R & D Japan, a pharmaceutical company. Of the drug development process depicted in Fig. 1.1, the author works on clinical trials-related tasks through designing development strategies, planning statistical analyses that fit the objectives of each protocol, sharing interpretations of study results with team members. The main disease areas the author has been involved in are oncology and rare disease from early to late phases of clinical trials.

While there are two statistical approaches, which are frequentist and Bayesian statistics, most clinical trials have relied on frequentist approaches. Although we usually follow conventional statistical designs based on frequentist strategies unless there are special circumstances, it is also well-known that clinical trials with small sample sizes owing to feasibility or ethical reasons have worth adopting Bayesian strategies to make the most of available information. For example, early-phase clinical trials and clinical trials whose target population is relatively small sometimes leverage Bayesian approaches to make their strategies more efficient. Although the history of utilization of Bayesian approaches in clinical trials is still short, there are enormous needs and expectations of effective applications and developments of Bayesian strategies for clinical trials. Gupta (2012) and Berry (2006), for instance, have summarized and discussed the advantages and challenges of the utilization of Bayesian approaches in drug development. The author has felt those growing needs and expectations via real projects on the job since the author started to work in the pharmaceutical industry. In addition, a way of more exploitation of Bayesian approaches in clinical trials has recently been sought as one of the efforts of research on

regulatory science of pharmaceuticals and medical devices supported by the Japan Agency for medical research and development (A. Hirakawa, Research representative, 2019-2021).

1.2 Motivation and contribution

The author has some experience of being in charge of dose-finding studies in Phase I clinical trials as a biostatistician. A primary objective of dose-finding studies is generally to find a maximum tolerated dose (MTD) of new drug candidates. The MTD is defined as the highest dose that does not cause an unacceptable level of toxicity according to the National Cancer Institute (NCI) dictionary (<https://www.cancer.gov/publications/dictionaries>). Dose-finding studies are one of the few clinical trials that positively exploit Bayesian strategies; therefore, those studies brought the author good opportunities to deal with Bayesian statistics in practical clinical trials. Dose-finding studies determine a safe dose that plays a key role pertaining to therapeutic dose range in subsequent trials; therefore, misspecification of the MTD in dose-finding studies leads to a low success rate of the drug development. Nevertheless, even statistical designs that are well-known and frequently used in dose-finding studies have some discussion points and standard designs have yet been established. The more the author learned statistical designs for dose-finding studies in order to understand the projects, the more the author recognized they have still open research questions. In particular, Bayesian approaches provide ever-increasing interest and challenge in the area of dose-finding studies because they can make the most of limited available data. Besides, waves of machine learning methodologies including novel topics such as artificial intelligence have been finally reaching the drug development field. Those waves have brought novel strategies that are worthy of consideration as new approaches for clinical trials. Given these circumstances, our study has cast light on a different Bayesian approach that has never been applied in clinical trials while attracting attention in recent years. We thought it could address the open research questions that conventional statistical designs have left.

In conventional statistical designs for dose-finding studies, few Bayesian designs adopt established frameworks incorporating dose-response estimations and targeted dose selections while having flexible modeling approaches. It often causes model misspecification, inflexible or suboptimal dose-allocation, and insufficient accuracy of final dose determinations. Bayesian optimization (Mockus et al., 1978; Mockus, 1975) has emerged as an efficient strategy to find a global optimizer (i.e., minimizer or maximizer) of an unknown function and been expanded its utilization to various areas. Shahriari et al. (2016) reported that Bayesian optimization has recently applied to a wide range of areas such as machine learning, sensor networks, and environmental monitoring. Given these backgrounds, we also paid attention to the advantage of Bayesian optimization. As a result, we considered Bayesian optimization would have sufficient features to address the current issues on statistical designs for dose-finding studies.

According to Fig. 1.2 depicting typical frameworks of Bayesian statistical designs in dose-finding studies and Bayesian optimization, we can see that those approaches have similar processes in terms of conducting repeated assessment based on Bayesian estimations. In addition, the ultimate goals are not exactly the same but closely similar to each other in terms of figuring out the best point over an unknown function or curve. The majority of conventional Bayesian designs in dose-finding studies are parametric model-based designs, although information about dose-response relationships are rarely available; therefore, prescribing strong assumptions to unknown dose-response relationships, which parametric model-based designs do, sometimes causes difficulties. Also, most conventional designs do not effectively utilize uncertainties in their dose selection process

even if there are nonnegligible uncertainties that possibly exist in particular at the beginning of the trial. Bayesian optimization estimates an unknown function of interest with a nonparametric model that realizes flexible modeling through sequential updates based on Bayes' rule. In addition, Bayesian optimization constructs a probabilistic model for an unknown function and then exploits the model via acquisition functions to make decisions where the most promising point is without ignoring uncertainties of posteriors. Selecting potential optimal points without relying on only local optimal points avoids concentration on a particular local region, which results in actively collecting data of other regions. Consequently, it allows us to identify a global optimal point within as few evaluations as possible. Optimization strategies based on acquisition functions in Bayesian optimization have been demonstrated to work effectively and realize as few evaluations as possible by such as Bull (2011). Because dose-finding studies usually evaluate a limited number of patients, it is ideal to adopt a design that returns more accurate results with fewer evaluations; therefore, it was considered that Bayesian optimization would become an ideal design for dose-finding studies as it could provide a sophisticated optimization strategy that had worked successfully in various areas. Although Bayesian optimization seemed to have the promising capability to be a better design for dose-finding studies, it had never been used in clinical trials before. It was a great challenge for our study to develop a new Bayesian design based on Bayesian optimization that can be an acceptable and alterable approach compared with conventional statistical designs because clinical trials have unique characteristics that are totally different from other fields such as ethical issues concerning patient safety. The biggest concern to apply a new design brought from the other fields to the clinical field would be related to whether it can appropriately incorporate ethical restrictions while having less impact on its original advantages.

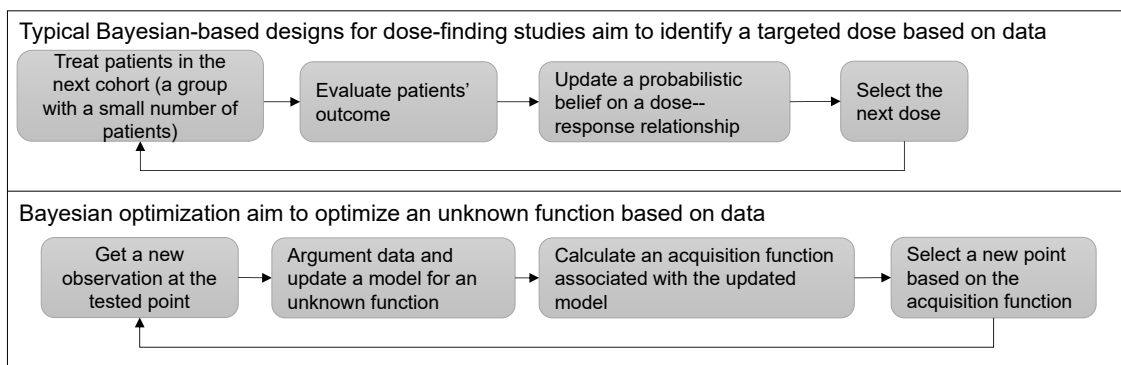


Fig. 1.2. Typical frameworks of Bayesian designs for dose-finding studies and Bayesian optimization

The development of Bayesian approaches to expand and accelerate their utilization in clinical trials is one of the recent challenges as a biostatistician in the pharmaceutical industry. There is still room for improvement in existing statistical designs for dose-finding studies where Bayesian approaches are positively utilized. Although we considered Bayesian optimization had some elements to solve the current issues, it was necessary to develop a new approach based on Bayesian optimization to fit clinical trials especially in terms of ethical aspects. Our goal of this thesis is to design and propose a new Bayesian design for dose-finding studies by utilizing the advantages of Bayesian optimization frameworks. We will introduce a Bayesian optimization design (BOD) that is a novel Bayesian design for dose-finding studies throughout this thesis. For the development of new statistical designs for dose-finding studies, the operation characteristics of proposed designs

are usually evaluated via simulation studies. Even if there is real clinical trial data, true dose–response curves are never revealed and all observations are obtained based on a specific statistical design in the trial. It is possible to use some real clinical trial data as one of true dose–response curves for simulations; however, simulation scenarios should cover a wide enough range of variety of scenarios in order to confirm the operating characteristics under a typical situation where no information about dose–response curves is available (Mayer et al., 2019). Thus, performance evaluation should include the majority of plausible dose–response scenarios including ones with various MTD locations and extreme ones that we may encounter. Accordingly, in the same way as standard steps for statistical design development for dose-finding studies, our study evaluates the performance of BOD by conducting simulation studies that compare it with existing representative designs under some plausible scenarios that it is likely we encounter in actuality. We believe that our study on a new Bayesian approach that could perform better than existing designs has contributed to expanding the limited choice of Bayesian approaches in clinical trials and showing the applicability of Bayesian designs from other fields that possibly lead to developing further effective applications in the future. Also, we believe that our work has brought further awareness for biostatistician communities of the usefulness of Bayesian strategies in clinical trials. Furthermore, our work has shown the usefulness of Bayesian optimization in clinical trial fields where ethical limitations exist. We believe our work will put some stimulation on applications of Bayesian approaches to clinical trial fields and finally contribute indirectly to an improvement on success rates of drug development.

1.3 Roadmap

Typical dose-finding studies to identify the MTD mainly assess safety aspects on investigational agents assuming mono-therapies with cytotoxic agents and sequentially evaluate patients divided into small groups called cohorts in single trials. Cytotoxic agents destroy cancer cells by inhibiting cell division based on the unique characteristics of cancer cells that often divide markedly faster than normal cells. Cancer therapies with cytotoxic agents are often called cytotoxic chemotherapy or just chemotherapy. Because cytotoxic agents do not distinguish tumors and healthy tissue, normal cells are also affected by the drug transported throughout the body from the bloodstream; therefore, careful safety management is necessary to ensure patient safety. Although the role of cytotoxic agents has decreased slightly with the development of therapies with modern mode-of-actions that are more effective and less toxic such as targeted therapies, cytotoxic chemotherapies remain fundamental options depending on cancer types (Twelves et al., 2016). In dose-finding studies, the information about tested dose levels and the number of patients who experienced specific toxicity events so-called dose limited toxicity (DLT) is analyzed for MTD estimation in addition to overall assessment using all available data. Statistical designs sequentially analyze the specific information of interest for MTD estimation and select an appropriate dose for the next test during the trial. The MTD becomes a candidate of the recommended dose determined at the end of the trial.

Chapter 2. Firstly, Chapter 2 explains several conventional statistical designs for MTD estimation assuming typical dose-finding studies. The underlying premise in typical dose-finding studies for cytotoxic agents is that toxicity increases monotonically with increasing dose levels. The statistical designs we introduce in Chapter 2 are well-known to biostatisticians and have been applied to various actual clinical trials. A major transition of trends on statistical designs for dose-finding studies can be seen by exhibiting those designs step by step.

Chapter 3. In accordance with the same underlying premise taken in the most typical dose-finding studies for cancer treatment, we have developed a new Bayesian design, BOD, for identifying the MTD for mono-therapies of cytotoxic agents. Chapter 3 introduces BOD for MTD estimation in detail such as statistical frameworks and its operating characteristics via simulation studies. Fig. 1.3 shows the relationship between chapters and targeted dose-finding studies. While we have developed BOD assuming dose-finding studies under three different situations as shown in Fig. 1.3, Chapter 3 introduces the basic components of BOD. In Chapters 4 and 5, we extend BOD introduced in Chapter 3 to more complex situations.

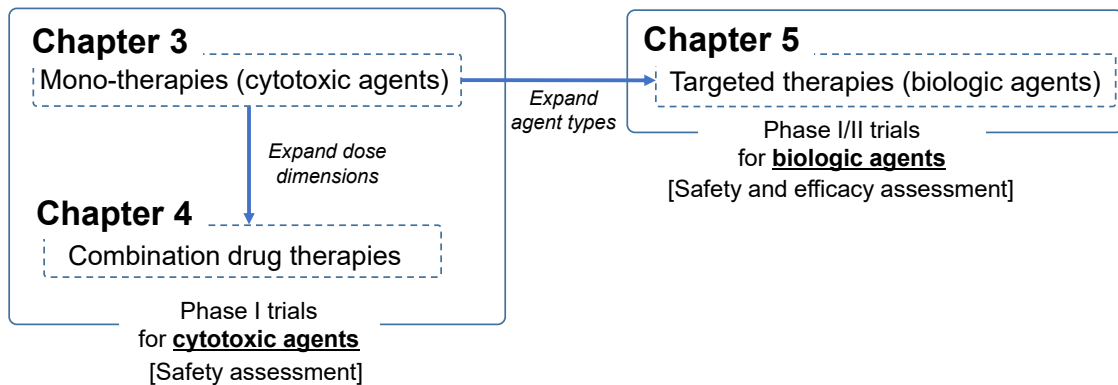


Fig. 1.3. Each chapter and the associated type in dose-finding studies where BOD applies

Chapter 4. It is still a great challenge to address one of the biggest issues, which is drug resistance, across cancer treatment. Combination therapy decreases the likelihood of occurrence of resistant cancer cells by combining drugs with different mechanisms. Also, combination therapy aims to incorporate the benefits from each drug; thereby, enhance therapeutic response in cancer. For the above background, combination therapy can be a key treatment option for cancer patients. Chapter 4 introduces BOD for MTD combination (MTDC) estimation assuming combination therapy of cytotoxic agents. BOD in Chapter 4 deals with two-dimensional doses as input data instead of one-dimensional dose that BOD in Chapter 3 assumes.

Chapter 5. Chapter 5 focuses on optimal dose (OD) estimation under dose-finding studies for targeted therapies instead of MTD estimation. Targeted therapies with such as biologic agents have been growing interest and increasingly developed in recent years. These therapies deliver a more focused treatment by directly interfering with a specific pathway involved in the targeted tumors. Arora et al. (2017) has mentioned that biologic agents have revolutionized therapy for a number of malignancies. Targeted therapies with biologic agents differ from conventional chemotherapies with the use of cytotoxic agents in terms of mode-of-action. For example, once the targeted pathway is inhibited with a certain amount of biologic agents, further administration no longer provides any benefit for patients. Dose-efficacy curves may exhibit non-monotone shapes such as unimodal and plateau in such a case. Cytotoxic agents show monotone relationships between dose levels and toxicity. Simultaneously, the efficacy of cytotoxic agents is assumed to increase monotonically with increasing dose levels. The MTD is clinically meaningful only if the drug shows monotonic increasing patterns with increasing dose levels on both toxicity and efficacy. Considering the difference of the mode-of-action between cytotoxic chemotherapies and targeted therapies, we can not put the monotonic assumption on dose-response relationships for targeted therapies unlike BOD in Chapters 3 and 4. Thus, instead of

MTD estimation, dose-finding studies for targeted therapies require identifying an OD that provides sufficient efficacy under an acceptable toxicity rate.

Organization of the thesis

This thesis is organized as follows:

- Chapter 2 is a brief literature review of statistical designs to introduce a major transition of statistical designs for typical dose-finding studies.
- Chapter 3 concerns BOD for mono-therapies with cytotoxic agents for finding the MTD. The related works on Chapter 3 are *A curve free method based on Bayesian optimization for oncology Phase I clinical trials*, A. Takahashi and T. Suzuki, Proceedings of the 2nd International Conference on Statistics: Theory and Applications (ICSTA'20), 2020 (Takahashi and Suzuki, 2020) and *Bayesian optimization for estimating the maximum tolerated dose in Phase I clinical trials*, A. Takahashi and T. Suzuki, Contemporary Clinical Trials Communications, 2021 (Takahashi and Suzuki, 2021a).
- Chapter 4 covers BOD for combination drug therapies in Phase I clinical trials for identifying a single MTDC and based on *Bayesian optimization design for finding a maximum tolerated dose combination in Phase I clinical trials*, A. Takahashi and T. Suzuki, International Journal of Biostatistics, 2021 (Takahashi and Suzuki, 2021b).
- Chapter 5 deals with BOD focusing on targeted therapies with biologic agents in Phase I/II clinical trials for finding a single OD. Chapter 5 is based on *Bayesian optimization design for dose-finding based on toxicity and efficacy outcomes in Phase I/II clinical trials*, A. Takahashi and T. Suzuki, Pharmaceutical Statistics, 2021 (Takahashi and Suzuki, 2021c).
- In Chapter 6, we conclude this thesis and discuss future work for our study along with some challenging points.

Term descriptions in this thesis

Glossary of important terms in this thesis is listed below:

Table 1.1. Glossary of important terms

No	Term	Brief descriptions for this thesis*
1	Dose-finding studies	One of Phase I clinical trials for identifying targeted doses depending on types of the cancer therapies while taking tolerability into account. They are typically conducted sequential tests of an investigational agent by escalating and de-escalating dose levels for successive cohorts.
2	Mono-therapies	Therapies using a single drug to treat a disease or condition.
3	Combination therapies	Therapies using multiple drugs to treat a disease or condition.
4	Targeted therapies	A type of cancer treatment that directly affects specific parts being involved in the growth, progression, and spread of cancer.
5	Cytotoxic agents	Anticancer agents that aim to destroy cancer cells by inhibiting cell division based on the unique characteristics of cancer cells.
6	Biologic agents	A substance that is made from a living organism or its products. Many targeted therapies for cancer treatment utilize biologic agents.
7	Cohort	Group composed of small number of patients (e.g., three patients per cohort) for treatment at a selected dose level.
8	Dose-response relationships	Depending on study objectives, dose-response relationships describe relationships between dose and the associated toxicity, efficacy, or both responses.
9	Dose limited toxicity (DLT)	Serious adverse effects that prevent an increase in dose level of the treatment.
10	Maximum tolerated dose (MTD)	The highest dose that does not cause unacceptable toxicity.
11	Maximum tolerated dose combination (MTDC)	A dose combination whose toxicity probability equals to the target toxicity rate.
12	Optimal dose (OD)	In this thesis, an optimal dose means a dose that satisfies the following two conditions: efficacy is greater than or equal to the target efficacy rate, toxicity does not exceed the target toxicity rate.
13	Target response rates	There are two types of target response rates in this thesis. One is the target toxicity rate that describes the maximum tolerated toxicity for MTD or MTDC determinations and the other is the target efficacy rate that describes the minimum acceptable efficacy rate for OD determinations.
14	Bayesian optimization	One of the sequential nonparametric Bayesian-based frameworks to seek a global optimization over unknown black-box functions while estimating the functions.
15	Bayesian optimization design (BOD)	Our proposed design that leverages the advantages of Bayesian optimization for dose-finding studies.
16	Parametric model-based designs	In this thesis, parametric model-based designs mean statistical designs that utilize parametric models to describe dose-response relationships.
17	Toxicity probability interval designs	Bayesian nonparametric designs that sequentially select appropriate doses based on posterior probabilities associated with pre-specified toxicity intervals for dose-finding studies.
18	Curve-free designs	Bayesian nonparametric designs that sequentially estimate dose-response relationships without specifying theoretical models and select and determine appropriate doses.
19	Gaussian process (GP)	A stochastic process that has Gaussian distributed finite dimensional marginal distributions (Quadrianto et al., 2011). Every finite collection of the random variables follow a multivariate normal distribution.
20	Acquisition function	Heuristics based on the latest model on unknown functions for drawing how desirable each point is for the next evaluation in order to reach the global optimal point.

* Note that the glossary partially relies on the NCI dictionary.

Chapter 2

Transition and problems of statistical designs for dose-finding studies

This chapter introduces the fundamental study procedures for dose-finding studies and popular statistical designs that provide essential frameworks that have brought various extended statistical designs. Besides, we briefly explain the transitions of statistical designs for dose-finding studies and the relationships among the designs. In preparation to apply Bayesian optimization to dose-finding studies, we also roughly describe typical Bayesian optimization.

2.1 Common procedures in dose-finding studies

The ultimate goal of oncology Phase I clinical trials for mono-therapies with cytotoxic agents is to identify the MTD defined as the highest dose that does not cause an unacceptable level of DLT. Specifically, the primary objective of dose-finding studies is to identify which dose x_j is the maximum dose that has the closest toxicity (i.e., DLT) probability to the target toxicity rate of θ among available dose levels ($j \in \{1, \dots, J\}$).

The evaluation period for DLT involved in the MTD determination is usually set as the first cycle of treatment (e.g., 28 days long) during a trial. In general, DLT is assessed according to the National Cancer Institute's Common Terminology Criteria for Adverse Events (CTCAE) classification and usually encompasses all grade 3 or higher severe toxicity except for grade 3 nonfebrile neutropenia and alopecia. According to Postel-Vinay (2015), a few recent trials have updated its DLT definition for instance by adding prolonged grade 2 toxicity or prolonging the DLT evaluation period in order to fit the agents with modern mode-of actions or to address late-onset toxicity (Postel-Vinay et al., 2011).

Dose-finding studies examine investigational agents by assigning patients to small groups called "cohort" from the lowest dose level based on sequential dose-escalation and de-escalation procedures. Cohort sizes are usually three patients. Once patients in the t th cohort are treated, all observations up to the t th test are carefully assessed. The next dose $x_{(t+1)}$ is then selected based on the evaluated observations. The series of treating patients with a selected dose, observing patient outcomes, and selecting the next dose are repeated until one of the predetermined stopping rules is met. The following steps are commonly used in dose-finding studies:

1. Patients in a cohort are treated with a selected dose. In general, patients in the first cohort are treated with the lowest dose level. Because dose-finding studies are first-in-human studies that are the first time to test the new treatment on people, patients in the first cohort are treated with a dose that is considered safe on human bodies based on extrapolation from preclinical data.

2. Patient outcomes are observed. The number of patients who experienced or did not experience DLT plays the most important role in dose selections for the next cohort.
3. After the t th test, the next dose $x_{(t+1)}$ is selected based on the latest data or cumulative data depending on statistical designs.
4. Steps 1 to 3 continue until one of the predetermined stopping rules is met. (e.g., All patients have been already enrolled in the trial.)
5. MTD determination is performed based on the last data or cumulative data.

Fig. 2.1 roughly draws the graphical depiction regarding common procedures in dose-finding studies.

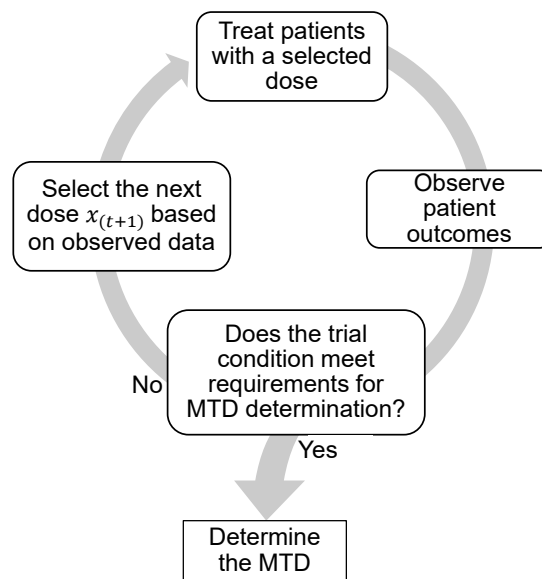


Fig. 2.1. Schematic of common procedures in dose-finding studies

2.2 The earliest statistical design

A number of statistical designs have been proposed for identifying the MTD under monotherapies with cytotoxic agents. Out of them, the earliest statistical designs is the traditional 3+3 design (Carter, 1973). It is categorized as rule-based designs, which is also referred to as algorithm-based designs, that require no modeling of dose-toxicity relationships. While there are some variations of the 3+3 design, the 3+3 design commonly sequentially treats 3-patient cohorts and dose-escalation (and de-escalation) for the next cohort is determined based on the number of patients with DLT at the current test dose. In accordance with the summary provided in Daimon (2012), the traditional 3+3 design imposes the following rule on dose-escalation:

1. Treat three patients with a selected dose level
 - If no patient experienced DLT, the next dose level is escalated by one dose level and the next step goes to Step 1.
 - If one patient experienced DLT in the cohort, the next step goes to Step 2.
 - If more than or equal to two patients experienced DLT in the cohort, the next step goes to Step 3.

2. Stay the same dose level for the next cohort and three patients in the next cohort are treated at the same dose level.
 - If no patient experienced DLT in the cohort, the next dose level is escalated by one dose level and the next step goes to Step 1.
 - If more than or equal to one patient experienced DLT in the cohort, the next step goes to Step 3.
3. Dose-escalation is stopped because the last tested dose is regarded as higher than the MTD.
4. The MTD is determined after Step 3. If MTD determination requires treating at least six patients with the MTD, the MTD will be the highest dose with 33% or less toxicity probability. When we reach Step 3, one dose level lower than the last tested dose will be the MTD; however, if the dose has not been administered to at least six patients, three patients will be additionally treated with the dose. These steps are repeated until the final condition for MTD determination meets.

Fig. 2.2 provides a graphical depiction of an example of possible dose allocations in a dose-finding study based on the traditional 3+3 design following the above rule. In Fig. 2.2, dose-escalation was stopped at the 7th cohort, and the trial selected the 4th dose level from the bottom as the MTD based on the patient outcomes at the last cohort.

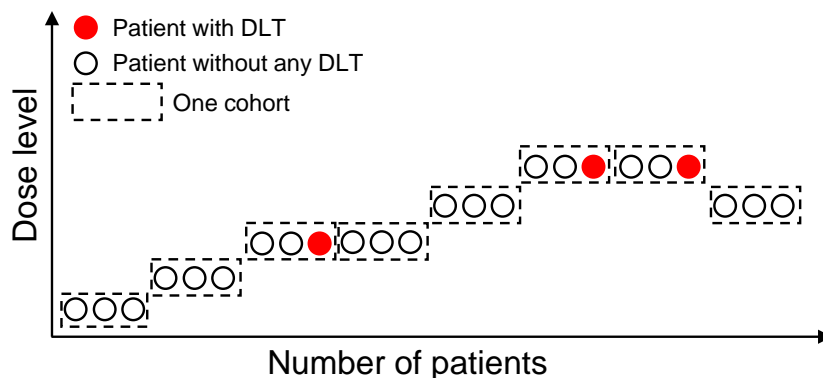


Fig. 2.2. Schematic of the traditional 3 + 3 design

If Steps 2 and 3 are removed, and the conditions of Steps 2 and 3 are alternated with

- If one patient experienced DLT in the cohort, the next step goes back to Step 1,
- If more than or equal to two patients experienced DLT in the cohort, the next dose level is de-escalated by one dose level ,

respectively, this variation equals to another 3+3 design proposed by Storer (1989), which is one of the most popular algorithms as the 3+3 designs.

Concerns to be highlighted

The 3+3 design remains the prevailing design for dose-finding studies because it is simple to implement and safe (Tourneau et al., 2009); however, it is well-known that the 3+3 design has substantial limitations such as treating a large portion of patients with potentially subtherapeutic doses due to infrequency of dose-escalation (Ratain et al., 1993; Reiner et al., 1999; Zohar and O'Quigley, 2009). It results in low accuracy of MTD identifications. Ivanova (2006) has numerically shown that the target toxicity rate of the 3+3 design locates between 0.16 and 0.27 on average, which is lower than the declared target

toxicity rate of 0.33. This is a significant disadvantage because the target toxicity rate in this design deviates from what we truly need to identify.

2.3 Parametric model-based designs

As alternative designs to rule-based designs represented by the traditional 3+3 design, parametric model-based designs have been emerged to improve the precision of MTD estimation. The continual reassessment method (CRM) proposed by O’Quigley et al. (1990) is the first parametric model-based design, which is one of the adaptive designs, for dose-finding studies. After the emergence of the CRM, a lot of variations of the CRM have been introduced such as dose escalation with overdose control (Bobb et al., 1998), the Bayesian logistic regression model (Neuenschwander et al., 2008), the Bayesian model averaging CRM (Yin and Yuan, 2009a) and others (e.g., Goodman et al., 1995; Faries, 1994; Leung and Wang, 2002). Many authors have demonstrated the CRM and its variations provide superior performance to the 3+3 design in terms of accuracy of MTD estimation, proper dose allocation, and flexibility (e.g., Onar et al., 2009; Iasonos et al., 2008; Onar-Thomas and Xiong, 2010; Ananthakrishnan et al., 2016; Boonstra et al., 2015; Rosenberger and Haines, 2002; Jaki et al., 2013; James et al., 2016). Accordingly, leading pharmaceutical companies have commonly applied parametric model-based designs represented by the CRM including its variations (Love et al., 2017).

Of the CRM family, Bayesian-based CRMs is often utilized in oncology Phase I clinical trials to identify the MTD. Compared with the 3+3 design, the CRM has roughly three advantages: Target toxicity rates are quantitatively clarified so that the CRM can identify the dose associated with the targeted toxicity rate. Dose–toxicity relationships are evaluated by exploiting not only the last observations but also all available data. Intuitive graphical illustrations on dose–toxicity relationships could be easily provided based on simple parametric models. The following describes the statistical frameworks of the CRM.

Dose-toxicity model and a prior distribution for the model parameter

The CRM models dose–toxicity relationships with simple one-parameter monotonically increasing functions. The popular models for the CRM are a power model, a one-parameter logistic model, or a hyperbolic tangent model. For example, a power model describes a toxicity probability at a dose level j ($j \in \{1, \dots, J\}$) as follows:

$$\pi_j(a) = (\pi_j^0)^{\exp(a)}, \quad (2.1)$$

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 provided by investigators and called skeletons for the model; however, it is often challenging to specify initial toxicity guesses because of a lack of information about dose–toxicity relationships.

Lee and Cheung (2009) has proposed a systematic approach to calibrate initial toxicity guesses for the CRM based on indifference intervals. The indifference interval for a dose level corresponding to the true MTD provides an interval of toxicity rates associated with the neighboring doses that may be selected instead of the true MTD because the difference of toxicity rates among doses within the interval is regarded as clinically negligible. The systematic approach provides an indifference interval to maximize the average percentage of correct MTD selections under the selected model for the CRM across possible scenarios of true toxicity probabilities.

As another representative model for the CRM, a one-parameter logistic model with a

fixed intercept is given by

$$\text{logit}\{\pi_j(a)\} = a_0 + x_j \times \exp(a), \quad (2.2)$$

where a_0 is a fixed constant that is usually 3 (O'Quigley and Chevret, 1991). While skeletons are the same values as re-scaled conceptual doses x_j in the power model, a logistic model needs additional calculations for conceptual doses associated with initial toxicity guesses. The values of x_j are calculated so that toxicity probabilities derived with a fixed parameter (e.g., $a = 1$) return π_j^0 at each j . We note the values of x_j do not need to take actual values of the investigational agents.

A prior distribution for a usually assumes a normal formulation (i.e., $P(a) = N(0, \sigma_a^2)$), where a prior variance for σ_a^2 can be calibrated by another systematic approach proposed by Lee and Cheung (2011) 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 dose-toxicity relationships at the start of a trial.

Model update and dose selection

We obtain binary responses regarding if a patient experienced DLT or not at every sequential test for cohorts. Once the dichotomous DLT response is observed in each patient in the t th cohort, a distribution on a is updated using all available observations up to the t th cohort $D_{1:t} = \{(n_{(1)}, x_{(1)}, y_{(1)}), \dots, (n_{(t)}, x_{(t)}, y_{(t)})\}$ according to Bayes' theorem as follows:

$$P(a | D_{1:t}) = \frac{L(D_{1:t} | a)P(a)}{\int_0^\infty P(u)L(D_{1:t}; u)du}, \quad (2.3)$$

$$L(D_{1:t} | a) = \prod_{z=1}^t \pi_{(z)}(a)^{y_{(z)}} \{1 - \pi_{(z)}(a)\}^{n_{(z)} - y_{(z)}}, \quad (2.4)$$

where $y_{(z)}$ denotes a number of patients experienced DLT out of $n_{(z)}$ patients treated with the z th tested dose $x_{(z)}$; $\pi_{(z)}$ denotes a toxicity probability corresponding to $x_{(z)}$. As a result, the CRM provides the following two possible point estimates:

- A plug-in mean

$$\hat{\pi}_j = \pi_j(a) |_{a=E(a|D_{1:t})}, \quad E(a|D_{1:t}) = \int_0^\infty aP(a | D_{1:t})da, \quad (2.5)$$

- A posterior mean

$$\hat{\pi}_j = E\{\pi_j(a)|D_{1:t}\} = \int_0^\infty \pi_j(a)P(a | D_{1:t})da. \quad (2.6)$$

Based on $\hat{\pi}_j$, a dose with the closest toxicity rate to a target toxicity rate θ is selected as the next dose $x_{(t+1)}$. Occasionally, a trial does not allow either to skip doses in dose-escalation or to escalate doses immediately after toxic outcomes in dose selections to ensure patient safety. Once the trial conditions meet pre-specified stopping rules, the MTD is determined as a dose with the closest toxicity to θ based on $\hat{\pi}_j$.

Open research questions

The CRM needs to specify a theoretical model describing dose–toxicity relationships before the beginning of a trial. Shen and O’Quigley (1996) has demonstrated that the CRM can achieve the primary goal with regards to the selection of the MTD even if the model is misspecified; however, Cheung and Chappell (2002) has reported that sufficient conditions to converge to the true MTD established by Shen and O’Quigley (1996) might be too restrictive. It means there are possible scenarios where the CRM cannot guarantee convergence to the true MTD. Paoletti and Kramar (2009) has reported that the choice of theoretical models in the CRM would affect the operating characteristics if sample sizes are small (e.g., 10 to 50 patients). As there is little information about true dose–toxicity relationships in practice, it may not always appropriate to apply strong restrictions for dose–toxicity shapes before the initiation of trials. Parametric model-based designs have usually the potential risk of model misspecification and the CRM is no exception.

Furthermore, most conventional designs including the CRM rely on point estimates without considering uncertainties of estimates during their dose selections. If there is non-negligible variability in estimated dose–toxicity relationships, dose selections without taking account of the uncertainty may lead to skewed dose allocations. For example, consider two different distributions with the same mean value of a toxicity probability. Even if these mean values are near to the target toxicity rate and the dose can be selected as an MTD candidate, the shapes of those distributions may be quite different in particular at the start of a trial (Neuenschwander et al., 2008; Mozgunov and Jaki, 2020). Thus, decision-making based on point estimates can be potentially risky since it cannot reflect such non-negligible variability.

Our study introduced from the next chapter will address issues on the model misspecification and on ignoring uncertainties in dose selections.

2.4 Nonparametric designs

2.4.1 Toxicity probability interval designs

As different approaches from parametric model-based designs like represented by the CRM, nonparametric Bayesian designs have been introduced to relax the effect of model selections. Toxicity probability interval designs might be major approaches of nonparametric Bayesian designs.

The modified toxicity probability interval (mTPI) proposed by Ji et al. (2010) is one of the most popular toxicity probability interval designs. Liu and Yuan (2015) has proposed the Bayesian optimal interval (BOIN) that 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 (Guo et al., 2017) 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. Because mTPI provides as simple implementation as the 3+3 design but much better performance than the 3+3 design, the popularity of the mTPI has been growing in both research and industry entities during the relatively short period since it was proposed (Ji and Wang, 2013). The following describes the statistical frameworks of the mTPI.

Toxicity probability distribution

The mTPI is a Bayesian adaptive design assisted by a beta-binomial model in spite of its simplicity on implementation without any logistic burden. Specifically, suppose π_j is a toxicity probability at a dose level j ($j \in \{1, \dots, J\}$). The prior distribution for π_j is given by

$$\pi_j \sim \text{Beta}(\alpha_j, \beta_j), \tag{2.7}$$

where α_j and β_j are hyperparameters for a beta distribution. In general, a non-informative prior is adopted for equation (2.7) (e.g., $\alpha_j = \beta_j = 1$, $\alpha_j = \beta_j = 0.005$). When y_j patients experienced DLT out of n_j patients at a dose level j after the t th test, a posterior distribution on a toxicity probability at j can be denoted by

$$\pi_j | D_{1:t} \sim \text{Beta}(\alpha_j + y_j, \beta_j + n_j - y_j), \tag{2.8}$$

based on the beta-binomial conjugacy.

Dose selection

The unit probability mass (UPM) under a posterior distribution of π_j at the current tested dose level j is calculated for three intervals corresponding to underdosing, proper dosing, and overdosing: $(0, \theta - \varepsilon_1)$, $(\theta - \varepsilon_1, \theta + \varepsilon_2)$, $(\theta + \varepsilon_2, 1)$. Ji and Wang (2013) and Ji et al. (2010) have recommended to assign small fractional values to ε_1 and ε_2 (e.g., 0.05) to consider uncertainties around a target toxicity rate θ . The UPM for a dose level j corresponding to an interval (p_1, p_2) can be calculated by $\{F(p_2) - F(p_1)\}/(p_2 - p_1)$, where F is a cumulative distribution function based on a beta distribution for π_j .

At each test during a trial, the mTPI takes either decision among ‘Stay at the same dose (\mathcal{S})’, ‘Escalation (\mathcal{E})’, or ‘De-escalation (\mathcal{D})’, which are associated with the interval $(\theta - \varepsilon_1, \theta + \varepsilon_2)$, $(0, \theta - \varepsilon_1)$ and $(\theta + \varepsilon_2, 1)$, respectively. Out of the three decisions, the mTPI selects a dose with the largest UPM as the next dose. This dose selection algorithm concerns minimizing the probabilities of incorrect decisions. The mTPI assumes an independent uniform prior for π_j and considers the six penalties as depicted in Fig. 2.3.

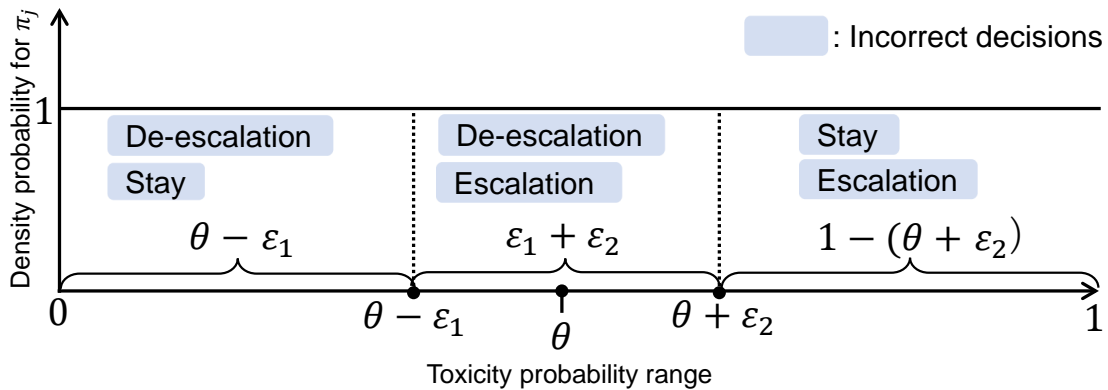


Fig. 2.3. Independent uniform prior for π_j and incorrect decisions where penalties are imposed when the true toxicity probability π_j falls in the divided range

Under the uniform prior, the mTPI sets the six penalties at $1/(\theta - \varepsilon_1)$, $1/(\varepsilon_1 + \varepsilon_2)$, and $1/(1 - \theta - \varepsilon_2)$ for each penalty pattern at the three divided range from left to right in Fig.

2.3, respectively. This possesses the property that the prior expected penalties for \mathcal{S} , \mathcal{E} , and \mathcal{D} are the same (=1). For example, the prior expected penalty on \mathcal{D} is given by

$$E\{P(\pi_j|\mathcal{D})\} = \frac{1}{\theta - \varepsilon_1}P(\pi_j < \theta - \varepsilon_1) + \frac{1}{\varepsilon_1 + \varepsilon_2}P(\theta - \varepsilon_1 \leq \pi_j \leq \theta + \varepsilon_2).$$

Given these penalties, the three posterior expected penalties on \mathcal{S} , \mathcal{E} , and \mathcal{D} equal (1 – UPMs) for the three intervals; therefore, choosing the largest UPM corresponds to choosing the decision that minimizes the posterior expected penalties.

Suppose the current tested dose level is j . If $P(\pi_{j+1} > \theta \mid D_{1:t}) > 0.95$, it implies that the upper adjacent dose level from the current tested dose shows excessive toxicity very likely. In this case, dose-escalation will be prohibited for the next dose selection (Ji et al., 2007).

Once either predetermined stopping rule is met, a single dose with the smallest difference between $\tilde{\pi}_j$ and the target toxicity rate θ is selected as the MTD from an admissible dose set \mathcal{A}_t , where $\tilde{\pi}_j$ denotes a sensible estimate of π_j and is usually derived from an isotonic regression. Specifically, after the posterior mean $\hat{\pi}_j$ is computed under the posterior beta distribution, the pooled adjacent violators algorithm (Barlow et al., 1972) on $\hat{\pi}_j$ is performed so that $\tilde{\pi}_j$ increases monotonically with increasing dose levels such as $\tilde{\pi}_1 \leq \dots \leq \tilde{\pi}_J$. The dose level at the MTD (j^*) is described by

$$j^* = \arg \min_{j \in \mathcal{A}_t} | \tilde{\pi}_j - \theta | . \quad (2.9)$$

Regarding an admissible dose set \mathcal{A}_t , it is composed of all tested dose levels that eliminate doses that satisfy a condition $P(\pi_j > \theta \mid D_{1:t}) \leq 0.95$; that is, any dose where the final posterior beta distribution shows high risks of excessive toxicity will not be used for the MTD determination.

Open research questions

Toxicity interval designs including the mTPI and other variations can be implemented easily and perform well under various scenarios. On the other hand, it focuses on only the current tested dose for the next dose selection and available doses for the dose selections are narrower than other types of designs like the CRM. Also, many toxicity interval designs perform the isotonic regression at the last stage of trials to provide dose–toxicity curves. It means dose–toxicity relationships across all possible doses are provided only for the final MTD determination. In general, dose level review meetings whose committee members include external clinicians assess all available data to decide the best action for the next dose selection; therefore, the latest statistical belief on dose–toxicity relationships possibly helps them understand the characteristics of investigational agents even if the estimate has a lot of uncertainties. There may still be room for improvement with respect to the above two points.

2.4.2 Curve-free designs

Curve-free designs are other nonparametric Bayesian designs that differ from toxicity interval designs. Curve-free designs can estimate dose–toxicity curves without any theoretical models that rely on parameters. Although the number of curve-free designs is much less than other types of designs, for example, Gasparini and Eisele (2000, 2001), Whitehead et al. (2010), and Tang et al. (2018) have proposed curve-free designs for dose-finding studies. Out of them, a nonparametric Bayesian design based on a product-of-beta-prior

(PBP) proposed by Gasparini and Eisele (2000, 2001) has been cited the most as a curve-free design from other authors. The following describes the statistical frameworks of the PBP.

Nonparametric modeling for toxicity probabilities

The PBP 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.

Let π_j denote a toxicity probability at a dose level j and reparameterize it with b_j as follows:

$$b_1 = 1 - \pi_1 \text{ and } b_j = \frac{1 - \pi_j}{1 - \pi_{j-1}} \text{ for } j = 2, \dots, J. \quad (2.10)$$

A beta prior distribution is independently assumed to each b_j ; that is, $b_j \sim \text{Beta}(o_j, \zeta_j)$. Equation (2.10) can be converted to the following expression:

$$\pi_j = 1 - \prod_{z=1}^j b_z. \quad (2.11)$$

Thus, the prior distribution for π_j induced by equation (2.11) 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.

