Nonparametric Bayesian Method for Drug Combination with Discrete Doses

Zahra S. Razaee* Galen Cook-Wiens * Mourad Tighiouart*

Abstract

We propose an adaptive design for early-phase drug-combination cancer trials with the goal of estimating the maximum tolerated dose (MTD). A nonparametric Bayesian model, using Beta priors truncated to the set of partially ordered dose combinations, is used to describe the probability of dose limiting toxicity (DLT). Dose allocation between successive cohorts of patients is estimated using a modified Continual Reassessment scheme. The updated probabilities of DLT are calculated with a Gibbs sampler that employs a weighting mechanism to calibrate the influence of data versus the prior. At the end of the trial, we recommend one or more dose combinations as the MTD based on our proposed algorithm. The design operating characteristics indicate that our method is comparable with existing methods.

Key Words: Cancer Phase I trials, Drug combination, Maximum tolerated dose, Nonparametric Bayesian method, Partial ordering

1. Introduction

The primary objective in conventional phase I clinical trials is to determine the maximum tolerated dose (MTD), defined as the dose with the probability of toxicity closest to a prespecified target. For safety and ethical concerns, most phase I trials are conducted adaptively, using the dose limiting toxicity (DLT) status of previously enrolled patients to determine the dose level for the next cohort of patients. The majority of such trials are designed for single agent, e.g., the conventional 3 + 3 design (Storer, 1989), the continual reassessment method(CRM) and its variants (O'Quigley, Pepe, and Fisher, 1990), (Goodman, Zahurak, and Piantadosi, 1995), (Korn and Simon, 1991), (Møller, 1995), (O'Quigley and Shen, 1996), (Leung and Wang, 2002), (O'Quigley and Paoletti, 2003), (Iasonos and O'Quigley, 2011), (Daimon, Zohar, and O'Quigley, 2011), (Liu, Yin, and Yuan, 2013), the efficient dose escalation with overdose control (EWOC) method and its variants (Babb, Rogatko, and Zacks, 1998; Tighiouart, Piantadosi, and Rogatko, 2014; Tighiouart and Rogatko, 2010; Wheeler, Sweeting, and Mander, 2017), the modified toxicity probability interval method (Ji and Wang, 2013), the Bayesian optimal design (Yuan et al., 2016), the nonparametric overdose control method (Lin and Yin, 2017), the semiparametric dose finding methods (Clertant and O'Quigley, 2017) and the Bayesian adaptive design using a flexible range of doses (Tighiouart, Cook-Wiens, and Rogatko, 2018).

Recent advances in drug discovery have intensified interest in using dual agents in phase I clinical trials. This interest is fueled by the fact that drug combinations may

*Cedars-Sinai Medical Center, Los Angeles, CA, USA.

induce a synergistic treatment effect by targeting multiple pathways simultaneously and inhibiting resistance mechanisms. A fundamental assumption for cytotoxic agents is monotonicity between toxicity and doses. For single agent, this assumption induces a complete ordering of the doses. However, in the case of drug combination treatment where the two agents are allowed to vary, it induces a partial ordering constraint on the probabilities of toxicities. The monotonicity assumption coupled with small sample size in phase I clinical trials and higher dimension of the dose space, make the design of combination trials challenging.

Various model-based designs for drug combinations, have been studied in the last decade. Thall et al., 2003 proposed a method using a six-parameter model to define the probability of toxicity as a function of the two doses with the requirement that each of the two agents had been studied previously as a single agent. Wang and Ivanova, 2005 used a two-stage design with regression model. Yin and Yuan, 2009a and Yin and Yuan, 2009b developed a design that models the probability of toxicity with a copula type model. Wages, Conaway, and O'Quigley, 2011 considered estimation of toxicity probabilities within a small number of simple orders. Tighiouart, Li, and Rogatko, 2017 and Tighiouart, Cook-Wiens, and Rogatko, 2018 used a reparametrized logistic model to describe the relationship between the doses of the two agents and the probability of dose limiting toxicity and extended the work of Tighiouart, Piantadosi, and Rogatko, 2014 by allowing the MTD curve to lie anywhere in the Cartesian plane of the dose levels of the two drugs and treating cohorts of two patients simultaneously with different dose combinations. The method was further extended to account for a baseline covariate by Diniz, Kim, and Tighiouart, 2018 and settings where an unknown fraction of DLTs is attributable to one or more agents by Jimenez, Tighiouart, and Gasparini, 2019.