The hyperparameters for π_j and b_j are determined as follows:

1. The hyperparameters for π_j are determined with prior inputs $(\pi_1^0, \pi_2^0, \dots, \pi_J^0)$ provided by investigators. One of the two parameters (A_j or B_j) is 1, then another is calculated by matching the prior median of π_j to the prior inputs of π_j^0 in order to make the variance largely subject to approximate unimodality. Specifically, if $\pi_j^0 \leq 0.5$, then $A_j = \log(0.5)/\log(\pi_j^0)$ and $B_j = 1$; whereas if $\pi_j^0 > 0.5$, then $A_j = 1$ and $B_j = \log(0.5)/\log(1 - \pi_j^0)$.
2. The parameters o_j and ζ_j are determined by matching the first two moments for b_j calculated by the first two moments of π_j based on A_j and B_j via equation (2.10). Gasparini and Eisele (2000, 2001) have provided detailed equations for the hyperparameters.

Dose selection

In the same manner in the other dose-finding procedures, binary DLT responses at each test are mainly contributed to the next dose selection. Once patient outcomes are observed at the t th test, posterior toxicity distributions for each dose level j are calculated. The next dose level for the $t + 1$ th test is then given by

$$j_{(t+1)} = \arg \min_{j \in \mathcal{A}_t} | \hat{\pi}_j - \theta |, \quad (2.12)$$

where $\mathcal{A}_t \subset \{1, \dots, J\}$ denotes an admissible dose set at the t th test whose conditions are predetermined before trials; $\hat{\pi}_j = E(\pi_j | D_{1:t})$ corresponds to the posterior expectations

of toxicity probabilities at each dose level. At the end of a trial, MTD determination is performed according to equation (2.12) so that the MTD has the closest toxicity rate to a target toxicity rate θ based on the final posterior mean of toxicity probabilities.

Open research questions

Gasparini and Eisele (2000, 2001) have shown the exact posterior distribution of b_j analytically as a mixture of beta distributions because binary toxicity outcomes follow a binomial distribution; therefore, they have also presented the exact computation for the posterior expectations for π_j . On the other hand, the matrix involved in the exact equation grows dynamically as the number of tests increases. Cheung (2011) has discussed the numerical problems of the exact computations and indicated that Markov Chain Monte Carlo (MCMC) is one of the options for avoiding the numerical problem. In addition, the PBP has been pointed that it might cause rigidity in situations where a low toxicity rate is targeted due to its vague priors. Cheung (2002) has discussed and given some solutions to the rigidity issue.

Jaki et al. (2013) has indicated that curve-free designs would prefer to the CRM if there is little evidence of high enough quality about dose–toxicity relationships. Nevertheless, the number of curve-free designs is overwhelmingly fewer than other types of designs so far. Furthermore, the existing design has some shortcomings such as numerical issues; therefore, it would have been yet a great challenge to develop sophisticated curve-free designs.

2.5 Preparation for leveraging Bayesian optimization

We introduced so far four major statistical designs for dose-finding studies by their types: the 3+3 design in rule-based designs, the CRM in parametric model-based designs, the mTPI in toxicity probability interval designs, and the PBP in curve-free designs. Fig. 2.4 briefly illustrates the design types we explained through sections 2.2 to 2.4.

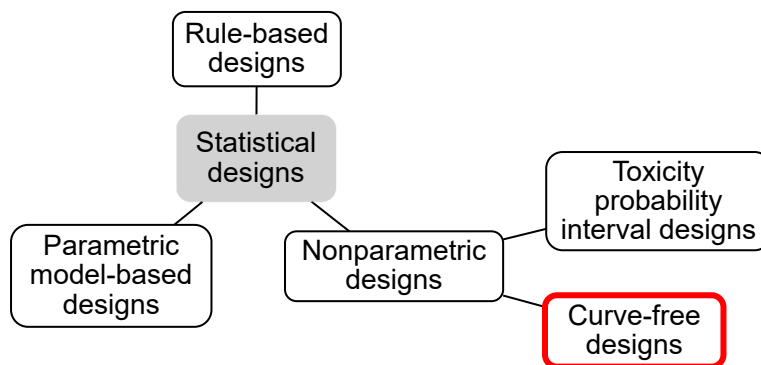


Fig. 2.4. Schematic of statistical design types for dose-finding studies

Given the open research questions we raised to the conventional statistical designs, we develop a new nonparametric Bayesian design that can be categorized in curve-free designs (the red box in Fig. 2.4) by utilizing Bayesian optimization frameworks. To the best of our knowledge, our study is the first to investigate the capability of Bayesian optimization under clinical trials. Bayesian optimization possibly offers the following advantages:

- Bayesian optimization can be applied to dose-finding studies in the same procedures

- as existing statistical designs (e.g., sequential evaluations with cohorts),
- Because of little information available about dose–response relationships, flexible modeling without relying on strong assumptions on the shapes is a reasonable approach. Bayesian optimization realizes flexible modeling on unknown dose–response relationships based on a nonparametric model,
 - Bayesian optimization can reflect non-negligible variability to dose selections based on acquisition functions to avoid skewed dose selections,
 - In general, sample sizes in dose-finding studies are small. Bayesian optimization can seek the MTD within as few sample sizes as possible.

We briefly introduce the mathematical background of typical Bayesian optimization based on Brochu et al. (2010) in the following section.

2.5.1 Typical Bayesian optimization

Bayesian optimization is a sophisticated framework to optimize expensive objective functions that take a lot of resources such as time and cost to evaluate (Shahriari et al., 2016). As it is called a black-box optimization algorithm, it aims to search for the optimum of an unknown objective function while estimating a model of the objective function through a nonparametric Bayesian regression.

Suppose g denotes an unknown objective function. The ultimate goal of Bayesian optimization is to solve the following problem:

$$x^* = \arg \min_{x \in \mathcal{X}} g(x), \quad (2.13)$$

where \mathcal{X} is input space of interest for searching. We are considering finding the global optimizer (minimizer or maximizer) of the unknown objective function g . Typical Bayesian optimization deals with continuous input x . We assume g can be evaluated at any arbitrary point x in the domain regardless of its shape. Although we do not know the shape of g , we observe g through point-wise observations y that might be noise-free or noisy depending on the situation. When treating noisy observations, typical Bayesian optimization usually assumes that observations follow a Gaussian distribution, and an arbitrary observation is described as $y \sim N(g(x), \sigma^2)$, where σ is the variance of noise injected into the function observations.

Model update

We follow Bayes' theorem to estimate the unknown objective function g . A Gaussian process prior is often put as a prior belief over g . The prior distribution of g is given by $g \sim \text{GP}(m, k)$, which is specified by a mean function $m(x)$ and a kernel function $k(x, x')$. Every time we observe y , g is sequentially refined through a Gaussian process regression based on the Bayesian posterior updating so that it can reflect observations. A posterior distribution based on up to t th observations is given by

$$P(\text{d}g \mid D_{1:t}) \propto P(D_{1:t} \mid g)P(\text{d}g), \quad (2.14)$$

where $D_{1:t} = \{x_{(1)}, \dots, x_{(t)}, y_{(1)}, \dots, y_{(t)}\}$. This Bayesian updating goes a conjugate analysis of a Gaussian distribution; therefore, the posterior distribution g at the next

arbitrary point $x_{(t+1)}$ can be analytically calculated and described as follows:

$$\begin{aligned} P(\text{dg}(x_{(t+1)}) \mid D_{1:t}, x_{(t+1)}) &= N(\text{dg} \mid \mu_{(t)}(x_{(t+1)}), \sigma_{(t)}^2(x_{(t+1)})), \\ \mu_{(t)}(x_{(t+1)}) &= \mathbf{k}^\top \mathbf{K}^{-1} \mathbf{g}_{1:t}, \\ \sigma_{(t)}^2(x_{(t+1)}) &= k(x_{(t+1)}, x_{(t+1)}) - \mathbf{k}^\top \mathbf{K}^{-1} \mathbf{k}, \end{aligned} \quad (2.15)$$

where $\mathbf{g}_{1:t} = g(x_{(1:t)})$, $\mathbf{k} = [k(x_{(t+1)}, x_{(1)}), k(x_{(t+1)}, x_{(2)}), \dots, k(x_{(t+1)}, x_{(t)})]$, and the kernel matrix \mathbf{K} is given by

$$\mathbf{K} = \begin{bmatrix} k(x_{(1)}, x_{(1)}) & \dots & k(x_{(1)}, x_{(t)}) \\ & \ddots & \\ k(x_{(t)}, x_{(1)}) & \dots & k(x_{(t)}, x_{(t)}) \end{bmatrix}.$$

Next point selection based on acquisition functions

Once the model is updated, an acquisition function that guides the next point to evaluate is calculated to provide a balance of information between exploitation and exploration on the posterior distribution for g .

There are various acquisition functions for Bayesian optimization. Based on Snoek et al. (2012), the most popular acquisition functions are roughly explained as follows:

- Probability of Improvement (PI) provides a strategy to maximize the probability of improving over currently the best value. This strategy tends to weigh on exploitation more than exploration.
- Expected improvement (EI) provides a strategy to maximize the expected improvement over the current best point (Mockus, 1975; Jones et al., 1998).
- Lower (upper, when considering maximization) confidence bound (LCB) directly balances between exploitation and exploration with the simple sum form and is a strategy to minimize regret over the course of their optimization (Srinivas et al., 2010).

The next point is selected so that it maximizes or minimizes an adopted acquisition function. Snoek et al. (2012) has also provided the form of each acquisition function when observations follow a Gaussian distribution. Our study will use the EI and the LCB that have shown to be better-behaved than PI. From the next chapter, we will explain how we leverage Bayesian optimization for dose-finding studies in detail along with explanations of concrete equations to calculate each acquisition function.

Chapter 3

Bayesian optimization design of MTD estimation for mono-therapies

Chapter 3 introduces Bayesian optimization design, BOD, that applies Bayesian optimization frameworks to MTD estimation problems in dose-finding studies investigating mono-therapies with cytotoxic agents. Chapter 3 covers typical dose-finding studies based on common situations frequently assumed such as assumptions for the mode-of-action owing to agent types, and the number of dose levels. Chapter 3 is constructed based on the works *A curve free method based on Bayesian optimization for oncology Phase I clinical trials*, A. Takahashi and T. Suzuki, Proceeding of 2nd International Conference on Statistics: Theory and Applications, 2020 (Takahashi and Suzuki, 2020); *Bayesian optimization for estimating the maximum tolerated dose in Phase I clinical trials*, A. Takahashi and T. Suzuki, Contemporary Clinical Trials Communications, 2021 (Takahashi and Suzuki, 2021a).

3.1 Basic structure of Bayesian optimization design

Section 3.1 introduces BOD that has the basic structure of BOD introduced throughout this thesis by the following subsections:

- Section 3.1.1 describes how BOD models dose–toxicity relationships,
- Section 3.1.2 shows a dose selection strategy on how BOD sequentially selects MTD candidates,
- Section 3.1.3 mentions overdose control imposed in BOD to address safety considerations, The procedures introduced here play an essential role as a statistical design for clinical trials from the perspective of ensuring patient safety,
- Section 3.1.4 specifically explains the implementation steps of BOD,
- Section 3.1.5 illustrates a brief example with some drawings that we can generate when implementing BOD.

BOD is expected to provide the following advantages; flexible nonparametric modeling in BOD relaxes model misspecification issues because dose–response relationships are expressed without strong assumptions found in parametric model-based designs; the current probabilistic belief on dose–response relationships across all doses can be sequentially refined via Bayes’ rule, which helps understand potential characteristics of investigational agents; uncertainties of estimates contribute to dose selections via an acquisition function that is one of the features of Bayesian optimization frameworks providing a balance of information between exploration and exploitation of the estimates, which would realize efficient searches for the targeted dose without disproportionately concentrating on a local

optimum.

First of all, we introduce BOD in detail from its statistical modeling at the following section.

3.1.1 Statistical model for dose–toxicity relationships

As described in Section 2.5, Bayesian optimization regards unknown objective functions as nonparametric models to seek the global optimizer. BOD also models dose–toxicity relationships based on the nonparametric approach. Let us express dose–toxicity relationships as follows:

$$f(x) = \text{logit}\{\pi(x)\} = \log \left\{ \frac{\pi(x)}{1 - \pi(x)} \right\}, \quad (3.1)$$

where $\pi(x)$ denotes toxicity probabilities at a corresponding conceptual dose x at an arbitrary dose level j ($j \in \{1, \dots, J\}$) in the finite dose range $\mathcal{X} = \{x \in \mathbb{R} \mid x_1 \leq x \leq x_J\}$ and $\pi(x) = [1 + \exp\{-f(x)\}]^{-1}$. The logit transformation for π guarantees that π bounds within the range 0 to 1 on a finite dose range. We note that conceptual doses x are not necessary to be actual dose values because dose–toxicity relationships taken into account in the estimation process only rely on distances between conceptual dose values but not actual values; however, conceptual doses should be as equally spaced with regards to toxicity probabilities as possible so that we could avoid skewed estimation of toxicity.

Since we do not know the exact form of f , we estimate it in the Bayesian manner. As shown in Section 2.5, Bayesian optimization puts prior distributions on objective functions to be optimized because the primary goal is to identify optimal points over the objective functions. BOD does not treat dose–toxicity functions as such objective functions; on the other hand, preliminary estimation on dose–toxicity relationships could be of interest and allow investigators to catch the characteristics of investigational agents. Although there is a limitation to estimate dose–toxicity relationships due to limited sample size in practice. Nonetheless, preliminary estimation could be useful information for subsequent clinical trials in addition to the current dose-finding study. Therefore, BOD firstly prescribes a prior belief over f and estimates it.

According to typical Bayesian optimization that often puts a Gaussian process prior over unknown functions, BOD describes a prior distribution for f by using a Gaussian process as follows:

$$f \sim \text{GP}(m, k), \quad (3.2)$$

where a mean function $m(x)$ and a covariance function $k(x, x')$ specify its stochastic process. A Gaussian process makes a model easy to treat because it leaves setting a prior distribution to designing a kernel function (Rasmussen and Williams, 2006). Although it is possible to apply other than a Gaussian process to prior distributions of f , Mockus et al. (1994) has already explicitly set the framework for a Gaussian process prior in Bayesian optimization. Also, a Gaussian process prior has proven successful in Bayesian optimization, for example, as mentioned in Shahriari et al. (2016). (On a different note, Russu et al. (2011) has utilized a Gaussian process for modeling population pharmacokinetics in dose-finding studies.)

The prior mean function for $m(x)$ can derive from pre-specified initial guesses for toxicity probabilities; however, because little is known about dose–toxicity relationships in general, one option for setting prior values for $m(x)$ might be utilizing an indifference interval determined by a systematic approach proposed for the CRM by Lee and Cheung (2009). Section 2.3 also describes indifference interval for the CRM. We will show specifics about

how to set prior values for $m(x)$ using the indifference interval approach at Step 1 for the model phase in Section 3.1.4.

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

$$k(x, x') = \sigma_f^2 \exp\left(-\frac{1}{2\rho^2} |x - x'|^2\right), \quad (3.3)$$

where σ_f^2 is a signal variance and generally set as 1, which determines the variation of function values from their mean; a scale parameter ρ controls the width of the kernel. The scale parameter expresses the distance between turning points in the explored dose space defined by a setting on conceptual doses. In general, ρ becomes a value within the range of conceptual doses. Because dose–toxicity relationships do not have many turning points, too small a value is inappropriate (e.g., a value of less than 0.1 is too small under the setting where each dose has a 0.2 increment interval per dose level as in our simulation study in Section 3.2). The squared exponential kernel is a very popular choice for Bayesian optimization and Schneider et al. (2017) has reported that the squared exponential kernel is suitable when the unknown function is sufficiently smooth. Dose–toxicity functions under mono-therapies are sufficiently smooth in general; therefore, we adopt the squared exponential kernel in BOD. In order to provide computational stability, a small value ξ , which is similar to noise in a regression model, is added on the diagonal elements of the covariance function (Neal, 1997). As a result, each element in the covariance matrix K for arbitrary conceptual doses x and x' is expressed as

$$K_{x,x'} = k(x, x') + \xi 1[x = x'], \quad (3.4)$$

where the indicator function $1[\cdot]$ returns 1 when the condition in the brackets is satisfied. Otherwise, it returns 0.

After patients are administered a specific dose x in accordance with the study protocol, investigators carefully observe and review patient outcomes. Out of all observations, the number of patients with DLT (y) will be directly involved in dose selections for the next cohort. The binary outcomes on DLT are described by a binomial distribution as $y \sim \text{Bin}(n, \pi)$, where $\text{Bin}(n, \pi)$ abbreviates a discrete probability distribution of the number of successes in a sequence of n independent experiments with a probability π . Then, the likelihood function of observed values up to the t th test denoted by $D_{1:t} = \{(n_{(1)}, x_{(1)}, y_{(1)}), \dots, (n_{(t)}, x_{(t)}, y_{(t)})\}$ is given by

$$L(D_{1:t} | f) = \prod_{z=1}^t \pi(x_{(z)})^{y_{(z)}} \{1 - \pi(x_{(z)})\}^{n_{(z)} - y_{(z)}}, \quad (3.5)$$

where $y_{(z)}$ is the number of patients with DLT in a cohort with $n_{(z)}$ patients treated at $x_{(z)}$ for the z th test.

Based on a Gaussian process prior and the likelihood function in equation 3.5, a posterior distribution for π along with f is updated through the Bayes' rule. In practice, posterior samples are generated for the functions π and f by MCMC.

3.1.2 Dose selection strategy

Suppose that we are interested in finding a dose that produces the closest toxicity to a target toxicity rate θ . So far, we have obtained an updated distribution for a dose–toxicity relationship according to Section 3.1.1. To transrate MTD estimation problems

into optimization problems so that we could utilize Bayesian optimization frameworks, we set an objective function to be optimized in BOD as follows:

$$g(x) = |\pi(x) - \theta|, \quad (3.6)$$

where $\pi(x)$ derives from $f(x)$ according to equation (3.1). Posterior samples for g are then calculated immediately after we get posterior samples for π based on those of f . The exact form of the function g is still unavailable; however, we can leverage those probabilistic beliefs in order to reach the next promising dose $x_{(t+1)}$ through designing an acquisition function \hat{g} that is an alternative of the true objective function g .

While we briefly introduced several acquisition functions in Section 2.5, we select the EI for \hat{g} in this chapter. Although there are many acquisition functions, the EI strategy is often a default choice in popular Bayesian optimization packages (Nguyen et al., 2017) as one of the most popular global optimization algorithms. It is because the EI strategy offers reasonable and consistent performance without the need to choose additional hyperparameters (Qin et al., 2017; Snoek et al., 2012). In addition, Bull (2011) has reported that the EI is efficient in the number of function evaluations required to find the global optimum. The EI strategy automatically balances exploitation and exploration and sequentially searches the point that offers the greatest expected improvement over the current best point. Exploitation means sampling where the prediction of the probabilistic model is high (e.g., posterior means). On the other hand, exploration means sampling at locations where the prediction uncertainty is high. Specifically, it considers both probability and magnitude of improvement over the current best point at candidate doses based on posterior distributions of g . 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

$$I(x) = \max\{0, g^+ - g(x)\}, \quad (3.7)$$

where $g^+ = \min_x [\mathbb{E}_g\{g(x)\} | D_{1:t}]$ that means the current best point providing the minimum value on $g(x)$ among all available doses (Gramacy and Lee, 2011). Accordingly, $I(x)$ provides a positive value if $g(x)$ turns out to be less than g^+ . Otherwise, $I(x)$ becomes zero. Secondly, the EI is calculated as the expectation of $I(x)$ as follows:

$$\text{EI}(x) = \mathbb{E}\{I(x) | D_{1:t}\} = \int_0^1 I(x) P\{g(x) | D_{1:t}\} dg, \quad (3.8)$$

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

$$x_{(t+1)} = \arg \max_{x \in \mathcal{A}_t} \{\text{EI}(x)\}, \quad (3.9)$$

where \mathcal{A}_t is an admissible dose set defined at Section 3.1.3. It imposes some dose-escalation restrictions for overdose control to ensure patient safety.

3.1.3 Overdose control for patient safety

Patients in the next cohort are treated with the selected dose according to equation (3.9) from an admissible dose set that limits candidate dose ranges based on such as toxicity

probabilities to ensure patient safety. Specifically, we impose overdose control on dose allocations during a trial through an admissible dose set \mathcal{A}_t that is refreshed at each test and includes doses that satisfy all conditions as follows:

1. All candidate doses in \mathcal{A}_t require to satisfy $P\{\pi(x) > \theta \mid D_{1:t}\} < c_2$. In addition, if the lowest dose does not satisfy $P\{\pi(x_1) > \theta \mid D_{1:t}\} \leq c_1$, \mathcal{A}_t includes only x_1 .
2. No dose skip is allowed; therefore, the highest dose level is up to one dose level higher than $x_{(t)}$. If $x_{(t)} = x_J$, the highest candidate dose is up to x_J .
3. 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 two or more patients experience DLT at x_1 , \mathcal{A}_t includes only x_1 .

For the first condition, we assume $c_1 \leq c_2$ in order not to miss unsafe situations under the assumption that toxicity increases monotonically with increasing dose levels. The third condition on dose de-escalation might be a typical setting when the CRM is implemented in actual clinical trials.

3.1.4 Implementation steps

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

Start-up phase

The start-up phase follows an algorithm-based procedure that is similar to the traditional 3+3 design 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 a dose higher by one level than the current dose.
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 described through Sections 3.1.1 to 3.1.3. The following list provides the implementation steps during trials as well as some preparations related to the model before the initiation of trials.

1. At the preparation stage of the study protocol, design parameters for a Gaussian process prior are determined as follows:
 - (a) If there is little information about dose–toxicity relationships, initial toxicity guesses generated by the indifference interval approach are available for setting a prior mean function for $m(x)$ as follows:
 - i. According to the systematic approach proposed by Lee and Cheung (2009), we calculate an indifference interval named the optimal δ while using the CRM based on a power model and an initial MTD at the center of the dose range (i.e., $J/2$),

which we could maximize the average percentage of correct MTD selections when implementing the CRM with the selected settings. The main purpose of utilizing the systematic approach is to obtain one of the good guesses on the slope of initial toxicity guesses pertaining to a prior mean function for $m(x)$ in BOD. Smaller δ makes the slope of initial toxicity guesses more gentle. As a result, dose-escalation becomes more aggressive than larger δ in BOD. In contrast, initial toxicity guesses based on a large δ tends to offer more conservative dose-escalation. Empirically, values close to the optimal δ (e.g., a range of $\delta \pm 0.02$) would provide good operating characteristics for BOD.

- ii. Aside from the above determination on δ , an initial MTD location ν on our initial toxicity guess for BOD could be determined based on the results of the start-up phase. Otherwise, we recommend locating an initial MTD 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. Once δ and ν are decided, we can generate the initial guesses whose indifference interval is δ and initial MTD location is the ν th dose level. Specifically, they can be generated with the `getprior` function in the R package `dfcrm` by specifying δ , ν , a power model, and the number of planned dose levels. We put logit transformed values of the generated initial guesses as a prior mean function for $m(x)$.

If there is an informative belief on dose–toxicity relationships, initial guesses should reflect it.

- (b) For a covariance function, the value of σ_f is set as 1 because it is a typical setting in Bayesian optimization and works well in most cases. The value of ρ indicates a typical distance between turning points; thus, it depends on conceptual dose ranges. An appropriate ρ provides at most two turning points in a range because dose–toxicity functions do not have many turning points. In our two stage-design (i.e., a combined design of the model phase after the start-up phase), a good value is often comparable to the length of a dose range.
2. Once we are ready for proceeding to the model phase evaluations, firstly posterior samples for f reflecting observed DLT outcomes are calculated with equations (3.1) through (3.5) based on the Bayes’ rule via MCMC. Secondly, they are transformed with the inverse logit function to obtain posterior samples of toxicity probabilities π for illustrating actual dose–toxicity relationships. A drawing of a posterior distribution on π can help understand the current belief on π at each dose level. After the above calculations, posterior samples for the function g are finally calculated with the posterior samples for π according to equation (3.6).
 3. Patients in the next cohort are treated with the selected dose according to the EI strategy described in equations (3.7) to (3.9).
 4. Steps 2 and 3 continue until the trial condition meets pre-specified stopping rules.
 5. Once the trial is terminated, the MTD (x^*) is determined based on the final posterior distribution for toxicity probabilities as follows:

$$x^* = \arg \max_{x \in \{x_j | \hat{\pi}(x_j) < (\theta + \epsilon_2)\}} P(\theta - \epsilon_1 < \pi(x) < \theta + \epsilon_1 | D_{1,t}), \quad (3.10)$$

where $\hat{\pi}(x_j)$ 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 an acceptable range for θ provided by ϵ_1 from the final MTD candidate set that does not retain doses with excessive toxicities according to the condition with ϵ_2 .

3.1.5 Example with illustrations

Fig. 3.1 illustrates 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.




Fig. 3.1. Posterior toxicity distributions and the associated EIs in BOD

The upper section having three figures obtained after each test provides posterior distributions for toxicity probabilities as its posterior mean (dashed line) and 10 to 90 percentiles area (filled area) along with sample statistics of toxicity probabilities at each tested dose level based on observed data (filled circle). They also include initial toxicity guesses (dotted line) assumed in the model and a true dose–toxicity curve (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 to the next dose level selected by equation (3.9) by arrow marks as well as the dotted vertical 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. In Fig. 3.1, the upper-left section illustrates the updated toxicity probability distribution at the end of the start-up phase, and the lower-left section provides the corresponding EI. The posterior mean function was still close to the initial toxicity guesses. Based on the EI 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 changed slightly after we obtained the patient outcome and then the EI selected x_5 . In the 7th test, there was one patient who experienced DLT. The upper-right section shows that the posterior mean function based on the updated toxicity probability distribution approached the true dose–toxicity curve. After the 7th test, the EI selected x_5 again as the next dose, which is the true

MTD in this example.

3.2 Performance evaluation

We conducted a simulation study to examine the performance of BOD for mono-therapies with cytotoxic agents. Zhou et al. (2018) that compared parametric model-based designs (CRM, dose-escalation with overdose control, and Bayesian logistic regression model) and toxicity probability interval designs (mTPI, BOIN, and keyboard design which operates in the same way as the mTPI-2) has reported that the CRM outperformed the other parametric model-based designs in terms of accuracy of identifying the MTD, and the BOIN outperformed the mTPI and provided comparable performance with the CRM. Similarly, Zhu et al. (2019) that compared the operating characteristics of the 3+3 design, the CRM, the BOIN, and the keyboard design has suggested that both the BOIN and the keyboard design provided comparable performance with the CRM. Horton et al. (2017) has also reported that the CRM tended to outperform the BOIN and mTPI as the number of dose levels increases. There are other studies that compared performance between the CRM and toxicity probability interval designs. Conclusions differ slightly among articles; however, the common point is that the CRM would consistently perform well in various scenarios and toxicity probability interval designs are attractive in their ease of use. Considering the literature review, we set the Bayesian-based CRM as a benchmark for the performance evaluation.

In addition, we included curve free designs that are the same categories of BOD into the evaluation. We selected the PBP and another curve free design proposed by Whitehead et al. (2010). The latter considers probabilities falling into pre-specified toxicity risk categories at each dose level (hereafter, we abbreviate this design WTW based on the initials of the authors). Appendix A.1 provides a brief explanation about the WTW.

Section 3.2.1 explains the simulation frameworks for each design. By following it, the simulation results are presented in Section 3.2.2.

3.2.1 Simulation settings

Suppose that we considered to find the MTD that has a toxicity probability closest to a target toxicity θ among eight dose levels ($J = 8$) under the maximum sample size of 36 and the cohort size of 3. The sample size of 36 was based on the average number of patients in model-guided dose-response studies reported by Iasonos and O'Quigley (2014). Table 3.1 shows fifteen scenarios used for the simulation study, which includes scenarios 1 to 6 that were excerpted from Tang et al. (2018). Fig. 3.2 also depicts those scenarios separately by θ .

Table 3.1. 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

The MTD is displayed as bold type by scenario and θ .

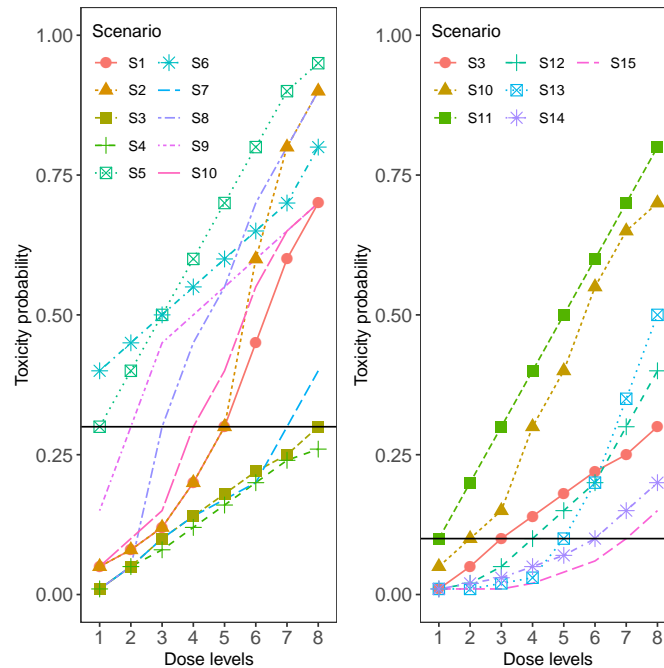


Fig. 3.2. The left figure shows toxicity scenarios for dose-finding studies whose target toxicity rate is 0.3. The right one shows toxicity scenarios for dose-finding studies whose target toxicity rate is 0.1.

We evaluated the operating characteristics of each design 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. Aside from the typical setting, we might encounter a lower target toxicity rate when DLT could be assumed to present severe symptoms. Assuming a lower toxicity rate case, we set $\theta = 0.1$ and evaluated scenarios 3, 10 to 15. Bold toxicity

probabilities in Table 3.1 denote the true MTD under each θ .

The trial started from the lowest dose in all designs. In addition, all designs terminated the trial when the maximum sample size was reached unless otherwise specified. The number of trials to evaluate the operating characteristics of each design was 1,000 in the simulation study.

For performance evaluation, besides descriptive summaries, we calculated mean squared errors (MSEs) and their 95% confidence intervals. MSEs are given by

$$\text{MSE} = M^{-1} \sum_{m=1}^M (\pi_{[m]}^* - \theta)^2, \quad (3.11)$$

where $\pi_{[m]}^*$ is a true toxicity probability at the recommended MTD in the m th trial under each scenario and M denotes the total number of simulations ($M = 1000$).

In the subsequent paragraphs, we detail simulation settings for each design: BOD, the CRM, the PBP, and the WTW in order.

Bayesian optimization design

We generated initial toxicity guesses using the indifference interval approach as described at Step 1 for the model phase in Section 3.1.4. We derived the optimal δ under each target toxicity rate by evaluating possible indifference intervals from 0.01 to 0.15 by 0.01 running 2,000 simulations, where these settings (i.e., the evaluation range and the number of iterations) were based on Lee and Cheung (2009) that is the original article for the systematic approach. 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 BOD. As a result, the equal value to the optimal δ (i.e., 0.05) and a lower value by 0.01 than the optimal δ (i.e., 0.02) seemed suitable for scenarios under $\theta = 0.3$ and $\theta = 0.1$, respectively.

For the initial MTD location ν in the initial guesses under $\theta = 0.3$, we relied on the last observations in the start-up phase as a similar approach to the 3+3 design. Based on the last observations in the start-up phase, we set ν at either the last tested dose level or one dose level lower than the last-tested dose level if more than two patients experience DLT at the last tested dose. Although the 3+3 based approach could be available for $\theta = 0.3$ to determine an initial thought of the MTD location, it would not apply to 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 conceptual doses of $(x_1, \dots, x_8) = (0.0, 0.2, 0.4, 0.6, 0.8, 1.0, 1.2, 1.4)$ 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, $c_1 = 0.5$ and $c_2 = 0.9$ were used under $\theta = 0.3$. The assumption of $c_1 = 0.5$ 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. For $\theta = 0.1$, overdose allocations tended to be occurred more easily than the settings for $\theta = 0.3$ because a relatively small δ was applied for $\theta = 0.1$ in addition to the lower toxicity rate setting; therefore, smaller values seemed to be appropriate for c_1 and c_2 when $\theta = 0.1$ than those for $\theta = 0.3$. As a result, we used $c_1 = c_2 = 0.4$ under $\theta = 0.1$ to avoid overdose allocations. We determined the values of c_2 under each θ based on a balance between MTD selections and overdose allocations.

For MTD determination described in equation (3.10), $\epsilon_1 = 0.05$ 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 estimates $\hat{\pi}_j$. Additionally, $\epsilon_1 = \epsilon_2 = 0.015$ was used for $\theta = 0.1$, where $\theta = 0.1$ and $\epsilon_1 = 0.015$ provided almost as small a ratio as that of $\theta = 0.3$ and $\epsilon_1 = 0.05$. When the target toxicity rate is lower than a typical setting, investigators might consider lowering overdose selections as much as possible from a clinical perspective; therefore, the value of $\epsilon_2 = 0.015$ conservatively set as the same value of ϵ_1 in addition to reducing the possible effect on overdose selections due to the small δ .

As exploratory analyses under $\theta = 0.3$, we evaluated BOD that provided only monotonically increasing dose-toxicity functions (hereafter, we call it BOD-mono) by assuming the following equation instead of equation (3.2):

$$f' \sim \text{GP}(m, k) \times \mathbb{1}[f' : \text{monotonically increasing function}]. \quad (3.12)$$

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 equation (3.12). As another exploratory analysis for BOD (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).

MCMC was implemented using the R package `rstan`. Appendix A.3 provides the R code used to implement BOD in this simulation study.

Continual reassessment method

According to the descriptions in Section 2.3, we calibrated design parameters for the CRM. In the simulation study, we evaluated a power model and a logistic model with an intercept of 3 that are commonly employed for the CRM. Based on Lee and Cheung (2009), the optimal δ was 0.05 for $\theta = 0.3$ and 0.03 for $\theta = 0.1$ when assuming the initial MTD location $\nu = J/2$ under each model. In addition, assuming the prior distribution of a model parameter a follows a normal distribution described as $a \sim \text{N}(0, \sigma_a^2)$, σ_a was calibrated based on Lee and Cheung (2011) to obtain a least informative prior variance σ_a^{LI} . 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 for the CRM can be implemented by R package `dfcrm`.

A product-of-beta prior design

According to Section 2.4.2, hyperparameters for π_j were calculated based on Gasparini and Eisele (2000, 2001). Because the hyperparameter calculations needed initial toxicity guesses, they were also generated by the 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 $\theta = 0.1$, which were the same δ values used in the CRM and BOD.

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

Toxicity risk approach

Following the setting that Whitehead et al. (2010) exemplified, we employed $\kappa = 5$ for a toxicity grid regardless of θ , and then set $\{h_1, \dots, h_5\} = \{0.1, 0.2, 0.3, 0.4, 0.6\}$ for $\theta = 0.3$ and $\{h_1, \dots, h_5\} = \{0.01, 0.05, 0.1, 0.2, 0.3\}$ for $\theta = 0.1$. We used a uniform joint prior for the joint distribution of r_j as shown in Appendix A.1 A dose maximizing the marginal posterior probability that the toxicity risk was equal to h_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. Specifically, only doses that satisfied a condition of $P(r_j = h_5 \mid D_{1:t}) < 0.2$ 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.

3.2.2 Simulation results

How BOD performed compared with the other designs?

- Typical target toxicity rate

Table 3.2 shows the operating characteristics under $\theta = 0.3$ that might be a typical target toxicity rate that we frequently encounter in practice. For observed toxicity, BOD shows the lowest toxicity percentages in all the designs under all scenarios due to lower overdose allocations. BOD treats fewer patients with overdose levels than the other designs in most scenarios. In particular, BOD successfully controls overdose allocations in scenarios 5 and 6 where the true MTD is at the lowest dose level compared with the other designs due to the effect of c_1 . BOD tends to select safer doses than the other designs owing to overdose control such as an admissible dose set. On the other hand, even if the MTD locates at the end of the explored dose range (i.e., scenarios 3 and 4), BOD identifies the MTD correctly approximately 10% more than the CRMs (CRM-p denoting the CRM based on the power model, and CRM-l denoting the CRM based on the logistic model).

As shown in Table 3.2, BOD shows higher correct MTD selection probabilities than the CRM-p in all scenarios. Compared between BOD and the CRM-l, BOD shows higher correct MTD selection probabilities in most scenarios, while scenarios 2 and 7 are comparable results between the two designs, and a lower correct selection probability is shown in BOD under scenario 1. Given the difference between the results of the CRMs in scenario 1, a logistic model might fit more to this scenario. Additionally, the initial MTD location of BOD depends on the last observations in the start-up phase when $\theta = 0.3$. 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. In addition, in scenarios 1 and 2 where the advantages of BOD are minimal, initial guesses for the CRMs at around MTD including adjacent doses overlap at or approach the true dose-toxicity curves closely when parallel-shifted in the vertical direction. In such a case, it is highly likely that an estimated curve around the MTD would be successfully approached to 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 BOD. The CRMs did not allow dose-escalation when patients in the current cohort experienced DLT. Due to this overdose control, the CRMs could not reach the MTD located on higher doses in the dose range as quickly as BOD.

A comparison of BOD and the PBP regarding correct MTD selection probabilities shows that BOD provides lower MTD correct probabilities 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 designs, the PBP shows a higher probability of overdose selection and higher observed toxicity than BOD in this scenario.

The WTW performs very well in one scenario but then provides much poorer performance than the other designs under scenarios 3, 4, and 5. 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 the dose-escalation in 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 the target toxicity rate. This result implies that the cutoff value on an admissible dose set should be carefully decided.

Table 3.2. Operating characteristics under a typical target toxicity rate ($\theta = 0.3$) by each design 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)

design	Scenario	MTD Determination		Dose allocation (%)		Toxicity (%)	Scenario	MTD Determination		Dose allocation (%)		Toxicity (%)
		Correct	Overdose	MTD	Overdose			Correct	Overdose	MTD	Overdose	
BOD	1	0.520	0.189	24.8	8.9	20.2	6	0.940	0.060	89.2	10.8	40.5
CRM-p		0.519	0.163	27.6	10.5	21.7		0.886	0.114	78.5	21.5	41.3
CRM-I		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
BOD	2	0.595	0.097	26.4	6.7	20.6	7	0.329	0.147	11.3	4.2	15.8
CRM-p		0.590	0.082	30.1	7.8	22.2		0.300	0.072	11.0	2.9	16.0
CRM-I		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
BOD	3	0.237	0.000	5.8	0.0	15.2	8	0.695	0.285	48.0	23.2	26.3
CRM-p		0.120	0.000	3.9	0.0	15.7		0.609	0.334	44.2	33.6	29.8
CRM-I		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
BOD	4	0.312	0.000	7.6	0.0	14.2	9	0.658	0.232	46.8	19.1	27.4
CRM-p		0.181	0.000	5.6	0.0	14.5		0.631	0.258	47.8	29.9	31.3
CRM-I		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
BOD	5	0.759	0.241	76.2	23.8	33.4	10	0.592	0.279	32.5	16.0	22.4
CRM-p		0.662	0.338	59.6	40.4	35.1		0.570	0.296	36.3	21.2	24.8
CRM-I		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

Figure 3.3 provides convergence results of MTD selections from the view of toxicity probabilities. The results generally support the MTD selection probabilities explained above. BOD provides comparable performance to CRMs in most scenarios and better convergence especially in scenarios 5, 6, and 8. The PBP and the WTW do not provide consistent results across the wide scenarios; that is, the best MSE is shown in one scenario but the poorest MSE is also shown in the other scenario compared with the other designs. Compared with the other designs, BOD tends to provide stable performance that is not inferior to the other designs in the variety of scenarios in terms of MTD selection accuracy measured by toxicity probabilities.

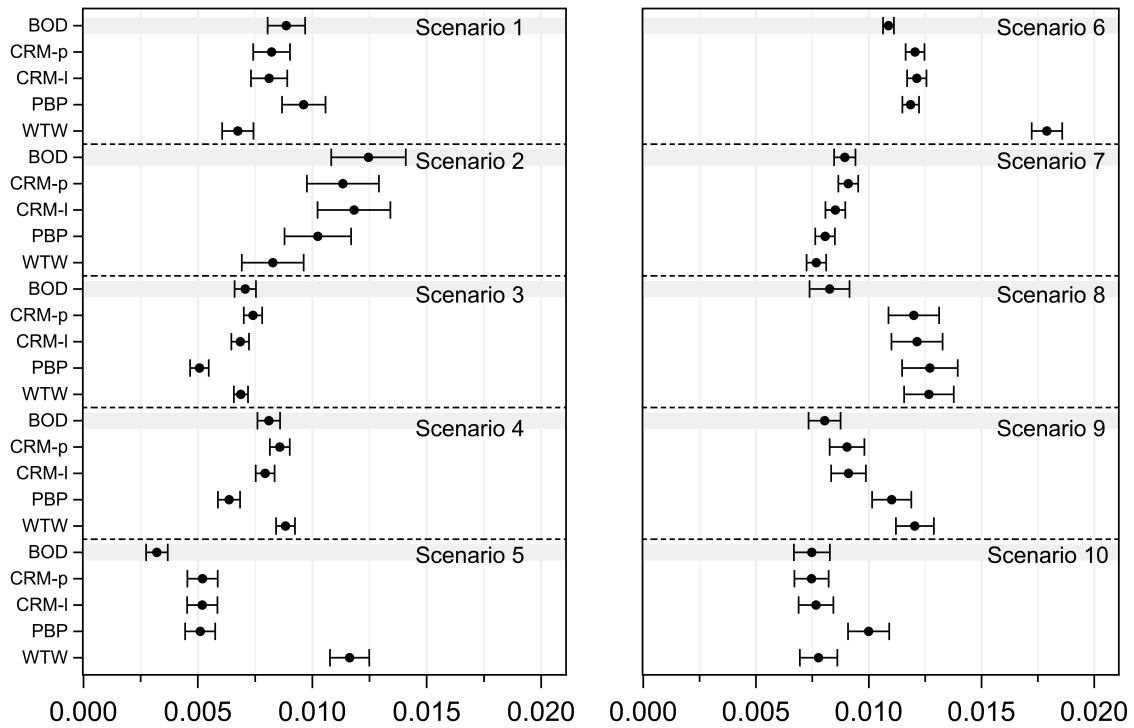


Fig. 3.3. Mean squared errors on toxicity probabilities at the recommended MTD

• Low target toxicity rate

Table 3.3 shows the operating characteristics under a lower target toxicity ($\theta = 0.1$). It supports what Table 3.2 shows. BOD 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. The same trend under scenarios 3, 4, and 7 assuming $\theta = 0.3$ is shown in scenario 15. The difference of dose selection strategies between the two designs might be another reason other than the effect of overdose control regarding why BOD arrives at the true MTD more than the CRMs when the MTD locates near the highest dose. BOD 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 BOD 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 3.3, BOD performs better than the PBPs in most

scenarios. Although concern has been raised that the PBP causes undesirable rigidity in situations where a low toxicity rate is targeted due to its vague priors (Cheung, 2002, 2011), the PBPs did not get stuck at suboptimal dose level under our simulation settings; however, BOD and the CRMs provide better results than the PBPs in most scenarios. In only scenario 11 where the lowest dose level is the true MTD, the PBPs outperform the other designs. Considering that a toxicity probability for 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 the lowest dose toxicity more accurately than the other doses. The WTW provides the best performance in scenarios 12, 13, and 14; however, the correct MTD selection probability at scenario 11 is half as high as the other designs and the overdose selections are shown with high percentages during and at the end of trials.

Table 3.3. Operating characteristics under a lower target toxicity rate ($\theta = 0.1$) by each design 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)

design	Scenario	MTD Determination		Dose allocation (%)		Toxicity (%)	Scenario	MTD Determination		Dose allocation (%)		Toxicity (%)
		Correct	Overdose	MTD	Overdose			Correct	Overdose	MTD	Overdose	
BOD	3	0.409	0.362	25.6	32.7	9.7	13	0.462	0.054	13.1	17.1	7.7
CRM-p		0.369	0.420	30.5	35.6	10.3		0.539	0.145	28.3	16.5	8.0
CRM-I		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
BOD	10	0.495	0.272	38.2	27.7	11.8	14	0.249	0.144	13.4	13.5	6.5
CRM-p		0.464	0.324	35.0	38.7	12.6		0.249	0.180	14.0	12.0	6.5
CRM-I		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
BOD	11	0.780	0.220	68.5	31.5	14.8	15	0.320	0.089	14.1	6.4	4.6
CRM-p		0.783	0.217	61.1	38.9	15.5		0.276	0.156	15.8	6.5	5.0
CRM-I		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	4.6	0.0	3.8	
BOD	12	0.421	0.212	21.7	23.3	8.7						
CRM-p		0.400	0.325	26.9	26.7	9.2						
CRM-I		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						

Figure 3.4 provides a similar relative trend among the designs to Figure 3.3. BOD provides comparable MSEs compared with CRMs in all scenarios and is more stable than the PBP and the WTW across the scenarios.




Fig. 3.4. Mean squared errors on toxicity probabilities at the recommended MTD under a lower toxicity target

How different settings affected BOD performance?

It might be natural to assume that a toxicity probability increases monotonically with increasing dose levels; however, Table 3.4 shows the monotonicity restriction to a Gaussian process prior (i.e., BOD-mono) does not improve the performance compared with BOD 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 of BOD based on equation (3.2) include partially non-monotonically increasing functions; however, posterior distributions, as well as their posterior mean functions, show monotonically increasing shapes owing to follow the trend that true dose–toxicity curves and initial guesses draw. On the other hand, toxicity probabilities in the posterior distributions provided by BOD-mono tend to go down in the lower dose range and go up in the higher dose range than BOD. That is because the posterior distributions are composed of only functions with positive slopes. As a result, the effect of the monotonicity restriction skews the posterior distributions but does not improve the performance of BOD.

Table 3.5 evaluates impact of initial toxicity guesses provided by a different δ on the operating characteristics for BOD that assumes equation (3.2). When the slope becomes gentle by using $\delta = 0.03$, initial toxicity guesses 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 the gentle slope. In contrast, increasing the

Table 3.4. Differences of the operating characteristics between BODs with the monotonicity restriction (BOD-mono) and without the restriction (BOD) under $\theta = 0.3$ (For BOD, the results of BOD in Table 3.2 are re-displayed here.)

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

slope by using $\delta = 0.07$ lowers correct selection probabilities in scenarios 3 and 4. On the other hand, BOD 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 correct selection probabilities and the average toxicity percentages are not affected by the value of δ . Sensitivity analyses for δ , for example around the optimal $\delta \pm 0.02$, might be needed considering the effect of δ .

Table 3.5. Impact of the slope of initial guesses (i.e., $\delta = 0.3, 0.7$) on the operating characteristics of BOD under $\theta = 0.3$ at each scenario

Scenario	δ	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

3.3 Summary and consideration

Chapter 3 has introduced BOD for MTD estimation for mono-therapies with cytotoxic agents assuming typical dose-finding studies. Additionally, performance evaluations of BOD have been presented through simulation studies compared with the CRM, the PBP, and the WTW.

If a theoretical model fits a true dose–toxicity relationship around the MTD through updating a model parameter, the CRM has good performance. Even in that case, BOD could have almost comparable performance in terms of correct MTD selections to the CRM while providing lower toxicity percentages than the CRM. On the other hand, if a theoretical model is far different from a true dose–toxicity curve around the MTD, the performance of the CRM gets lower than that in the other cases. Section 3.2 shows that BOD provides better performance than the CRM, especially in the latter cases. In general, little is known about dose–toxicity relationships; therefore, BOD has the 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 designs, BOD provides more stable results in terms of correct MTD selections and MSEs of toxicity probabilities at the recommended MTD. Additionally, overdose control imposed in BOD works successfully because BOD provides lower or comparable observed toxicity percentages than the other designs in most scenarios.

One of the features in BOD is to select doses based on posterior distributions on dose–toxicity relationships without neglecting their uncertainties. 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 to other areas Bayesian optimization usually applies to. BOD achieves to allocate fewer patients to overdose levels than the CRMs in most scenarios under the restrictions for overdose control while

providing correct MTD selections that are better than or comparable to the CRMs.

BOD has five design parameters in the model (i.e., δ , ν , σ_f , ρ , ξ). Although it is not mandatory to generate initial guesses by using the indifference interval approach, it might be familiar with biostatisticians who belong to pharmaceutical industries. In this case, δ that determines the slope of initial toxicity guesses can refer to the 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 Section 3.2, the initial MTD location ν could be calibrated based on the last observations in the start-up phase or the center of the dose range. For the kernel parameters, σ_f is fixed as a value of 1 as mentioned in Section 3.1.4. 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 equal value to 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., c_1 , c_2 , ϵ_1 , ϵ_2). These values are adjusted based on a balance between correct MTD selections and safe dose allocations during the preparation stage for the study. We should also take into account clinical perspective as well as a typical setting (e.g., $\epsilon_1 = 0.05$ is often used when $\theta = 0.3$.) for the adjustment. The calibration approaches for the design parameters are still open discussion, and they need further evaluations in future work.

In typical Bayesian optimization, observed values are generally assumed to follow a normal distribution as mentioned in Section 2.5. Although the simulation results under the normal approximation assumption are not provided in this thesis, they generally suggested higher doses than the true MTD with high probability and did not work well in various simulation patterns consistently. It implies such a normal approximation for binary 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). According to our exploratory simulation results, the effect of acquisition functions is minimal at least between the expected improvement and the lower confidence bound criteria. We will introduce the LCB strategy for dose-finding studies in Chapter 5.

Although we put monotonically increasing constants on a prior mean function for $m(x)$ and estimate dose–toxicity functions 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 BOD, the model misspecification issue leaves in this case; however, it might be sometimes a reasonable approach when a particular theoretical model is likely to fit true dose–toxicity relationships. We might extend BOD by collaborating with such a parametric approach in future work.

While further discussions and evaluations in particular about how to calibrate design parameters are required, our simulation study suggests that BOD 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 are rarely conducted by only model information. All observations are carefully reviewed by clinical experts for dose selections. Illustrations of posterior distributions of dose–toxicity relationships in BOD such as Fig. 3.1 will be useful information in the comprehensive review by clinical experts.

The scenarios we addressed in this chapter were only monotonically increasing dose–toxicity relationships with one-dimensional dose; thus, the advantages of BOD seem to be minimal in terms of improvement of the MTD selections probabilities although BOD has a benefit to provide safer dose selections with comparable and stable MTD selection

results compared with the other designs. If what we would like to address lies under more complex situations such as two-dimensional inputs or outputs where it is unlikely that the monotonicity assumption can be applied on dose–toxicity relationships, it is expected that the advantages of nonparametric approaches appear more. Chapters 4 and 5 will address applications of BOD to such complex situations.

A

Appendices for Chapter 3

A.1 Toxicity risk approach

Whitehead et al. (2010) has proposed a toxicity risk approach that is one of the curve-free designs for MTD estimation and we call it WTW in this thesis. The WTW deals with a toxicity risk r_j that is a probability with which a patient experiences DLT at a dose level j ($j \in \{1, \dots, J\}$). The r_j is modeled directly and the model assumes r_j is equal to one of a grid of h values ($h_1 < \dots < h_\kappa$). For example, $\{h_1, \dots, h_5\} = \{0.1, 0.2, 0.3, 0.4, 0.6\}$ under a target toxicity rate $\theta = 0.3$ with $\kappa = 5$ provide that a dose level j is interpreted as very safe, safe, ideal, risky, and toxic if r_j is equal to h_1 to h_5 , respectively.

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. A uniform joint prior with monotonicity restrictions is given by a joint distribution of r_j as follows:

$$\pi^0(q_1, \dots, q_J) = \text{P}(r_1 = q_1, \dots, r_J = q_J) = e_0, \quad (\text{A.1})$$

where $q_1, \dots, q_J \in \{h_1, \dots, h_\kappa\}$ and $q_1 \leq \dots \leq q_J$, with $\pi^0(q_1, \dots, q_J) = 0$ otherwise. The value of e_0 can be calculated based on the fact that the joint probability π^0 must sum to 1 over all combinations of q_j .

The posterior joint distribution π corresponding to the uniform joint prior π^0 is described as follows:

$$\pi(q_1, \dots, q_J) \mid D_{1:t} \propto L(r_1, \dots, r_J) = \prod_{z=1}^t r_{(z)}^{y_{(z)}} (1 - r_{(z)})^{n_{(z)} - y_{(z)}}, \quad (\text{A.2})$$

where $n_{(z)}$ patients in the z th cohort have been treated at a dose level (z) with a risk of $r_{(z)}$ resulting in $y_{(z)}$ patients with DLT. According to π , the posterior marginal distribution for each risk at each dose level $\pi_j = \pi_j(q_{\hat{\kappa}}) = \text{P}(r_j = q_{\hat{\kappa}})$ is calculated, where $\hat{\kappa} = 1, \dots, \kappa$.

A dose maximizing the marginal posterior probability that the toxicity risk is equal to the “ideal” risk is 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 is taken into account in the dose selection to judge whether the dose is admissible or not. For example, only doses that satisfy a condition of $\text{P}(r_j = h_\kappa \mid D_{1:t}) < 0.2$ are included in an admissible dose set for the next dose selection, where h_κ means the highest risk (i.e., “toxic” risk) in the toxicity grid.

A.2 Initial idea for expanding BOD to address continuous toxicity variables

BOD focuses on a binary variable as toxicity outcomes in its modeling. On the other hand, it is worth considering altering it to continuous variables in the BOD modeling for MTD estimation. Continuous variables might include for instance biomarker values, pharmacokinetic parameters, and toxicity scores that are a single measurement reflecting various toxicity types or grades.

Appendix A.2 addresses a possible modeling approach for continuous variables. Although all descriptions here are just our initial thought that has not been validated nor confirmed by specific performance evaluations, we hope this will bring some ideas for future work.

Modeling for continuous variables following a Gaussian distribution

A continuous variables for toxicity outcomes could more accurately take patient toxicity burden on cytotoxic agents into account for MTD estimation. One idea to extend BOD to deal with continuous variables could be the following approach: Suppose f denotes an unknown function describing a dose–toxicity relationship on a continuous variable Y . An observation y at a dose level j might be described as follows:

$$y_j = f(x_j) + v_j, \quad (\text{A.3})$$

$$v_j \sim \text{N}(0, \sigma^2), \quad (\text{A.4})$$

where v_j is a Gaussian noise (i.e., $y \sim \text{N}(f(x), \sigma^2)$). According to a similar approach provided in Wang and Ivanova (2015), the variance σ^2 , for example, follows a prior distribution specified by an Inverse Gamma distribution like $\text{IG}(v_0, \sigma^{0^2})$, or could be updated sequentially based on a marginal likelihood maximization approach.

In the same manner in BOD for a binary variable, we set a prior distribution over f using a Gaussian process as $f \sim \text{GP}(m, k)$ and estimate f through Bayes' rule. A prior mean function for $m(x)$ reflects an initial guess on the continuous variable at each dose level. The covariance function, which is specified by a kernel function, could be equation (3.3). According to Section 2.5, a posterior distribution for f at the next arbitrary point can be analytically calculated when y follows a Normal distribution as described in equation (2.15).

The MTD is defined as a dose that returns a pre-specified ideal value y^* for the MTD determination on the continuous variable. Accordingly, we seek a dose that minimizes an absolute value of the difference between $f(x)$ and y^* as follows:

$$x^* = \arg \min_{x \in \mathcal{A}} |f(x) - y^*|. \quad (\text{A.5})$$

Modeling for continuous variables following a non-Gaussian distribution

As a different possible example, another type of continuous variables could be a normalized equivalent toxicity score that has been designed to take toxicity grades into account by Chen et al. (2010). Suppose that Y denotes a normalized equivalent toxicity score for identifying the MTD, it is dealt as a quasi-continuous variable ranging from 0 to 1 and modeled as follows:

$$f(x_j) = \text{logit}(y_j) \quad (\text{A.6})$$

The likelihood function given data denoted as $L(f | D)$ (D is composed of all observations) is described by a quasi-Bernoulli likelihood whose basic formulation has been shown in Chena et al., 2012. Again, a posterior distribution over f having a Gaussian process prior is calculated based on Bayes' rule through MCMC.

In the same manner as the above, we seek a dose that minimizes an absolute value of the difference between y and a pre-specified ideal value y^* as follows:

$$x^* = \arg \min_{x \in \mathcal{A}} | [1 + \exp\{-f(x)\}]^{-1} - y^* | . \quad (\text{A.7})$$

Dose selection

According to objective functions corresponding to equations (A.5) and (A.7), an acquisition function such as the EI or the LCB guides the next appropriate dose during a trial in the same manner in Section 3.1.2. An admissible dose set \mathcal{A} should impose some appropriate conditions to ensure patient safety based on the posterior distribution on f . Once the trial meets one of pre-specified stopping criteria, the MTD is declared as a dose that falls within a certain interval centering y^* on the final posterior distribution for f with a maximum probability for example.

A.3 BOD simulation code for MTD estimation

We implemented BOD using R for the simulation study described in Section 3.2. Appendix A.3 provides R codes for implementing BOD to identify the MTD whose target toxicity rate is 0.3.

- Listing A.1 is R code that constructs the simulation body.
- Listing A.2 presents a `stan` model that is compiled from the R package `rstan` in Listing A.1 to implement the MCMC for obtaining posterior samples.

Listing A.1. Simulation body of BOD for MTD estimation

```

1 #BOD_single_simulation_body.R
2 #####
3 # Bayesian optimization design for identifying the MTD #
4 #####
5
6 #####
7 # Initial Setting #
8 #####
9
10 ### environment settings ###
11 ## change working directory
12 setwd("XXX") # Set your current directory
13
14 ## Load R packages
15 library(rstan) # for MCMC
16 library(dfcrm) # for generating initial toxicity guesses using the
    getprior function
17
18 # We recommend calling the following rstan options
19 rstan_options(auto_write = TRUE) #To avoid recompilation of unchanged Stan
    programs
20 options(mc.cores = parallel::detectCores()) # For execution on a local,
    multicore CPU with excess RAM
21 Sys.setenv(LOCAL_CPPFLAGS = '-march=native') # For improved execution time
22
23
24 ### Function settings ###
25 # Function for calculating absolute values
26 fun2 <- function(fx, p){
27   g1 <- abs(fx - p)
28   return(g1)
29 }
30
31 # Function for calculating logit values
32 fun3 <- function(a){
33   b <- log(a/(1-a))
34   return(b)
35 }
36
37 # Function that modified the rbind function to be able to combine data
    frames with different column names
38 rbindCOrder <- function(...) {
39   n <- length(list(...))
40   temp <- list(...)[[1]]
41   names(temp) <- NA
42   for (i in 2:n) {
43     tmp<-list(...)[[i]]

```

```

44     names(tmp) <- NA
45     temp <- rbind(temp, tmp)
46   }
47   names(temp) <- names(list(...)[[1]])
48   return(temp)
49 }
50
51 #####
52 # Specific settings for the simulation #
53 #####
54
55 ### Scenarios ###
56 ## For example, Scenarios 1~3 are loaded as follows:
57 S1 <- c(0.05, 0.08, 0.12, 0.20, 0.30, 0.45, 0.60, 0.70)
58 S2 <- c(0.05, 0.08, 0.12, 0.20, 0.30, 0.60, 0.80, 0.90)
59 S3 <- c(0.01, 0.05, 0.10, 0.14, 0.18, 0.22, 0.25, 0.30)
60 true.table <- cbind(S1, S2, S3)
61
62 ### Initial settings ###
63 iter <- 1000 # number of simulations
64 cohort <- 3 # number of patients in a cohort
65 N <- length(S1) # number of dose levels
66 dose <- seq(0, 1.4, length = N) # conceptual dose values
67 max_sub <- 36 # maximum number of patients in a trial
68 theta <- 0.3 # a target toxicity rate
69 til <- 0.001 # small values to be used to calculate posterior density
        functions  $P(g(x)|D)$  for the EI (equation (3.8))
70 x <- seq(0, 1, by = (til)) # dummy data of  $g(x)$  for the EI calculations
71
72 # Stan model is compiled for MCMC
73 # Suppose that a stan file 'BOD_model_single.stan' provided after this
        simulation body is stored in the current directory.
74 model_first <- rstan::stan_model(file = 'BOD_model_single.stan')
75
76 # Design parameters for BOD
77 delta <- 0.05 # an indifference interval width for generating initial
        toxicity guesses
78 sigma_f <- 1 # a signal variance in the squared exponential kernel
79 rho <- 1.4 # a scale parameter in the squared exponential kernel
80 xi <- 0.08 # a small value to be added in a covariance function
81 c1 <- 0.5 # overdose control in Section 3.1.3
82 c2 <- 0.9 # overdose control in Section 3.1.3
83 eps1 <- 0.05 # varepsilon_1 in equation (3.10)
84 eps2 <- 0.1 # varepsilon_2 in equation (3.10)
85
86 # Covariance function in the squared exponential kernel
87 K <- matrix(rep(0,N*N), nrow=N, ncol=N)
88 for (i in 1:(N - 1)) {
89   K[i, i] <- (sigma_f)^2
90   for (j in (i + 1):N) {
91     K[i, j] = (sigma_f)^2*exp(-0.5/(rho)^2 * ((dose[i] - dose[j])^2))
92     K[j, i] = K[i, j]
93   }
94 }
95 K[N, N] <-(sigma_f)^2
96 K2 <- K+diag(xi, N) # Add a small value to diagonal elements in K
97 L <- t(chol(K2)) # Cholesky decomposition
98
99
100 #####
101 # Simulation iteration part #

```

```

102 #####
103
104 ### Select Scenarios ###
105 for(S in 1:1){
106
107   true <- true.table[, (S)] # true toxicity probabilities
108
109   for(sim in 1:iter){
110
111     # Initial values for output vectors
112     tox_n <- numeric(N) # number of patients with DLT by dose
113     tox_p <- rep(NA, N) # sample statistics of observed toxicity by dose
114     sum_n <- numeric(N) # number of patients already treated in a trial by
      dose
115     next_l <- 1 # the next dose (start from the lowest dose level 1)
116     total_n <- 0 # total number of patients treated in a trial
117     p_DLT <- 0 # sample statistics of observed toxicity in a trial
118     num <- 0 # number of tests in a trial
119     first <- 0 # If the start-up phase terminated, first=1. Otherwise,
      first=0.
120
121
122     #####
123     # Conduct one trial #
124     #####
125
126     while(total_n < max_sub){ # repeat tests until a trial reaches the
      maximum sample size
127
128       #####
129       # Start-up phase #
130       #####
131
132       while(first==0){
133
134         test_level <- next_l # dose level for the next cohort
135         tox_num <- rbinom(1, cohort, true[test_level]) # number of
          patients with DLT in a cohort
136         tox_n[test_level] <- tox_n[test_level] + tox_num
137         sum_n[test_level] <- sum_n[test_level] + cohort
138         tox_p[test_level] <- tox_n[test_level]/sum_n[test_level]
139         total_n <- sum(sum_n)
140         num <- total_n/cohort
141
142         # retain outputs
143         if(num == 1){obs <- cbind(test_level, tox_num)}else{
144           obs <- rbind(obs, cbind(test_level, tox_num))
145         }
146
147         if(next_l == N || sum(tox_n) >= 2){ # end the start-up phase
148           first <- 1
149           t<-1
150
151           # initial MTD location for a target toxicity rate of 0.3
152           if(tox_num > 2){ # considering only the last DLT observations
153             if(next_l == 1){
154               locate <- 1
155             }else{
156               locate <- next_l - 1
157             }
158           }else{

```

```

159     locate <- next_l
160   }
161
162   # generate initial toxicity guesses
163   guess <- getprior(delta, theta, locate, N)
164
165   # The following code was used instead when theta = 0.1.
166   # locate <- N/2
167   # guess <- getprior(delta, theta, N/2, N)
168
169   logit_guess <- fun3(guess)
170
171   # dose-escalation in the start-up phase
172 }else if(tox_num == 0){
173   next_l <- test_level+1
174 }else if(tox_num == 1){
175   next_l <- test_level
176 }
177 }
178
179 #####
180 # Model phase #
181 #####
182
183 if(t > 1){
184   test_level <- next_l # dose level for the next cohort
185   tox_num <- rbinom(1, cohort, true[test_level])
186   tox_n[test_level] <- tox_n[test_level] + tox_num
187   sum_n[test_level] <- sum_n[test_level] + cohort
188   tox_p[test_level] <- tox_n[test_level]/sum_n[test_level]
189   total_n <- sum(sum_n)
190   num <- total_n/cohort
191   obs <- rbind(obs,cbind(test_level, tox_num))
192 }
193
194 t<-t+1
195
196 cat("\n
197   *****\n
198   Scinario: ",S,"( model phase )\n
199   Simulation number - Test number: ",sim,"-", num,"\n
200   *****\n")
201
202 ### Update toxicity distributions ###
203 y.obs <- as.array(obs[, 2]) # number of DLT at each test
204 t.dose <- as.array(obs[, 1]) # test levels at each test
205
206 # load data for MCMC
207 data <- list(N = N, L = L, logit_guess = logit_guess, Y = y.obs,
208             tested = t.dose, Num = num, Cohort = cohort)
209
210 # run mcmc sampling
211 fit <- rstan::sampling(model_first,
212                       data=data,
213                       pars=c('Ep'),
214                       iter=9000, warmup=3500, chains=1, thin=10
215                       )
216
217 ms <- rstan::extract(fit)
218
219 # posterior samples for toxicity probabilities

```



```

219     d_mcmc <- exp_stars <- data.frame(Ep = ms$Ep)
220
221     # posterior mean for toxicity probabilities
222     exp_star <- apply(exp_stars, 2, mean)
223
224     # calculate cumulative distributions
225     ove <- rep(0, N)
226     pro <- rep(0, N)
227
228     for(n in 1: N){
229         exp_stars1 <- sort(exp_stars[, n])
230         Fn <- ecdf(exp_stars1) # empirical cumulative distribution
                functions
231         ove[n] <- 1 - Fn(theta) # probabilities that each dose is higher
                than the target toxicity rate
232         pro[n] <- Fn(theta + eps1) - Fn(theta - eps1) # probabilities
                that each dose falls within the proper interval including the
                target toxicity rate
233     }
234
235     ### overdose control at bullet # 3 in Section 3.1.3 ###
236     accep_maxd2<-N
237     if(tox_num > 1 ){
238         accep_maxd2 <- test_level - 1
239         if(accep_maxd2 == 0) accep_maxd2 <- 1
240     }
241
242     ### calculate the expected improvement ###
243     EI <- numeric(N)
244     gx <- fun2(fx = exp_stars, p = theta) # posterior samples for g
245     g_best <- min(apply(gx, 2, mean), na.rm = TRUE) # the current best
                value for g(x)
246
247     # improvement function described as equation (3.7)
248     I <- g_best - x
249     Imp0 <- which(I < 0)
250     I[Imp0] <- 0
251
252     for(j in 1:N){
253         # fit a spline regression to obtain a density function for g at
                dose level j
254         gh <- splinefun(density(x = as.vector(t(gx[j]))), n = 1000,
                from = 0, to = 1), method = "natural")
255
256         # probability density function (pdf) of g at dose level j
257         start <- -til/2
258         til_n <- 1 # element number when dividing a pdf of g by small
                widths
259
260         p_gx <- numeric(length(x)) # vector for the pdf
261
262         while(start < 1){
263             p_gx[til_n] <- integrate(gh, start, start + til)$value
264             if(p_gx[til_n] < 0){ p_gx[j] <- 0} # In order to correspond
                very small values near to 0 (there is no negative value in
                practice)
265             start <- start + til
266             til_n <- til_n + 1
267         }
268
269         # calculate expected improvement values by equation (3.8)
270         EI[j] <- I%*%p_gx

```

```

271     }
272
273     next_l <- which(EI == max(EI)) # equation (3.9)
274     if(length(next_l) > 1) next_l <- next_l[1]
275
276     ### overdose control at bullet #1 in Section 3.1.3 ###
277     eli_dose<-0
278     for(j in 1:N){
279         if(ove[j] < c2){eli_dose <- c(eli_dose,i)}
280     }
281     eli_dose <- eli_dose[-1]
282
283     if(length(eli_dose) == 0){
284         accep_maxd <- 1
285     }else{
286         accep_maxd <- max(eli_dose)
287     }
288
289     accep_maxd <- min(accep_maxd2, accep_maxd)
290
291     ### overdose control at bullet #2 in Section 3.1.3 ###
292     if(next_l > (test_level + 1 )) {
293         next_l <- test_level+1
294     }
295
296     if(next_l > accep_maxd){
297         next_l <- accep_maxd
298     }
299
300     ### overdose control at bullet #1 in Section 3.1.3 ###
301     if(ove[1] > c1){
302         next_l <- 1
303     }
304 } # end while loop
305
306 ### Select the MTD ###
307 # final MTD candidate set
308 admiss2_<- which(exp_star < (theta + eps2))
309 if(length(admiss2_) == 0){
310     MTD <- 1
311 }else{
312     MTD <- which(pro[admiss2_] == max(pro[admiss2_])) # equation (3.10)
313 }
314
315 ### Output Results ###
316 p_DLT <- sum(tox_n)/sum(sum_n)*100
317
318 # make a list of outputs
319 Result_box <- cbind(theta, max_sub, S, sim, MTD, S, total_n, data.frame
320 (t(sum_n),total_DLT=sum(tox_n), p_DLT, data.frame(t(tox_n)), t(ove
321 ), t(pro), t(EI), t(exp_star), delta, locate, rho, sigma_f, xi, c1,
322 c2, eps1, eps2)
323
324 # write results
325 if(sim == 1){Result_box1 <- Result_box}
326 else{Result_box1 <- rbindCOrder(Result_box1, Result_box)}
327
328 write.table(Result_box1, "output.txt", row.names = F, quote = F,append
329 = F)

```

```

328
329 } # end iterations of one trial
330
331 } # end scenarios' loop

```

A stan model file ‘BOD_model_single.stan’ is composed of the following contents.

Listing A.2. Stan code for BOD of the MTD estimation

```

1 //BOD_model_single.stan
2 data {
3     int<lower = 1> N; // number of dose levels
4     int<lower = 1> Num; // number of tests in a trial
5     int<lower = 1> Cohort; //cohort size
6     vector[N] logit_guess; // initial toxicity guesses (logit scale)
7     int<lower = 1> tested[Num]; // tested dose levels
8     int<lower = 0> Y[Num]; // number of patients with DLT at each
          test
9     matrix[N,N] L; // the squared exponential kernel (Cholesky
          decomposition)
10 }
11
12 transformed data {
13     vector[N] mu_zero = rep_vector(0,N);
14 }
15
16 parameters {
17     vector[N] zero_m;
18 }
19
20 // Specify a parameter we need to obtain as an output
21 transformed parameters {
22     vector[N] Ep;
23     Ep = inv_logit(logit_guess + zero_m); # posterior toxicity
          probabilities
24 }
25
26 model { // Gaussian process prior
27     zero_m ~ multi_normal_cholesky(mu_zero, L);
28
29     // likelihood function
30     for (i in 1:Num){
31         target += binomial_lpmf(Y[i] | Cohort, Ep[tested[i]]);
32     }
33 }

```

Chapter 4

Bayesian optimization design of MTDC estimation for combination drug therapies

Chapter 4 introduces BOD for estimating a single MTD combination (MTDC) for dose-finding studies of combination drug therapies. Mono-therapies we treated in Chapter 3 have one-dimensional dose as input data for MTD estimation. BOD for mono-therapies in Chapter 3 has one-dimensional doses as input data for MTD estimation. On the other hand, Chapter 4 deals with two different agents assuming combination drug therapies that result in two-dimensional input data although output data from patients (i.e., toxicity outcomes) are the same as Chapter 3. We organize this chapter as follows:

- Section 4.1 briefly reviews conventional statistical designs of dose-finding studies for combination drug therapies to clarify the current issues we should address. Additionally, we explain two representative existing statistical designs that are evaluated in our simulation study in Section 4.3.
- Section 4.2 introduces BOD for MTDC estimation in terms of its statistical modeling, implementation steps, and some illustrations based on a brief example to make readers understand easily.
- Section 4.3 explains settings for the simulation study for performance evaluation and provides the results. We evaluate the operating characteristics of BOD for MTDC estimation compared with the BOIN's and CRM's extension designs.
- Section 4.4 summarizes and discusses performance evaluation results on BOD along with some topics of interest for further work.

We introduce this chapter based on the work *Bayesian optimization design for finding a maximum tolerated dose combination in Phase I clinical trials*, A. Takahashi and T. Suzuki, The International Journal of Biostatistics, 2021 (Takahashi and Suzuki, 2021b).

4.1 Overview of drug development and statistical designs for dose-finding studies of combination drug therapies

4.1.1 Dose-finding studies for combination drug therapies

As Mokhtari et al. (2017) has mentioned that combination drug therapies that combine two or more agents are a cornerstone of cancer therapies, the development of combination drug therapies for cancer treatment has become more commonplace because combination drug

therapies can offer synergistic and additive anti-cancer benefits, unlike mono-therapies. In particular, resistant patients to the treatment of traditional mono-therapies can reap such additional benefits. Additionally, combination drug therapies potentially reduce drug resistance.

In the same manner in Phase I clinical trials for mono-therapies, those for combination therapies primarily aim to identify an MTDC. The underlying premise in conventional statistical designs for dose-finding studies of mono-therapies with cytotoxic agents is that toxicity increases monotonically with increasing dose levels. In combination drug therapies, however, we cannot simply apply such monotonicity assumptions because complex drug–drug interactions make complete toxicity orders unclear in the dose combination matrices. Dose–toxicity relationships in combination drug therapies are much more complex than those for mono-therapies. The complexity of dose-finding studies for combination drug therapies is caused not only by the unknown dose–toxicity relationships but also by the possible existence of multiple MTDCs in the dose combination matrices.

4.1.2 MTDC definition

Before introducing statistical designs, let us clarify the MTDC definition and the difference between MTD and MTDC. Suppose that we investigate a combination therapy of two agents (agent 1 and agent 2) with $I \times J$ combinations and the primary objective of the trial is to identify a single MTDC.

For dose-finding studies for mono-therapies with cytotoxic agents, the goal is to identify the MTD (x^*) that will produce DLT with a target toxicity rate of θ . The relationship between x^* and θ is described as

$$P(Y = 1 \mid x = x^*) = \theta, \quad (4.1)$$

where x denotes a variable of doses and a random variable of Y returns 1 when a patient experiences DLT, otherwise it returns 0. The MTD x^* is usually a particular dose at a single dose level that is the maximum dose with acceptable toxicity because toxicity monotonically increases with increasing dose levels.

In the same way, an MTDC \mathbf{x}^* for combination drug therapies can be described as follows (Jimenez et al., 2019):

$$P(Y = 1 \mid \mathbf{x} = \mathbf{x}^*) = \theta, \quad (4.2)$$

where $\mathbf{x} \in \{\mathbf{x}_{1,1}, \dots, \mathbf{x}_{I,J}\}$; $\mathbf{x}_{i,j} = (x_i^{(1)}, x_j^{(2)})$ expresses a dose combination composed of two discrete conceptual doses of a dose level i for agent 1 and a dose level j for agent 2, which is picked up from two-dimensional dose combination space ($i \in \{1, \dots, I\}$, $j \in \{1, \dots, J\}$).

One of the different points between “MTD” and “MTDC” is the latter could be a subset of the possible dose combinations that have a probability of DLT for a patient equal to θ in the explored dose combination matrices. A single MTDC in this chapter means a dose that achieves equation (4.2); therefore, we note that it is not necessary to be the maximum dose combination in the possible subset.

4.1.3 Bayesian optimal interval design

As a strategy for finding a single MTDC or MTDCs, a simple way is to transform the two-dimensional dose-finding problem into a one-dimensional dose-finding problem. Zhang and Yuan (2016) has introduced the waterfall design to find an MTD contour by dividing a trial into a series of sub-trials and applying a one-dimensional approach to each sub-trial. The

authors utilized, for example, the BOIN design that is one of toxicity probability interval designs like the mTPI in Section 2.4.1 and has been introduced by Liu and Yuan (2015) and Yuan et al. (2016) for MTD estimation as a one-dimensional approach. Decreasing dimensions is straightforward, but on the other hand, such designs generally need more patients for a trial.

Lin and Yin (2017a) has extended the BOIN design directly to drug combination situations without dividing a trial into sub-trials to find a single MTDC. Hereafter, we call this extended version ‘BOIN’. As mentioned in the previous paragraph, the BOIN is categorized in toxicity probability interval designs, which is the same category as the mTPI; accordingly, basic assumptions on toxicity probabilities are similar to the mTPI shown in Section 2.4.1. The following are specifics on the BOIN for MTDC estimation.

Toxicity probability distribution

Similar to the mTPI, the BOIN also constructs a beta-binomial model on toxicity probabilities at each dose combination. Suppose $\pi_{i,j}$ is a toxicity probability at a dose combination (i, j) where $i \in \{1, \dots, I\}$ and $j \in \{1, \dots, J\}$. In the same way as equation (2.7), the prior distribution of $\pi_{i,j}$ is given by

$$\pi_{i,j} \sim \text{Beta}(\alpha_{i,j}, \beta_{i,j}), \quad (4.3)$$

where $\alpha_{i,j}$ and $\beta_{i,j}$ are hyperparameters for a beta distribution. A vague beta prior is usually assumed.

Once patient outcomes are obtained, a posterior toxicity probability for the current dose combination level (i, j) is updated based on a beta-binomial model such that

$$\pi_{i,j} \mid D_{1:t} \sim \text{Beta}(\alpha_{i,j} + y_{i,j}, \beta_{i,j} + n_{i,j} - y_{i,j}), \quad (4.4)$$

where $y_{i,j}$ patients experienced DLT out of $n_{i,j}$ patients at the dose combination (i, j) .

Dose selection

Unlike the mTPI, the next dose combination is selected by a sample statistic of the toxicity probability at the current dose combination based on optimal escalation and de-escalation boundaries instead of the upper probability mass.

Optimal boundaries are obtained by minimizing the chance of making incorrect decisions. Firstly, investigators specify ϕ_1 denoting the highest DLT rate requiring dose escalation and ϕ_2 denoting the lowest DLT rate requiring dose de-escalation. Liu and Yuan (2015) recommended default values of $\phi_1 = 0.6\theta$ and $\phi_2 = 1.4\theta$ for general use. Alternatively, we can adjust those values to set conservative or aggressive dose selections as needed; however, making a range between $[\phi_1, \phi_2]$ too narrow leads to be difficult to differentiate θ from the probabilities that are close to it especially under small sample size trials. The BOIN formulates the following three hypotheses: $H_{0i,j} : \pi_{i,j} = \theta$; $H_{1i,j} : \pi_{i,j} = \phi_1$; and $H_{2i,j} : \pi_{i,j} = \phi_2$. H_0 indicates the current dose combination is the MTD; H_1 indicates the current dose combination is subtherapeutic; H_2 indicates the current dose combination is overly toxic.

The correct decisions under the three hypotheses H_0, H_1 , and H_2 are ‘Stay at the same dose combination (\mathcal{S})’, ‘Escalation (\mathcal{E})’, and ‘De-escalation (\mathcal{D})’ respectively. Correspondingly, the incorrect decisions under H_0, H_1 , and H_2 are other than the correct decisions, which are described as $\bar{\mathcal{S}}$, $\bar{\mathcal{E}}$, and $\bar{\mathcal{D}}$ respectively. The probability of making incorrect

decisions at the current dose combination based on the boundary (λ_1, λ_2) is given by

$$\begin{aligned}
\mathcal{C}^T(\lambda_1, \lambda_2) &= P(H_{0i,j})P(\bar{\mathcal{S}} | H_{0i,j}) + P(H_{1i,j})P(\bar{\mathcal{E}} | H_{1i,j}) + P(H_{2i,j})P(\bar{\mathcal{D}} | H_{2i,j}) \quad (4.5) \\
&= P(H_{0i,j})P(\hat{\pi}_{i,j} \leq \lambda_1 \text{ or } \hat{\pi}_{i,j} \geq \lambda_2 | H_{0i,j}) + \\
&\quad P(H_{1i,j})P(\hat{\pi}_{i,j} > \lambda_1 | H_{1i,j}) + P(H_{2i,j})P(\hat{\pi}_{i,j} < \lambda_2 | H_{2i,j}) \\
&= P(H_{0i,j})P(y_{i,j} \leq n_{i,j}\lambda_1 \text{ or } y_{i,j} \geq n_{i,j}\lambda_2 | H_{0i,j}) + \\
&\quad P(H_{1i,j})P(y_{i,j} > n_{i,j}\lambda_1 | H_{1i,j}) + P(H_{2i,j})P(y_{i,j} < n_{i,j}\lambda_2 | H_{2i,j}) \\
&= P(H_{0i,j})\{F(n_{i,j}\lambda_1; n_{i,j}, \theta) + 1 - F(n_{i,j}\lambda_2 - 1; n_{i,j}, \theta)\} + \\
&\quad P(H_{1i,j})\{1 - F(n_{i,j}\lambda_1; n_{i,j}, \phi_1)\} + P(H_{2i,j})\{F(n_{i,j}\lambda_2 - 1; n_{i,j}, \phi_2)\},
\end{aligned}$$

where $F(y; n, \theta)$ is a cumulative distribution function of a binomial distribution based on the parameters of size n and a probability θ at an observation y ; $\hat{\pi}_{i,j} = y_{i,j}/n_{i,j}$ is a sample statistic of a toxicity probability at the current dose combination (i, j) . Based on equation (4.5), optimal boundaries based on λ_1 and λ_2 can be calculated by minimizing the probability of incorrectly assigning patients to subtherapeutic or overly toxic doses (i.e., $\mathcal{C}^T(\lambda_1, \lambda_2)$), thereby optimizing for patient ethics. In the simplest case where we assign non-informative prior to three hypotheses (i.e., $P(H_{0i,j}) = P(H_{1i,j}) = P(H_{2i,j}) = 1/3$), the boundaries always satisfy the property of $\phi_1 < \lambda_1 < \theta$ and $\theta < \lambda_2 < \phi_2$. Liu and Yuan (2015) shows optimal boundaries corresponding to some commonly encountered target toxicity rates under assuming the non-informative prior on the hypotheses.

Based on the optimal boundary, the next dose combination is decided as follows:

1. When a sample statistic at the current dose combination falls within the optimal boundaries $[\lambda_1, \lambda_2]$, the next dose combination stays at the current dose combination.
2. When a sample statistic at the current dose combination is equal to or greater than the upper boundary λ_2 , the next candidate is a de-escalated dose combination (i.e., $(i-1, j)$ or $(i, j-1)$).
3. Otherwise, the decision for the next dose combination is to escalate a dose combination (i.e., $(i+1, j)$ or $(i, j+1)$).
4. In the above escalation and de-escalation decisions, either dose combination is selected whichever has the highest posterior probability that the next candidate falls within the boundaries.

As shown above, the BOIN is transparent and accessible to non-statisticians because decision-making relies on only sample statistics and simple to implement.

Overdose control is usually performed in the dose selection so that the current dose combination and higher dose combinations with potentially excessive toxicity are eliminated from the set of the next candidates because of safety concerns. Specifically, if a posterior toxicity probability greater than θ at the current dose combination is greater than 0.95 when $n_{i,j}$ is at least three patients, the current dose combination and the higher dose combinations are eliminated from a trial. The cutoff value of 0.95 is the recommended value for general use.

The cycle of observing patient outcomes, updating toxicity distributions of the current dose combination, and selecting the next candidate is repeated until the trial meets either of pre-specified stopping rules. Once the trial condition meets pre-specified stopping rules, a single MTDC with toxicity closest to θ is selected based on isotonic estimates described in a matrix form of toxicity probabilities. If there is a tie, the higher dose combination is selected when the isotonic estimate is lower than the target toxicity rate.

Open research questions

We have raised open research questions for toxicity interval designs for MTD estimation in Section 2.4.1. We can also apply the same to the BOIN. One is the limited allowance window for dose selections in the BOIN. The BOIN offers only the adjacent dose combinations of the current dose combination in its dose selection process. This works successfully from the view of overdose control; however, the limited range of dose selection might make it difficult to reach the higher dose combinations quickly depending on true dose-toxicity scenarios. The other is the BOIN provides the dose-toxicity relationships across tested dose combinations by isotonic estimates at the last stage of trials; thus, we do not obtain the updating beliefs on dose-toxicity relationships across a dose combination matrix during a trial. These points can be addressed in our work.

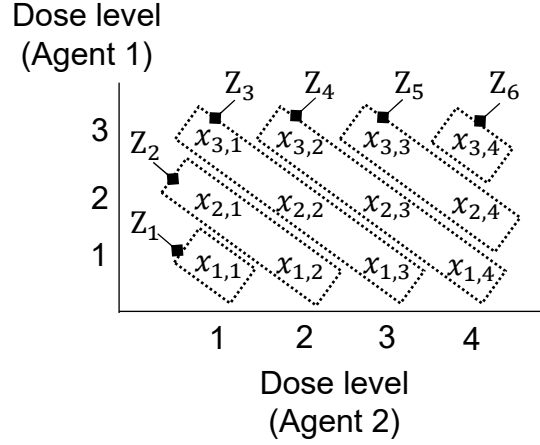
4.1.4 Partial ordering continual reassessment method

Wages et al. (2011) and Wages and Conaway (2013) have developed the partial ordering continual reassessment method (POCRM) that also addresses the two-dimensional problem directly without dividing a trial into sub-trials. The POCRM has extended the most popular parametric model-based approach for mono-therapies called the CRM explained in Section 2.3 to MTDC estimation for combination drug therapies.

Besides the POCRM, Yin and Yuan (2009b), Riviere et al. (2014), and Diniz et al. (2017) have proposed other parametric model-based designs using a copula regression model, a logistic regression model, and a model re-parameterized in terms of the probability of toxicity under a known link function, respectively. Of parametric model-based designs that are a major class in statistical designs for dose-finding studies for MTDC estimation, the POCRM has been frequently evaluated as a competitor in several articles. Riviere et al. (2015a) evaluated the performance of rule-based designs (up and down design for combination drug therapy and up and down with T-statistics) and parametric model-based designs (the POCRM, the copula model, and the Gumbel model). They reported that the parametric model-based designs performed better than the rule-based designs in terms of correct selections of a single MTDC and there was no major difference among the compared parametric model-based designs. Additionally, Hirakawa et al. (2015) has compared five parametric model-based designs, which were the copula regression model, the hierarchical Bayesian model, the shrinkage logistic model, the POCRM, and the order-restricted interface method. According to their simulation results, although each of the designs performed well for various scenarios, on average the hierarchical Bayesian model and the POCRM yielded the highest percentage of correct selections of a single MTDC.

Start-up phase

Wages et al. (2011) has introduced the POCRM based on a two-stage design composed of a start-up phase and a model phase. Other proposed designs for dose-finding studies for combination drug therapies such as Wang and Ivanova (2005); Yin and Yuan (2009b); Wages and Conaway (2013); Riviere et al. (2014); Zhang and Yuan (2016); Wages (2017) also include a start-up phase for MTDC estimation, while the strategy varies among designs. The start-up phase in the POCRM is based on the "toxicity zone" that guides in choosing orders that specify a potential dose-toxicity relationship before a model phase is conducted. Fig. 4.1 describes an example of toxicity zone described as Z . The POCRM considers that toxicity increases among zones in ascending order as $Z_1 < \dots < Z_6$, while we do not know the toxicity order in the internal of each zone.