In contrast to the parametric models, the nonparametric models do not impose any functional form on the dose–toxicity relationship. Parametric models suffer from potential model misspecification, which may lead to unsafe dose escalation. On the other hand, nonparametric models can capture more subtle aspect of the data, hence they are more flexible. Several nonparametric models for dual agent clinical trials have been studied in the past. Lin and Yin, 2016 estimated the toxicity order of two drugs by two-dimensional isotonic regression and reduced the two-dimensional drug combination searching space into a one dimension and used a parametric CRM model based on the updated toxicity order. Mander and Sweeting, 2015 considered a product of independent Beta probabilities escalation strategy allowing the prior distributions for each dose combination to be unconstrained and imposing the monotonicity assumption when escalating by choosing only monotonic contours.

In this paper, we propose a nonparametric Bayesian method for combinations of drugs (NBCD) by modeling the joint prior probabilities of DLTs on the space of all dose combinations with independent beta distributions truncated to the set of combinations that satisfy the partial order. Unlike the PIPE algorithm proposed by Mander and Sweeting, 2015, our approach guarantees that the joint posterior distribution of the probabilities of DLT estimated with Gibbs sampler satisfies the partial order constraint. A weighted mechanism is introduced when allocating doses to successive cohorts of

patients in order to calibrate the influence of data and that of the prior.

The rest of this paper is organized as follows. In Section 3, we propose a non-parametric Bayesian model with beta priors truncated to the set of partially ordered dose combinations. In Section 4, we present our trial design and an algorithm for dose recommendation for phase II. We study the performance of our model with extensive simulations studies in Section 5.

Notation. We use n_{ij} and z_{ij} to denote the number of patient assigned to the dose combination d_{ij} and the number of DLTs observed at $(i,j), respectively$. We use $\boldsymbol{\alpha} = (\alpha_{ij}) \in \mathbb{R}^K$ and $\boldsymbol{\beta} = (\beta_{ij}) \in \mathbb{R}^K$ that collect all α_{ij} and β_{ij} parameters, respectively. For any positive integer I, we write $[I] = \{1, \ldots, I\}$. We denote by $\tilde{p}_{ij}$ and $\hat{p}_{ij}$, the prior and posterior median of the probability of DLT at dose level d_{ij}, respectively.

2. Problem formulation

Consider two drugs A and B with ordered dose levels $\mathcal{A} = \{d_i^A, i = 1, \ldots, I\}$ and $\mathcal{B} = \{d_j^B, j = 1, \ldots, J\}$, respectively. Let $\mathcal{D} = \mathcal{A} \times \mathcal{B}$ be the set of dose combinations available in the trial. We denote a typical element of $\mathcal{D}$ as $d_{ij} = (d_i^A, d_j^B)$. To each dose combination d_{ij}, there is a true probability of toxicity, or DLT, which we denote as p_{ij}^*. That is,

$$p_{ij}^* = \mathbb{P}\big(\text{Dose combination } d_{ij} \text{ causes a DLT}\big). \tag{1}$$

Later in our model, we use p_{ij} when modeling these probabilities as random.

Given a target probability of DLT $\theta \in (0,1)$, we are interested in dose combinations whose p_{ij}^* are close to θ. In order to make this notion precise, we define δ-approximate MTD as

$$\text{MTD}(\delta) := \{d_{ij} \mid (i,j) \in \mathcal{M}_\delta\} \tag{2}$$

where,

$$\mathcal{M}_\delta = \big\{(i,j) \in [I] \times [J] : \ |p_{ij}^* - \theta| \leq \delta\big\} \tag{3}$$

Our goal is to recover any dose combination in MTD(δ) for prescribed values of δ and θ. In the sequel, for simplicity, we often say dose combination (i,j) instead of d_{ij}.

2.1 The model

When modeling the drug trials, we consider the probabilities of DLT as random variables p_{ij} for which we will specify a prior. Let $\boldsymbol{p} = (p_{11}, p_{12}, p_{13}, \ldots, p_{IJ})$ be the random vector obtained by collecting all the p_{ij}s. During the trial, the patients are assigned to various dose combinations and their toxicity response is recorded. Assume that at a given stage

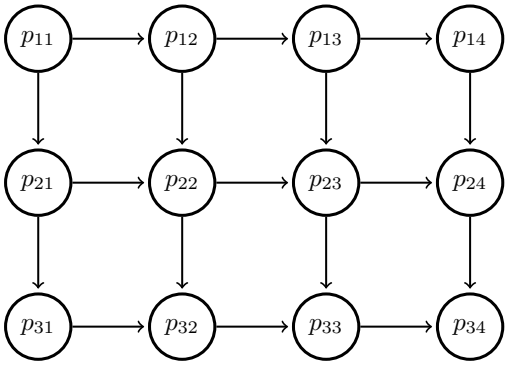

Figure 1: Hasse diagram for a 3×4 lattice

in the trial, we have assigned a total of n_{ij} patients to dose combination (i, j). The number of these patients who experienced DLT, denoted as z_{ij} is distributed as