Fig. 4.1. Toxicity zoning for 2 agents with 3×4 combinations

In the start-up phase, the initial escalation scheme proceeds according to zones, testing dose combinations by ascending order of Z until the first DLT is observed. When a particular zone has more than one dose combination, either dose combination is chosen randomly from the zone. Each zone is not allowed to be skipped until all dose combinations in the previous zone have been allocated to a cohort of patients (e.g., Dose combinations in Z_3 can be tested only after all dose combinations in Z_1 and Z_2 are tested.). Once the start-up phase is terminated due to the first DLT, the model phase starts. As an exception, the highest dose combination is declared the MTDC without the model phase if all patients are treated with no DLT at the start-up phase.

Model phase

After the start-up phase is terminated and trial conditions have not reached either stopping criteria yet, the design switches to the model phase based on the power model:

$$\pi(x_{i,j}) \approx \psi_s(x_{i,j}, a) = \{\pi_{i,j}^0(s)\}^a, \quad (4.6)$$

where $\pi_{i,j}^0(s)$ is an initial guess for a toxicity probability at a conceptual dose combination value $x_{i,j}$ ($x_{i,j} \in \mathbb{R}$) under an ordering s ; a denotes a model parameter.

For the ordering s , Wages and Conaway (2013) recommended to use six possible orderings named across rows, up columns, up diagonals, down diagonals, alternating down-up diagonals, and alternating up-down diagonals (i.e., $s = 1, \dots, 6$ in this case). Prior probabilities for the orderings are decided and described as $P(s) = \{P(1), \dots, P(S)\}$, where $P(s) \geq 0$, $\sum_s P(s) = 1$, and $s \in \{1, \dots, S\}$. The simplest case that works well in practice is $P(s) = 1/S$.

Skeleton values for $\pi_{i,j}^0(s)$ are selected by investigators. Alternatively, they could be calculated with the indifference interval approach that the CRM often uses as described in Section 2.3. If the indifference interval approach is used for generating skeleton values, one set of possible toxicity probabilities can be generated; therefore, initial guesses for each ordering s are decided by rearranging the skeleton values to correspond s after the generation.

After patients in the t th cohort are treated, the likelihood function on all observations $D_{1:t}$ is given by

$$L_s(D_{1:t} | a) = \prod_{z=1}^t \psi_s(x_{(z)}, a)^{y_{(z)}} \{1 - \psi_s(x_{(z)}, a)\}^{n_{(z)} - y_{(z)}}, \quad (4.7)$$

where $y_{(z)}$ patients out of $n_{(z)}$ patients in the z th cohort who were treated with a dose combination denoted by $x_{(z)}$ experienced DLT.

For selecting the next dose combination, a single ordering s^* is chosen based on ordering weights given by

$$\varpi(s) = \frac{\{L_s(D_{1:t} | a) |_{a=\hat{a}_s}\}P(s)}{\sum_{s=1}^S \{L_s(D_{1:t} | a) |_{a=\hat{a}_s}\}P(s)}, \quad (4.8)$$

$$s^* = \arg \max_{s \in \{1, \dots, S\}} \varpi(s), \quad (4.9)$$

where \hat{a}_s denotes the maximum likelihood estimate for each ordering based on equation (4.7). Given s^* that indicates a plausible choice for orderings, the next dose combination is selected such that

$$x_{(t+1)} = \arg \min_{x \in \{x_{1,1}, \dots, x_{I,J}\}} |\psi_{s^*}(x, \hat{a}_{s^*}) - \theta|. \quad (4.10)$$

Patients in the next cohort are treated with the selected dose combination corresponding to $x_{(t+1)}$.

According to equation (4.10), an MTDC is determined at the end of a trial from dose combinations that have already been administered to patients.

Open research questions

The open research questions on parametric model-based designs for MTD estimation in Section 2.3 are again applicable here. Parametric model-based designs have been reported to provide better performance than rule-based designs in various scenarios; however, they always have the potential risk of model misspecification. As it is more challenging to know true dose–toxicity relationships for combination drug therapies than mono-therapies, it is highly likely that strong assumptions for dose–toxicity shapes are difficult to apply before the initiation of trials. Additionally, we can apply the same concerns pointed in Section 2.3 about the risk of ignoring uncertainties of estimates in dose selections. As one of the options to avoid suboptimal dose allocations during trials, we would propose to take non-negligible variability on estimated dose–toxicity relationships into account in dose selections.

4.2 Modeling frameworks and practical procedures of Bayesian optimization design for MTDC estimation

Given the raised points in Section 4.1, we have developed BOD for combination drug therapies, which seeks a single MTDC. The modeling approach of BOD for MTDC estimation is almost similar to that of Chapter 3 except for the dimension of input data.

Firstly, an unknown dose–toxicity function is nonparametrically modeled. A Gaussian process is then put as a prior belief over it so as to treat the model as a Bayesian non-parametric regression model. Once patient outcomes are observed, we refine our current belief on the model via Bayes' rule. Simultaneously, a probabilistic belief on an objective function that we will optimize to find an MTDC derives from the estimated dose–toxicity function. With the probabilistic belief, we sequentially induce an acquisition function that is an alternative to the true objective function and guides where the most promising point is while leveraging uncertainties of the objective function. The next candidate of dose combinations is selected by maximizing the acquisition function. We repeat observing

patient outcomes, updating the model, and selecting a dose combination until either of the pre-specified stopping rules is met. Finally, a single MTDC is determined based on the final updated dose–toxicity function.

In this section, we introduce BOD for MTDC estimation in detail with the following subsections:

- BOD for MTDC estimation is composed of a two-stage design like the POCRM. Section 4.2.1 explains the modeling part in BOD. Because most components in the dose selection strategy are the same as Section 3.1.2, we describe only short explanations while skipping overlap contents with Section 3.1.2.
- Section 4.2.2 introduces all steps to implement BOD for MTDC estimation including the start-up phase for BOD.
- Section 4.2.3 illustrates a brief example and depicts some drawings BOD could provide.

4.2.1 Statistical modeling of dose-toxicity relationships and dose selection strategy

Dose–toxicity relationship

Suppose that each agent we focus on here monotonically increases toxicity with increasing dose levels when we evaluate each one individually. Consider a trial investigating a combination drug therapy of two agents (agent 1 and agent 2) with $I \times J$ combinations. The nonparametric model for the dose matrix is given by

$$f(\mathbf{x}) = \text{logit}\{\pi(\mathbf{x})\} = \log \left\{ \frac{\pi(\mathbf{x})}{1 - \pi(\mathbf{x})} \right\}, \quad (4.11)$$

where $\pi(\mathbf{x})$ denotes toxicity probabilities at a corresponding re-scaled combination dose value \mathbf{x} in the finite range $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^2 \mid \mathbf{x}_{1,1}, \dots, \mathbf{x}_{I,J}\}$; $\mathbf{x}_{i,j} = (x_i^{(1)}, x_j^{(2)})$ expresses a dose combination with two discrete conceptual values combining a dose level i for agent 1 and a dose level j for agent 2, which is picked up from the two-dimensional dose combination space ($i \in \{1, \dots, I\}$, $j \in \{1, \dots, J\}$). The only difference between equations (3.1) and (4.11) is whether input data is 1-dimensional dose or 2-dimensional doses.

Although there is a limitation to estimate dose–toxicity relationships on whole dose combination space due to limited sample size in practice, preliminary estimation might help to interpret the characteristics of investigated agents for subsequent clinical trials. Thus, in the same manner in Section 3.1, we estimate f in a Bayesian manner by putting a Gaussian process prior on an unknown function space; that is, $f \sim \text{GP}(m, k)$ specified by a mean function $m(\mathbf{x})$ and a covariance function $k(\mathbf{x}, \mathbf{x}')$. A prior mean function for $m(\mathbf{x})$ can derive from pre-specified initial guesses of dose–toxicity relationships if there is historical data or data of previous early studies; however, little information is available about dose–toxicity relationships particularly in more complex situations such as combination drug therapies. Bayesian optimization often uses a flat mean prior with a value of 0 as a prior setting for objective functions. For dose-finding studies, if there is no information on dose–toxicity relationships, one option is to apply a mean prior with small values that have a slight monotonically increasing trend in order to have less impact on dose selections during the trial. For a covariance function, We apply the squared

exponential kernel that substitutes two-dimensional doses to equation (3.3) as follows:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp \left\{ -\frac{1}{2\rho^2} \sum_{a=1}^2 (x^{(a)} - x^{(a')})^2 \right\}. \quad (4.12)$$

Although there are other options for kernel functions, the squared exponential kernel is also suitable for MTDC estimation for combination therapy because dose–toxicity relationships are still typically smooth. In order to improve the matrix computations and the efficiency of sampling, we apply a modification with a small value ξ that is the same manner in equation (3.4). As a result, an element in a covariance matrix K between arbitrary dose combinations \mathbf{x} and \mathbf{x}' is given by

$$K_{\mathbf{x}, \mathbf{x}'} = k(\mathbf{x}, \mathbf{x}') + \xi 1[\mathbf{x} = \mathbf{x}']. \quad (4.13)$$

We will explain how to set parameters for a Gaussian process prior assuming that there is no prior information about dose–toxicity relationships in Section 4.2.2 (Step 2 in the model phase) and provide practical settings in Section 4.3.2.

The likelihood function is the same as equation (3.5) but cumulative data up to t th test are described as $D_{1:t} = \{(n_{(1)}, \mathbf{x}_{(1)}, y_{(1)}), \dots, (n_{(t)}, \mathbf{x}_{(t)}, y_{(t)})\}$. Once patient outcomes are observed, posterior distributions of toxicity probabilities π along with f are sampled through the MCMC.

Dose selection strategy

We have addressed how to model and estimate dose–toxicity relationships at the above. Our goal is to find a single MTDC that provides toxicity closest to the target toxicity rate θ according to equation (4.2). Accordingly, the MTDC estimation problem can be transformed into an optimization problem by defining the following objective function:

$$g(\mathbf{x}) = |\pi(\mathbf{x}) - \theta|. \quad (4.14)$$

Equation (4.14) is almost the same as equation (3.6) except for the dimension of input space. A posterior distribution for g is simultaneously sampled when the corresponding posterior distributions of π and f are sampled.

In the same manner in Section 3.1, we apply the EI as an acquisition function \hat{g} to decide the next dose combination $\mathbf{x}_{(t+1)}$. Dose selection follows the steps described in equations (3.7), (3.8) and (3.9) while substituting \mathbf{x} instead of x as input data. The definition of an admissible dose set \mathcal{A}_t for the dose selection will be given in the model phase Step 4 in Section 4.2.2.

4.2.2 BOD implementation steps for MTDC estimation

We introduce the entire sequence of implementation steps of BOD for MTDC estimation in this section. Firstly, we will introduce the start-up phase in detail. After that, specific steps will be described for the model phase.

Start-up phase

The start-up phase is conducted as a part of BOD before the model phase introduced in Section 4.2.1 because the information available at the beginning of the trial may be too

limited to rely entirely on posterior estimates based on the model. The start-up phase is implemented according to the following rules:

1. Patients in the first cohort are treated at the lowest dose combination (i.e., $\mathbf{x}_{(1)} = \mathbf{x}_{1,1}$). If the first DLT is observed at the first cohort, the next dose combination is again the lowest dose combination.
2. If no patient has experienced DLT in the trial yet or if patients in the second cohort following the first cohort with DLT did not experience DLT, the dose combination for the next cohort increases in one direction by one level; that is, agent 1 remains fixed at the lowest level and only agent 2 is escalated by one level.
3. If $\mathbf{x}_{1,J}$ has been tested, agent 1 is escalated by one dose level and a dose level for agent 2 is set to the lowest level until DLT is observed.
4. If DLT is observed after the first cohort or $\mathbf{x}_{I,1}$ has been tested, then the model phase begins.

Our start-up phase could be familiar steps based on a simple algorithm-based escalation approach because dose combinations are escalated only in one direction under the assumption that toxicity increases monotonically with increasing dose levels. In addition, it could often allow more available dose combinations at the initiation of the model phase than a zone strategy that escalates dose combinations in diagonal zone order in a dose combination matrix (e.g., the POCRM utilizes a zone strategy as described in Section 4.1.4). For the sake of shortening time, it might be possible to simultaneously start from both vertical direction and horizontal direction because dose escalations by each direction do not affect each other; however, the start-up phase introduced here assumes that each direction is executed in sequence.

Model phase

After the start-up phase, the model phase described in Section 4.2.1 is implemented as follows:

1. Conceptual dose values \mathbf{x} are specified. Conceptual doses should have equal intervals between doses because the correlation of toxicity between dose combinations is determined by the distance between conceptual doses. As long as this condition is satisfied, \mathbf{x} can take any values.
2. Parameters for a Gaussian process prior $f \sim \text{GP}(m, k)$ are decided as follows:
 - (a) Like the POCRM, initial toxicity guesses associated with a prior mean function for $m(\mathbf{x})$ could be generated by the indifference interval approach with `getprior` function in the R package `dfcrm` in the same manner in Section 3.1.4. As we stated in Section 3.1.4, the `getprior` function requires the indifference interval half-width (δ), the target toxicity value (θ), the initial MTDC location (ν), and the number of combinations ($I \times J$). Unlike BOD for MTD estimation that utilizes the systematic approach (Lee and Cheung, 2009) to calibrate δ , BOD for MTDC estimation uses relatively small fixed values for δ because dose-toxicity surfaces for two agents are out of scope in the systematic approach. By specifying a small value of δ , initial toxicity guesses have a gentle increasing trend with increasing dose levels and the effect of the initial values could be minimized in the dose selections. The simplest setting for a toxicity ordering on initial toxicity guesses is ‘across rows’ which is one of the partial orderings described in Wages and Conaway (2013). For BOD, an ‘across rows’ means toxicity shows monotonicity across rows that can be described as $\pi(\mathbf{x}_{1,1}) \leq \dots \leq \pi(\mathbf{x}_{1,J}) \leq \pi(\mathbf{x}_{2,1}) \leq \dots \leq \pi(\mathbf{x}_{2,J}) \leq \dots \leq \pi(\mathbf{x}_{I,J})$. According to our experiences, it is worth evaluating δ within (0.01, 0.05) if the `getprior`

function is used for generating initial guesses while the appropriate range for δ might depend on the number of combinations. A value of ν should be on the center of the dose matrix to avoid skewed dose selections affected by the initial guesses as much as possible. With these tunable parameters (δ and ν), initial toxicity guesses are generated by setting a power model in the `getprior` function. The logit-transformed values of the generated initial guesses are placed into the prior mean function for $m(\mathbf{x})$.

- (b) For the kernel parameters, σ_f is generally set to 1 and ρ is varied to fit the model. Considering that it might be expected that a dose–toxicity function has potentially 2 or 3 turning points in the dose combination space at most, ρ is specified according to the relationship between the expected tuning points and conceptual dose values. A value of ξ is decided according to the calculation speed; however, the smaller ξ is, the less the operating characteristics are affected.
 - (c) If something useful data regarding dose–toxicity relationships such as historical data is available, the above design parameters should reflect them as needed.
3. Based on all observed data so far, posterior samples for f are generated by MCMC. Simultaneously, posterior samples for π and g are sequentially calculated.
 4. An admissible dose combination set \mathcal{A}_t is refreshed at each test and defined as follows:
 - (a) Among the dose combinations already administered to patients, let \bar{i}_j denote the highest tested dose level of agent 1 for each dose level of agent 2 ($j \in \{1, \dots, J\}$). In the same manner, let \bar{j}_i denote the highest tested dose level of agent 2 for each dose level of agent 1 ($i \in \{1, \dots, I\}$). \mathcal{A}_t includes all dose combinations (i, j) that meet $(i \leq \bar{i}_j + 1, j)$ or $(i, j \leq \bar{j}_i + 1)$. Accordingly, we prohibit the diagonal escalations from the highest dose combinations already administered to patients on each row and column in the dose combination matrix.
 - (b) In addition, we impose that all dose combinations in \mathcal{A}_t satisfy $\hat{\pi}(\mathbf{x}) < (\theta + \hat{\epsilon}_1)$ for overdose control, where $\hat{\pi}(\mathbf{x})$ is a posterior mean at a dose combination \mathbf{x} ; $\hat{\epsilon}_1$ is a small value providing an acceptable width that is preliminarily determined in terms of safety perspective and uncertainties of estimates. A value of $\hat{\epsilon}_1$ could take such as 0.05 and 0.1 for general use.
 - (c) If we need to avoid early termination due to an empty admissible dose set, leaving the lowest dose combination (i.e., $x_{1,1}$) in \mathcal{A}_t might be a possible option when $\mathcal{A}_t = \emptyset$.
 5. After the acquisition function EI is calculated according to equations (3.7) and (3.8), the next dose combination is selected by maximizing the EI based on equation (3.9).
 6. Patients in the next cohort are treated with the selected dose combination.
 7. Steps 3 through 6 are repeated until a pre-specified stopping rule is met, e.g., the maximum sample size is reached or the admissible dose set is empty.
 8. After Step 7 is completed, the dose combination falling within the range $\theta \pm \hat{\epsilon}_1$ with the highest probability among a subset of the tested dose combinations is determined as an MTDC \mathbf{x}^* . The subset for the final MTDC determination satisfies $\hat{\pi}(\mathbf{x}_{i,j}) < (\theta + \hat{\epsilon}_1)$ within the tested dose combinations ($n_{i,j} > 0$) for overdose control. On the basis of the final posterior distribution of π , an MTDC \mathbf{x}^* is given by

$$\mathbf{x}^* = \underset{\mathbf{x} \in \{\mathbf{x}_{i,j} | \hat{\pi}(\mathbf{x}_{i,j}) < (\theta + \hat{\epsilon}_1), n_{i,j} > 0, i \in \{1, \dots, I\}, j \in \{1, \dots, J\}\}}{\arg \max} \mathbb{P}(\theta - \hat{\epsilon}_1 \leq \pi(\mathbf{x}) \leq \theta + \hat{\epsilon}_1 | D_{1:t}). \quad (4.15)$$

4.2.3 Illustration and example

In this section, we show some illustrations of BOD for MTDC estimation about updates of dose–toxicity relationships and dose selections. Table 4.1 shows outputs observed up to the 8th test in a trial assuming scenario 7 in Table 4.2 that was prepared for the simulation study.

Table 4.1. A trial example

Phase	Test number	cumulative number of patients	(i, j)	DLT response	True toxicity
Start-up	1	1	(1,1)	0	0.07
	2	2	(1,2)	0	0.10
	3	3	(1,3)	1	0.12
Model	4	4	(1,3)	0	0.12
	5	5	(2,2)	0	0.30
	6	6	(2,2)	1	0.30
	7	7	(1,3)	0	0.12
	8	8	(2,2)	0	0.30

According to procedures of the start-up phase in Section 4.2.2, 3 tests were conducted until the first DLT was observed in this example. In Fig. 4.2, the upper figures on upper and lower sections display the posterior means and 10 and 90 percentiles of the posterior distributions on toxicity probabilities. Additionally, the corresponding EI after each test is shown in lower figures. In practice, we obtain posterior distributions among all dose combinations through posterior samples rather than just posterior means and 10 and 90 percentiles. The upper left section in Fig. 4.2 displays posterior estimates after the 3rd test and the corresponding EI at all dose combinations $(i, j)(i \in \{1, 2, 3\}, j \in \{1, 2, 3, 4, 5\})$. According to equation (3.9), the dose combination of (1,3) pointed by an arrow symbol in the figure of the EI was selected for the 4th test. In the same manner, next dose combinations are selected after each test as shown in Fig. 4.2. We note that admissible dose sets temporally exclude some dose combinations from the next dose selection from the view of patient safety as described in the model phase Step 4 in Section 4.2.2.

Fig. 4.3 depicts posterior density probability functions that were used to calculate the EI after the 8th test. As shown in equation (3.7), $I(\mathbf{x})$ takes either 0 or $g^+ - g(\mathbf{x})$, whichever is greater. Filled area in Fig. 4.3 shows one where values of $g(\mathbf{x})$ are lower than g^+ defined with data up to the 8th test; thus, the corresponding $I(\mathbf{x})$ returns positive values. On the other hand, the other areas provide greater values of $g(\mathbf{x})$ than g^+ ; therefore, $I(\mathbf{x})$ becomes 0. The EI is calculated based on equation (3.8) using the probability density functions and $I(\mathbf{x})$ derived from $g(\mathbf{x})$. As a result, for example, the EI after the 8th test draws the line chart displayed in the lower right section in Fig. 4.2.

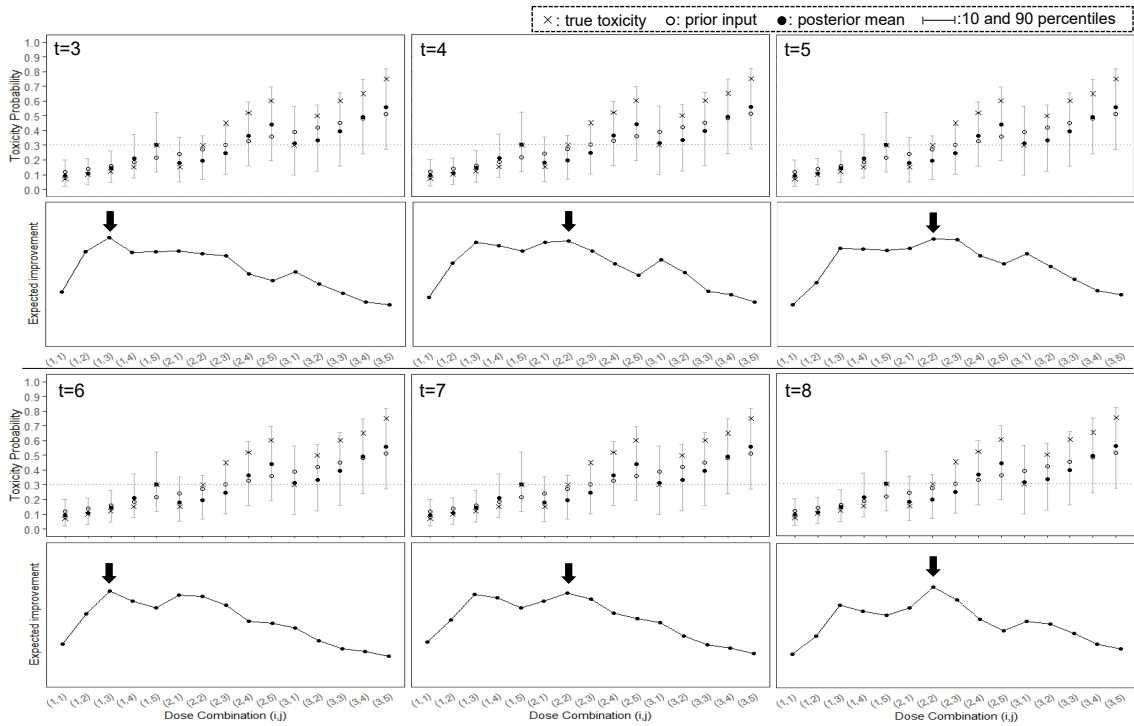


Fig. 4.2. An example of BOD for MTDC estimation shows updates of dose–toxicity relationships and the corresponding EI during a trial.

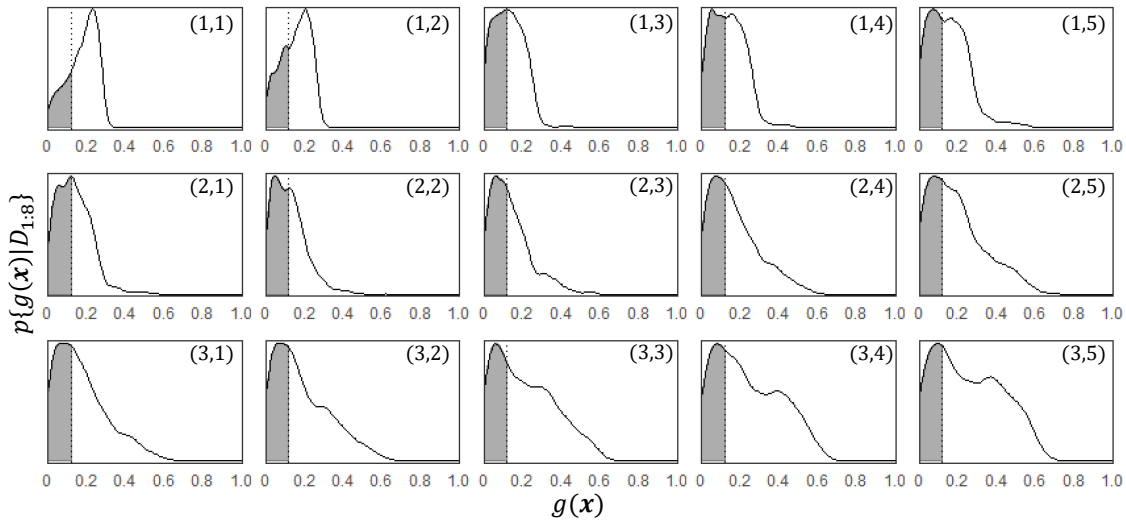


Fig. 4.3. Probability density functions after the 8th test at each dose combination (i, j)

These steps are repeated until the end of the trial. After the trial meets one of the pre-specified stopping rules, a single MTDC is determined based on equation (4.15).

4.3 Performance evaluation

We conducted a simulation study to examine the performance of BOD for MTDC estimation by comparing it with the BOIN based on Lin and Yin (2017a) and the POCRM based on Wages and Conaway (2013). In the following sections, we present our simulation

framework and results.

4.3.1 General settings for the simulation study

We considered dose-finding studies for two agents having 3×5 dose combinations and the target toxicity rate θ was 0.3. In this simulation study, we used the same toxicity scenarios as the simulation settings in Riviere et al. (2015a). Those ten scenarios, which seem to cover a wide variety of underlying realities, are shown in Table 4.2. We conducted our simulation study under the trial with cohorts including a single patient (i.e., the cohort of size 1) and a small sample size of 30. Regarding the sample size of 30, we referred to Hirakawa et al. (2015) whose simulation study was based on four actual examples of dose-finding studies for combination drug therapies that had been reviewed by the FDA, approved by the Institutional Review Board, or involved initial planning discussions with clinicians for a study; therefore, this small sample size might be the trial designs that we might encounter in practice. In most cases, feasibility has a great impact on designing Phase I clinical trials.

As a stopping rule, all designs except for the BOIN with elimination rule terminated the trials when the number of patients reaches the maximum sample size so that we could compare the different designs under as consistent conditions as possible. ‘The BOIN with elimination rule’ that aligned with the original proposal of the BOIN possibly caused early termination before reaching the maximum sample size as an exception; therefore, we also evaluated the BOIN that did not cause early termination. The specifics for the BOIN in the simulation study will be explained in Section 4.3.3.

We supposed that the lowest dose combination was the initial test dose combination in all the designs. All simulations were computed for 2,000 trials. As performance evaluations, we provided descriptive summaries, MSEs calculated by equation (3.11) that substituted $M = 2000$, and 95% confidence intervals for the MSEs.

Table 4.2. Toxicity scenarios for 3×5 dose combinations

		Dose level of agent 2											
		1	2	3	4	5	1	2	3	4	5		
		Scenario 1					Scenario 2						
		3	0.15	0.30	0.45	0.50	0.60	3	0.45	0.55	0.60	0.70	0.80
		2	0.10	0.15	0.30	0.45	0.55	2	0.30	0.45	0.50	0.60	0.75
Dose level of agent 1		1	0.05	0.10	0.15	0.30	0.45	1	0.15	0.30	0.45	0.50	0.60
		Scenario 3					Scenario 4						
		3	0.10	0.15	0.30	0.45	0.55	3	0.50	0.60	0.70	0.80	0.90
		2	0.07	0.10	0.15	0.30	0.45	2	0.45	0.55	0.65	0.75	0.85
		1	0.02	0.07	0.10	0.15	0.30	1	0.30	0.45	0.60	0.70	0.80
		Scenario 5					Scenario 6						
		3	0.07	0.09	0.12	0.15	0.30	3	0.15	0.30	0.45	0.50	0.60
		2	0.03	0.05	0.10	0.13	0.15	2	0.09	0.12	0.15	0.30	0.45
		1	0.01	0.02	0.08	0.10	0.11	1	0.05	0.08	0.10	0.13	0.15
		Scenario 7					Scenario 8						
		3	0.30	0.50	0.60	0.65	0.75	3	0.08	0.15	0.45	0.60	0.80
		2	0.15	0.30	0.45	0.52	0.60	2	0.05	0.12	0.30	0.55	0.70
Dose level of agent 1		1	0.07	0.10	0.12	0.15	0.30	1	0.02	0.10	0.15	0.50	0.60
		Scenario 9					Scenario 10						
		3	0.15	0.30	0.45	0.55	0.65	3	0.70	0.75	0.80	0.85	0.90
		2	0.02	0.05	0.08	0.12	0.15	2	0.45	0.50	0.60	0.65	0.70
		1	0.005	0.01	0.02	0.04	0.07	1	0.05	0.10	0.15	0.30	0.45

4.3.2 Simulation settings of BOD for MTDC estimation

We followed all the steps described in Section 4.2.2; however, $\mathcal{A}_t = \{\mathbf{x}_{1,1}\}$ when $\mathcal{A}_t = \emptyset$ in the model phase Step 4 to avoid early termination in the simulation study. The conceptual doses for agent 1 and agent 2 were $\{0.2, 0.4, 0.6\}$ and $\{0.2, 0.4, 0.6, 0.8, 1.0\}$, respectively; therefore, the available dose combinations of the 3×5 matrix consisted of all the combinations of these conceptual doses in increasing order of dose level for each agent.

We generated initial toxicity guesses with the `getprior` function using $\delta = 0.015$ and $\nu = 8$. Based on preliminary evaluations, $\delta = 0.015$ seemed to work well for 15 dose combinations from the view of safety dose allocations and correct MTDC selection probabilities. The setting of $\nu = 8$ indicates that an initial MTDC was the center of the dose matrix (i.e., $\mathbf{x}_{2,3}$). The parameters related to the squared exponential kernel were $\sigma_f^2 = 1$, $\rho = 0.4$, and $\xi = 0.08$. The value of ρ indicates that the turning point was every 0.4 lengths of conceptual doses. It was in a multiple of 0.2 used for the conceptual doses and was a slightly smaller value than a half-length of conceptual doses for agent 2. We determined the value of ξ only based on computational speed. Step 4 in the model phase and equation (4.15) used $\hat{\epsilon}_1 = 0.10$.

As evaluation for the impact of δ and ρ , we also simulated with different settings for them under specific scenarios. Although posterior distributions for f as well as π provide a unique partial-ordering trend, individual functions of them do not always follow the feature; therefore, we also evaluated BOD treating only toxicity probability functions that did not reverse toxicity at least over 0.15 when the dose level of one agent increased while the dose level of the other was fixed.

Appendix B.1 provides the codes for BOD that were used in the simulation study.

4.3.3 Simulation settings for the competitors

Bayesian optimal interval design

We followed the procedures in Section 4.1.3 regarding the mathematical background of the BOIN. Suppose that the current dose combination is composed of dose levels (i, j) . Dose selections are conducted based on the sample statistic given by $y_{i,j}/n_{i,j}$ during a trial. For the optimal boundaries for the dose selections, we used $(0.236, 0.358)$ based on the recommended values for $\theta = 0.3$ in Liu and Yuan (2015). Posterior toxicity probabilities at each dose combination were calculated by a beta-binomial model assuming a vague beta prior $\text{Beta}(1, 1)$.

We evaluated two types of the BOIN. One was the BOIN with dose elimination rule that imposed the following rule to avoid assigning patients to overly toxic dose combinations: If a posterior toxicity probability greater than θ at the current dose combination was greater than 0.95 when $n_{i,j}$ was at least three patients, the current dose combination and the higher dose combinations $\{(i', j') \mid (i' \geq i) \wedge (j' \geq j)\}$ were eliminated from the trial. If the lowest dose combination was eliminated, the trial was terminated because of safety concerns; that is, it possibly decreased the number of patients in a trial due to early termination depending on scenarios. The cutoff value of 0.95 is the recommended value for general use in the BOIN (Yan et al., 2020). The other BOIN evaluated in this simulation was the BOIN without the elimination rule that did not apply the elimination rule; therefore, the latter BOIN never terminated trials early.

At the end of the trial, a single MTDC with toxicity closest to θ was selected based on isotonic estimates described in a matrix form of the toxicity probabilities. If there was a tie, the higher dose level was selected when the isotonic estimate was lower than θ .

We implemented the BOIN with the R package `BOIN`. As a side note, the BOIN can be implemented via a user-friendly cloud software and desktop software provided from <https://biostatistics.mdanderson.org/SoftwareDownload/> as alternatives of the R package.

Partial ordering continual reassessment method

As shown in Section 4.1.4, the two-stage POCRM was evaluated in this simulation study. As a start-up phase of the two stage POCRM, the initial escalation scheme proceeded according to zones, testing (i, j) , in order, $(1, 1)$, $(1, 2)$, $(2, 1)$, $(1, 3)$, $(2, 2)$, $(3, 1)$, $(1, 4)$, $(2, 3)$, $(3, 2)$, $(1, 5)$, $(2, 4)$, $(3, 3)$, $(2, 5)$, $(3, 4)$, and $(3, 5)$ until the first DLT was observed. After the start-up phase finished, the design switched to the model-based procedure based on the power model. As recommended in Wages and Conaway (2013), six possible orderings—named across rows, up columns, up diagonals, down diagonals, alternating down-up diagonals, and alternating up-down diagonals—were used in the simulation (i.e., $s = 1, \dots, 6$).

Riviere et al. (2015a) put an initial MTDC on $(3, 3)$ for the POCRM, which was near the highest dose combination in the explored dose combination space. On the other hand, Wages and Conaway (2013) recommended that an initial MTDC location should be placed in the center of the working model to ensure enough space within the range both below and above the dose. This suggestion was also discussed in Riviere et al. (2015a). In this simulation, therefore, we adopted the center of the dose combination ranges, $(2, 3)$, as an initial MTDC location, rather than $(3, 3)$. Skeleton values $\pi_{i,j}^0(s)$ to configure the orderings were generated by using the `getprior` function in R package `dcrm` with $\delta = 0.03$ and $\nu = 8$, where the value of δ was the same setting as in Riviere et al. (2015a). The six possible orderings were provided through the rearrangement of the generated values using

the `getwm` function in R package `pocrm` that can help rearrange the skeleton values. We also examined the POCRM using a different initial guess with $\delta = 0.05$ in the `getprior` function as a sensitivity analysis. We implemented the POCRM using the `pocrm` package in R.

4.3.4 Simulation results

In this section, BOD for MTDC estimation is displayed as ‘BOD (all)’. As mentioned in Section 4.3.2, we evaluated BOD with a partial monotonicity restriction on toxicity in addition to BOD (all). BOD with the monotonicity restriction is described as ‘BOD (mono)’. The BOIN and the POCRM have also two patterns in the simulation study as explained in Section 4.3.3. They are displayed as ‘BOIN with elimination’, ‘BOIN w/o elimination’ (i.e., the BOIN without the elimination rule), and ‘POCRM (a specified value for δ)’.

- MTDC selection

Table 4.3 shows selection probabilities under the ten scenarios. Good performance is suggested by the correct selection probabilities of MTDC shown by BOD (all). Simultaneously, BOD (all) shows lower selection probabilities of dose combinations with higher toxicity than θ (Overdoses) compared with the other designs in most scenarios.

In only scenario 5, BOD (all) provides a lower correct selection probability than that of the BOINs (with and w/o elimination) and POCRM (0.03). This could be due to the unique properties that are effective especially in this scenario. For the POCRM, the correct selection probability under scenario 5 drops dramatically as increasing δ , whereas those in the other scenarios are not so affected by δ . If only the highest dose combination is an MTDC, the POCRM with a small δ that makes an initial toxicity guess at the highest dose combination be close to θ could allow reaching the true MTDC more rapidly than BOD (all). The BOIN allows only ± 1 dose level changes from the current dose combination in its dose selection. In contrast, BOD selects a candidate from an admissible dose set that typically provides a wider selection range than the candidates of the BOIN. BOD tries to seek plausible dose combinations from the wider space without concentrating on a local optimum by the acquisition function. It allows BOD to realize more flexible dose selections than the BOIN and worked well in most scenarios; however, the BOIN’s simple dose selection procedure could provide a more efficient approach to reach the highest dose combination rapidly than BOD in the case of scenario 5.

The relatively poor results shown by the POCRMs indicate that pre-specified assumptions about the dose–toxicity relationships might hinder the efficient search for the true MTDC in two-dimensional complex space because, in most scenarios, nonparametric approaches provide better performance than the POCRMs in the simulation study. On the other hand, the POCRMs work better than the BOINs in scenarios 9 and 10 in terms of the correct selection probability of MTDC. For one thing, one of their pre-specified orderings could fit the true dose-toxicity relationships because the “across rows” ordering seems to be suitable for these scenarios. With regard to scenario 9, the zone strategy the POCRM applied could have worked well. Scenario 9 has very low toxicities at dose combinations below the zone containing an MTDC locating on $(i, j) = (3, 2)$. Therefore, it is highly likely that the first DLT is observed at an MTDC in the start-up phase. A synergistic effect among these multiple factors—scenarios, zone strategies, prior settings, and orderings—would enable the POCRM to reach an MTDC more rapidly than in other situations. On the other hand, even if an initial MTDC is located on the true MTDC, the POCRMs do not always provide excellent results, as exemplified by scenario 8. Ad-

Table 4.3. Operating characteristics of selection probabilities (true MTDC or overdoses) by each design and scenario

Selected dose	Scenario	BOD (all)	BOD (mono)	BOIN with elimination	BOIN w/o elimination	POCRM (0.03)	POCRM (0.05)
MTDC	1	0.62	0.63	0.50	0.51	0.23	0.23
	2	0.59	0.57	0.52	0.55	0.39	0.43
	3	0.57	0.60	0.53	0.52	0.34	0.39
	4	0.66	0.51	0.48	0.65	0.73	0.70
	5	0.49	0.48	0.61	0.64	0.62	0.45
	6	0.52	0.53	0.42	0.42	0.30	0.31
	7	0.60	0.67	0.56	0.55	0.44	0.52
	8	0.30	0.24	0.26	0.29	0.03	0.03
	9	0.44	0.49	0.24	0.24	0.32	0.35
	10	0.40	0.36	0.29	0.26	0.28	0.35
	Average		0.52	0.51	0.44	0.46	0.37
Overdose	1	0.19	0.18	0.28	0.29	0.40	0.39
	2	0.32	0.38	0.36	0.36	0.46	0.44
	3	0.39	0.34	0.42	0.43	0.53	0.52
	4	0.34	0.50	0.34	0.35	0.27	0.30
	5	0.00	0.00	0.00	0.00	0.00	0.00
	6	0.18	0.19	0.25	0.27	0.29	0.27
	7	0.29	0.25	0.29	0.31	0.42	0.35
	8	0.29	0.34	0.30	0.34	0.46	0.51
	9	0.23	0.23	0.19	0.24	0.27	0.25
	10	0.34	0.46	0.44	0.49	0.45	0.44
	Average		0.26	0.29	0.29	0.31	0.36

ditionally, a correct MTDC can be selected if the estimated toxicity probability at a true MTDC is close to θ on the estimated plausible ordering even if the ordering is wrong. As the POCRM does not provide excellent performances under many scenarios in this simulation, however, it would be difficult to reach a single correct MTDC at the end of the trial when toxicity orderings are misspecified. Regarding prior settings for the POCRM, Riviere et al. (2015a) have reported that definitions for reasonable and unreasonable prior settings for the POCRM have remained unclear so far.

The BOIN with elimination shows early termination in some trials under scenarios 2 and 4 with the average numbers of patients of 29.5 and 26.2, respectively. Because of the elimination rule, the correct selection probability drops by 17% under scenario 4 (Table 4.3). It indicates the dose elimination rule has a high impact on the performance of the BOIN when the dose combination matrix has an MTDC at only the lowest dose combination. The BOIN is free from rigid assumptions, unlike the POCRM, while there is a limitation on the search space compared with BOD. It could be why the performance of the BOINs was intermediate between BOD (all) and the POCRM.

MSEs with precision information support Table 4.3. In Fig. 4.4, we interpret the performance as more favorable when MSEs are smaller. We can see consistent trends to the above descriptions as BOD (all) has the smallest MSEs in most scenarios.




Fig. 4.4. Mean squared errors on toxicity probabilities at the recommended MTDC

- Dose allocations and observed toxicity

BOD (all) shows the lowest observed toxicity percentage among all the designs (Table 4.4) and allocates more patients to the true MTDC and fewer patients to overdoses compared with the other designs in most scenarios (Table 4.5). Low MTDC allocation under scenario 5 shown in BOD (all) may be explained by the same reason mentioned in the second paragraph in the descriptions for MTDC selection results. In scenario 4, BOD (all) selects more overdoses during the trial than the other designs (Table 4.5), although the observed toxicity percentage of BOD (all) is similar to that of the other designs (Table 4.4). This can be explained by Table 4.6 that provides the average actual number of patients treated with each dose combination under scenario 4. Although BOD (all) treats fewer patients on the dose combination of (1,1) than the other designs, BOD (all) selected doses near the dose combination of (1,1). The higher selections on the horizontal area of dose level 1 of agent 1 seem to be the effect of the start-up phase. The other designs reached farther areas of dose combinations in the matrix; for example, the other designs select the combinations of $i = 3$, while BOD (all) never selects them. For scenario 8, BOD (all) allocates as many patients to the true MTDC as the BOINs, while the POCRM ($\delta = 0.03$ and 0.05) show the lowest allocation to the true MTDC.

Table 4.4. Average percentage of the total number of patients experienced with DLT for each design under 10 scenarios

Scenario	BOD (all)	BOD (mono)	BOIN with elimination	BOIN w/o elimination	POCRM (0.03)	POCRM (0.05)
1	26.3	27.2	31.7	32.0	27.5	27.1
2	35.5	37.5	37.0	36.3	37.8	37.7
3	21.5	22.2	29.0	29.3	22.2	20.9
4	40.9	42.8	44.7	39.3	40.3	41.1
5	13.7	14.1	20.3	20.7	15.6	13.7
6	22.3	23.0	30.0	30.7	23.7	22.5
7	27.4	28.2	33.0	34.0	30.4	29.6
8	28.4	28.7	30.7	32.3	28.8	27.8
9	20.9	22.5	27.7	30.7	22.1	20.3
10	32.4	34.7	35.3	36.0	36.2	35.7
Average	26.9	28.1	31.9	32.1	28.4	27.6

Table 4.5. Dose allocation probabilities to MTDCs and overdoses under 10 scenarios

Allocated dose	Scenario	BOD (all)	BOD (mono)	BOIN with elimination	BOIN w/o elimination	POCRM (0.03)	POCRM (0.05)
MTDC	1	0.44	0.45	0.33	0.32	0.22	0.23
	2	0.45	0.43	0.32	0.32	0.28	0.30
	3	0.35	0.38	0.34	0.32	0.25	0.27
	4	0.48	0.33	0.56	0.55	0.61	0.56
	5	0.19	0.19	0.47	0.50	0.30	0.20
	6	0.30	0.31	0.26	0.25	0.20	0.21
	7	0.43	0.50	0.32	0.29	0.33	0.37
	8	0.16	0.12	0.15	0.16	0.05	0.04
	9	0.25	0.29	0.15	0.14	0.17	0.18
	10	0.28	0.26	0.13	0.12	0.15	0.20
	Average	0.33	0.33	0.30	0.30	0.26	0.26
Overdose	1	0.19	0.19	0.37	0.39	0.31	0.30
	2	0.43	0.49	0.46	0.47	0.53	0.51
	3	0.45	0.42	0.56	0.56	0.54	0.51
	4	0.52	0.67	0.44	0.45	0.39	0.44
	5	0.00	0.00	0.00	0.00	0.00	0.00
	6	0.13	0.14	0.37	0.39	0.23	0.20
	7	0.31	0.27	0.45	0.49	0.40	0.37
	8	0.34	0.37	0.41	0.43	0.40	0.40
	9	0.16	0.18	0.33	0.41	0.23	0.18
	10	0.40	0.49	0.50	0.52	0.51	0.51
	Average	0.29	0.32	0.39	0.41	0.35	0.34

Table 4.6. Average number of patients treated with each dose combination in scenario 4

		Dose level of agent 2											
		1	2	3	4	5	1	2	3	4	5		
		BOD (all)					BOD (mono)						
		3	0.0	0.0	0.0	0.0	0.0	3	0.6	0.1	0.0	0.0	0.0
Dose level of agent 1		2	1.3	0.7	0.1	0.0	0.0	2	5.3	1.6	0.1	0.0	0.0
		1	14.3	9.8	2.9	0.7	0.2	1	9.9	8.8	2.7	0.8	0.1
		BOIN with elimination					BOIN w/o elimination						
		3	1.1	0.7	0.3	0.1	0.0	3	1.1	0.9	0.4	0.1	0.0
		2	3.7	1.6	0.5	0.1	0.0	2	3.8	1.8	0.5	0.1	0.0
		1	13.2	3.9	0.8	0.2	0.0	1	16.6	3.9	0.8	0.1	0.0
		POCRM (0.03)					POCRM (0.05)						
		3	0.3	0.6	0.3	0.0	0.0	3	0.3	0.3	0.0	0.0	0.0
		2	1.8	0.9	0.3	0.3	0.3	2	2.4	1.2	0.3	0.3	0.3
		1	18.3	3.3	2.1	1.2	0.3	1	16.8	3.9	2.4	1.5	0.3

- BOD with different settings

As shown in Table 4.3, BOD (mono) provides comparable correct selection probability to BOD (all); however, the selection probability of overdoses is higher than the other designs in several scenarios. In addition, observed toxicity percentages in BOD (mono) slightly increase compared with BOD (all) in most scenarios (Table 4.4). In particular, the difference is exhibited in scenario 4. BOD (mono) offers posterior distributions of dose–toxicity relationships with almost a partial ordering feature while the estimated toxicity probabilities tend to be lower at the low dose combinations and higher at the high dose combinations than those of BOD (all). As a result, BOD (mono) shows higher overdose selections and higher observed toxicity percentages than BOD (all). For BOD (all), posterior distributions of π follow a partial ordering feature even without the monotonicity restriction because of the following reasons; a prior mean function belongs to one of the partial orderings; an admissible dose combination set takes the feature into considerations by increasing dose selection range in order from the tested dose combinations; posterior distributions necessarily reflect the partial ordering that the true dose–toxicity relationship has. Given these facts and simulation results, the monotonicity restriction would not be necessary for the BOD.

As an evaluation of tunable parameters (δ , ρ), we simulated with scenarios 1, 4, 5, and 6. When ρ is a very small value (i.e., 0.1), the performance decreases significantly (Table 4.7). It is unlikely that dose–toxicity relationships have many turning points as mentioned in Section 4.2.2. This result demonstrates that too small ρ compared with the length of the conceptual dose is inappropriate for BOD. On the other hand, if $\rho = 1.0$ that is the same as the length of the conceptual dose for agent 2, the results are comparable to BOD (all) with the original settings under these scenarios. It implies that it is sufficient to evaluate ρ from the value that provides 2 or 3 turning points up to a longer length of conceptual dose in two agents. The larger value of δ (e.g., $\delta = 0.030$) affects the performance in scenarios 4 and 5 that have an MTDC on the lowest or highest dose combinations, while the results in the other scenarios retain. The larger δ is, the steeper the slope of initial guesses is. It means that an initial toxicity probability at the lowest or highest dose combinations becomes far from θ when δ increases. As a result, δ of 0.030 lower the performance in scenarios 4 and 5. The results under $\delta = 0.010$ show comparable performance to BOD (all) with the original settings in terms of correct selection probabilities of MTDC and

observed toxicity percentages. On the other hand, the selection probability of overdoses is 6% increased in scenario 6 compared to BOD (all) with the original settings. The flatter prior mean function might increase overdose selections depending on scenarios; therefore, δ needs to be carefully decided based on simulations or an alternative approach. We also discuss the design parameters in Section 4.4.

Table 4.7. Selection and allocation probabilities and percentage of the total number of toxicities when different parameters (δ, ρ) were used in BOD (all)

Scenario	(δ, ρ)	Selection probabilities		Toxicity (%)
		MTDC	Overdose	
1	(0.015, 0.1)	0.40	0.29	25.4
	(0.015, 1.0)	0.57	0.17	26.3
	(0.010, 0.4)	0.58	0.20	26.3
	(0.030, 0.4)	0.60	0.16	24.9
4	(0.015, 0.1)	0.13	0.87	51.5
	(0.015, 1.0)	0.72	0.28	35.7
	(0.010, 0.4)	0.72	0.28	37.9
	(0.030, 0.4)	0.21	0.79	48.2
5	(0.015, 0.1)	0.00	0.00	8.9
	(0.015, 1.0)	0.66	0.00	16.2
	(0.010, 0.4)	0.54	0.00	15.3
	(0.030, 0.4)	0.18	0.00	10.2
6	(0.015, 0.1)	0.32	0.11	16.7
	(0.015, 1.0)	0.50	0.22	24.3
	(0.010, 0.4)	0.46	0.24	23.1
	(0.030, 0.4)	0.53	0.13	19.9

4.4 Summary and discussion

BOD for identifying a single MTDC in a two-dimensional dose combination space has been presented in this chapter. Although the basic strategy of BOD in Chapters 3 and 4 has a lot in common on settings related to Bayesian optimization frameworks, detail conditions pertaining to cancer therapies are unique in each chapter (e.g., the start-up phase design, initial toxicity guesses, dose selection conditions, and overdose control). For MTDC estimation, BOD (all) performs better than the competitors. It is more challenging to know about dose-toxicity relationships on dose combination matrices than mono-therapies. The simulation results indicate nonparametric designs could have more advantage than parametric model-based designs for MTDC estimation. The good performance shown by BOD (all) indicates that modeling flexibility of BOD works effectively in the MTDC estimation problem. The simulation results also show that BOD (all) offering more flexible dose selection works better than the BOIN that has the rule to limit its dose selection range depending on the dose combination currently being tested. Considering the favorable results shown by BOD (all), BOD features—flexible modeling on dose-toxicity relationships and flexible dose selections without ignoring uncertainties of posterior dose-toxicity distributions—could bring some benefits to MTDC estimation.

Our simulation study aimed to evaluate the operating characteristics among BOD and the competitors under the same conditions as much as possible. Of two common cohort sizes, one and three, our simulation study employed cohorts including a single patient to align with Wages et al. (2011) and Wages and Conaway (2013). Designs with a cohort size

of three can collect more toxicity data for a given dose combination on a shorter study duration while Zhou (2004) has reported that designs with a cohort size of one achieve more accurate estimates and fewer toxicities. Hirakawa et al. (2016) has reported that a cohort size of one may be favorable to have the best chance to identify a true MTDC and lowering selection rates for unacceptable toxic dose combinations. Tighiouart and Rogatko (2010) has reported the difference in cohort sizes did not affect the safety of a trial and therapeutic doses based on a Bayesian adaptive design due to small clinical trials; however, they have recommended using a cohort size of one to avoid seeing simultaneous serious toxic events when a group of patients is treated at the same dose level. For actual clinical trials, it is crucial to evaluate the effect of cohort sizes on estimation accuracy and safety under assumed trial conditions, and carefully determine a cohort size while considering the balance with feasibility (e.g., study duration).

For dose-finding studies for combination drug therapies, it is challenging to fully know which dose combinations are non-toxic due to drug-drug interactions even after MTDC determination; however, we can assume toxicity increases monotonically with increasing dose levels when we escalate the dose levels of one agent while fixing the other agent's dose level. Once we determine an MTDC, lower dose combinations on the same diagonal line than the MTDC are regarded as non-toxic; dose combinations in which dose levels of one agent are the same as the MTDC while setting the other agent sets lower levels than the MTDC are also regarded as non-toxic; otherwise, there are still limitations to conclude the other specific dose combinations to be non-toxic. On the other hand, BOD gives some insight into general dose-toxicity relationships for explored dose combinations even though toxicity estimation across all dose combinations is not the specific aim in dose-finding studies. This would help interpret which dose combinations can be used safely in future studies.

While we focused on finding a single MTDC, BOD estimates toxicity probabilities of all dose combinations at each test; therefore, it is feasible in theory to extend BOD to the selection of multiple MTDCs. We leave this extension as our future work. The computational aspect may be needed improvement in future work. The POCRM is based on a likelihood estimation. Additionally, the BOIN calculates posterior distributions of the current dose combinations analytically through the beta-binomial model. On the other hand, BOD needs MCMC to estimate posterior distributions. While the time difference is not relevant for the use in conducting one trial, BOD takes more time than the other designs due to MCMC when simulation evaluation is needed. GPGPU-based parallel computing might allow us to make the calculation much faster. Additionally, calibration approaches for design parameters are still an open discussion. As almost similar to BOD for MTD estimation in Chapter 3, BOD for MTDC estimation has the following parameters: δ and ν for mean prior functions, ρ , σ_f and ξ for kernel parameters, and $\hat{\epsilon}_1$ for the model phase Step 4. Of these, ν , σ_f and ξ are automatically decided; the value of ν should be the center of the dose combination matrix; σ_f is usually 1 as described in the model phase Step 2 in Section 4.2.2; and the value of ξ can be determined only based on computational speed. Therefore, the number of parameters to be considered is three (i.e., δ , ρ and $\hat{\epsilon}_1$). The value of δ provides the slope of dose-toxicity relationships when one of the partial orderings of across rows is applied. On the basis of our experiences, we provided the recommendation range for evaluating δ in the model phase Step 2 in Section 4.2.2. The parameter of ρ expresses the number of turning points in the dose combination matrix and there could be up to 2 or 3 turning points in the dose combination space. It might be good to try values from the minimum of one-fourth to the maximum length of the conceptual dose. For the value of $\hat{\epsilon}_1$, 0.05 or 0.10 could be generally accepted in terms of clinical perspective. Additionally, it would be possible to obtain typically valid hyperparameters using various available data based on empirical Bayes methods instead

of searching the values for the design parameters. It is very challenging to assess the risk when the design parameters are misspecified; however, we think at least we can mitigate and minimize the risk by evaluating the impact of design parameters according to the above approaches.

One of our motivations to apply Bayesian optimization frameworks was brought from that most conventional statistical designs utilize the best guess in their dose selections although relying on the best guess might lead to suboptimal dose allocations during the trial. In this chapter, we employed the EI as an optimization strategy in BOD for MTDC estimation. Acquisition functions like the EI are one of the key components of the Bayesian optimization frameworks and return trade-offs between exploitation and exploration so that we can avoid concentrate on local optimum points and search for global optimum as few evaluations as possible. Chapter 9.2 in Fedorov and Leonov (2013) has discussed largely the problems of the best intention designs that rely on only the best point of a stochastic model for the next dose selection in dose-finding studies. It has indicated that the best intensive designs may converge to wrong doses and can result in a false prediction that differs from what we need to identify. Given the concerns about the best intention designs, BOD is one of the reasonable approaches that select dose combinations based on the balance of exploitation and exploration that does not only rely on the current best point.

Unlike statistical designs for MTD estimation, no standard solution has yet been developed for the MTDC estimation problem. Considering the current situation and the comprehensively good performance exhibited by BOD (all) under most scenarios in terms of correct MTDC selection probabilities and safer dose allocations, BOD for MTDC estimation that is displayed as BOD (all) in Section 4.3 has a high potential to be a powerful tool for use in dose-finding studies of combination drug therapies.

B

Appendices for Chapter 4

B.1 BOD simulation code for MTDC estimation

We implemented BOD for MTDC estimation using R for the simulation study described in Section 4.3. Appendix B.1 provides R codes for implementing BOD (all) to identify a single MTDC for a combination drug therapy.

- Listing B.1 is R codes that construct the simulation body.
- Listing B.2 presents a `stan` model that is compiled from the R package `rstan` in Listing B.1 to implement the MCMC for obtaining posterior samples.

Listing B.1. Simulation body of BOD for MTDC estimation

```

1 #BOD_combination_simulation_body.R
2 #####
3 # Bayesian optimization Design for MTDC estimation #
4 #####
5
6 ### Initial setting ###
7 setwd("XXX") # set your working directly
8 library(dfcrm) # for the getprior function
9 library(rstan) # for MCMC estimation
10
11 ### Compile stan model for MCMC ###
12 rstan_options(auto_write = TRUE) # rstan option that allows you to
    automatically save a bare version of a compiled Stan program to the
    hard disk so that it does not need to be recompiled.
13 stanmodel <- stan_model(file = 'BOD_model_combination.stan') # a stan file
    in the working directly
14
15 ### Define functions ###
16 # logit function
17 logit_f <- function(a){
18   b<- log(a/(1-a))
19   return(b)
20 }
21
22 # For making a listing of outputs to avoid errors regarding variable names
23 rbindCOrder <- function(...) {
24   n <- length(list(...))
25   temp <- list(...)[[1]]
26   names(temp) <- NA
27   for (i in 2:n) {
28     tmp <- list(...)[[i]]
29     names(tmp) <- NA
30     temp <- rbind(temp, tmp)

```

```

31 }
32 names(temp) <- names(list(...)[[1]])
33 return(temp)
34 }
35
36 ### Trial and simulation settings ###
37 sim_min <- 1 # the minimum number of the simulation number
38 iter <- XXX # set number of iterations
39 cohort <- 1 # cohort size
40 max_sub <- 30 # sample size
41 Pt <- 0.3 # target toxicity rate
42 A <- c(0.2,0.4,0.6) # conceptual doses for agent 1
43 B <- c(0.2,0.4,0.6,0.8,1) # conceptual doses for agent 2
44 N <- length(A)*length(B) # number of the dose combinations
45 X <- expand.grid(B,A) # all conceptual doses of agents 1 and 2
46 X <- X[,c(2,1)] # swap the first and second columns
47 X$grid <- row(X)[,1] # add the grid column to X
48 names(X)[1] <- paste("A") # change the column name
49 names(X)[2] <- paste("B") # change the column name
50 X2 <- X[,-3] # X without the grid column
51 slope <- 0.015 # a design parameter delta used in the getprior function
52 eps1 <- 0.1 # a design parameter for overdose control in the model phase
    Step 4 and equation (4.15) for MTDC determination
53
54 #For expected improvement
55 til <- 0.001 # increment width of the following dummy values
56 g_t1 <- seq(from=0, to=1,by=(til)) # dummy values of g(x)
57
58
59 ### a Gaussian process prior ###
60 # a prior mean function
61 g0 <- logit_f(getprior(slope, Pt, 8 N)) # 8 indicates the center of the
    dose matrix of 3x5.
62
63 # a covariance function based on the squared exponential kernel
64 rho <- 0.4 # a scale parameter
65 sigma_f <- 1 # a signal variance that is generally fixed as 1.
66 K <- matrix(rep(0,N*N), nrow=N, ncol=N) # dummy matrix for a covariance
    matrix K
67
68 # calculate K
69 for (i in 1:(N - 1)) {
70   K[i, i] <- (sigma_f)^2
71   for (j in (i + 1):N) {
72     K[i, j] = (sigma_f)^2*exp(-0.5/(rho)^2 * ((X2[i,1] - X2[j,1])^2+(X2
        [i,2] - X2[j,2])^2))
73     K[j, i] = K[i, j]
74   }
75 }
76 K[N, N] <-(sigma_f)^2
77 xi <- 0.08 # a small value to be added to diagonal elements in K
78 K2 <- K + diag(xi,N) # for improving computational stability
79 L <- t(chol(K2)) # Cholesky decomposition
80
81 ### True toxicity scenarios ###
82 # These scenarios are introduced in (Riviere et al., 2015a)
83 S1 <- c(0.05, 0.10, 0.15, 0.30, 0.45, 0.10, 0.15, 0.30, 0.45, 0.55,
    0.15, 0.30, 0.45, 0.50, 0.60)
84 S2 <- c(0.15, 0.30, 0.45, 0.50, 0.60, 0.30, 0.45, 0.50, 0.60, 0.75,
    0.45, 0.55, 0.60, 0.70, 0.80)
85 S3 <- c(0.02, 0.07, 0.10, 0.15, 0.30, 0.07, 0.10, 0.15, 0.30, 0.45,

```

```

      0.10, 0.15, 0.30, 0.45, 0.55)
86 S4 <- c(0.30, 0.45, 0.60, 0.70, 0.80, 0.45, 0.55, 0.65, 0.75, 0.85,
      0.50, 0.60, 0.70, 0.80, 0.90)
87 S5 <- c(0.01, 0.02, 0.08, 0.10, 0.11, 0.03, 0.05, 0.10, 0.13, 0.15,
      0.07, 0.09, 0.12, 0.15, 0.30)
88 S6 <- c(0.05, 0.08, 0.10, 0.13, 0.15, 0.09, 0.12, 0.15, 0.30, 0.45,
      0.15, 0.30, 0.45, 0.50, 0.60)
89 S7 <- c(0.07, 0.10, 0.12, 0.15, 0.30, 0.15, 0.30, 0.45, 0.52, 0.60,
      0.30, 0.50, 0.60, 0.65, 0.75)
90 S8 <- c(0.02, 0.10, 0.15, 0.50, 0.60, 0.05, 0.12, 0.30, 0.55, 0.70,
      0.08, 0.15, 0.45, 0.60, 0.80)
91 S9 <- c(0.005, 0.01, 0.02, 0.04, 0.07, 0.02, 0.05, 0.08, 0.12, 0.15,
      0.15, 0.30, 0.45, 0.55, 0.65)
92 S10 <- c(0.05, 0.10, 0.15, 0.30, 0.45, 0.45, 0.50, 0.60, 0.65, 0.70,
      0.70, 0.75, 0.80, 0.85, 0.90)
93 Sc <- cbind(S1, S2, S3, S4, S5, S6, S7, S8, S9, S10)
94
95 #####
96 # Simulation part #
97 #####
98 for(S in 1:1){ # Select scenarios
99
100   # true toxicity probabilities
101   X$True <- Sc[,S] # true toxicity at the selected scenario
102   near_mtd <- which((X$True <= (Pt + 0.05)) & (X$True >= (Pt - 0.05)))
103
104   for(sim in sim_min:iter){
105     # Initial values for a trial
106     next_l <- 1 # the first dose combination in the trial is the lowest
      dose combination
107     t <- 0 # number of tests
108     p_DLT <- 0 # percentage of the number of patients who experienced
      DLT
109     num <- rep(0,N) # number of patients treated in a trial
110     num_t <- rep(0,N) # number of patients who experienced DLT
111     DLT <- 0 # DLT flag for the start-up phase
112
113     #####
114     # Start-up phase #
115     #####
116     while(DLT==0){
117       t <- t + 1 # update the test number
118
119       if(t == 1){ # for the initial test
120         Test <- cbind(X[X$grid == next_l,], c = cohort) # a listing of
          (test doses, cohort size, the number of patients with DLT)
121         y <- rbinom(1, cohort, X$True[next_l]) # number of patients with
          DLT in a cohort
122         Test$y <- y
123         DLT <- y
124       }else{
125         y <- rbinom(1, cohort, X$True[next_l])
126         DLT <- y
127         obs <- cbind(X[X$grid == next_l,], c = cohort, y)
128         Test <- rbind(Test, obs)
129       }
130
131       # update the number of patients and patients with DLT at each grid
132       num[next_l] <- num[next_l] + cohort
133       num_t[next_l] <- num_t[next_l] + y
134

```

```

135     if(next_l == 11){ # The last possible dose combination for the
136       start-up phase when 3 x 5 dose combinations
137     DLT <- 1 # end the start-up phase
138     }
139     if(DLT == 0){
140       if(next_l < 5){ # fix the lowest dose of agent 1 and only
141         escalate agent 2
142         next_l <- (next_l + 1)
143       }else if(next_l == 5){ # fix the lowest dose of agent 2 and
144         only escalate agent 1
145         next_l <- 6
146       }else if(next_l == 6){ # fix the lowest dose of agent 2 and
147         only escalate agent 1
148         next_l <- 11
149       }
150     }
151     # if patients in the first cohort experienced DLT, the lowest dose
152     combination is re-tested.
153     if(DLT > 0 && sum(num) == cohort){
154       DLT <- 0
155       next_l <- 1
156     }
157   }
158   #####
159   # Model phase #
160   #####
161   start_t <- t
162   t <- start_t -1
163   while(sum(num) + cohort <= max_sub){
164     t <- t+1
165     # cat("Scinario / Sim - Test : " ,S ,"/sim= " , sim , " - t=" ,
166       t, "\n")
167     if(t > start_t){
168       y <- rbinom(1, cohort, X$True[next_l])
169       obs <- cbind(X[X$grid == next_l,], c = cohort, y)
170       Test <- rbind(Test, obs)
171       num[next_l] <- num[next_l] + cohort
172       num_t[next_l] <- num_t[next_l] + y
173     }
174     ### Sampling from the posterior distribution (MCMC part) ###
175     y <- as.array(Test$y) # toxicity outcome
176     c <- as.array(Test$c) # cohort size at each test
177     t.grid <- as.array(Test$grid) # tested grid
178     # input output data
179     data <- list(t = t, c = c, N = N, y = y, test_x = t.grid, logitg
180       = g0, L = L)
181     # generate posterior samples by rstan
182     fit <- sampling(stanmodel,
183       data = data,
184       pars = c('Ep'),
185       iter = 7000, warmup = 3000, thin = 4, chains = 1
186     )
187   }
188 }

```

```

189 # Check the MCMC results graphically
190 # ggmcmc(ggs(fit,inc_warmup=TRUE), file = 'MCMC.pdf')
191
192 ms <- rstan::extract(fit) # outputs of the sampling
193 post <- data.frame(Ep = ms$Ep) # posterior samples of toxicity
      probabilities
194
195 # calculate probabilities from a cumulative distribution
196 ove <- rep(0,N) # probabilities over the target toxicity rate
197 pro2 <- rep(0,N) # probabilities falling within the proper
      interval including the target toxicity rate
198
199 for(n in 1: N){
200   exp_stars1 <- sort(post[,n])
201   Fn <- ecdf(exp_stars1) # empirical cumulative distribution
      function
202   ove[n] <- 1 - Fn(Pt)
203   pro2[n] <- Fn(Pt + 0.1) - Fn(Pt - 0.1)
204 }
205
206 ##### Calculate the acquisition function (Expected improvement) ###
207 EI <- numeric(N)
208 gg <- abs(post - Pt) # posterior samples of the objective function
209 g_best <- min(apply(gg, 2, mean), na.rm = TRUE) # the current
      best point
210
211 # improvement function:  $I(x) = \max(g\_best - g(x), 0)$ 
212 I <- g_best - gg # calculate  $I(x)$  using all possible values of  $g(x)$ 
      ranging  $[0,1]$ .
213 Imp0 <- which(I < 0)
214 I[Imp0] <- 0
215
216 for(i in 1:N){
217   # estimate a probability density function at the grid i
218   f <- splinefun(density(x = as.vector(t(gg[i]))), n = 1000,
      from = 0, to = 1), method = "natural")
219
220   # probability density between "start and start+til"
221   start <- -til/2
222   til_n <- 1
223   p_gx <- numeric(length(g_t1)) # vector for the probability
      density function
224
225   while(til_n < (length(g_t1) + 1)){
226     p_gx[til_n] <- integrate(f, start, start + til)$value
227     if(p_gx[til_n] < 0){p_gx[til_n] <- 0}
228     start <- start + til
229     til_n <- til_n + 1
230   }
231
232   #expected improvement
233   EI[i] <- I %*% p_gx # the expected improvement function
234 }
235
236 r <- -(rank(EI)-N) + 1 # rank grid numbers in descending order
      with the EI
237 pos_tox <- apply(post, 2, mean) # posterior means of toxicity
      probabilities
238 At <- cbind(X, EI, r, pos_tox)
239 At$C <- rep(0, nrow(At))
240

```



```

241     # Admissible dose combination set (At) for the next dose
242     combination selections
243     # the model phase Step 4 requirement (a) in Section 4.2.2
244     if(length(which(Test$grid == 1)) > 0) At$C[1] <- At$C[2] <-
245     At$C[6] <- 1
246     if(length(which(Test$grid == 2)) > 0) At$C[1] <- At$C[2] <-
247     At$C[3] <- At$C[7] <- 1
248     if(length(which(Test$grid == 3)) > 0) At$C[2] <- At$C[3] <-
249     At$C[4] <- At$C[8] <- 1
250     if(length(which(Test$grid == 4)) > 0) At$C[3] <- At$C[4] <-
251     At$C[5] <- At$C[9] <- 1
252     if(length(which(Test$grid == 5)) > 0) At$C[10] <- At$C[5] <-
253     At$C[4] <- 1
254     if(length(which(Test$grid == 6)) > 0) At$C[1] <- At$C[6] <-
255     At$C[11] <- At$C[7] <- 1
256     if(length(which(Test$grid == 7)) > 0) At$C[1] <- At$C[6] <-
257     At$C[7] <- At$C[8] <- At$C[12] <- At$C[2] <- 1
258     if(length(which(Test$grid == 8)) > 0) At$C[2] <- At$C[7] <-
259     At$C[8] <- At$C[9] <- At$C[13] <- At$C[3] <- 1
260     if(length(which(Test$grid == 9)) > 0) At$C[3] <- At$C[8] <-
261     At$C[9] <- At$C[10] <- At$C[14] <- At$C[4] <- 1
262     if(length(which(Test$grid == 10)) > 0) At$C[4] <- At$C[9] <-
263     At$C[10] <- At$C[15] <- At$C[5] <- 1
264     if(length(which(Test$grid == 11)) > 0) At$C[1] <- At$C[12] <-
265     At$C[11] <- At$C[6] <- 1
266     if(length(which(Test$grid == 12)) > 0) At$C[6] <- At$C[2] <-
267     At$C[11] <- At$C[7] <- At$C[13] <- At$C[12] <- 1
268     if(length(which(Test$grid == 13)) > 0) At$C[1] <- At$C[7] <-
269     At$C[3] <- At$C[12] <- At$C[13] <- At$C[8] <- At$C[14] <-
270     1
271     if(length(which(Test$grid == 14)) > 0) At$C[2] <- At$C[8] <-
272     At$C[4] <- At$C[13] <- At$C[14] <- At$C[9] <- At$C[15] <-
273     1
274     if(length(which(Test$grid == 15)) > 0) At$C[3] <- At$C[9] <-
275     At$C[5] <- At$C[14] <- At$C[15] <- At$C[10] <- 1
276
277     # the model phase Step 4 requirement (b) in Section 4.2.2
278     for(i in 1:nrow(At)){
279         if(At$pos_tox[i] >= (Pt + eps1)){
280             At$C[i] <- 0 # if the grid is possibly too toxic, At$C is
281             0.
282         }
283     }
284
285     At <- At[order(At$r),]
286     At <- At[(At$C == 1),] # retain only dose combinations in At
287
288     # select the next dose combination
289     if(nrow(At) > 0){
290         # select a grid that maximizes the EI within dose combinations
291         in At (r is the minimum number)
292         At <- At[At$r == min(At$r), ]
293         next_l <- At$grid
294     }else{
295         # if all dose combination is excluded from At, the next
296         candidate is the lowest dose level.
297         next_l <- 1
298     }
299 } ### end one trial

```

```

281     ### Output Results ###
282     # final candidates for MTDC determination
283     select.dose1 <- c(which(pos_tox < (Pt + eps1)), which(num > 0))
284     select.dose <- select.dose1[duplicated(select.dose1)]
285
286     # determine a single MTDC
287     if(length(select.dose) > 0){
288     MTD <- which(pro2 == max(pro2[select.dose])) # equation (4.15)
289     }else{
290     MTD <- 1
291     }
292
293     p_DLT <- sum(num_t)/sum(num) # percentage of patients with DLT in a
      trial
294     C <- is.element(MTD, near_mtd) # correct selection of MTDC (as
      reference)
295
296     # Output the results
297     Result_box<- cbind(S, sim, MTD, total = sum(num), p_DLT, data.frame
      (t(num)), tox = sum(num_t), data.frame(t(num_t)), data.frame(t(
      pos_tox)), t(EI), t(pro2),t(ove), C, slope, rho, sigma_f)
298
299     if(sim == sim_min){
300     Result_box1 <- Result_box
301     }else{
302     Result_box1 <- rbindCOrder(Result_box1, Result_box)
303     }
304
305     write.table(Result_box1, "XXX/output.txt", row.names = F, quote = F,
      append = F)
306   }### end iteration of the simulation
307 }### end one scenario

```

A stan model file ‘BOD_model_combination.stan’ is composed of the following contents.

Listing B.2. Stan model for BOD of MTDC estimation

```

1 // BOD_model_combination.stan
2 // data inputs
3 data {
4   int<lower = 1> t; // cumulative number of tests at the trial (t =
      1,2,3,...)
5   int<lower = 1> c[t]; // number of patients in a cohort
6   int<lower = 1> N; // total number of dose combinations (total grids) (N
      = 3x5 = 15)
7   int<lower = 0, upper = 3> y[t]; // number of patients with DLT at each
      test
8   int<lower = 1, upper = N> test_x[t]; // grid number where patients are
      tested at each test
9   vector[N] logitg; // logit transformed initial guesses
10  matrix[N, N] L; //a covariance function based on the squared exponential
      kernel
11 }
12
13 transformed data{
14   vector[N] mu_zero = rep_vector(0, N); // make a vector with dummy
      values of 0
15 }
16
17 // parameter section (a parameter declared here means what we need to
      output)
18 parameters {
19   vector[N] zero_m;
20 }
21
22 transformed parameters{
23   // Estimate toxicity probability based on the Gaussian process prior
      with pre-specified initial guesses
24   vector[N] Ep;
25   Ep = inv_logit(zero_m + logitg); // toxicity probabilities (Ep)
26 }
27
28 // model section
29 model {
30   //----- declare a Gaussian process prior -----//
31   zero_m ~ multi_normal_cholesky(mu_zero, L);
32
33   //----- likelihood function-----//
34   for (j in 1:t){
35     y[j] ~ binomial(c[j], Ep[test_x[j]]);
36   }
37 }

```

Chapter 5

Bayesian optimization design of OD estimation for targeted therapies

Chapter 5 introduces BOD for OD estimation in Phase I/II clinical trials. BOD for the MTD or a single MTDC in Chapters 3 and 4 deals with toxicity outcomes. On the other hand, BOD for OD estimation incorporates not only toxicity outcomes but also efficacy outcomes, as introduced in Chapter 1. We organize Chapter 5 by the following sections:

- Section 5.1 covers a literature review about statistical designs for OD estimation along with brief explanations about some conventional statistical designs.
- Section 5.2 introduces the mathematical framework on BOD for OD estimation.
- Section 5.3 shows performance evaluation of BOD for OD estimation compared with some existing statistical designs by a simulation study.
- Section 5.4 summarizes the simulation results and introduces some possible future work.

We introduce this chapter based on the work *Bayesian optimization design for dose - finding based on toxicity and efficacy outcomes in Phase I/II clinical trials*, A. Takahashi and T. Suzuki, Pharmaceutical Statistics, 2021 (Takahashi and Suzuki, 2021c).

5.1 Overview of drug development and statistical designs for dose-finding studies of targeted therapies

5.1.1 Difference between mono-therapies with cytotoxic agents and targeted therapies with biologic agents in dose-finding studies

In cancer treatment, the primary goal of Phase I clinical trials is to identify the MTD based on toxicity outcomes. Although Phase I clinical trials typically collect preliminary efficacy information about investigational agents, the MTD is usually determined independently from these efficacy data. After Phase I clinical trials, Phase II clinical trials are conducted to investigate efficacy at or around the MTD. Traditional dose-finding studies assume treatment with cytotoxic agents that directly inhibit the growth of tumors. The MTD is expected to provide the greatest therapeutic benefit under the monotonicity assumption that toxicity increases with increasing dose levels; therefore, it is reasonable to conduct Phase II clinical trials to evaluate efficacy of a specific dose after the MTD determination in Phase I clinical trials when both toxicity and efficacy have monotonicity assumptions on dose-response relationships. On the other hand, targeted therapies cannot apply the monotonicity assumption to their dose-response relationships, unlike mono-therapies with

cytotoxic agents.

According to the NCI dictionary, targeted therapies are defined as a type of treatment that uses drugs or other substances to identify and attack specific types of cancer cells with less harm to normal cells. It also describes that targeted therapies block the action of certain enzymes, proteins, or other molecules involved in the growth and spread of cancer cells, or help the immune system kill cancer cells or deliver toxic substances directly to cancer cells and kill them. Biologic agents created from living organisms by biological processes greatly contribute to cancer treatment as major targeted therapies. Accordingly, biologic therapies have increasingly drawn attention in drug development (Andrews et al., 2015). Biologic agents are designed to directly target specific pathways related to specific tumors. Targeted therapies with biological agents are expected to repair, stimulate, or enhance the immune response that pertains to the cause of cancer such as malfunction of the immune system. In terms of the relationship between the amount of dose and response to tumors, once a targeted pathway is inhibited with a certain amount of biologic agents, further administration will no longer provide any benefit for patients. Additionally, targeted therapy may have fewer side effects because not attacking healthy cells, unlike chemotherapy based on cytotoxic agents.

Considering the mode of action, it is inappropriate to assume that, for biologic agents, both toxicity and efficacy will increase monotonically with increasing dose levels. For example, non-monotone patterns on dose-efficacy relationships may be drawn as Fig. 5.1.




Fig. 5.1. Dose–response relationships with non-monotone patterns

Thus, instead of an MTD, biologic agents require the estimation of an optimal dose, OD, that provides sufficient efficacy under an acceptable toxicity rate as exemplified in Fig. 5.2. Suppose the target response rates for toxicity and efficacy denote θ_T and θ_E , respectively. The OD is defined as a dose providing the efficacy of greater than or equal to θ_E and the toxicity of not over θ_T . Given the commonality of trial designs that incorporate both toxicity and efficacy data, and the fact that new drugs with a variety of mode of actions (e.g., biologic agents) are being increasingly developed, statistical strategies to find ODs have been getting more and more of interest and importance; thus, it is worthwhile to improve statistical designs for dose-finding studies of OD estimation.




Fig. 5.2. Optimal dose estimation for a targeted therapy with biologic agents (All curves except for a red dot line are dose-efficacy curves)

There are several approaches for patient outcomes that statistical designs for OD estimation deal with. Mandrekar et al. (2010) and Zhong et al. (2012) handle trivariate outcomes that are composed of acceptable toxicity without efficacy, acceptable toxicity with efficacy, and severe toxicity. There are a few statistical designs that treat continuous variables instead of binary variables especially on efficacy outcomes for OD estimation. However, overall, most statistical designs for OD estimation utilize bivariate outcomes composed of binary responses for both toxicity and efficacy.

In the following, we will introduce some popular statistical designs dealing with binary outcomes of toxicity and efficacy for OD estimation while considering non-monotone patterns on dose–efficacy relationships.

5.1.2 Parametric model-based designs

For parametric model-based designs, the CRM described in Section 2.3 has been extended for MTDC estimation as mentioned in Section 4.1.4, and OD estimation is no exception. Braun (2002) has introduced the bivariate CRM to handle bivariate outcomes as an extended version of the CRM; however, the bivariate CRM does not account for non-monotone patterns in dose–efficacy relationships. As other parametric model-based designs, there are a conditional auto-regressive model (Muenz et al., 2019), a toxicity-efficacy odds ratio trade-off (Yin et al., 2006), a Gumbel bivariate logistic regression (Dragalin and Fedorov, 2006), and a logistic design (Zang et al., 2014). A Clayton model (Yuan and Yin, 2009) would be a more advanced approach that incorporates toxicity outcomes with a Cox proportional hazards model and efficacy outcomes with a cure rate model in order to consider the correlation between onset times and outcomes. Of parametric model-based designs, the efficacy-toxicity trade-off (EffTox) design introduced by Thall and Cook (2004) is one of the most popular one because it has been utilized in comparative analyses in several articles such as Takeda et al. (2018); Li et al. (2017); Liu and Johnson (2016); Bekel and Shen (2005). Thall et al. (2006) and Brock et al. (2017) have also introduced how to implement the EffTox with practical considerations. The EffTox incorporates trade-off contours between probabilities of toxicity and efficacy in its dose selection algorithm, after providing a joint probability function based on a Gumbel distri-

bution (Murtaugh and Fisher, 1990) composed of marginal probabilities for dose–toxicity and dose–efficacy relationships derived from logit models. The EffTox covers dose–efficacy relationships having non-monotonic patterns through a quadratic term in its logit model. We introduce the EffTox more detail from the following.

Dose-response model for the EffTox

The EffTox models marginal probabilities of efficacy and toxicity as follows:

$$\text{logit}\{\pi^E(x)\} = \mu_E + \gamma_{E_1}x + \gamma_{E_2}x^2, \quad (5.1)$$

and

$$\text{logit}\{\pi^T(x)\} = \mu_T + \gamma_Tx, \quad (5.2)$$

where x is a conceptual dose at an arbitrary dose level. A conceptual dose corresponding to an actual dose (d) at a dose level j ($j \in \{1, \dots, J\}$) is given by $x_j = \log d_j - J^{-1} \sum_{j=1}^J \log d_j$. In equation (5.1), non-monotonic patterns are considered through γ_{E_2} with a quadratic term.

In the same manner as Murtaugh and Fisher (1990), a joint probability function of toxicity and efficacy is given by a Gumbel distribution as follows:

$$\begin{aligned} P(\mathbf{Y} = \{y^T, y^E\}) &= (\pi^T)^{y^T} (1 - \pi^T)^{1-y^T} (\pi^E)^{y^E} (1 - \pi^E)^{1-y^E} \\ &\quad + (-1)^{y^T+y^E} \pi^T (1 - \pi^T) \pi^E (1 - \pi^E) \left(\frac{e^\eta - 1}{e^\eta + 1} \right), \end{aligned} \quad (5.3)$$

where \mathbf{Y} is composed of random variables Y^T on a toxicity outcome and Y^E on an efficacy outcome in a patient treated with a dose x and takes $\{0, 1\}$; $\pi^T = \pi^T(x)$; $\pi^E = \pi^E(x)$; η denotes an association parameter between toxicity and efficacy outcomes. Thus, the EffTox has six parameters in its models: μ_E , γ_{E_1} , γ_{E_2} , μ_T , γ_T and η . These parameters with often normal priors are estimated through Bayes' rule.

Utility contour and dose selection

Once patient outcomes are observed, posterior sampling for π^T and π^E is performed by MCMC. After posterior samples on π^T and π^E are generated, the EffTox incorporates them through the following utility contour:

$$U_{\text{EffTox}}(x) = 1 - \left[\left\{ \frac{1 - \pi^E(x)}{1 - \pi^{E_1}} \right\}^\iota + \left\{ \frac{\pi^T(x)}{\pi^{T_2}} \right\}^\iota \right]^{\frac{1}{\iota}}, \quad (5.4)$$

where ι is determined to have the utility contour go through pre-specified three points $(\pi^{E_1}, 0)$, $(1, \pi^{T_2})$ and (π^{E_3}, π^{T_3}) . The value of π^{E_1} returns the smallest desirable efficacy probability that is expected when a toxicity probability is 0. The maximum value of toxicity probability when an efficacy probability is 1 denotes π^{T_2} that provides the same desirability as $(\pi^{E_1}, 0)$. The third point (π^{E_3}, π^{T_3}) is equally desirable having the intermediate position between the other two pre-specified points.

The next dose is selected from an admissible dose set \mathcal{A}_t by maximizing either a plug-in mean or a posterior mean on equation 5.4 as follows:

$$x_{(t+1)} = \arg \max_{x \in \mathcal{A}_t} \hat{U}_{\text{EffTox}}(x). \quad (5.5)$$

Note that Thall and Cook (2004) introduced plug-in means on posterior probabilities of efficacy and toxicity into equation (5.4) to obtain $\hat{U}_{\text{EffTox}}(x)$; however, the R package `trialr`, which is an open-source to implement the EffTox, adopts posterior means of the utility contour for the dose selection. The admissible dose set \mathcal{A}_t at the t th test is composed of all doses that satisfy following criteria:

$$\mathbb{P}\{\pi^E(x) > \underline{\lambda}^E \mid D_{1:t}\} > c_E, \quad (5.6)$$

and

$$\mathbb{P}\{\pi^T(x) < \bar{\lambda}^T \mid D_{1:t}\} > c_T, \quad (5.7)$$

where $\underline{\lambda}^E$ and $\bar{\lambda}^T$ are acceptable response rates on a lower bound of efficacy and an upper bound of toxicity, respectively. The tunable design parameters of $\underline{\lambda}^E$, $\bar{\lambda}^T$, c_E and c_T are provided by investigators before the trial.

If $\mathcal{A}_t = \emptyset$, then the trial is terminated and the EffTox declares that there is no OD in the explored doses. When the trial is terminated according to pre-specified stopping criteria and $\mathcal{A}_t \neq \emptyset$, an OD is determined as a dose maximizing utility score based on equation (5.5).

Open research questions

The EffTox can be implemented with both desktop software and the R package; however, it requires numerous parameters to be specified before trials. Dose–response models need at least six parameters, and at least three additional parameters are needed for trade-off utility functions. In the practical implementation, the EffTox requires twelve hyperparameters related to the six model parameters; these hyperparameters are calculated by eliciting means from both dose–response curves, and by referencing the target mean effective sample size. Although the desktop software automatically generates those numerous hyperparameters when the elicited means and target mean effective sample sizes are designated, complex parameter settings might make implementation difficult.

While numerous authors have introduced parametric model-based designs, those strategies requiring strict assumptions for dose–response relationships do not always offer the best solution because there is little or no information about toxicity or efficacy on investigational agents in practice, which is a common concern for dose-finding studies including MTD and MTDC estimation. Thus, parametric model-based designs are faced with a potential risk of model misspecification, particularly in small sample size trials. In addition, for at least OD estimations based on bivariate binary outcomes, several nonparametric designs have been shown to outperform certain parametric model-based designs (Zang et al., 2014; Li et al., 2017). These facts imply that parametric model-based designs may not be always the best choice for addressing OD estimation problems.

5.1.3 Optimal interval designs

Although a majority of statistical designs for OD estimations are parametric model-based designs, optimal interval designs that are similar to toxicity probability interval designs in Sections 2.4.1 and 4.1.3 may be currently the most prominent of nonparametric approaches. They utilize divided probability intervals related to certain levels of response probabilities for both toxicity and efficacy. For example, Lin and Yin (2017b) has introduced a simple toxicity and efficacy interval design (STEIN) that is an optimal interval design and has extended the BOIN for MTD estimation (Liu and Yuan, 2015) to address

the OD estimation problem. In the STEIN, the dose selection algorithm is based on optimal intervals for toxicity and efficacy minimizing probabilities of misclassifying observed values against its pre-specified hypotheses to guide whether the next dose should remain at the same dose level or move from the current dose level by one dose level. Similar to the STEIN, Takeda et al. (2018, 2020) have also introduced the extended version of the BOIN. Additionally, Li et al. (2017) has introduced an optimal interval design incorporating a joint unit probability mass for toxicity and efficacy probability intervals. Of optimal interval designs, we focus on the STEIN that would offer the simple strategy in terms of its implementation as we can see from the design name. The specifics of the STEIN are introduced in the following.

Optimal intervals on toxicity and efficacy probabilities in the STEIN

The STEIN's optimal intervals are determined as follows. Suppose that the current dose level is j ($j \in \{1, \dots, J\}$). For a toxicity probability, the STEIN applies three simple hypotheses at each dose level j ; $H_{1j}^T : \pi_j^T = \phi_1^T$ versus $H_{0j}^T : \pi_j^T = \phi_0^T$ versus $H_{2j}^T : \pi_j^T = \phi_2^T$. In addition, threshold values of ϕ_1^T and ϕ_2^T that express overly safe and excessively toxic construct a local indifference interval for an acceptable toxicity ϕ_0^T ($0 < \phi_1^T < \phi_0^T < \phi_2^T < 1$). Let $\hat{\pi}_j^T$ denote a sample statistic of toxicity proportions that divide the number of response outcomes at j by the number of patients treated at j . The probability that $\hat{\pi}_j^T$ is incorrectly classified under three hypotheses is given by

$$\begin{aligned} \mathcal{C}^T(\phi_L^T, \phi_U^T) = & \text{P}(H_{0j}^T) \text{P}(\hat{\pi}_j^T \leq \phi_L^T \text{ or } \hat{\pi}_j^T \geq \phi_U^T | H_{0j}^T) \\ & + \text{P}(H_{1j}^T) \text{P}(\hat{\pi}_j^T > \phi_L^T | H_{1j}^T) + \text{P}(H_{2j}^T) \text{P}(\hat{\pi}_j^T < \phi_U^T | H_{2j}^T), \end{aligned} \quad (5.8)$$

where prior probabilities to take each hypothesis are usually set a uniform prior as $\text{P}(H_{0j}^T) = \text{P}(H_{1j}^T) = \text{P}(H_{2j}^T) = 1/3$. The lower and upper critical threshold values of toxicity probability expressed as ϕ_L^T and ϕ_U^T are determined to minimize misclassification probability \mathcal{C}^T and treated as an indifferent tolerance interval of the target toxicity rate.

For efficacy, the optimal interval is determined by the similar framework to that for toxicity. Firstly, two hypotheses are given for efficacy; $H_{1j}^E : \pi_j^E = \phi_1^E$ versus $H_{2j}^E : \pi_j^E = \phi_2^E$, where pre-specified threshold values of ϕ_1^E and ϕ_2^E are a clinically uninteresting response rate and a clinically desired response rate, respectively ($\phi_1^E < \phi_2^E$). The misclassification probability of the efficacy outcome $\hat{\pi}_j^E$ denoting a sample statistic of proportion of observed response at j is given by

$$\mathcal{C}^E(\phi^E) = \text{P}(H_{1j}^E) \text{P}(\hat{\pi}_j^E \leq \phi^E | H_{1j}^E) + \text{P}(H_{2j}^E) \text{P}(\hat{\pi}_j^E < \phi^E | H_{2j}^E), \quad (5.9)$$

where prior probabilities for each hypothesis are usually $\text{P}(H_{1j}^E) = \text{P}(H_{2j}^E) = 1/2$. The sample critical value ϕ^E is determined by minimizing \mathcal{C}^E .

The STEIN's dose allocation procedures

The STEIN starts from the lowest dose level to treat patients in the first cohort. After the t th test, dose selection is based on the optimal intervals for toxicity and efficacy outcomes as follows:

1. If $\hat{\pi}_j^T \geq \phi_U^T$, the next dose level is $j - 1$.
2. If $\hat{\pi}_j^T < \phi_U^T$ and $\hat{\pi}_j^E \geq \phi^E$, the next dose level remains j .
3. If $\hat{\pi}_j^T < \phi_U^T$ and $\hat{\pi}_j^E < \phi^E$, the next dose is selected to maximize the posterior

efficacy probability greater than ϕ^E given by

$$j_{(t+1)} = \arg \max_{j \in \mathcal{A}_t} P(\pi_j^E > \phi^E | D_{1:t}). \quad (5.10)$$

In the calculation of posterior efficacy probabilities, a uniform vague prior is assigned to each π_j^E . The admissible dose set \mathcal{A}_t is $\{j-1, j, j+1\}$ when $\hat{\pi}_j^T \leq \phi_L^T$; $\{j-1, j\}$ when $\phi_L^T < \hat{\pi}_j^T < \phi_U^T$; and otherwise, $\{j-1\}$.

Fig. 5.3 depicts partitions for the dose selections.




Fig. 5.3. Dose selection partitions in the STEIN

The trial keeps treating patients with a selected dose based on the above rules until the maximum sample size is reached unless other stopping criteria are specified. When all patients are treated, an OD is selected by a utility function given by

$$U_{\text{STEIN}}(\tilde{\pi}_j^T, \tilde{\pi}_j^E) = \tilde{\pi}_j^E - w_1 \tilde{\pi}_j^T - w_2 \tilde{\pi}_j^T 1[\tilde{\pi}_j^T > \phi_0^T], \quad (5.11)$$

where w_1 and w_2 are nonnegative pre-specified values. The isotonicly transformed values $\tilde{\pi}_j^T$ construct a dose-toxicity curve after performing an isotonic regression on observed toxicity probabilities $\hat{\pi}_j^T$. For a dose-efficacy curve, unimodal isotonic regressions (Turner and Wollan, 1997), which incorporate all possible models in dose-efficacy curves, are performed J times using observed efficacy probabilities $\hat{\pi}_j^E$ to estimate the unimodal isotonicly transformed values $\tilde{\pi}_j^E$. An OD is determined as a dose that maximizes the utility function.

Open research questions

Optimal interval designs including the STEIN are generally simple to implement because they do not rely on complex models. The same improvement points raised in Sections 2.4.1 and 4.1.3 can be applied again because they are a common concern among nonparametric interval designs even if we change our ultimate goal. Optimal interval designs for OD estimation also focus on the current dose and judge whether the next dose should be escalated or de-escalated ± 1 dose levels from the current dose; therefore, they offer limited dose candidates in dose selections. In addition, they do not focus on the relationships of toxicity and efficacy across dose levels until the end of a trial. If the current probabilities

pertaining to toxicity and efficacy relationships among all dose levels are provided step by step, it may help capture and anticipate the dose–response relationships.

5.1.4 Isotonic design

As another nonparametric approach that is different from optimal interval designs, we introduce the isotonic design proposed by Zang et al. (2014). The isotonic design utilizes double-sided isotonic regressions for estimating dose–efficacy relationships. In the double-sided isotonic regression, two separate isotonic regressions are conducted by dividing them at each branch dose level sequentially to make dose–efficacy curves representing non-monotone patterns. A user-friendly cloud software we can access from <https://biostatistics.mdanderson.org/shinyapps/MTADF/> helps implement the isotonic design.

In the isotonic design, toxicity at each dose level are independently modeled by a beta-binomial model, that is, $Y_j \sim \text{Bin}(n_j, \pi_j^T)$ and $\pi_j^T \sim \text{Beta}(\alpha_j, \beta_j)$, where α_j and β_j are the hyperparameters. The isotonic design starts evaluations from the lowest dose. Once patient outcomes are observed, an admissible dose set \mathcal{A}_t after the t th cohort is decided as follows:

$$\mathcal{A}_t = \{j : \tilde{\text{P}}(\pi_j^T > \theta_T | D_{1:t}) < \acute{c}_T, j = 1, \dots, J\}, \quad (5.12)$$

where $\tilde{\text{P}}(\pi_j^T > \theta_T | D_{1:t})$ is isotonicallly transformed values of posterior probabilities $\text{P}(\pi_j^T > \theta_T | D_{1:t})$. A toxicity threshold of \acute{c}_T pre-specified for overdose control has been recommended 0.8 for general use.

Dose selection

The next dose is selected to provide the highest isotonic estimate on efficacy probabilities in \mathcal{A}_t . The isotonic estimates of efficacy probabilities are briefly calculated as follows:

- The double-sided isotonic regression that fits two separate standard isotonic regressions using the pooled adjacent violators algorithm (Barlow et al., 1972) is performed J times to generate all J possible isotonic estimate curves $\{\tilde{\pi}_j^{E_\varrho}; j = 1, \dots, J\}$, $\varrho = 1, \dots, J$.
- Dose selection utilizes the isotonic estimate curve $\{\tilde{\pi}_j^{E_{\varrho^*}}\}$ based on the following ϱ^* that provides the smallest sum of the squared errors:

$$\varrho^* = \arg \min_{\varrho \in \{1, \dots, J\}} \sum_{j=1}^J (\tilde{\pi}_j^{E_\varrho} - \hat{\pi}_j^E)^2. \quad (5.13)$$

Once the maximum sample size is reached, the dose with the highest efficacy is selected as an OD from an admissible dose set. In the above dose selection, the isotonic regression is conducted using only already tried doses; therefore, if the highest tested dose level has the maximum efficacy and the next higher dose is included in \mathcal{A}_t , the next dose becomes automatically an escalated dose to further explore the dose–efficacy relationship and search the maximum point.

Open research questions

The isotonic design decides the next dose on the basis of 1) an admissible dose set determined by isotonicallly transformed posterior probabilities for toxicity and 2) the highest

isotonic estimate from the double-sided isotonic regression for efficacy. Most existing designs rely their dose selections on the current best point without considering uncertainties, and the isotonic design is no exception. Such decision-making might cause suboptimal dose allocations because estimates might have high uncertainties owing to a lack of information about investigational agents, particularly at the beginning of a trial.

5.2 Modeling frameworks and practical procedures of Bayesian optimization design for OD estimation

Given the open research questions raised in Section 5.1, Section 5.2 introduces BOD for identifying single ODs based on bivariate binary outcomes assuming Phase I/II clinical trials for targeted therapies with biologic agents. BOD for OD estimation is also a similar approach to BOD in Chapters 3 and 4 except for bivariate outcomes on efficacy and toxicity. Firstly, we separately model dose–toxicity and dose–efficacy relationships as unknown nonlinear functions. Additionally, we place prior distributions on both function spaces to estimate the unknown functions through observation data. In order to solve OD estimation problems through Bayesian optimization frameworks, we design a utility function that becomes an objective function to be optimized and incorporates both dose–response relationships on toxicity and efficacy so that ODs represent the minimizer of the utility function among available dose levels. Once we observe patient outcomes, posterior distributions for both responses are updated according to the Bayes’ rule. Simultaneously, posterior distributions for the utility function are computed by the posterior distributions for each response probabilities. An acquisition function based on posterior distributions of the utility function guides dose selections and designates the next candidate for an OD. We repeat the cycle of obtaining data, updating unknown functions for each response, and selecting the next candidate dose until pre-specified stopping rules are met. Finally, we determine a single OD based on the final posterior distribution for the utility function.

In this section, we introduce modeling frameworks on BOD for OD estimation. This section is composed of the following:

- Section 5.2.1 presents BOD’s nonparametric modeling on dose–response relationships.
- Section 5.2.2 explains how to incorporate toxicity and efficacy estimations for OD selections.
- Section 5.2.3 introduces all steps to implement BOD for OD estimation with additional restrictions for overdose control. In this section, we explain another acquisition function that is also one of the most popular ones aside from the EI that BODs in Chapters 3 and 4 utilize.

5.2.1 Statistical modeling on dose–response relationships

BOD for OD estimation expresses unknown dose–response relationships for both toxicity and efficacy as follows:

$$f^R(x) = \text{logit}\{\pi^R(x)\} = \log \left\{ \frac{\pi^R(x)}{1 - \pi^R(x)} \right\}, \quad (5.14)$$

where $\pi^R(x)$ denotes response probabilities at a corresponding rescaled dose x in the finite dose range $\mathcal{X} = \{x \in \mathbb{R} \mid x_1 \leq x \leq x_J\}$. The structure of equation (5.14) is the same as BODs for MTD and MTDC estimation in Chapters 3 and 4. Regarding the description

R , we assign T to R when the focus is on toxicity responses. Likewise, we assign E to R when focusing on efficacy responses. Because both toxicity and efficacy responses can be expressed in the same statistical framework, we henceforth use R as a common variable denoting either toxicity or efficacy. In addition, our nonparametric approach with flexible modeling enables us to apply the same concept to each dose–response relationship via nonlinear functions. Regarding patient outcomes, the number of patients with a response (Y_j^R) at an arbitrary dose level j ($j \in \{1, \dots, J\}$) follows a binomial distribution $Y_j^R \sim \text{Bin}(n, \pi_j^R)$, where $\pi_j^R = \text{logit}^{-1}\{f^R(x_j)\} = [1 + \exp\{-f^R(x_j)\}]^{-1}$. On the basis of the binomial distribution, the likelihood function up to the t th test is given by

$$L(D_{1:t} \mid f^T, f^E) = \prod_{R \in \{T, E\}} \prod_{z=1}^t \pi^R(x_{(z)})^{y_{(z)}^R} \{1 - \pi^R(x_{(z)})\}^{(n_{(z)} - y_{(z)}^R)}, \quad (5.15)$$

where $D_{1:t} = \{(n_{(1)}, x_{(1)}, y_{(1)}^T, y_{(1)}^E), \dots, (n_{(t)}, x_{(t)}, y_{(t)}^T, y_{(t)}^E)\}$; $y_{(z)}^R$ denotes the number of patients who experienced the response R out of $n_{(z)}$ patients treated with a dose $x_{(z)}$ at the z th test.

We note that responses for toxicity and efficacy are treated independently in the model. If we consider the correlation, for example, by using a joint distribution of toxicity and efficacy responses, it could improve the efficiency of the estimates of π^R in theory. The ignorability of the dependence between toxicity and efficacy responses in OD estimation problems has been discussed and evaluated in Liu and Johnson (2016); Muenz et al. (2019). They reported that including an association between toxicity and efficacy responses does not improve operating characteristics, particularly when the sample size is small. Thus, assumptions of dependence between toxicity and efficacy responses within a subject are not applied in BOD for OD estimation regardless of associations between them to retain the simplicity of the estimation process in early phase trials with small sample sizes.

A Gaussian process is adopted as a prior distribution for f^R to estimate it in a Bayesian manner in the same way as BODs for MTD and MTDC estimation, that is, $f^R \sim \text{GP}(m^R, k)$, where $m^R(x)$ is a mean function corresponding to either of responses R and $k(x, x')$ is a covariance function. Prior inputs on response probabilities for π^R associated with $m^R(x)$ can be derived from pre-specified initial guesses by investigators; however, there is little information regarding investigational agents in many cases in practice. A concrete example for setting prior inputs under those cases will be introduced in Section 5.3.2. For a covariance function, we use the squared exponential kernel given by equation (3.3). As mentioned in Section 3.1.1, it is still expected that the squared exponential kernel can provide sufficient solutions for OD estimation because both dose–response relationships are assumed to be sufficiently smooth functions in general. Additionally, we apply the same covariance function expressed by equation (3.4) having a small value (ξ) to improve computational stability.

Under the statistical models and assumptions described in this section, posterior distributions for f^R and π^R are computed based on the Bayes’ rule by the MCMC after patient outcomes are observed.

5.2.2 Integrating toxicity and efficacy information

Utility functions integrate estimated information about toxicity and efficacy probabilities to identify ODs that satisfy pre-specified toxicity and efficacy target rates. For example, Li et al. (2017); Liu and Johnson (2016); Asakawa and Hamada (2013); Tao et al. (2013); Zhong et al. (2012); Bekel and Shen (2005); Thall and Cook (2004); Braun (2002) utilize a variety of utility functions in their proposed designs. We employ the following utility

function for BOD:

$$U_{\text{BOD}}(x) = w_T 1[\pi^T(x) > \theta_T] \{\pi^T(x) - \theta_T\}^2 + w_E 1[\pi^E(x) < \theta_E] \{\pi^E(x) - \theta_E\}^2, (5.16)$$

where w_R is any nonnegative real number used to weight a penalty when estimated posterior response probabilities π^R deviate from the pre-specified target rate θ_R toward unfavorable conditions. The balance between the weight values for safety and efficacy is important; that is, the ratio between w_T and w_E affects operating characteristics, while the absolute value itself is negligible. As previously mentioned, the utility function of BOD is regarded as objective functions to be optimized in Bayesian optimization frameworks. On the basis of Bayesian optimization frameworks, we seek a dose that minimizes the utility function.

The target rate θ_R denotes the maximum tolerated toxicity rate as θ_T and the minimum acceptable efficacy rate as θ_E . In general, the values for θ_R are provided by clinicians based on their knowledge and expectation, although useful information from such as historical data may be involved in the consideration as needed. In any case, the values need to be set considering the following background. In equation (5.16), we assume doses offer clinically meaningful effect as ODs when they satisfy both conditions that toxicity is less than or equal to θ_T and efficacy is greater than or equal to θ_E . Doses with less toxic and higher efficacy than each θ_R are more ideal; however, in practice, it might be the most realistic situation to seek one dose that meets the above-required conditions as quickly as possible within the limited evaluations and sample size.

As a side note, the utility function focuses on worsen effects of both responses by adding the distances from each target response rate θ_R in the form of a sum. If both beneficial and worsen effects were considered in such a simple sum function, the utility score would be affected by the beneficial or worsen effect that has greater magnitude when one of the toxicity or efficacy responses shows a benefit while the other shows some worsen the effect. Thus, no discrimination is made in equation (5.16) among the dose levels that are within the safe and efficacious region where the dose levels with toxicity rates less than or equal to θ_T and efficacy rates greater than or equal to θ_E . If discrimination is required within those dose levels, different approaches including different types of utility functions will be needed.

The posterior distribution of $U_{\text{BOD}}(x)$ is calculated with the updated distributions of π^R . A single candidate OD is selected as a next dose based on posterior samples of $U_{\text{BOD}}(x)$, as described in the following section.

5.2.3 BOD implementation steps for OD estimation

In this section, we introduce BOD implementation steps along with some restrictions for overdose control.

Overdose control

In BOD for OD estimation, we define an admissible dose set $\mathcal{A}_t \subset \mathcal{X}$ for overdose control so that overly toxic doses are eliminated temporally from the dose selection. \mathcal{A}_t is determined before the next dose selection and refreshed at each test. \mathcal{A}_t satisfies following all conditions:

- Dose skipping for dose escalation is restricted; candidate doses are up to one dose level higher than the already-tested maximum dose level.
- If two or more patients experience toxicity at $x_{(t)}$ (assuming a cohort size of 3), the maximum candidate dose in \mathcal{A}_t is one dose level lower than $x_{(t)}$.

- All dose levels in \mathcal{A}_t should meet the following condition:

$$\hat{\pi}^T(x) < (\theta_T + \check{\epsilon}_1), \quad (5.17)$$

where $\hat{\pi}^T(x)$ is a posterior mean based on posterior distributions of $\pi^T(x)$ and $\check{\epsilon}_1$ is a pre-specified value for overdose control. Since $\hat{\pi}^T(x)$ is estimated based on very limited data, $\check{\epsilon}_1$ is determined from the viewpoint of a balance of the additional allowance from θ_T and patient safety. The appropriate value for $\check{\epsilon}_1$ can remove potentially toxic doses and does not control dose selections overly. Although the value of $\check{\epsilon}_1$ is dependent on the target toxicity rate θ_T , it would be worth evaluating values in the range of $[0.05, 0.20]$ for $\check{\epsilon}_1$.

We note that efficacy is not considered in \mathcal{A}_t while BOD can be generalized to incorporate it. For example, $\hat{\pi}^E(x) > (\theta_E - \check{\epsilon}_1)$ might be an additional condition to reflect efficacy results in \mathcal{A}_t , where $\hat{\pi}^E(x)$ is a posterior mean based on posterior distributions of $\pi^E(x)$.

Dose selection strategy

Recall that we aim to achieve our ultimate goal (i.e., identifying a dose providing the minimum value of $U_{\text{BOD}}(x)$) through an acquisition function which is an alternative of the true objective function $U_{\text{BOD}}(x)$ because the exact form of $U_{\text{BOD}}(x)$ is still unavailable.

After the t th cohort is tested, the next candidate is selected from \mathcal{A}_t using the lower confidence bound criteria, LCB (Srinivas et al., 2010). Although BODs for MTD and MTDC estimation in Chapters 3 and 4 utilize the EI, we employ the LCB in this chapter, currently, one of the most popular acquisition functions aside from the EI and empirically known to have sound operating characteristics. Additionally, the strategy based on the LCB has been shown to provably converge (Srinivas et al., 2010). Although we do not show the evaluation results using the EI in BOD for OD estimation, it is expected that the performance is almost similar between those two acquisition functions based on our preliminary simulation evaluations.

The LCB returns exploitation and exploration trade-off without additional complex calculations. It might work well even we only rely on the current best point on average without any exploration, depending on the situation. On the other hand, there is a possibility of another better point than the current best point. A trade-off between exploitation and exploration plays an important role in searching such a better point faster especially when the number of evaluations is limited.

The LCB is typically defined as $\text{LCB}(x) = \mu_u(x) + u\sigma_u(x)$, where $\mu_u(x)$ and $\sigma_u(x)$ is a mean function and a square root variance at a dose x on the posterior distribution of an objective function, respectively; the tunable arbitrary constant $u > 0$ balances between exploitation and exploration directly. Instead, we alternate the LCB derived by $\sigma_u(x)$ to the one defined as the q th percentile ($0 < q < 50$) of the posterior distribution of $U_{\text{BOD}}(x)$ because it is more straightforward to use available posterior samples directly without additional calculations. Hereafter, we denote the LCB based on the q th percentile as $\text{LCB}(x)$. $\text{LCB}(x)$ directly balances exploitation and exploration of $U_{\text{BOD}}(x)$ to determine where next to test from the objective function. The next candidate dose is given by

$$x_{(t+1)} = \arg \min_{x \in \mathcal{A}_t} [\text{LCB}(x)]. \quad (5.18)$$

The closer q is 50, the stronger exploitation of the estimated $U_{\text{BOD}}(x)$ becomes. In contrast, $\text{LCB}(x)$ with a smaller value of q returns values with larger uncertainties on the estimated $U_{\text{BOD}}(x)$; therefore, it includes more explorations. We utilize 10th percentile

for q in this thesis. The value of q is tunable; however, 10th percentile would work successfully according to our experiences applying BOD to various simulation scenarios for dose-finding studies. If stronger exploitation or exploration are needed, the value of q is changed to larger or smaller, respectively (e.g., a value in the range within [5, 20] may be the candidate.).

Implementation steps

Dose selection algorithm in BOD for OD estimation is as follows:

1. Patients in the first cohort are treated with the lowest dose level.
2. Once patient outcomes are observed, posterior distributions for $f^R(x)$ as well as $\pi^R(x)$ are computed. Simultaneously, a posterior distribution for $U_{\text{BOD}}(x)$ is updated.
3. Based on the cumulative data so far, \mathcal{A}_t is determined for the next dose selection.
4. The next dose is selected in accordance with equation (5.18).
5. Steps 2 through 4 are repeated until the total number of patients reaches a maximum sample size, or other trial conditions trigger any pre-specified stopping rules.
6. At the end of a trial, a single OD x^* is determined based on the final posterior distribution of $U_{\text{BOD}}(x)$ as follows:

$$x^* = \min_x [\mathbb{E}_{U_{\text{BOD}}} \{U_{\text{BOD}}(x) \mid D_{1:t}, x \in \{x \mid \hat{\pi}^T(x) < (\theta_T + \check{\epsilon}_2), \hat{\pi}^E(x) > (\theta_E - \check{\epsilon}_2)\}\}], \quad (5.19)$$

where $\check{\epsilon}_2$ is equal to $\check{\epsilon}_1$ or is obtained by adding around 0.05 to $\check{\epsilon}_1$.

- 6' Here is an optional condition. If we wish to add a criterion to the OD estimation in order to assess the appropriateness of the final estimate, one option is to consider whether posterior probabilities for $U_{\text{BOD}}(x)$ are greater than a threshold c_{u1} :

$$P(U_{\text{BOD}}(x) > c_{u1} \mid D_{1:t}) < c_{u2}. \quad (5.20)$$

A positive value c_{u1} represents a threshold value to define undesirable regions based on the setting of equation (5.16). The value c_{u2} ranges from 0 to 1. If an arbitrary dose x meets desirable conditions more convincingly as ODS, $U_{\text{BOD}}(x) \mid D_{1:t}$ draws a distribution concentrated near 0. In contrast, $U_{\text{BOD}}(x) \mid D_{1:t}$ returns greater values than 0 when an arbitrary dose x deviates further from the desirable conditions. Therefore, if $P(U_{\text{BOD}}(x) > c_{u1} \mid D_{1:t})$ is larger, there is a higher probability that the dose x provides undesirable utility scores. If the final OD estimation meets the condition in equation (5.19) along with equation (5.20), the dose is declared as the final estimation in the trial. Otherwise, the final estimation becomes NA as no appropriate OD is found in the trial. A smaller c_{u2} value generates a more conservative decision.

5.3 Performance evaluation

As a performance evaluation for BOD, the operating characteristics of BOD were evaluated and compared with three different designs via a simulation study. Section 5.3 covers simulation frameworks for the simulation study and the results, which is composed of the following subsections:

- Section 5.3.1 explains common settings across a simulation study,
- Sections 5.3.2 and 5.3.3 introduce specific settings for each design compared in the simulation study,
- Sections 5.3.4 shows simulation results.

While there is currently no standard design for OD estimation, we compared BOD against the EffTox (Thall and Cook, 2004), the STEIN (Lin and Yin, 2017b), and the isotonic design (Zang et al., 2014); the EffTox explained in Section 5.1.2 is one of the popular parametric model-based designs; the STEIN explained in Section 5.1.3 offers one of the major nonparametric designs for OD estimation on the basis of the BOIN; the isotonic design explained in Section 5.1.4, which is another nonparametric design other than optimal interval designs, has been reported to outperform one of the parametric model-based designs.

5.3.1 General settings for the simulation study

Suppose that we desired to identify a single OD defined as a dose providing the efficacy of greater than or equal to 0.5, and toxicity that did not exceed 0.3. Thus, target response rates for toxicity and efficacy in the simulation study were $\theta_T = 0.3$ and $\theta_E = 0.5$, respectively.

We selected twelve true dose–response relationships associated with six doses (x_1, \dots, x_6), as shown in Fig. 5.4. Table 5.1 provides the true probability combinations of toxicity and efficacy responses by scenario and displays true OD in bold font. In scenario 1, toxicity and efficacy both monotonically increase with increasing dose levels, and higher doses have unacceptable toxicity. In scenario 2, the dose–efficacy curve exhibits a unimodal shape. In scenario 3, efficacy remains near the target rate and varies only slightly between doses; on the other hand, toxicity increases linearly and is unacceptable from the middle of the dose range. In scenario 4, efficacy decreases monotonically with negligible toxicity in all dose levels. In scenario 5, efficacy increases substantially with increasing dose levels under acceptable toxicity, and there are two OD candidates. Scenario 6 produces monotonic increases for both toxicity and efficacy with increasing dose levels, similar to scenario 1; however, it exhibits more gradual shapes than scenario 1. In scenario 7, intermediate dose levels 3 and 4 are the most desirable. The OD in scenario 8 is at the highest dose level because both toxicity and efficacy show very moderate slopes with increasing dose levels. In scenario 9, all dose levels exhibit excessive toxicity compared with the target rate, and there is no OD. In scenario 10, the lowest dose is the only one that does not exceed the target toxicity rate, but its efficacy is lower than the target; therefore, there is no OD in the explored dose range. In scenario 11, toxicity monotonically increases but the highest toxicity is lower than θ_T , and efficacy shows a bell shape. Scenario 12 has also monotonically increasing toxicity and the highest toxicity is lower than θ_T while the dose–efficacy relationship shows a plateau shape.

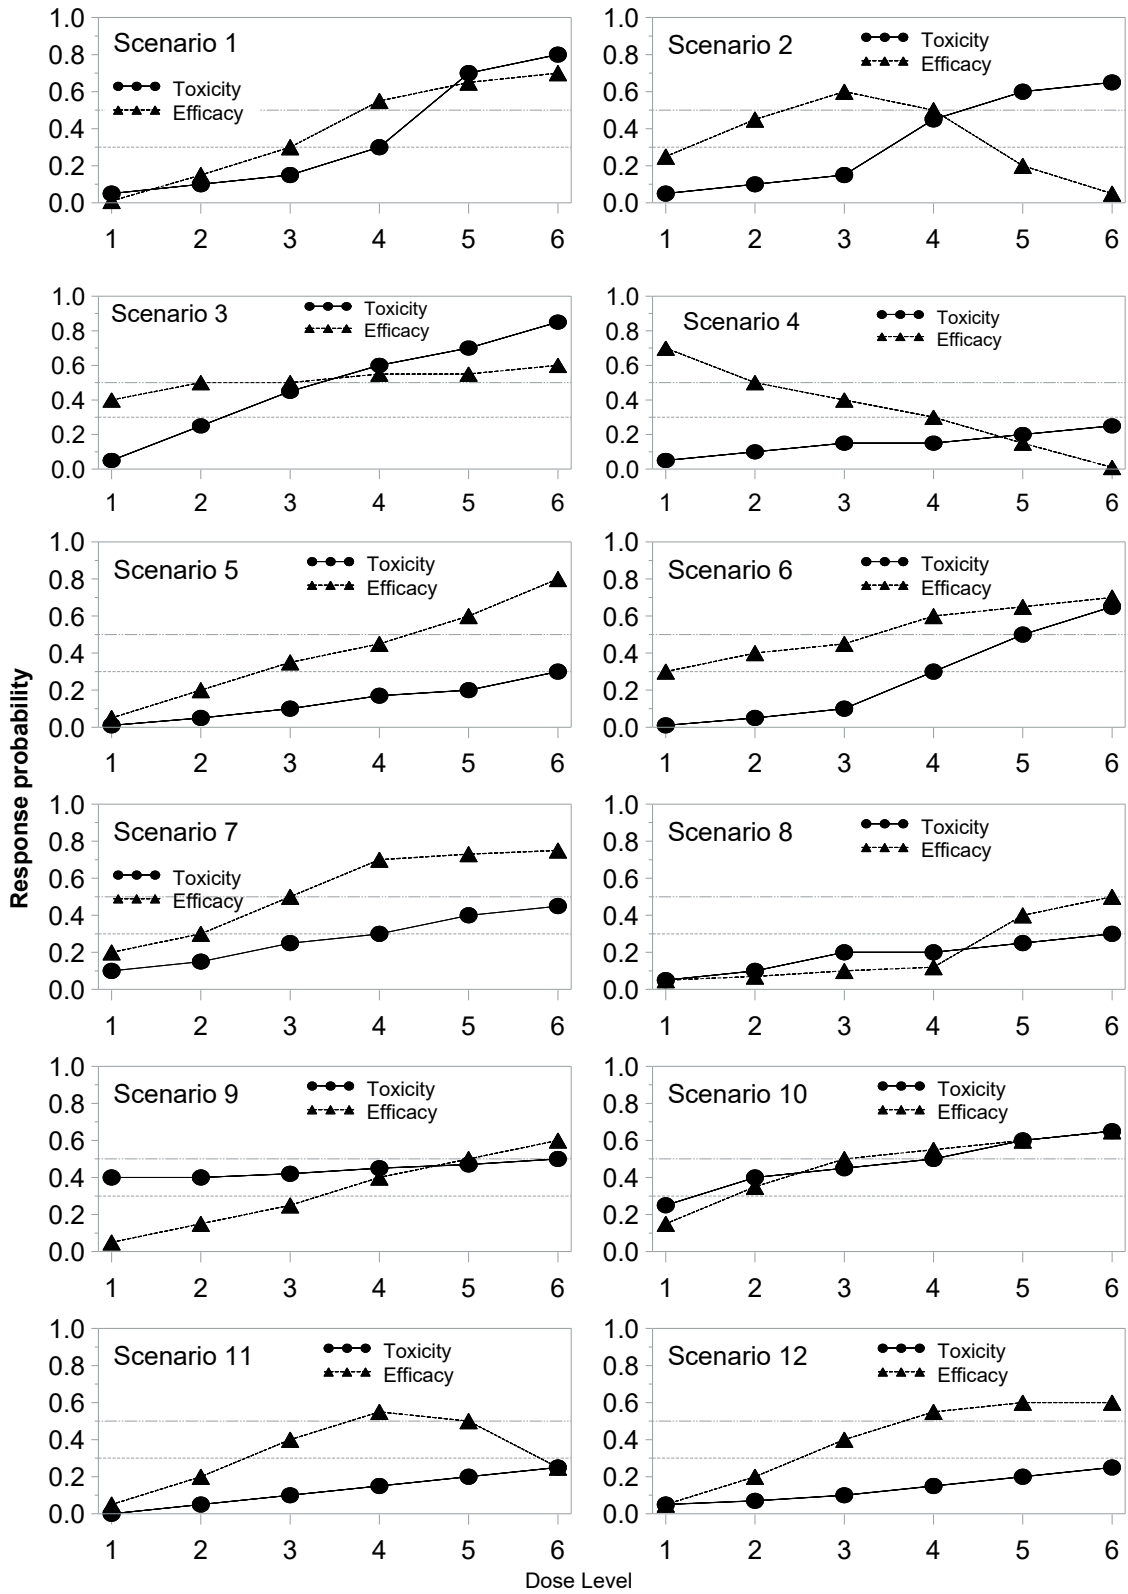


Fig. 5.4. Dose-response curves by scenario

Table 5.1. Scenarios on true toxicity and efficacy probability combinations

Scenario	True probability at each dose level (Toxicity, Efficacy)					
	1	2	3	4	5	6
1	(0.05,0.01)	(0.10,0.15)	(0.15,0.30)	(0.30,0.55)	(0.70,0.65)	(0.80,0.70)
2	(0.05,0.25)	(0.10,0.45)	(0.15,0.60)	(0.45,0.50)	(0.60,0.20)	(0.65,0.05)
3	(0.05,0.40)	(0.25,0.50)	(0.45,0.50)	(0.60,0.55)	(0.70,0.55)	(0.85,0.60)
4	(0.05,0.70)	(0.10,0.50)	(0.15,0.40)	(0.15,0.30)	(0.20,0.15)	(0.25,0.01)
5	(0.01,0.05)	(0.05,0.20)	(0.10,0.35)	(0.17,0.45)	(0.20,0.60)	(0.30,0.80)
6	(0.01,0.30)	(0.05,0.40)	(0.10,0.45)	(0.30,0.60)	(0.50,0.65)	(0.65,0.70)
7	(0.10,0.20)	(0.15,0.30)	(0.25,0.50)	(0.30,0.70)	(0.40,0.73)	(0.45,0.75)
8	(0.05,0.05)	(0.10,0.07)	(0.20,0.10)	(0.20,0.12)	(0.25,0.40)	(0.30,0.50)
9	(0.40,0.05)	(0.40,0.15)	(0.42,0.25)	(0.45,0.40)	(0.47,0.50)	(0.50,0.60)
10	(0.25,0.15)	(0.40,0.35)	(0.45,0.50)	(0.50,0.55)	(0.60,0.60)	(0.65,0.65)
11	(0.00,0.05)	(0.05,0.20)	(0.10,0.40)	(0.15,0.55)	(0.20,0.50)	(0.25,0.25)
12	(0.05,0.05)	(0.07,0.20)	(0.10,0.40)	(0.15,0.55)	(0.20,0.60)	(0.25,0.60)

In the simulation study, the maximum sample size was 36 with a cohort size of 3. For all designs, trials were started from the lowest dose level. The simulations were executed 1,000 times for each design.

Besides descriptive summaries such as OD selection probabilities, dose allocations, and observed response percentages, we present MSEs along with their 95% confidence intervals as a precision index to evaluate simulation outcomes. The MSE for OD estimation is given by

$$\text{MSE} = M^{-1} \sum_{m=1}^M \{1[\pi_{[m]}^{*T} > \theta_T](\pi_{[m]}^{*T} - \theta_T)^2 + 1[\pi_{[m]}^{*E} < \theta_E](\pi_{[m]}^{*E} - \theta_E)^2\}, \quad (5.21)$$

where $\pi_{[m]}^{*R}$ ($R \in \{T, E\}$) denotes the true response probabilities at recommended ODs in the m th trial under the scenarios. M is the total number of simulations ($M = 1000$).

5.3.2 Simulation settings of BOD for OD estimation

We set the following parameters for BOD. In the same manner of BODs for MTD and MTDC estimation in Sections 3.2 and 4.3, we generated prior mean functions for Gaussian process priors by the `getprior` function in R package `dfcrm` assuming that we do not have any informative prior. As another reason to use the `getprior` function, many biostatisticians are likely familiar with using this function to generate initial guesses for the CRM. The required parameters in the `getprior` function are δ , θ_R , initial OD location ν , and J . For ν , we recommend using the center of J unless there is a specific informative prior. It ensures not to concentrate unnecessarily on extreme dose levels in the dose selection process. Bayesian optimization often puts flat mean functions with 0 on its prior mean functions for objective functions; however, in dose-finding studies, it is reasonable to use prior mean functions with a slightly monotonically increasing trend. For dose–efficacy relationships, a relatively flat curve is appropriate to minimize the effect of the prior because it may follow a non-monotone pattern. In order to create such priors, we generated initial guesses of π^E using `getprior` ($\delta = 0.01$, $\theta_R = \theta_E$, $\nu = 3$, $J = 6$). The smaller the value of δ is, the flatter the slope of initial guesses becomes; therefore, the generated initial guesses of π^E were the flattest when δ was changed in 0.01 increments and had a slightly monotonically increasing trend. In contrast, it is unlikely that dose–toxicity curves follow a non-monotone pattern; therefore, in the same manner as BOD for MTD estimation in Section 3.2, we calibrated δ by the systematic approach (Lee and Cheung, 2009) based on the indifference interval. According to the systematic approach with 2,000

simulations, the optimal δ was 0.06 when $\theta_R = \theta_T$, $\nu = 3$, $J = 6$, and the maximum number of patients of 36. By utilizing this result, we generated the initial guesses on π^T using `getprior`($\delta = 0.06, \theta_R = \theta_T, \nu = 3, J = 6$). Although we used the optimal δ for dose–toxicity relationships, we recommend conducting sensitivity analyses using several values of δ (e.g., values around the range of an optimal $\delta \pm 0.02$) from the view of safe dose allocations and correct OD selections.

The conceptual doses used in the covariance functions were $(x_1, \dots, x_6) = (0.0, 0.2, 0.4, 0.6, 0.8, 1.0)$. Regarding kernel parameters in the covariance function in equation (3.3), σ_f was 1 following one of the common settings in Bayesian optimization frameworks because there is no specific reason to increase the maximum value of the covariance; a scale parameter was examined by two patterns, namely $\rho = 0.4$ and $\rho = 0.6$. For a scale parameter, the two patterns with turning points of every 0.4 or 0.6 provide 2 or 1 turning points in the dose range of $[0, 1]$. Since the dose–response curves are generally smooth and do not have many turning points, we consider that the appropriate number for turning points would be less than 3. Additionally, $\xi = 0.05$ was used for equation (3.4). While the value of ξ was simply determined by the appropriateness of the calculation speed, a smaller value of ξ is generally better in terms of less impact on the operating characteristics.

For the utility function described in equation (5.16), safety and efficacy were weighted equally during the dose selection process; that is, the ratio between w_T and w_E was 1 using the actual values of 1 as the simplest setting. Fig. 5.5 shows the example of utility scores calculated by equation (5.16) when substituted for $w_T = w_E = 1$ and the discrete probabilities of responses. As shown in Fig. 5.5, the utility scores become 0 when both toxicity and efficacy satisfy ideal conditions. Otherwise, the utility scores become greater than 0. The parameters for the admissible dose set in equation (5.17) were $\tilde{\epsilon}_1 = 0.1$. It would be generally acceptable in terms of safety perspective that doses with lower toxicity than 0.4 were included in the admissible dose set when $\theta_T = 0.3$.

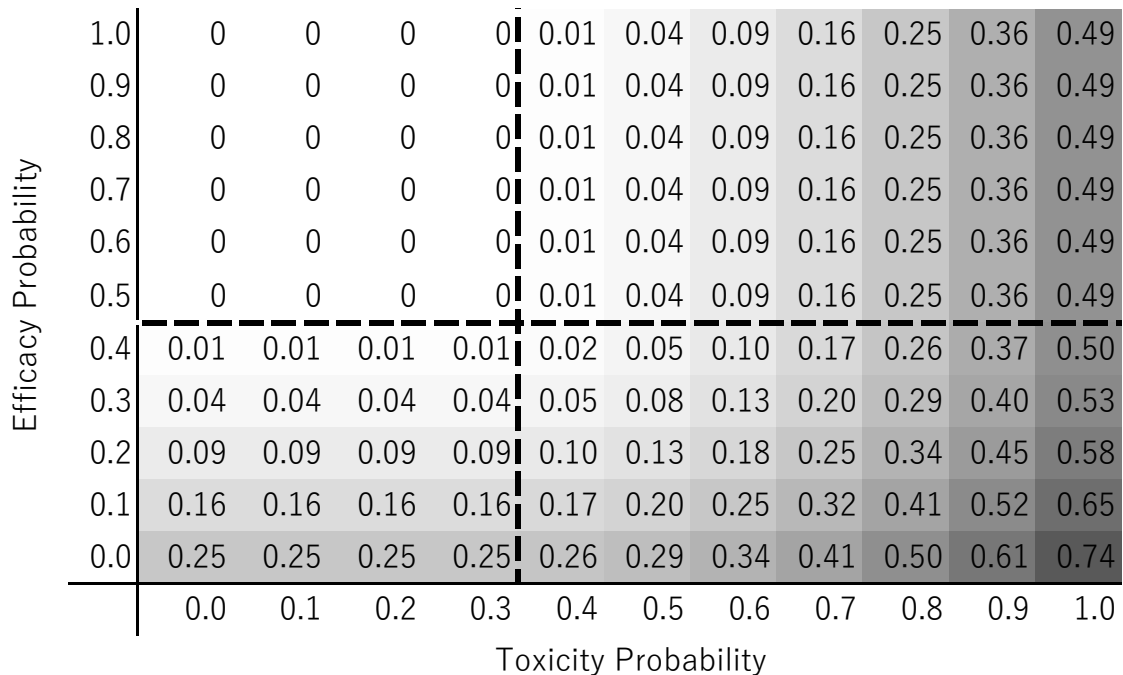


Fig. 5.5. Utility scores under the specified weight parameters

For the final OD selection based on equation (5.19), $\tilde{\epsilon}_2$ was set to $\tilde{\epsilon}_1 + 0.05$ that was slightly looser than $\tilde{\epsilon}_1$ so as to limit unnecessary NA decisions. If there was no candidate dose for the final OD selection, the final decision was NA in the trial. The threshold value of c_{u1} in equation (5.20) was 0.02, which was provided when both efficacy and toxicity simultaneously deviated to an undesirable direction by 0.1 as shown in Fig. 5.5. We investigated two patterns pertaining to the final determination of the OD using $c_{u2} = 1.0$ and 0.7. We note that $c_{u2} = 1.0$ indicates an OD is determined only by equation (5.19), and no additional conditions are provided for OD determination; that is, equation (5.20) is ignored.

Posterior sampling was done using Stan. In the simulation results, BOD is displayed with the scale parameter ρ for the squared exponential kernel and the value of c_{u2} : $\text{BOD}(\rho, c_{u2})$. We present simulation results of $\text{BOD}(0.4, c_{u2})$ in Section 5.3.4 and those of $\text{BOD}(0.6, c_{u2})$ are presented in Appendix C.1. Simulation codes for BOD are presented in Appendix C.2.

5.3.3 Simulation settings for the competitors

In order to ensure comparability among all designs, similar settings were established for each design as much as possible.

EffTox

The simulations for the EffTox were conducted by the R package `trialr`. For the parameter settings, we referred to desktop software that was available via <https://biostatistics.mdanderson.org/>. All hyperparameters were determined by the hyperparameter calculator in the desktop software using prior mean probabilities ranging from 0.2 to 0.45 for toxicity and from 0.4 to 0.65 for efficacy, and a target mean effective sample size of 1.0. The prior mean probabilities for both toxicity and efficacy exhibited target response rates in the center of the dose range (x_3).

The EffTox need to specify three points from the efficacy-toxicity plain to set utility contours described as equation (5.4); we used (0.1,0), (1,0.9), and (0.5,0.3) for $(\pi^{E_1}, 0)$, $(1, \pi^{T_2})$, and (π^{E_3}, π^{T_3}) , respectively, such that the contours pass the target toxicity rates $(\pi^{E_3}, \pi^{T_3}) = (\theta_E, \theta_T)$ and would become sufficiently steep as recommended in this design.

Admissible dose sets included all doses that satisfied the conditions $\text{P}\{\pi^E(x) \mid D_{1:t} > \theta_E\} > 0.1$ and $\text{P}\{\pi^T(x) \mid D_{1:t} < \theta_T\} > 0.1$. The cutoff value of 0.1 was obtained from Thall et al. (2014) and was the default setting in the desktop software. At the end of a trial, an OD that had the maximum utility was selected from the final admissible dose set. If no dose met the admissible dose set criteria during a trial, the trial terminated early and the final declaration was NA.

STEIN

We implemented the STEIN on the basis of the original article (Lin and Yin, 2017b). We used $\phi_1^E = 0.3$ and $\phi_2^E = 0.7$ being within the recommended ranges in the original article and made the optimal value of ϕ^E equal to θ_E . Additionally, the parameters for the toxicity probability interval were derived from the recommended calculations by using $\phi_0^T = \theta_T$, $\phi_1^T = 0.75\phi_0^T$ and $\phi_2^T = 1.25\phi_0^T$. The lower and upper critical threshold values of toxicity probability, which constructed an optimal interval, were (0.261, 0.337). Admissible dose sets for the next dose selection were determined by comparing a sample statistic y_j^T/n_j with the toxicity interval (0.261, 0.337) at each test.

In accordance with equation (5.11) with pre-specified weight parameters, the final OD was a dose maximizing the following STEIN's utility function:

$$U_{\text{STEIN}}(\tilde{\pi}_j^T, \tilde{\pi}_j^E) = \tilde{\pi}_j^E - 0.33\tilde{\pi}_j^T - 1.09\tilde{\pi}_j^T 1[\tilde{\pi}_j^T > \theta_T].$$

The weight parameter values of 0.33 and 1.09 were the same as the setting in the simulation study of the original article. On the basis of the practical implementation proposed in the original article, an additional overdose control was imposed by eliminating overly toxic doses that met $P(\pi_j^T > \theta_T \mid D_{1:t}) > 0.95$ from an admissible dose set, where the prior distribution was $\pi_j^T \sim \text{Unif}(0, 1)$. Trials were terminated if no admissible dose was identified.

Isotonic

The isotonic design was implemented using cloud software located at <https://biostatistics.mdanderson.org/shinyapps/MTADF/>. For safety monitoring, we used $\acute{c}_T = 0.8$ in equation (5.12); that is, unacceptable toxic dose levels that met $\tilde{P}(\pi_j^T \mid D_{1:t} > \theta_T) < 0.8$ were eliminated from a trial. The threshold value of 0.8 has been recommended to achieve overdose control in the cloud software. We note that the isotonic design does not require the value of θ_E because it selects dose levels with the highest efficacy probability within an admissible dose set.

5.3.4 Simulation results

• OD selection probabilities and MSE results

Fig. 5.6 shows the precision results generated by equation (5.21). MSEs shown in BOD tend to be consistently small compared with the other designs. In contrast, the other designs perform very well in one scenario but then provide OD estimations that deviate from the true OD in other scenarios. The impact of scale parameters in the kernel function was minimal between $\text{BOD}(0.4, \cdot)$ and $\text{BOD}(0.6, \cdot)$; however, appropriate values must be carefully determined through such as a grid search before a trial. For the criterion in equation (5.20), the amount of the impact on MSEs is small.

The correct OD selection probabilities shown in Table 5.2 closely align to the MSE results shown in Fig. 5.6. In most scenarios, BOD provides the highest correct OD selection probabilities than the other designs. In scenarios 9, the lowest dose level (with 40% toxicity) already exceeds the target toxicity rate, and toxicity increases 10% from the lowest to the highest dose level. The final determination of BOD is correctly NA with high probabilities in this scenario. Additionally, BOD allocates most patients to the lowest dose and shows the lowest observed toxicity percentage in all the designs as shown in Table 5.3. BOD narrows down the final OD candidates based on the posterior mean of the toxicity and efficacy probabilities to remain appropriate candidates and to exclude inappropriate candidates as well as the overdose control through admissible dose sets. The results at scenario 9 indicate that BOD can successfully exclude overly toxic doses from both the final OD declaration and admissible dose sets. Compared between $\text{BOD}(0.4, 1.0)$ and $\text{BOD}(0.4, 0.7)$, the latter selects NA more correctly in scenario 10 with no true OD. The criterion in equation (5.20) improves OD selection probabilities when the scenario does not have a true OD; however, OD selection without using the criterion works well in all scenarios regardless of the existence of a true OD compared with the other designs.

The STEIN tends to select one dose level lower than the true OD more than BOD in most scenarios. This would be because of its limited dose-escalation algorithm. The

STEIN selects dose escalation decision only when a sample statistic at the current dose level is lower than the lower boundary and the escalated dose has the highest posterior efficacy in an admissible dose set. The lower boundary in the simulation appears to be a strict limitation because it only allows toxicity probabilities lower than 26.1% when determining whether to include an escalated dose in the admissible dose set; therefore, dose escalation is limited more strictly when toxicity occurs.

In scenario 4, the EffTox shows a slightly higher correct OD selection probability than BOD, while the MSEs are comparable between the two designs. In scenarios 9 and 10 where there are no true OD, the EffTox provides a high probability of NA as well as a large stopping percentage. However, it allocates patients to a wider range of doses than the other designs (Table 5.3). As a result, the EffTox has higher toxicity percentage in those scenarios than the other designs. In the other scenarios, the EffTox does not perform well compared with the other designs in terms of the correct OD selection. In addition, probabilities of NA selection and early termination are frequently occurred due to the criteria for admissible dose sets with the cutoff value of 0.1 even if the scenario has a true OD, as shown in scenarios 1, 8, and 11.

The isotonic design applies doubly isotonic regressions to estimate dose–efficacy relationships; therefore, it was expected that the doubly isotonic approach provides more efficient estimation results if the true dose–efficacy curve does not exhibit a monotonically increasing shape. In our simulation study, however, the isotonic design does not show the highest correct OD selection probability even if the dose–efficacy curve exhibits a unimodal shape such as scenarios 2 and 11.

In scenario 12 where x_4 and x_5 are ODs, all the designs except for BOD select x_4 more than x_5 . In contrast, BOD tends to select x_5 whose efficacy is higher than that of x_4 while the selection probability of x_4 is higher than the EffTox and the isotonic designs. It would reflect the property of the utility function in BOD that x_4 more likely returns a positive utility score (i.e., penalty) than x_5 given the differences between the true response probabilities on these doses and the target response rates.




Fig. 5.6. Mean squared errors and the 95% confidence intervals for OD estimation by designs and scenarios

Table 5.2. OD selection probabilities at each dose level by designs and scenarios

Method	Dose level						NA	%Stop	Dose level						NA	%Stop
	1	2	3	4	5	6			1	2	3	4	5	6		
	Scenario 1															
BOD(0.4,1.0)	0.000	0.000	0.101	0.776	0.001	0.000	0.122	-	0.022	0.108	0.831	0.010	0.005	0.000	0.024	
BOD(0.4,0.7)	0.000	0.000	0.099	0.743	0.001	0.000	0.157	-	0.022	0.108	0.831	0.010	0.005	0.000	0.024	
EffTox	0.000	0.000	0.070	0.412	0.032	0.001	0.485	0.470	0.027	0.328	0.449	0.049	0.003	0.000	0.144	
STEIN	0.010	0.053	0.170	0.606	0.161	0.000	-	0.001	0.030	0.251	0.651	0.067	0.001	0.000	0.000	
Isotonic	0.019	0.165	0.330	0.445	0.041	0.000	-	-	0.128	0.226	0.545	0.098	0.003	0.000	-	
	Scenario 2															
BOD(0.4,1.0)	0.270	0.587	0.110	0.015	0.001	0.000	0.017	-	0.852	0.113	0.022	0.009	0.001	0.000	0.003	
BOD(0.4,0.7)	0.270	0.587	0.110	0.010	0.001	0.000	0.022	-	0.852	0.113	0.022	0.009	0.001	0.000	0.003	
EffTox	0.319	0.419	0.060	0.007	0.000	0.002	0.193	0.177	0.883	0.080	0.012	0.005	0.001	0.000	0.019	
STEIN	0.369	0.479	0.137	0.014	0.001	0.000	-	0.000	0.832	0.094	0.060	0.014	0.000	0.000	0.000	
Isotonic	0.368	0.490	0.122	0.020	0.000	0.000	-	-	0.749	0.171	0.062	0.017	0.001	0.000	-	
	Scenario 3															
BOD(0.4,1.0)	0.000	0.003	0.052	0.243	0.675	0.007	0.020	-	0.049	0.064	0.349	0.513	0.000	0.000	0.025	
BOD(0.4,0.7)	0.000	0.003	0.052	0.243	0.672	0.007	0.023	-	0.049	0.064	0.349	0.513	0.000	0.000	0.025	
EffTox	0.001	0.001	0.066	0.278	0.175	0.338	0.141	0.141	0.068	0.200	0.409	0.187	0.020	0.005	0.111	
STEIN	0.007	0.051	0.173	0.321	0.335	0.113	-	0.000	0.083	0.173	0.283	0.388	0.072	0.001	0.000	
Isotonic	0.081	0.159	0.202	0.139	0.221	0.198	-	-	0.226	0.251	0.229	0.224	0.066	0.004	-	
	Scenario 4															
BOD(0.4,1.0)	0.007	0.099	0.454	0.412	0.002	0.000	0.026	-	0.000	0.000	0.001	0.013	0.266	0.230	0.490	
BOD(0.4,0.7)	0.007	0.099	0.454	0.409	0.002	0.000	0.029	-	0.000	0.000	0.001	0.003	0.239	0.186	0.571	
EffTox	0.005	0.076	0.332	0.186	0.084	0.161	0.156	0.147	0.000	0.000	0.000	0.000	0.030	0.200	0.764	
STEIN	0.088	0.147	0.424	0.319	0.021	0.001	-	0.003	0.189	0.184	0.161	0.079	0.260	0.127	0.001	
Isotonic	0.237	0.158	0.167	0.273	0.149	0.016	-	-	0.123	0.173	0.178	0.082	0.268	0.176	-	
	Scenario 5															
BOD(0.4,1.0)	0.000	0.000	0.001	0.021	0.013	0.001	0.964	-	0.008	0.196	0.169	0.035	0.000	0.000	0.592	
BOD(0.4,0.7)	0.000	0.000	0.001	0.004	0.005	0.001	0.989	-	0.004	0.094	0.135	0.030	0.000	0.000	0.737	
EffTox	0.001	0.000	0.001	0.012	0.012	0.102	0.872	0.856	0.005	0.055	0.098	0.016	0.006	0.005	0.815	
STEIN	0.794	0.148	0.049	0.007	0.002	0.000	-	0.522	0.487	0.357	0.146	0.010	0.000	0.000	-	
Isotonic	0.512	0.208	0.137	0.098	0.044	0.001	-	-	0.505	0.279	0.156	0.053	0.007	0.000	-	
	Scenario 6															
BOD(0.4,1.0)	0.000	0.001	0.075	0.603	0.269	0.007	0.045	-	0.000	0.002	0.062	0.367	0.521	0.023	0.025	
BOD(0.4,0.7)	0.000	0.001	0.075	0.603	0.269	0.006	0.046	-	0.000	0.002	0.062	0.367	0.519	0.023	0.027	
EffTox	0.000	0.002	0.116	0.358	0.141	0.116	0.267	0.244	0.000	0.000	0.086	0.239	0.164	0.429	0.082	
STEIN	0.002	0.043	0.223	0.513	0.207	0.01.2	-	0.000	0.001	0.037	0.174	0.475	0.262	0.051	0.000	
Isotonic	0.079	0.124	0.194	0.393	0.199	0.01.1	-	-	0.067	0.131	0.177	0.268	0.237	0.120	-	

- Dose allocation

Regarding the average percentages of patients with toxicity responses (%tox) in Table 5.3, BOD shows less than or around 20% in most scenarios. BOD shows slightly higher %tox than the other designs depending on scenarios, but the actual values of %tox are less than 20% for all scenarios where BOD shows the highest %tox. Under scenarios 9 and 10 where most doses have excessive toxicity, BOD shows the lowest %tox. These results indicate that overdose control in BOD works successfully.

In contrast, BOD provides the highest efficacy response percentages (%eff) than the other designs in most scenarios while the EffTox shows the highest %eff in scenario 4. It indicates the utility function works to seek and allocates efficacious doses during a trial. The EffTox seems to fit scenario 4 because the true OD provides a much higher score in all doses based on the utility function of the EffTox. In scenarios 9 and 10, BOD does not show the highest %eff; however, it is due to overdose control. In scenario 11, the STEIN allocates patients to the true OD the most in all the designs while it does not show the highest %eff due to higher allocation percentages to lower dose levels than the OD. In scenario 12, BOD allocates patients to both of two ODs at comparable percentages. In contrast, the other designs allocate more patients to x_4 than x_5 . These allocation percentages at x_5 are around half amount of BOD. BOD considers a trade-off of exploitation and exploration; therefore, this result would also reflect such an exploration effect.

Overall, BOD with the current settings selects efficacious dose with controlling overdose allocations, and the employed dose selection criteria (e.g., admissible dose set) work successfully.

Table 5.3. Average number of patients treated with each dose level and average percentages of observed responses

Method	Dose level						Total	%tox	%eff
	1	2	3	4	5	6			
Scenario 1									
BOD(0.4)	4.1	3.9	10.3	15.6	2.1	0.0	36.0	23.1	37.7
EffTox	3.1	3.1	6.6	11.6	3.0	0.8	28.2	27.7	36.6
STEIN	5.0	6.6	10.2	12.9	1.3	0.0	36.0	19.9	33.4
Isotonic	3.7	7.4	9.4	10.7	3.4	1.4	36.0	25.0	36.4
Scenario 2									
BOD(0.4)	11.6	15.8	7.5	1.1	0.0	0.0	36.0	23.7	47.0
EffTox	16.2	12.4	3.6	0.5	0.2	0.1	32.9	19.8	44.0
STEIN	14.0	17.7	4.0	0.3	0.0	0.0	36.0	19.8	46.7
Isotonic	13.1	13.9	6.1	2.1	0.7	0.1	36.0	24.2	46.1
Scenario 3									
BOD(0.4)	3.4	3.4	5.4	9.0	9.8	5.0	36.0	16.0	46.2
EffTox	3.4	3.4	6.3	8.5	4.6	6.6	32.9	14.8	42.6
STEIN	4.3	6.0	8.7	9.9	5.8	1.2	36.0	12.6	36.5
Isotonic	5.5	7.5	7.9	5.4	5.0	4.8	36.0	12.8	38.3
Scenario 4									
BOD(0.4)	4.9	7.0	11.4	9.0	3.7	0.0	36.0	23.9	49.3
EffTox	5.0	6.4	10.1	5.4	2.6	3.9	33.4	25.1	47.8
STEIN	7.6	10.2	12.7	5.2	0.2	0.0	35.9	20.1	41.1
Isotonic	10.3	7.5	6.1	6.5	4.0	1.5	36.0	21.9	44.4
Scenario 5									
BOD(0.4)	24.4	6.0	3.7	1.5	0.3	0.0	36.0	40.4	10.3
EffTox	3.3	2.9	3.0	2.5	1.7	4.5	17.8	49.2	24.8
STEIN	17.7	4.6	1.2	0.2	0.0	0.0	23.6	49.4	7.4
Isotonic	15.7	9.2	5.5	3.6	1.7	0.3	36.0	41.4	16.9
Scenario 6									
BOD(0.4)	3.5	3.4	5.6	11.9	9.8	1.8	36.0	13.7	41.3
EffTox	3.4	3.5	6.5	9.5	5.1	4.4	32.5	12.8	35.5
STEIN	3.8	5.4	9.5	12.3	4.3	0.7	36.0	11.5	38.9
Isotonic	5.4	6.6	7.4	9.5	5.6	1.5	36.0	10.8	35.8
Scenario 7									
BOD(0.4)	4.9	7.0	11.4	9.0	3.7	0.0	36.0	23.9	49.3
EffTox	5.0	6.4	10.1	5.4	2.6	3.9	33.4	25.1	47.8
STEIN	7.6	10.2	12.7	5.2	0.2	0.0	35.9	20.1	41.1
Isotonic	10.3	7.5	6.1	6.5	4.0	1.5	36.0	21.9	44.4
Scenario 8									
BOD(0.4)	4.9	7.0	11.4	9.0	3.7	0.0	36.0	23.9	49.3
EffTox	5.0	6.4	10.1	5.4	2.6	3.9	33.4	25.1	47.8
STEIN	7.6	10.2	12.7	5.2	0.2	0.0	35.9	20.1	41.1
Isotonic	10.3	7.5	6.1	6.5	4.0	1.5	36.0	21.9	44.4
Scenario 9									
BOD(0.4)	24.4	6.0	3.7	1.5	0.3	0.0	36.0	40.4	10.3
EffTox	3.3	2.9	3.0	2.5	1.7	4.5	17.8	49.2	24.8
STEIN	17.7	4.6	1.2	0.2	0.0	0.0	23.6	49.4	7.4
Isotonic	15.7	9.2	5.5	3.6	1.7	0.3	36.0	41.4	16.9
Scenario 10									
BOD(0.4)	24.4	6.0	3.7	1.5	0.3	0.0	36.0	40.4	10.3
EffTox	3.3	2.9	3.0	2.5	1.7	4.5	17.8	49.2	24.8
STEIN	17.7	4.6	1.2	0.2	0.0	0.0	23.6	49.4	7.4
Isotonic	15.7	9.2	5.5	3.6	1.7	0.3	36.0	41.4	16.9
Scenario 11									
BOD(0.4)	3.5	3.4	5.6	11.9	9.8	1.8	36.0	13.7	41.3
EffTox	3.4	3.5	6.5	9.5	5.1	4.4	32.5	12.8	35.5
STEIN	3.8	5.4	9.5	12.3	4.3	0.7	36.0	11.5	38.9
Isotonic	5.4	6.6	7.4	9.5	5.6	1.5	36.0	10.8	35.8
Scenario 12									
BOD(0.4)	3.5	3.4	5.6	11.9	9.8	1.8	36.0	13.7	41.3
EffTox	3.4	3.5	6.5	9.5	5.1	4.4	32.5	12.8	35.5
STEIN	3.8	5.4	9.5	12.3	4.3	0.7	36.0	11.5	38.9
Isotonic	5.4	6.6	7.4	9.5	5.6	1.5	36.0	10.8	35.8

- Operating characteristics of the existing designs

In the STEIN algorithm, dose escalations occur only when 1) a sample statistic is less than or equal to the lower boundary (0.261) and 2) one dose level higher than the current dose exhibits the highest posterior efficacy probability in the admissible dose set. For example, if one out of three patients experienced toxicity, the dose-escalation decision is not selected. The results in the STEIN in Tables 5.2 and 5.3 reflect this strict limitation for dose-escalation decisions. An OD is estimated by isotonic estimates and unimodal isotonic estimates among only tested doses. If a true OD is located at the highest dose level, such as scenario 8, the STEIN shows a lower correct OD selection probability than the other scenarios owing to the limited dose-escalation decisions. In addition, the STEIN tends to select lower doses than the true OD with relatively high frequencies compared with BOD in most scenarios.

The EffTox provides ideal determinations with high probabilities of NAs and early terminations under scenarios 9 and 10. NA determinations are provided when there is no dose that meets conditions $P\{\pi^E(x) \mid D_{1:t} > \theta_E\} > 0.1$ and $P\{\pi^T(x) \mid D_{1:t} < \theta_T\} > 0.1$. The EffTox also provides more NA determinations than the other designs under the other scenarios with true ODs (Table 5.2). This indicates that the cutoff value of 0.1 should be refined for each scenario although it may be difficult to determine an ideal cutoff value when little information is known about dose–response relationships. Brock et al. (2017) has discussed the cutoff value settings in the EffTox. A parameter value for determining the extent of the curvature of the utility contours in the EffTox (i.e., ι in equation (5.4)) was approximately 0.85 based on the three parameters in Section 5.3.3. Compared between the simulation results and the true utility scores based on 0.85, the EffTox shows good results under the scenarios where the utility score at ODs is the maximum value or markedly different from that of the other doses. Otherwise, the performance of the EffTox tends to be low. For example, the true utility scores of the EffTox in scenario 6 are (0.20, 0.24, 0.22, 0.12, -0.07, and -0.18), where a true OD of x_4 returns a lower score than x_1, x_2 , and x_3 . As a result, the performance in scenario 6 is lower than that of the other designs. There might be more suitable values for the three parameters in scenario 6.

For isotonic design, the threshold value for overdose control would have a considerable impact on the operating characteristics. We used a threshold value of 0.8 in the simulation as recommended in the cloud software. Fig. 5.7 shows additional simulation results to examine the impact of the threshold values under scenario 1. If the threshold value is 0.7, OD recommendations in all trials are only x_1 . This indicates that dose elimination based on a threshold value of 0.7 is too conservative. On the other hand, if no dose elimination is applied with a threshold value of 1, x_5 and x_6 that have higher efficacy probabilities are selected more often than the other doses. The results with values of 0.8 and 0.9 show higher selection probabilities for x_4 although they are lower than the correct OD selection probability shown in BOD. A threshold value of 0.8 is recommended for general use; however, a slightly larger value than 0.8 would be suitable as a threshold value in scenario 1.

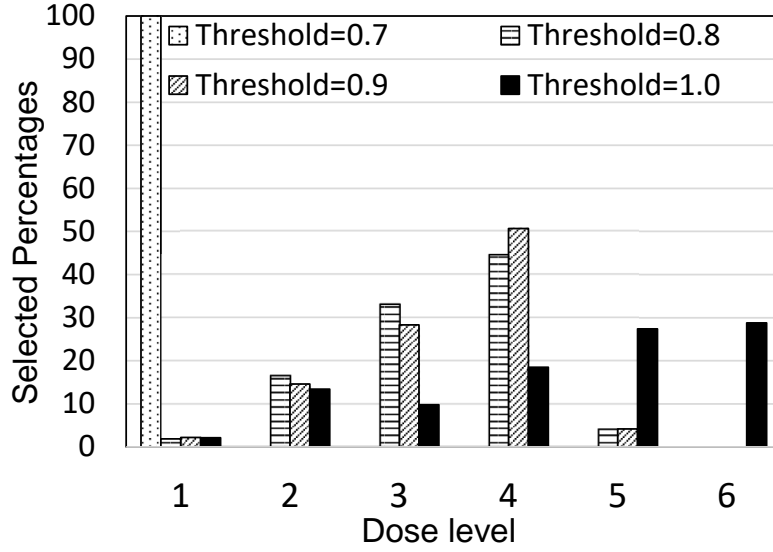


Fig. 5.7. Optimal dose selections in the isotonic design with different threshold values

5.4 Summary and discussion

Chapter 5 introduced BOD for OD estimation. In our simulation study with twelve scenarios, BOD consistently produced smaller MSEs and more accurate OD selections in most scenarios compared with the other designs.

For acquisition functions, $LCB(x)$ is very simple without any complex calculations. We selected 10 percentile for $LCB(x)$, however, what percentile is the most suitable as $LCB(x)$ and how to calibrate the balance are still an open discussion. Considering the difficulty of parameter calibrations, it might be worth evaluating the operating characteristics of the EI strategy because it does not need any parameter tuning as shown in Section 3.1.2.

We assumed that toxicity and efficacy outcomes were independent of each other in the model. On the other hand, some investigators might be interested in incorporating the impact of the association between them. One option could be including an association parameter η with a vague prior distribution (e.g., a beta distribution with parameters (2,2)). The likelihood function takes the association parameter into account such as

$$L(D_{1:t} | f^T, f^E, \eta) \propto \prod_{z=1}^t \pi^T(x_{(z)})^{y_{(z)}^T} \{1 - \pi^T(x_{(z)})\}^{(n_{(z)} - y_{(z)}^T)} \\ \pi^E(x_{(z)})^{y_{(z)}^E} \{1 - \pi^E(x_{(z)})\}^{(n_{(z)} - y_{(z)}^E)} \eta^{y_{(z)}^T y_{(z)}^E} (1 - \eta)^{1 - y_{(z)}^T y_{(z)}^E},$$

instead of equation (5.15) according to Asakawa and Hamada (2013). We note that it might be challenging to estimate such the association parameter because the sample size is small in general. Another option could be addressing the association through a utility function by adding an association term to equation (5.16). For example, an additional term related to the association between responses might be expressed as

$$\omega_{ET} 1[(\pi^T > \theta_T) \wedge (\pi^E < \theta_E)] (\pi^T - \theta_T) (\theta_E - \pi^E).$$

The association term provides an additional penalty that corresponds to the strength of the association if both response probabilities on toxicity and efficacy do not satisfy the

target response rates. Because we have not evaluated the feasibility of these extensions related to the association in this thesis, further evaluations will be needed as future work.

We focused on binary variables for toxicity and efficacy outcomes. On the other hand, BOD could be extended to situations where toxicity is binary and efficacy is continuous. Mandrekar et al. (2010) reported a case where the dichotomous efficacy outcomes were inaccurate and suboptimal. Several authors have applied their proposed designs for continuous outcomes; Bekel and Shen (2005) has treated continuous variables such as biomarker evaluations in place of binary efficacy outcomes in OD estimations; Tao et al. (2013) has applied Archimedean copula regression to bivariate continuous outcomes. If we treat a continuous efficacy variable instead of a binary efficacy variable, at least the following will be changed:

- Observed actual values of a continuous efficacy variable Y^{Ec} are modeled with an unknown function f^{Ec} ;
- A prior distribution for f^{Ec} is determined based on initial information or knowledge about the efficacy endpoint;
- Posterior distributions of f^{Ec} can be analytically calculated if Y^{Ec} is assumed to follow a normal distribution (i.e., $Y_j^{Ec} \sim N(f^{Ec}(x_j), \sigma^2)$), otherwise posterior samples for f^{Ec} are computed by MCMC;
- A utility function could be expressed by, for example,

$$U_{\text{BOD}_c}(x) = w_T 1[\pi^T(x) > \theta_T] \left\{ \frac{\pi^T(x) - \theta_T}{\theta_T} \right\}^2 + w_{Ec} 1[f^{Ec}(x) < y^{E*}] \left\{ \frac{y^{E*} - f^{Ec}(x)}{y^{E*}} \right\}^2,$$

where y^{E*} is a target value for the efficacy endpoint.

We have left such extensions regarding the continuous efficacy variable in future work.

As a further extended approach, it might be possible to incorporate both parametric and nonparametric approaches in BOD as hybrid designs. Although hybrid designs leave model misspecification issues, it is reasonable to assume that dose–toxicity relationships follow monotonically increasing curves; therefore, parametric model-based designs might provide effective dose findings especially for dose–toxicity relationships. Some authors have introduced hybrid designs for Phase I or Phase II clinical trials. For example, a hybrid design introduced by Yuan and Yin (2011a) relies on nonparametric approaches in the dose selection when toxicity outcomes at the current dose show strong enough information in terms of decision-making, otherwise relies on parametric model-based designs such as the CRM. Additionally, Yuan and Yin (2011b) has introduced another hybrid design that makes a weighted average estimate to contribute both parametric and nonparametric estimates to the estimation of a single dose–response curve. While these hybrid designs apply parametric and nonparametric approaches to a single dose–response curve, hybrid designs we assume here incorporate two dose–response curves estimated separately by parametric and nonparametric approaches. As such a design, Yeung et al. (2017) models relationships between doses and binary toxicity responses with a logistic regression model and mean efficacy responses with a random walk model. A gain function that explicitly trades-off efficacy and harm by incorporating the estimated values based on the two models then guides its dose-escalation procedure. As our initial thought, one of the simplest approaches for toxicity estimation as a hybrid design for BOD might be to utilize the CRM to estimate a dose providing the closest toxicity to the target toxicity rate because the CRM is one of the most popular parametric model-based designs and is well known to provide good operating characteristics in dose-finding studies. Simultaneously, we could apply the approach that BOD for OD estimation offers to estimate dose–efficacy relationships. In this case, one option of dose selection is selecting the most efficacious

dose among an admissible dose set based on the CRM through an acquisition function for dose–efficacy relationships. Another option is constructing a new utility function to incorporate toxicity and efficacy information. The simplest approach to do this might be keeping the second term in equation (5.16) as it is and including a penalty term based on toxicity probability estimated by the CRM instead of the first term.

For the utility function, the ratio between the weight parameters for toxicity and efficacy was set to 1 in our simulation study. On the other hand, various situations could be assumed with regard to a balance between safety and efficacy in practice. The utility function consists of squared differences between estimated posterior probabilities and target response rates; however, using absolute values instead of squared values might be an alternative. Utilizing a utility function similar to that of the STEIN might be another alternative. Additionally, a dose selection process based on separate evaluations for each dose–response relationship without utility functions may be a different alternative approach. Because there are a variety of candidates for dose selection processes including utility functions, they should be further evaluated.

BOD requires parameters for a Gaussian process prior and the utility function. A flat mean prior function for dose–efficacy relationship can be generated by the `getprior` function with a relatively small value of δ . Initial toxicity guesses are generated with δ derived by the systematic approach proposed for the CRM. A covariance function based on the squared exponential kernel has three parameters; σ_f, ρ, ξ . As the same setting as BODs in Chapters 3 and 4, σ_f is 1 and ξ is determined according to the desired computation speed (a smaller value is recommended). For the parameter ρ , a value providing 1 or 2 turning points would usually provide good operating characteristics. The weight parameters (w_t, w_e) for the utility function should be determined based on their balances (i.e., ratios) while the ratio of 1 worked successfully in our simulation study. If the weight for toxicity must be increased owing to safety concerns, w_t should be greater than w_e . The optional criterion for OD determinations in equation (5.20) would be better to use just as a reference because it does not always improve OD determinations. If prior knowledge based on historical data or real-world evidence on investigational agents is available, Gaussian process priors (i.e., prior mean functions and scale parameters for each covariance function) for toxicity and efficacy could be derived from such informative data. Additionally, BOD assumes that the discrete test doses are equally spaced concerning the probabilities of responses in the explored dose range; therefore, it is desirable that the investigated doses are determined so that they are equally spaced with respect to initial response probabilities using those data.

There is still no standard design for OD estimation even when most analyzed situations have binary outcomes for both toxicity and efficacy. The STEIN can be implemented easily and provided satisfactory performance in our simulation although there is concern about the possibility of significant deviation in the final OD determination depending on scenarios. The EffTox shows high rates of NA determinations and early terminations regardless of the existence of true ODs in scenarios because of the overdose control criteria. The isotonic design did not show outstanding results in our simulation study. We need further evaluations for BOD including possible extensions along with an evaluation on robustness; however, given the above situation and our simulation results, we conclude that BOD has a high potential to provide outstanding performance in terms of correct OD selections while providing proper overdose control.

C

Appendices for Chapter 5

C.1 Supplemental simulation results

Section C.1 provides simulation results of BOD (0.6, c_{u2}) that contributed to the performance evaluation but was omitted in Section 5.3.4.

Table C.1. Mean squared errors in BOD (0.6)

Scenario	BOD(0.6, 1.0)			BOD(0.6, 0.7)		
	MSE	95% CI*		MSE	95% CI*	
		Lower	Upper		Lower	Upper
1	0.0049	0.0040	0.0059	0.0049	0.0040	0.0058
2	0.0033	0.0023	0.0043	0.0031	0.0022	0.0040
3	0.0064	0.0056	0.0072	0.0061	0.0054	0.0069
4	0.0006	0.0003	0.0009	0.0006	0.0003	0.0009
5	0.0014	0.0011	0.0018	0.0014	0.0011	0.0018
6	0.0027	0.0023	0.0032	0.0027	0.0023	0.0032
7	0.0045	0.0036	0.0053	0.0045	0.0036	0.0053
8	0.0048	0.0040	0.0056	0.0043	0.0039	0.0048
9	0.0363	0.0271	0.0454	0.0289	-	-
10	0.0317	0.0298	0.0336	0.0316	0.0290	0.0342
11	0.0030	0.0022	0.0039	0.0030	0.0021	0.0038
12	0.0006	0.0003	0.0009	0.0006	0.0003	0.0009

* CI: Confidence interval

Table C.2. OD selection probabilities at each dose level in BOD (0.6) (%)

Method	Dose level						%Stop	NA	%Stop	Dose level						NA	%Stop			
	1	2	3	4	5	6				1	2	3	4	5	6					
	Scenario 1												Scenario 2							
BOD(0.6,1.0)	0.000	0.001	0.104	0.793	0.001	0.000	0.101	-	0.028	0.104	0.802	0.019	0.004	0.000	0.043	-				
BOD(0.6,0.7)	0.000	0.001	0.104	0.769	0.000	0.000	0.126	-	0.028	0.104	0.802	0.018	0.003	0.000	0.045	-				
	Scenario 3												Scenario 4							
BOD(0.6,1.0)	0.237	0.623	0.110	0.016	0.000	0.000	0.014	-	0.853	0.114	0.026	0.005	0.001	0.000	0.001	-				
BOD(0.6,0.7)	0.237	0.623	0.110	0.013	0.000	0.000	0.017	-	0.853	0.114	0.026	0.005	0.001	0.000	0.001	-				
	Scenario 5												Scenario 6							
BOD(0.6,1.0)	0.000	0.002	0.030	0.217	0.724	0.016	0.011	-	0.028	0.055	0.403	0.491	0.000	0.000	0.023	-				
BOD(0.6,0.7)	0.000	0.002	0.030	0.217	0.723	0.016	0.012	-	0.028	0.055	0.403	0.490	0.000	0.000	0.024	-				
	Scenario 7												Scenario 8							
BOD(0.6,1.0)	0.007	0.095	0.500	0.379	0.006	0.000	0.013	-	0.000	0.000	0.000	0.002	0.236	0.315	0.447	-				
BOD(0.6,0.7)	0.007	0.095	0.500	0.375	0.006	0.000	0.017	-	0.000	0.000	0.000	0.000	0.211	0.278	0.511	-				
	Scenario 9												Scenario 10							
BOD(0.6,1.0)	0.000	0.000	0.002	0.004	0.009	0.000	0.985	-	0.010	0.158	0.131	0.021	0.000	0.000	0.680	-				
BOD(0.6,0.7)	0.000	0.000	0.000	0.000	0.001	0.000	0.999	-	0.007	0.081	0.093	0.017	0.000	0.000	0.802	-				
	Scenario 11												Scenario 12							
BOD(0.6,1.0)	0.000	0.005	0.052	0.451	0.360	0.028	0.104	-	0.000	0.002	0.042	0.240	0.599	0.090	0.027	-				
BOD(0.6,0.7)	0.000	0.005	0.052	0.451	0.359	0.027	0.106	-	0.000	0.002	0.042	0.240	0.595	0.090	0.031	-				

Table C.3. Average number of patients treated with each dose level and average percentages of observed responses in BOD (0.6)

Scenario	Dose level						Total	%tox	%eff
	1	2	3	4	5	6			
1	3.8	4.0	9.3	16.3	2.6	0.0	36.0	24.4	39.3
2	5.4	6.0	16.2	7.6	0.9	0.1	36.0	20.2	48.8
3	10.2	16.3	8.1	1.4	0.1	0.0	36.0	25.2	47.3
4	24.9	5.8	3.2	1.6	0.5	0.1	36.0	7.4	61.3
5	3.4	3.3	4.5	7.9	9.1	7.8	36.0	17.2	49.2
6	5.4	4.6	7.7	13.3	4.7	0.2	36.0	21.1	50.6
7	4.9	6.7	10.4	9.3	3.8	0.9	36.0	24.5	50.4
8	5.2	4.1	6.4	8.2	7.5	4.7	36.0	19.1	21.1
9	25.2	6.5	3.0	1.1	0.2	0.0	36.0	40.3	9.7
10	17.6	12.5	4.6	1.1	0.1	0.0	36.0	33.7	27.8
11	3.6	3.5	4.7	9.5	10.0	4.7	36.0	14.7	39.6
12	3.6	3.6	4.7	7.5	8.7	7.9	36.0	16.1	46.9

C.2 BOD simulation code for OD estimation

We implemented BOD for OD estimation using R for the simulation study described in Section 5.3. Appendix C.2 provides R codes for implementing BOD to identify an OD for a targeted therapy.

- Listing C.1 is R code that constructs the simulation body.
- Listing C.2 presents a `stan` model that is compiled from the R package `rstan` in Listing C.1 to implement the MCMC for obtaining posterior samples.
- Listing C.3 shows a part of toxicity and efficacy scenarios used in the simulation study.

Listing C.1. Simulation body of BOD for OD estimation

```

1 #BOD_od_simulation_body.R
2 #####
3 ### Initial settings before running simulations ###
4 #####
5
6 ### Set a working directly ###
7 setwd("XXX") # set a location of directly for the simulation
8
9 #### packages ###
10 library(rstan) # for MCMC
11 rstan_options(auto_write = TRUE) # rstan option
12 options(mc.cores = parallel::detectCores()) # rstan option
13 #Sys.setenv(LOCAL_CPPFLAGS = '-march=native') # rstan option
14 library(dfcrm) #for using getprior function
15 #library(ggmc) # for checking convergence of MCMC graphically
16
17 ### Create functions ###
18 # logit function
19 logit_f <- function(a){
20   b <- log(a/(1-a))
21   return(b)
22 }
23
24 # modified rbind() that incorporates dataframes with different row names
25 rbindCOrder <- function(...) {
26   n <- length(list(...))
27   temp <- list(...)[[1]]
28   names(temp) <- NA
29   for (i in 2:n) {
30     tmp <- list(...)[[i]]
31     names(tmp) <- NA
32     temp <- rbind(temp, tmp)
33   }
34   names(temp) <- names(list(...)[[1]])
35   return(temp)
36 }
37
38 ### Compile a Stan file ###
39 stanmodel <- stan_model(file='BOD_model_od.stan') # BOD_model_od.stan.stan
40   is located at the working directly.
41
42 ### Trial Settings ###
43 N <- 6 # number of dose levels
44 Pt <- 0.3 # maximum acceptable toxicity rate
45 Pe <- 0.5 # target response rate

```

```

45 max_sub <- 36 # maximum number of patients in a trial
46 cohort <- 3 # cohort size
47 X <- c(seq(0, 1, length = N)) #conceptual doses
48
49 ### Scenarios for the simulation study ###
50 sc <- as.data.frame((read.table("Scenario_OD.txt", header = T))) #
    Scenario.txt is located at the working directly.
51 S_N <- length(sc[,-1])/2 # number of scenarios
52
53
54 ### For Gaussian process priors ###
55 # Prior mean functions for toxicity and efficacy
56 Guess_T <- logit_f(getprior(0.06, Pt, 3, N)) # for toxicity
57 Guess_E <- logit_f(getprior(0.01, Pe, 3, N)) # for efficacy
58
59 # Calculate covariance functions for toxicity and efficacy
60 # The simulation study applies the same covariance functions for both
    toxicity and efficacy; however, a Stan file can treat different
    covariance functions by using L and L2, as needed.
61
62 # calculate a fixed covariance function K for toxicity
63 sigma_f <- 1 # signal variance
64 rho <- 0.4 # scale parameter
65 xi <- 0.05 # a small value added to diagonal elements in K
66
67 K <- matrix(rep(0,N*N), nrow = N, ncol = N)
68 for (i in 1:(N - 1)) {
69   K[i, i] <- (sigma_f)^2
70   for (j in (i + 1):N) {
71     K[i, j] = (sigma_f)^2*exp(-0.5/(rho)^2 * ((X[i] - X[j])^2))
72     K[j, i] = K[i, j]
73   }
74 }
75 K[N, N] <- (sigma_f)^2
76 K2 <- K + diag(xi, N) # add small values for diagonal elements in K
77 L <- t(chol(K2)) # Cholesky decomposition for K2
78
79 # calculate a fixed covariance function K for efficacy
80 sigma_f <- 1 # signal variance
81 rho2 <- 0.4 # scale parameter
82
83 K <- matrix(rep(0,N*N), nrow = N, ncol = N)
84 for (i in 1:(N - 1)) {
85   K[i, i] <- (sigma_f)^2
86   for (j in (i + 1):N) {
87     K[i, j] = (sigma_f)^2*exp(-0.5/(rho2)^2 * ((X[i] - X[j])^2))
88     K[j, i] = K[i, j]
89   }
90 }
91 K[N, N] <- (sigma_f)^2
92 K2 <- K + diag(xi, N) # add small values for diagonal elements in K
93 L2 <- t(chol(K2)) # Cholesky decomposition
94
95
96 ### parameters for the utility function ###
97 e_p1 <- 1 # weight parameter for efficacy in equation (5.16)
98 s_p1 <- 1 # weight parameter for toxicity in equation (5.16)
99 # For the weight parameters, a Stan file with the current setting below (
    Listing C.2) can load only integer values.
100
101 ### a parameter for an admissible dose set ###

```

```

102 vareps1 <- 0.1 # a parameter in equation (5.17)
103
104 ### parameters for OD determination ###
105 vareps2 <- 0.15 # a parameter in equation (5.19)
106
107 # Here are optional parameters in equation (5.20)
108 c_u1 <- e_p1*(0.1)^2 + s_p1*(0.1)^2 # an example of c_u1
109 c_u2 <- 0.5 # an example of c_u2
110
111 ### Iterations and random seeds ###
112 sim_min <- 1 # minimum number in iterations
113 sim_max <- 1000 # maximum number of iterations
114 rand0 <- matrix(1:(100*1000*S_N), nrow = (100*S_N), ncol = 1000) #
    random seeds matrix
115
116 #####
117 # Start a simulation study #
118 #####
119 for(S in 1:2){ # loop selected scenarios
120   Sn <- 2+2*(S-1) # identify a location of the scenario in the loaded
    data
121   True_e <- sc[,Sn] # true efficacy probability in the scenario
122   True_t <- sc[,Sn+1] # true toxicity probability in the scenario
123   Table <- cbind(X, True_e, True_t)
124   tox_over <- which( True_t > Pt ) # doses with over toxicity
125
126 #####
127 # Iterate trials #
128 #####
129 for(sim in sim_min: sim_max){
130
131   ### Reset outcome values at every trial###
132   next_l <- 1 # start from the lowest dose level
133   t <- 0 # test number
134   over_a <- 0 # overdose allocation
135   p_DLT <- 0 #percentage of patients with toxicity
136   p_Res <- 0 #percentage of patients with response
137   num <- rep(0,N) # number of patients in each dose level
138   num_t <- rep(0,N) # number of patients with toxicity by dose
139   num_e <- rep(0,N) # number of patients with response by dose
140
141   #####
142   # Simulate one trial #
143   #####
144   # Test until total number of patients reaches a maximum number of
    patients
145   while(sum(num) + cohort <= max_sub){
146     t <- t + 1 # update a test number
147     seeda <- rand0[sim, (max_sub/cohort*(S-1) + t)] # random seed
148
149     cat("Scinario / Sim - Test : " ,S , " / " , sim , " - " , t,"\n")
150
151     # observed toxicity response
152     set.seed(seeda)
153     y_t <- rbinom(1,cohort , True_t[next_l])
154
155     # observed efficacy response
156     set.seed(seeda)
157     y_e <- rbinom(1, cohort , True_e[next_l])
158
159     # observed data

```

```

160     obs <- as.data.frame(cbind(t_level = next_l, y_e, y_t))
161
162     if(t == 1){
163       Test <- obs
164     }else{
165       Test <- rbind(Test, obs)
166     }
167
168     #update number of patients and responses at each dose level
169     num[next_l] <- num[next_l] + cohort
170     num_t[next_l] <- num_t[next_l] + y_t
171     num_e[next_l] <- num_e[next_l] + y_e
172
173     ### Sampling from the posterior distribution ###
174     # data loaded in the Stan file
175     y_t.out <- as.array(Test$y_t) # toxicity outcomes
176     y_e.out <- as.array(Test$y_e) # efficacy outcomes
177     t.x <- as.array(Test$t_level) # tested dose levels
178     data <- list(cohort = cohort, N = N, pt = Pt, pe = Pe, e_p1 = e_p1,
179                 s_p1 = s_p1, t = t, y_t = y_t.out, y_e = y_e.out, test_x = t.x,
180                 E_logg = Guess_E, T_logg = Guess_T, L = L, L2 = L2)
181
182     # Sampling from the posterior distributions
183     # [toxicity (Ep_t), efficacy (Ep_e) and utility (Utility)]
184     fit <- sampling(stanmodel,
185                   data = data,
186                   pars = c('Ep_t', 'Ep_e', 'Utility'),
187                   iter = 9000, warmup = 3000, thin = 4, chains = 1,
188                   seed = (seeda + t)
189                   #,control = list(adapt_delta = 0.9999999999999999,
190                   # max_treedepth = 20)
191                   )
192
193     ### MCMC parameters ###
194     # Check MCMC results graphically by using following code
195     # ggmcmc(ggs(fit,inc_warmup=TRUE),file='output2.pdf')
196     # Values of iter, warmup, thin and chains in the Stan sampling
197     # should be determined based on the outputs of 'fit'
198     #(e.g., Conditions that Rhat is less than 1.1 and n_eff is greater
199     # than 100) in addition to graphs (e.g., trace plots and auto
200     # correlation).
201     #The recommended number of chains by Stan development team is 4;
202     # however, we set 1 in the simulation study due to low computing
203     # speed. Given the above settings, 1500 posterior functions are
204     # sampled from each posterior distribution.
205     #####
206
207     ms <- rstan::extract(fit) # Get Stan outputs
208
209     # Posterior samples
210     post_tox <- ms$Ep_t # toxicity
211     post_eff <- ms$Ep_e # efficacy
212     post_u <- ms$Utility # utility
213
214     ### Calculate probabilities based on posterior distributions ###
215     # initial vectors for toxicity and utility functions
216     prob_t <- rep(0, N)
217     prob_u <- rep(0, N)
218
219     for(n in 1: N){

```

```

212     # posterior toxicity functions
213     exp_stars1 <- sort(post_tox[, n])
214
215     # posterior utility functions
216     exp_stars2 <- sort(post_u[, n])
217
218     # empirical cumulative distribution function
219     Fn1 <- ecdf(exp_stars1)
220     Fn2 <- ecdf(exp_stars2)
221
222     prob_t[n] <- 1 - Fn1(Pt) # probability that toxicity at each dose
223     prob_u[n] <- 1 - Fn2(c_u1) # probability that utility value at
224     # each dose level is greater than the pre-specified cutoff value
225     # (c_u1)
226 }
227
228 ### Determine an admissible dose set for overdose control ###
229 max_d1 <- N
230 max_d2 <- N
231 max_d3 <- N
232 mean_tox <- apply(post_tox, 2, mean) # posterior toxicity mean
233
234 # Skipping dose is restricted (1st bullet required for overdose
235 # control in Section 5.2.3)
236 if(max(Test$t_level) < N){
237     max_d1 <- max(Test$t_level) + 1
238 }else{
239     max_d1 <- N
240 }
241
242 # One dose level lower than the current test dose when 2 or more
243 # patients experienced DLT at the current tested dose (2nd bullet
244 # required for overdose control in Section 5.2.3)
245 if(y_t > 1){
246     if(next_l == 1){
247         max_d2 <- 1
248     }else{
249         max_d2 <- next_l - 1
250     }
251 }
252
253 # Doses with potentially overly toxic are eliminated (3rd bullet
254 # required for overdose control in Section 5.2.3)
255 if(length( which(mean_tox > (Pt + vareps1)) ) > 0 ){
256     max_d3 <- min( which(mean_tox > (Pt + vareps1)) ) -1
257     if(max_d3 < 1) max_d3 <- 1
258 }else{
259     max_d3 <- N
260 }
261
262 ##### Acquisition function (LCB) for the utility function ###
263 # 10-percentile on the posterior distribution of the utility
264 # function is calculated
265 q10 <- apply(post_u, 2, quantile, probs = c(0.1))
266
267 # LCB in the admissible dose set
268 q10a <- q10[1:(min(max_d1, max_d2, max_d3))]
269
270 # Which dose does provide the minimum value of LCB?

```

```

265     cand <- which(q10a == min(q10a))
266
267     mean_eff <- apply(post_eff, 2, mean) # posterior efficacy mean
268     # If the number of doses selected by LCB is more than 1, a dose with
        the highest efficacy among the admissible dose set is selected.
269     if(length(cand) > 1){
270         cand <- cand[which(mean_eff[cand] == max(mean_eff[cand]))]
271     }
272
273     # The next dose
274     next_l <- cand
275
276 } ### end a trial loop
277
278 #####
279 # Optimal dose is determined #
280 #####
281 utility_mean <- apply(post_u, 2, mean) # posterior utility mean
282
283 # Candidate doses for optimal dose selection
284 admis_f_t <- which(mean_tox < (Pt + vareps2))
285 admis_f_e <- which(mean_eff > (Pe - vareps2))
286 admis_f <- intersect(admis_f_e, admis_f_t)
287
288 # Optimal dose is selected from the candidate dose set
289 if(length(admis_f) == 0){
290     OD <- 0 #OD does not exist
291 }else{
292     OD <- admis_f[which(utility_mean[admis_f] == min(utility_mean[
        admis_f]))]
293 }
294
295 #####
296 # Save output of the trial #
297 #####
298 #number of patients treated with overdoses
299 if(length(tox_over) == 0){
300     over_a <- 0
301 }else{
302     over_a <- sum(num[tox_over])
303 }
304 # percentage of patients with toxicity
305 p_DLT <- sum(num_t)/sum(num)
306 # percentage of patients with response
307 p_Res <- sum(num_e)/sum(num)
308
309 Result_box <- cbind(S, sim, rho, rho2, vareps1, vareps1, OD, total=sum
        (num), data.frame(t(num)), tox = sum(num_t), data.frame(t(num_t))
        , res = sum(num_e), data.frame(t(num_e)), p_Res, p_DLT, over_a, t
        (prob_t), t(prob_u), c_u1, c_u2, t(mean_tox), t(mean_eff), t(
        utility_mean), t(q10))
310
311 # Combine all outputs
312 if(sim == sim_min){
313     Result_box1 <- Result_box
314 }else{
315     Result_box1 <- rbindCOrder(Result_box1, Result_box)
316 }
317 write.table(Result_box1, "output.txt", row.names = F, quote = F, append
        = F)
318 }

```

319 }

A stan model file ‘BOD_model_od.stan’ is composed of the following contents.

Listing C.2. Stan model for BOD of OD estimation

```

1 // BOD_model_od.stan
2 // data input
3 data {
4
5   int<lower = 1> cohort; //number of patients in a cohort
6   int<lower = 1> N; //total number of dose level
7   real pt; //target toxicity rate
8   real pe; //target efficacy rate
9   int<lower = 1> s_p1; //weight parameter for safety
10  int<lower = 1> e_p1; //weight parameter for efficacy
11  int<lower = 1> t; //test number in the trial (t = 1,2,3,...)
12  int<lower = 0, upper = cohort> y_t[t]; //number of patients with
    toxicity at each test
13  int<lower = 0, upper = cohort> y_e[t]; //number of patients with
    efficacy at each test
14  int<lower = 0, upper = N> test_x[t]; //tested doses
15  vector[N] T_logg; //logit transformed initial guess curves for toxicity
16  vector[N] E_logg; //logit transformed initial guess curves for efficacy
17  matrix[N, N] L; //covariance function for toxicity (Cholesky)
18  matrix[N, N] L2; //covariance function for efficacy (Cholesky)
19
20 }
21
22 transformed data{
23   vector[N] mu_zero = rep_vector(0, N);
24 }
25
26 parameters {
27   vector[N] zero_m1;
28   vector[N] zero_m2;
29 }
30
31 transformed parameters{
32
33   vector[N] Ep_t; //dose-toxicity function
34   vector[N] Ep_e; //dose-efficacy function
35   vector[N] Utility; //Utility function
36
37   //--- Define parameters to be output---//
38   Ep_t = inv_logit(zero_m1 + T_logg);
39   Ep_e = inv_logit(zero_m2 + E_logg);
40
41   for(i in 1:N){
42     Utility[i] = e_p1*(Ep_e[i] < pe)*square(Ep_e[i] - pe)+
43                 s_p1*(Ep_t[i] > pt)*square(Ep_t[i] - pt);
44   }
45 }
46 }
47
48 model {
49   //-----declare GPs for the priors-----//
50   zero_m1 ~ multi_normal_cholesky(mu_zero, L);
51   zero_m2 ~ multi_normal_cholesky(mu_zero, L2);

```



```

52
53 //--- likelihood based on a binomial distribution ---//
54 for (j in 1:t){
55     target += binomial_lpmf(y_t[j] | cohort, Ep_t[test_x[j]])+
56               binomial_lpmf(y_e[j] | cohort, Ep_e[test_x[j]]);
57 }
58
59 }

```

The following are toxicity and efficacy probabilities in scenarios 1 and 2.

Listing C.3. Scenarios used in the simulation study for OD estimation (Only scenarios 1 and 2 are listed below.)

1	Level	E1	T1	E2	T2
2	1	0.01	0.05	0.25	0.05
3	2	0.15	0.10	0.45	0.10
4	3	0.30	0.15	0.60	0.15
5	4	0.55	0.30	0.50	0.45
6	5	0.65	0.70	0.20	0.60
7	6	0.70	0.80	0.05	0.65

Chapter 6

Conclusion and future work

In this thesis, we have introduced a new dose-finding statistical design, BOD, that leverages advantages brought from Bayesian optimization while avoiding dose allocations on both overdoses and subtherapeutic doses as much as possible and simultaneously aiming to identify a targeted dose more accurately than existing statistical designs. BOD provides safe and accurate dose-finding under dose-finding studies for investigating mono-therapies or combination therapies in Phase I clinical trials, and targeted therapies of biologic agents in Phase I/II clinical trials.

Although we have examined the operating characteristics of BOD with various scenarios as experimental comparisons throughout this thesis and confirmed that BOD works well compared with the conventional designs, further numerical investigations could reveal more detailed properties for BOD. Additionally, although we have provided some examples for practical settings of design parameters for BOD, we have not developed systematic approaches for them; thus, they are still an open discussion. Statisticians in pharmaceutical industries usually conduct some assessment on the operating characteristics for a statistical design under some scenarios where we may encounter in the real world. Considering the application of BOD in clinical trials, the computational aspect related to MCMC is one of the challenging points because BOD requires more time to evaluate its operating characteristics in simulation studies.

BOD introduced in Chapter 3 covers the simplest and the most typical dose-finding studies that aim to find the MTD for mono-therapies. The operating characteristics are almost comparable with the compared designs and the advantages shown in BOD seem minimal compared with BOD for MTDC and OD estimation. Although the simulation results suggest BOD for MTD estimation provides safer dose allocations than the existing designs, parametric model-based designs that strictly follow the monotonicity may be a reasonable choice rather than nonparametric model-based designs even though little is available about dose-toxicity relationships because it is highly likely that dose-toxicity relationships have monotonicity trends against dose levels. In terms of the benefit to assume the monotonicity on dose-toxicity relationships, a similar discussion is provided in Section 5.4 for BOD for OD estimation. Because some investigators may think it is reasonable to assume toxicity increases monotonically with increasing dose levels even if investigational agents are biologic ones, they may be interested in a hybrid approach that combines a parametric model-based design for dose-toxicity relationships and a nonparametric model-based design based on BOD for dose-efficacy relationships. Therefore, it may be expected to attempt such a hybrid approach and investigate how much its operating characteristics show improvement in future work.

BOD for OD estimation and MTDC estimation in Chapters 4 and 5 provides appropriate dose allocations and more accurate dose selections compared with the other designs. It implies that BOD with flexible modeling and dose selections considering uncertainties

can successfully address more complex situations showing non-monotonic relationships between dose levels and response probabilities due to such as nonconventional mode-of-actions or drug interactions. Biologic agents have been increasingly investigated in recent Phase I clinical trials and the development of targeted therapies has progressed significantly to date (Braun, 2014; Seebacher et al., 2019). Many Phase I clinical trials are designed to investigate combinations of two or more agents (Tourneau et al., 2009). Additionally, as Yuan et al. (2019) has described the treatment landscape for non-small-cell lung cancer, development strategies especially in oncology areas for new treatment including combination therapies have been getting wider options and more complex than before. Considering such increasing diversity of drug development, it is worth being involved in the development of new statistical approaches like BOD providing a more efficient strategy that is different from conventional statistical designs assuming only the simplest MTD estimation problems.

Nevertheless, attractive dose-finding designs that have recently emerged for drug combination therapies are hardly used in practice so far. Instead, conventional statistical designs developed for mono-therapies have been applied (Riviere et al., 2015b). Even worse, traditional rule-based designs such as the 3+3 design whose inferiority has already been demonstrated compared with other conventional statistical designs are still the most used ones among clinicians and oncologists. As Rogatko et al. (2007) has reported that there are notable time lags between the application in practical clinical trials and the publication of novel statistical designs in terms of Phase I cancer clinical trials. They have shown that even if novel statistical designs that outperform conventional statistical designs are recognized in public, it takes a long time to bring statistical improvements into clinical practice. Such a skewed situation between actual applications and scientific evidence comes from barriers to the implementation of adaptive designs. For example, Love et al. (2017) has reported the most prominent barriers to implement model-based designs are lack of suitable training, investigators' preference for traditional rule-based designs due to their simplicity and rich experiences, and limited resources for study designs before funding.

Statistical designs for dose-finding studies had been broadly fallen into two classes for decades: the rule-based designs, the model-based designs. On the other hand, toxicity interval designs categorized in nonparametric designs have joined recently as a new class owing to the simplicity of implementation while having a Bayesian structure. Although Jaki et al. (2013) has reported the advantage of nonparametric model-based designs compared with other design types, the number of researches on nonparametric model-based designs for dose-finding studies is still very limited, which might attribute to the implementation complexity (Braun, 2014). Nonparametric model-based designs are still a minority in this field, however, something triggers (e.g., resolutions of implementation barriers and a change in the environment of regulatory agencies who are the gatekeeper of statistical designs for clinical trials) may accelerate researches of nonparametric model-based designs in the future. Developing a user-friendly interface is one of the mandatory approaches for relaxing the implementation barriers on BOD, which should be addressed in future work. In fact, most statistical designs utilized frequently in clinical trials have provided user-friendly interfaces such as R packages, Shiny applications, and cloud software. Although we have provided R codes to implement BOD, those who are not familiar with R or non-statisticians require user-friendly interfaces to understand what BOD does in its modeling and dose selections in practice. These improvements might result in the sophistication of nonparametric designs including BOD in terms of its implementation and numerical aspects. Still, improving the current skewed situation where better designs have not been actively employed in dose-finding studies will remain a challenging issue. For a statistician in the pharmaceutical industry, it is important to take the proper time

to actively consider and incorporate more efficient and safer designs and negotiate carefully with regulatory authorities. Additionally, as Rogatko et al. (2007) has discussed, it is essential for biostatisticians to collaborate with clinical colleagues efficiently as two-way communication so that more appropriate designs could be widely adopted in clinical trials.

Drug development has recently been actively attempting to incorporate modeling and simulation called model-informed drug development (MIDD), which is a wide range of applications of leveraging models on exposure and biological data to provide quantitative evaluations for improving efficiency in drug development (Wang et al., 2019). So far, this movement has strongly affected the pharmacometrics area. However, it is likely that it will also associate with statistical designs in dose-finding studies. Biological data such as Pharmacokinetics (PK) and Pharmacodynamics (PD) data are usually collected in dose-finding studies for exploratory assessments to understand investigated agents. Nevertheless, most dose-finding studies primarily utilize dichotomous outcomes on toxicity and/or efficacy in terms of MTD, MTDC, or OD determinations. It is theoretically possible to extend BOD so that it can deal with continuous variables on such as not only alternative toxicity/efficacy data but also PK/PD data and biomarker data instead of binary variables; therefore, those extensions adding other values without losing information could increase success rates of drug development as one of the MIDD approaches. We leave those possible extensions in future work.

We also left an expectation on the application of Bayesian optimization in areas that are not only clinical trials but also other fields. Typical Bayesian optimization rarely applies to optimization problems where locations of optimal points have strong limitations. On the other hand, dose-finding studies require to identify doses that are located right next to the area where we should refrain from aggressively searching to ensure patient safety. This thesis indicates that Bayesian optimization can apply to optimization problems with unique limitations on point selections over input space. There might be similar problem settings in other fields such as where optimization problems should take safety concerns into account. Therefore, we have high expectations for the Bayesian optimization community not only to cover typical cases but also to consider such as the above special ones in the future.

Acknowledgment

I would like to express my deepest gratitude to my supervisor Professor Taiji Suzuki for accepting me into his laboratory and providing me thoughtful guidance and continuous support. His insightful advice has brought me valuable unforgettable experiences that I could study a cutting-edge approach in my field of expertise. Although this research journey was much harder than I expected, thanks to his continuous encouragement, I could keep my motivation and to pursue my research interest. I also express my sincere appreciation to Professor Takafumi Kanamori and Professor Naoto Miyoshi for kindly accepting me into their laboratories after the transfer of my supervisor and providing me continuous support even during the leave of absence. I appreciate Professor Takafumi Kanamori for his thoughtful support and insightful advice at both presentations at seminars and off-line communications. Additionally, I would like to thank all my dissertation committee members including Professor Taiji Suzuki, Professor Takafumi Kanamori, Professor Naoto Miyoshi, Professor Sumio Watanabe, Professor Misako Takayasu, and Professor Yumiharu Nakano for taking the time to review my thesis and giving me valuable advice for improving the thesis. I appreciate all members of the School of Computing, Department of Mathematical and Computing Science in Tokyo institute of technology for their support. I would like to appreciate the Tokyo institute of technology for providing an excellent environment for studying. Thanks to the warm support and encouragement of my colleagues who have worked together in Pfizer R & D Japan, I could spend a very comfortable and motivated research life at the Tokyo Institute of Technology. I would like to appreciate Pfizer R & D Japan for allowing me to be a working student at the Tokyo Institute of Technology and for partially supporting the financial aspect. I pay my heartfelt gratitude to those who have taken care of my daughters while being away from home due to research work and my job. I finally would like to thank my dear husband and two cute daughters, and parents for all their support and encouragement.

Bibliography

- A. Hirakawa, Research representative. AMED 医薬品等規制調和・評価研究事業, 希少疾患領域の医薬品開発を効率化するための小規模臨床試験のデザイン・統計解析法の研究開発とその適正利用のための基本的考え方の策定. <https://research-er.jp/projects/view/1074219>, 2019-2021. (Last visit on the website: 14 January 2021).
- R. Ananthkrishnan, S. Green, M. Chang, G. Doros, J. Massaro, and M. LaValley. Systematic comparison of the statistical operating characteristics of various Phase I oncology designs. *Contemporary Clinical Trials Communications*, 5:34–48, 2016.
- L. Andrews, S. Ralston, E. Blomme, and K. Barnhart. A snapshot of biologic drug development: Challenges and opportunities. *Human and Experimental Toxicology*, 34(12):1279–1285, 2015.
- N. Arora, A. Gupta, and P. Paul. Biological agents in gastrointestinal cancers: Adverse effects and their management. *Journal of Gastrointestinal Oncology*, 8(3):485–98, 2017.
- T. Asakawa and C. Hamada. A pragmatic dose-finding approach using short-term surrogate efficacy outcomes to evaluate binary efficacy and toxicity outcomes in Phase I cancer clinical trials. *Pharmaceutical Statistics*, 12(5):315–27, 2013.
- R. E. Barlow, D. J. Bartholomew, J. M. Bremner, and H. D. Brunk. *Statistical Inference Under Order Restrictions*. John Wiley & Sons, New York, 1972.
- B. N. Bekele and Y. Shen. A Bayesian approach to jointly modeling toxicity and biomarker expression in Phase I/II dose-finding trial. *Biometrics*, 61(2):344–54, 2005.
- D. A. Berry. Bayesian clinical trials. *Nature Reviews Drug Discovery*, 5:27–36, 2006.
- J. Bobb, A. Rogatko, and S. Zacks. Cancer Phase I clinical trials: Efficient dose escalation with overdose control. *Statistics in Medicine*, 17(10):1103–20, 1998.
- P. S. Boonstra, J. Shen, J. M. G. Taylor, T. M. Braun, K. A. Griffith, S. Daignault, G. P. Kalemkerian, T. S. Lawrence, and M. J. Schipper. A statistical evaluation of dose expansion cohorts in Phase I clinical trials. *Journal of the National Cancer Institute*, 107(3):dju429, 2015.
- T. M. Braun. The bivariate continual reassessment method: Extending the CRM to Phase I trials of two competing outcomes. *Controlled Clinical Trials*, 23(3):240–56, 2002.
- T. M. Braun. The current design of oncology Phase I clinical trials: Progressing from algorithms to statistical models. *Chinese Clinical Oncology*, 3(1):1–11, 2014.
- E. Brochu, V. M. Cora, and N. D. Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *ArXiv*, 2010. 1012.2599.
- K. Brock, L. Billingham, M. Copland, S. Siddique, M. Sirovica, and C. Yap. Implementing the EffTox dose-finding design in the matchpoint trial. *BMC Medical Research Methodology*, 17(1):112, 2017.
- A. D. Bull. Convergence rates of efficient global optimization algorithms. *Journal of Machine Learning Research*, 12(88):2879–904, 2011.
- S. K. Carter. Study design principles for the clinical evaluation of new drugs as developed by the chemotherapy program of the National Cancer Institute. In *M.J. Staquet (Ed.), The Design of Clinical Trials in Cancer Therapy*, pages 242–89, Brussels, 1973. Editions Scientifiques Européennes.

- Z. Chen, M. D. Krailo, S. P. Azen, and M. Tighiouart. A novel toxicity scoring system treating toxicity response as a quasi-continuous variable in Phase I clinical trials. *Contemporary Clinical Trials*, 31(5):473–82, 2010.
- Z. Chena, M. Tighiouart, and J. Kowalski. Dose escalation with overdose control using a quasi-continuous toxicity score in cancer Phase I clinical trials. *Contemporary Clinical Trials*, 33(5):949–58, 2012.
- Y. K. Cheung. On the use of nonparametric curves in Phase I trials with low toxicity tolerance. *Biometrics*, 58(1):237–40, 2002.
- Y. K. Cheung. *Dose Finding by the Continual Reassessment Method*. CRC Press, Boca Raton, FL, 2011.
- Y. K. Cheung and R. Chappell. A simple technique to evaluate model sensitivity in the continual reassessment method. *Biometrics*, 58(3):671–4, 2002.
- T. Daimon. The continual reassessment method (CRM) and related topics – A review. Part I: Understanding the CRM for dose-finding studies. *Japanese Journal of Biometrics*, 33(1):1–29, 2012.
- M. A. Diniz, Q. Li, and M. Tighiouart. Dose finding for drug combination in early cancer Phase I trials using conditional continual reassessment method. *Journal of Biometrics & Biostatistics*, 8(6):381, 2017.
- V. Dragalin and V. Fedorov. Adaptive designs for dose-finding based on efficacy-toxicity response. *Journal of Statistical Planning and Inference*, 136(6):1800–23, 2006.
- D. Faries. Practical modifications of the continual reassessment method for Phase I cancer clinical trials. *Journal of Biopharmaceutical Statistics*, 4(2):147–64, 1994.
- V. V. Fedorov and S. L. Leonov. *Optimal Design for Nonlinear Response Models*. CRC Press, Boca Raton, FL, 2013.
- M. Gasparini and J. Eisele. A curve-free method for Phase I clinical trials. *Biometrics*, 56(2):609–15, 2000.
- M. Gasparini and J. 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):659–60, 2001.
- S. Goodman, M. L. Zahurak, and S. Piantadosi. Some practical improvements in the continual reassessment method for Phase I studies. *Statistics in Medicine*, 14(11):1149–61, 1995.
- R. B. Gramacy and H. K. H. Lee. Optimization under unknown constraints. In *Bayesian Statistics 9*, pages 229–56, New York, 2011. Oxford University Press.
- W. Guo, S. J. Wang, S. Yang, H. Lynn, and Y. Ji. A Bayesian interval dose-finding design addressing Ockham’s razor: mTPI-2. *Contemporary Clinical Trials*, 58:23–33, 2017.
- S. K. Gupta. Use of Bayesian statistics in drug development: Advantages and challenges. *International Journal of Applied and Basic Medical Research*, 2(1):3–6, 2012.
- A. Hirakawa, N. A. Wages, H. Sato, and S. Matsui. A comparative study of adaptive dose-finding designs for Phase I oncology trials of combination therapies. *Statistics in Medicine*, 34(24):3194–213, 2015.
- A. Hirakawa, H. Sato, and M. Goshio. Effect of design specifications in dose finding trials for combination therapies in oncology. *Pharmaceutical Statistics*, 15(6):531–40, 2016.
- B. J. Horton, N. A. Wages, and M. R. Conaway. Performance of toxicity probability interval based designs in contrast to the continual reassessment method. *Statistics in Medicine*, 36(2):291–300, 2017.
- A. Iasonos and J. O’Quigley. Adaptive dose-finding studies: A review of model-guided Phase I clinical trials. *Journal of Clinical Oncology*, 32(23):2505–11, 2014.
- A. Iasonos, A. S. Wilton, E. R. Riedel, V. E. Seshan, and D. R. Spriggs. 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):465–77, 2008.
- A. Ivanova. Escalation, group and A + B designs for dose-finding trials. *Statistics in*

- Medicine*, 25(21):3668–78, 2006.
- T. Jaki, S. Clive, and C. J. Weir. Principles of dose finding studies in cancer: A comparison of trial designs. *Cancer Chemotherapy and Pharmacology*, 71(5):1107–14, 2013.
- G. D. James, S. N. Symeonides, J. Marshall, J. Young, and G. Clack. Continual reassessment method for dose escalation clinical trials in oncology: A comparison of prior skeleton approaches using AZD3514 data. *BMC Cancer*, 16(1):703, 2016.
- Y. Ji and S. J. Wang. 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):1785–91, 2013.
- Y. Ji, Y. Li, and B. N. Bekele. Dose-finding in Phase I clinical trials based on toxicity probability intervals. *Clinical Trials*, 4(3):235–44, 2007.
- Y. Ji, P. Liu, Y. Li, and B. N. Bekele. A modified toxicity probability interval method for dose-finding trials. *Clinical Trials*, 7(6):653–63, 2010.
- J. L. Jimenez, M. Tighiouart, and M. Gasparini. Cancer Phase I trial design using drug combinations when a fraction of dose limiting toxicities is attributable to one or more agents. *Biometrical Journal*, 61(2):319–32, 2019.
- D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13:455–92, 1998.
- S. M. Lee and Y. K. Cheung. Model calibration in the continual reassessment method. *Clinical Trials*, 6(3):227–38, 2009.
- S. M. Lee and Y. K. Cheung. Calibration of prior variance in the bayesian continual reassessment method. *Statistics in Medicine*, 30(17):2081–89, 2011.
- D. H. Y. Leung and Y. G. Wang. An extension of the continual reassessment method using decision theory. *Statistics in Medicine*, 21(1):51–63, 2002.
- D. H. Li, J. B. Whitmore, W. Guo, and Y. Ji. Toxicity and efficacy probability interval design for Phase I adaptive cell therapy dose-finding clinical trials. *Clinical Cancer Research*, 23(1):13–20, 2017.
- R. Lin and G. Yin. Bayesian optimal interval design for dose finding in drug-combination trials. *Pharmaceutical Statistics*, 26(5):2155–67, 2017a.
- R. Lin and G. Yin. STEIN: A simple toxicity and efficacy interval design for seamless Phase I/II clinical trials. *Statistics in Medicine*, 36(26):4106–20, 2017b.
- S. Liu and V. E. Johnson. A robust Bayesian dose-finding design for Phase I/II clinical trials. *Biostatistics*, 17(2):249–63, 2016.
- S. Liu and Y. Yuan. Bayesian optimal interval designs for Phase I clinical trials. *Journal of the Royal Statistical Society Series C (Applied Statistics)*, 64(3):507–23, 2015.
- S. B. Love, S. Brown, C. J. Weir, C. Harbron, C. Yap, B. Gaschler-Markefski, J. Matcham, L. Caffrey, C. Mckevitt, S. Clive, C. Craddock, J. Spicer, and V. Cornelius. Embracing model-based designs for dose-finding trials. *British Journal of Cancer*, 117(3):332–9, 2017.
- S. J. Mandrekar, R. Qin, and D. J. Sargent. Model-based Phase I designs incorporating toxicity and efficacy for single and dual agent drug combinations: Methods and challenges. *Statistics in Medicine*, 29(10):1077–83, 2010.
- C. Mayer, I. Perevozskaya, S. Leonov, V. Dragalin, Y. Pritchett, A. Bedding, A. Hartford, P. Fardipour, and G. Cicconetti. Simulation practices for adaptive trial designs in drug and device development. *Statistics in Biopharmaceutical Research*, 11(4):1–20, 2019.
- A. Mockus, J. Mockus, and L. Mockus. Bayesian approach adapting stochastic and heuristic method of global and discrete optimization. *Informatika*, 5(1-2):123–66, 1994.
- J. Mockus. On Bayesian methods for seeking the extremum. In *Optimization Techniques IFIP Technical Conference*, pages 400–4. Springer, Berlin, Heidelberg, 1975.
- J. Mockus, V. Tiesis, and A. Zilinskas. The application of Bayesian methods for seeking the extremum. In *Towards Global Optimisation 2*, pages 117–29. Elsevier Science Ltd,

- North Holland, Amsterdam, 1978.
- R. B. Mokhtari, T. S. Homayouni, N. Baluch, E. Morgatskaya, S. Kumar, B. Das, and H. Yeger. Combination therapy in combating cancer. *Oncotarget*, 8(23):38022–38043., 2017.
- P. Mozgunov and T. Jaki. Improving safety of the continual reassessment method via a modified allocation rule. *Statistics in Medicine*, 39(7):906–22, 2020.
- D. G. Muenz, J. M. G. Taylor, and T. M. Braun. Phase I-II trial design for biologic agents using conditional auto-regressive models for toxicity and efficacy. *Journal of the Royal Statistical Society Series C (Applied Statistics)*, 68(2):331–45, 2019.
- P. A. Murtaugh and L. D. Fisher. Bivariate binary models of efficacy and toxicity in dose-ranging trials. *Communications in Statistics Part A - Theory and Methods*, 19(6):2003–20, 1990.
- R. M. Neal. Monte carlo implementation of Gaussian process models for Bayesian regression and classification. Technical report, Department of Statistics, University of Toronto, 1997.
- B. Neuenschwander, M. Branson, and T. Gsponer. Critical aspects of the Bayesian approach to Phase I cancer trials. *Statistics in Medicine*, 27(13):2420–39, 2008.
- V. Nguyen, S. Gupta, S. Rana, C. Li, and S. Venkatesh. Regret for expected improvement over the best-observed value and stopping condition. In *Proceedings of Machine Learning Research*, volume 77, pages 279–94, 2017.
- A. Onar, M. Kocak, and J. 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. *Journal of Biopharmaceutical Statistics*, 19(3):437–55, 2009.
- A. Onar-Thomas and Z. 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. *Contemporary Clinical Trials*, 31(3):259–70, 2010.
- J. O’Quigley and S. Chevret. Methods for dose finding studies in cancer clinical trials: A review and results of a monte carlo study. *Statistics in Medicine*, 10(11):1647–64, 1991.
- J. O’Quigley, M. Pepe, and L. Fisher. Continual reassessment method: A practical design for Phase I clinical trials in cancer. *Biometrics*, 46(1):33–48, 1990.
- J. Orloff, F. Douglas, J. Pinheiro, S. Levinson, M. Branson, P. Chaturvedi, E. Ette, P. Gallo, G. Hirsch, C. Mehta, N. Patel, S. Sabir, S. Springs, D. Stanski, M. R. Evers, E. Fleming, N. Singh, T. Tramontin, and H. Golub. The future of drug development: advancing clinical trial design. *Nature Reviews Drug Discovery*, 8:949–57, 2009.
- X. Paoletti and A. Kramar. A comparison of model choices for the continual reassessment method in Phase I cancer trials. *Statistics in Medicine*, 28(24):3012–28, 2009.
- S. Postel-Vinay. Redefining dose-limiting toxicity. *Clinical Advances in Hematology & Oncology*, 13(2):87–9, 2015.
- S. Postel-Vinay, C. Gomez-Roca, L. R. Molife, B. Anghan, A. Levy, I. Judson, J. D. Bono, J.-C. Soria, S. Kaye, and X. Paoletti. Phase I trials of molecularly targeted agents: Should we pay more attention to late toxicities? *Journal of Clinical Oncology*, 29(13):1728–35, 2011.
- C. Qin, D. Klabjan, and D. Russo. Improving the expected improvement algorithm. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 5387–97, 2017.
- N. Quadrianto, K. Kersting, and Z. Xu. Gaussian process. In C. Sammut and G. I. Webb, editors, *Encyclopedia of Machine Learning*. Springer, Boston, MA, 2011.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, London, England, 2006.

- M. J. Ratain, R. Mick, R. L. Schilsky, and M. Siegler. Statistical and ethical issues in the design and conduct of Phase I and II clinical trials of new anticancer agents. *Journal of the National Cancer Institute*, 85(20):1637–43, 1993.
- E. Reiner, X. Paoletti, and J. O’Quigley. Operating characteristics of the standard Phase I clinical trial design. *Computational Statistics & Data Analysis*, 30(3):303–15, 1999.
- M.-K. Riviere, Y. Yuan, F. Dubois, and S. Zohar. A Bayesian dose-finding design for drug combination clinical trials based on the logistic model. *Pharmaceutical Statistics*, 13(4):247–57, 2014.
- M.-K. Riviere, F. Dubois, and S. Zohar. Competing designs for drug combination in Phase I dose-finding clinical trials. *Statistics in Medicine*, 34(1):1–12, 2015a.
- M.-K. Riviere, C. L. Tourneau, X. Paoletti, F. Dubois, and S. Zohar. Designs of drug-combination Phase I trials in oncology: A systematic review of the literature. *Annals of Oncology*, 26:669–74, 2015b.
- P. R. Robuck and J. I. Wurzelmann. Understanding the drug development process. *Inflammatory Bowel Diseases*, 11(suppl 1):S13–S6, 2005.
- A. Rogatko, D. Schoeneck, W. Jonas, M. Tighiouart, F. R. Khuri, and A. Porter. Translation of innovative designs into Phase I trials. *Journal of Clinical Oncology*, 25(31):4982–86, 2007.
- W. F. Rosenberger and L. M. Haines. Competing designs for Phase I clinical trials: A review. *Statistics in Medicine*, 21(18):2757–70, 2002.
- A. Russu, I. Poggesi, R. Gomeni, and G. D. Nicolao. Bayesian population modeling of Phase I dose escalation studies: Gaussian process versus parametric approaches. *IEEE Transactions on Biomedical Engineering*, 58(11):3156–64, 2011.
- P.-I. Schneider, X. G. Santiago, C. Rockstuhl, and S. Burger. Global optimization of complex optical structures using Bayesian optimization based on Gaussian processes. In *SPIE Digital Optical Technologies, 2017*, page 103350O, 2017.
- N. A. Seebacher, A. E. Stacy, G. M. Porter, and A. M. Merlo. Clinical development of targeted and immune based anti-cancer therapies. *Journal of Experimental & Clinical Cancer Research*, 38(156):1–39, 2019.
- B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. D. Freitas. Taking the human out of the loop: A review of Bayesian optimization. In *Proceedings of the IEEE*, volume 104, pages 148–75, 2016.
- L. Z. Shen and J. O’Quigley. Consistency of continual reassessment method under model misspecification. *Blometrika*, 83(2):395–405, 1996.
- J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In *NIPS’12: Proceedings of the 25th International Conference on Neural Information Processing Systems*, volume 2, pages 2951–9, 2012.
- N. Srinivas, A. Krause, S. M. Kakade, and M. 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*, pages 1015–22, 2010.
- B. E. Storer. Design and analysis of Phase I clinical trials. *Biometrics*, 45(3):925–37, 1989.
- A. Takahashi and T. Suzuki. A curve free method based on Bayesian optimization for oncology Phase I clinical trials. In *2nd International Conference on Statistics: Theory and Applications (ICSTA’20)*, page ICSTA 106, 2020.
- A. Takahashi and T. Suzuki. Bayesian optimization for estimating the maximum tolerated dose in Phase I clinical trials. *Contemporary Clinical Trials Communications*, 21:100753, 2021a. doi: 10.1016/j.conctc.2021.100753.
- A. Takahashi and T. Suzuki. Bayesian optimization design for finding a maximum tolerated dose combination in Phase I clinical trials. *The International Journal of Biostatistics*, page to appear, 2021b. doi: 10.1515/ijb-2020-0147.

- A. Takahashi and T. Suzuki. Bayesian optimization design for dose-finding based on toxicity and efficacy outcomes in Phase I/II clinical trials. *Pharmaceutical Statistics*, 20(3):422–39, 2021c.
- K. Takeda, M. Taguri, and S. Morita. BOIN-ET: Bayesian optimal interval design for dose finding based on both efficacy and toxicity outcomes. *Pharmaceutical Statistics*, 17(4):383–95, 2018.
- K. Takeda, S. Morita, and M. Taguri. TITE-BOIN-ET: Time-to-event Bayesian optimal interval design to accelerate dose-finding based on both efficacy and toxicity outcomes. *Pharmaceutical Statistics*, 19(3):335–49, 2020.
- N. Tang, S. Wang, and G. Ye. A nonparametric Bayesian continual reassessment method in single-agent dose-finding studies. *BMC Medical Research Methodology*, 18(1):172, 2018.
- Y. Tao, J. Lin, Z. Li, J. Lin, T. Lu, and F. Yan. Dose-finding based on bivariate efficacy-toxicity outcome using Archimedean copula. *PLOS One*, 8(11):e78805, 2013.
- P. F. Thall and J. D. Cook. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics*, 60(3):684–93, 2004.
- P. F. Thall, J. D. Cook, and E. Estey. Adaptive dose selection using efficacy-toxicity trade-offs: Illustrations and practical considerations. *Journal of Biopharmaceutical Statistics*, 16(5):623–38, 2006.
- P. F. Thall, R. C. Herrick, H. Q. Nguyen, H. Q. Nguyen, J. J. Venier, and J. C. Norris. Effective sample size for computing prior hyperparameters in Bayesian Phase I-II dose-finding. *Clinical Trials*, 11(6):657–66, 2014.
- M. Tighiouart and A. Rogatko. Dose finding with escalation with overdose control (EWOC) in cancer clinical trials. *Statistical Science*, 25(2):217–26, 2010.
- C. L. Tourneau, J. J. Lee, and L. L. Siu. Dose escalation methods in Phase I cancer clinical trials. *Journal of the National Cancer Institute*, 101(10):708–720, 2009.
- T. R. Turner and P. C. Wollan. Locating a maximum using isotonic regression. *Computational Statistics & Data Analysis*, 25(3):305–320, 1997.
- C. Twelves, M. Jove, A. Gombos, and A. Awad. Cytotoxic chemotherapy: Still the mainstay of clinical practice for all subtypes metastatic breast cancer. *Critical Reviews in Oncology/Hematology*, 100:74–87, 2016.
- N. A. Wages. Identifying a maximum tolerated contour in two-dimensional dose finding. *Statistics in Medicine*, 36(2):242–53, 2017.
- N. A. Wages and M. R. Conaway. Specifications of a continual reassessment method design for Phase I trials of combined drugs. *Pharmaceutical Statistics*, 12(4):217–24, 2013.
- N. A. Wages, M. R. Conaway, and J. O’Quigley. Continual reassessment method for partial ordering. *Biometrics*, 67(4):1555–63, 2011.
- K. Wang and A. Ivanova. Two-dimensional dose finding in discrete dose space. *Biometrics*, 61(1):217–22, 2005.
- Y. Wang and A. Ivanova. Dose finding with continuous outcome in Phase I oncology trials. *Pharmaceutical Statistics*, 14(2):102–7, 2015.
- Y. Wang, H. Zhu, R. Madabushi, Q. Liu, S. Huang, and I. Zineh. Model-informed drug development: Current US regulatory practice and future considerations. *Clinical Pharmacology and Therapeutics*, 105(4):899–911, 2019.
- J. Whitehead, H. Thygesen, and A. Whitehead. A Bayesian dose-finding procedure for Phase I clinical trials based only on the assumptions of monotonicity. *Statistics in Medicine*, 29(17):1808–24, 2010.
- F. Yan, L. Zhang, Y. Zhou, H. Pan, S. Liu, and Y. Yuan. BOIN: An R package for designing single-agent and drug-combination dose-finding trials using Bayesian optimal interval designs. *Journal of Statistical Software*, 94(13):527, 2020.
- W. Y. Yeung, B. Reigner, U. Beyer, and C. Diack. Bayesian adaptive dose-escalation

- designs for simultaneously estimating the optimal and maximum safe dose based on safety and efficacy. *Pharmaceutical Statistics*, 16(6):396–413, 2017.
- G. Yin and Y. Yuan. Bayesian model averaging continual reassessment method in Phase I clinical trials. *Journal of the American Statistical Association*, 104(487):954–68, 2009a.
- G. Yin and Y. Yuan. Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society Series C (Applied Statistics)*, 58(2):211–24, 2009b.
- G. Yin, Y. Li, and Y. Ji. Bayesian dose-finding in Phase I/II clinical trials using toxicity and efficacy odds ratios. *Biometrics*, 62(3):777–87, 2006.
- M. Yuan, L.-L. Huang, J.-H. Chen, J. Wu, and Q. Xu. The emerging treatment landscape of targeted therapy in non-small-cell lung cancer. *Signal Transduction and Targeted Therapy*, 4(61), 2019.
- Y. Yuan and G. Yin. Bayesian dose finding by jointly modeling toxicity and efficacy as time-to-event outcomes. *Journal of the Royal Statistical Society Series C (Applied Statistics)*, 58(5):719–36, 2009.
- Y. Yuan and G. Yin. Bayesian hybrid dose-finding design in Phase I oncology clinical trials. *Statistics in Medicine*, 30(17):2098–108, 2011a.
- Y. Yuan and G. Yin. Dose-response curve estimation: A semiparametric mixture approach. *Biometrics*, 67(4):1543–54, 2011b.
- Y. Yuan, K. R. Hess, S. G. Hilsenbeck, and M. R. Gilbert. Bayesian optimal interval designs: A simple and well-performing design for Phase I oncology trials. *Clinical Cancer Research*, 22(17):4291–301, 2016.
- Y. Zang, J. J. Lee, and Y. Yuan. Adaptive designs for identifying optimal biological dose for molecularly targeted agents. *Clinical Trials*, 11(3):319–27, 2014.
- L. Zhang and Y. Yuan. A practical Bayesian design to identify the maximum tolerated dose contour for drug combination trials. *Statistics in Medicine*, 35(27):4924–36, 2016.
- W. Zhong, J. S. Koopmeiners, and B. P. Carlin. A trivariate continual reassessment method for Phase I/II trials of toxicity, efficacy, and surrogate efficacy. *Statistics in Medicine*, 31(29):3885–95, 2012.
- H. Zhou, Y. Yuan, and L. Nie. Accuracy, safety, and reliability of novel Phase I trial designs. *Clinical Cancer Research*, 24(18):4357–64, 2018.
- Y. Zhou. Choice of designs and doses for early phase trials. *Fundamental & Clinical Pharmacology*, 18(3):373–8, 2004.
- Y. Zhu, W.-T. Hwang, and Y. Li. Evaluating the effects of design parameters on the performances of Phase I trial designs. *Contemporary Clinical Trials Communications*, 15:100379, 2019.
- S. Zohar and J. O’Quigley. Re: Dose escalation methods in Phase I cancer clinical trials. *Journal of the National Cancer Institute*, 101(24):1732–3, 2009.