$$z_{ij} \mid \boldsymbol{p} \sim \text{Bin}(n_{ij}, p_{ij}).$$

Letting $\mathcal{N} = \{(i, j) \mid n_{ij} > 0\}$, the likelihood of the model is

$$L(\boldsymbol{z} \mid \boldsymbol{p}) = \prod_{(i,j) \in \mathcal{N}} \binom{n_{ij}}{z_{ij}} p_{ij}^{z_{ij}} (1 - p_{ij})^{n_{ij} - z_{ij}}. \tag{4}$$

where $\boldsymbol{z} = (z_{ij})$ collects all the z_{ij}s. In Section 3.1, we specify a prior for $\boldsymbol{p}$ which allows to obtain the posterior estimate of $\boldsymbol{p}$ given $\boldsymbol{z} = (z_{ij})$. Given these estimates, we update our estimate of $\text{MTD}(\delta)$, assign more patients, and so on. Before specifying the prior on $\boldsymbol{p}$, we need to better understand the constraints on $\boldsymbol{p}$.

2.2 Lattice constraints

We assume that the dose combination are ordered so that $p_{ij} \leq p_{i'j'}$ if $i \leq i'$ and $j \leq j'$. These constraints define a partial order on the collection $\{p_{ij}, i = 1, \ldots, I, j = 1, \ldots, J\}$ which is illustrated using a Hesse diagram in Figure 1.

Recall that $\boldsymbol{p} = (p_{11}, p_{12}, p_{13}, \ldots, p_{IJ})$ collect all the probabilities of DLT. The partial ordering constraints on $\boldsymbol{p}$ can be encoded as the intersection of the following sets:

$$\begin{aligned}
\Omega_1 = \Omega_{11} &= \{\, \boldsymbol{p} \mid \; 0 \qquad\qquad\qquad\qquad\quad < p_{11} < \min(p_{12}, p_{21}) \,\}, \\
&\vdots \\
\Omega_k = \Omega_{ij} &= \{\, \boldsymbol{p} \mid \; \max(p_{i-1,j}, p_{i,j-1}) \; < p_{ij} \; < \min(p_{i,j+1}, p_{i+1,j}) \,\}, \\
&\vdots \\
\Omega_K = \Omega_{IJ} &= \{\, \boldsymbol{p} \mid \; \max(p_{I-1J}, p_{I,J-1}) < p_{IJ} < 1 \,\}.
\end{aligned} \tag{5}$$

where $K = IJ$. Note that we are using the bijection $\eta : [I] \times [J] \to [IJ]$ given by $\eta(i, j) = (i - 1)J + j$ to transform a two-dimensional indexing to a one-dimensional

index. For example, when $J = 4$ we have $\Omega_5 = \Omega_{2,1}$ and so on. We will use these two indexings interchangeably throughout the paper. In particular, we often write the elements of $\boldsymbol{p}$ in the one-dimensional index as well $\boldsymbol{p} = (p_1, p_2, \ldots, p_K)$.

The partial ordering constraints can be summarized as requiring

$$\boldsymbol{p} \in \Omega, \quad \text{where} \quad \Omega := \bigcap_{k=1}^{K} \Omega_k \subset [0,1]^K.$$

In the sequel, we refer to Ω as the lattice. We note that there is redundancy in the specification of Ω_k's in (5) in that the same constraint might be enforced by multiple Ω_k's. This redundancy is helpful in deriving the Gibbs sampler of Section 3.2.

3. Nonparametric Bayesian model for dual agents

We start by specifying our prior on $\boldsymbol{p}$ and then discuss how we can sample from the prior and the posterior.

3.1 Nonparametric prior

Perhaps the most basic nonparametric prior on $\boldsymbol{p}$ is the uniform distribution on Ω. The uniform distribution on the lattice has density

$$f_u(\boldsymbol{p}) \propto 1_\Omega(\boldsymbol{p}), \quad \text{where} \quad 1_\Omega(\boldsymbol{p}) = \prod_{k=1}^{K} 1_{\Omega_k}(\boldsymbol{p}) \tag{6}$$

is the indicator of the lattice. By the uniform distribution being nonparametric, we mean that one is not assuming a specific functional form for p_{ij} based on a lower-dimensional parameter. We can extend (6) to a model with more general marginals. Assume that we want a prior on $\boldsymbol{p}$ that is obtained as follows: Draw the coordinates of $\boldsymbol{p}$ independently with p_{ij} having density $b_{ij}(p_{ij})$, then truncate the joint distribution of $\boldsymbol{p}$ to the set Ω. The density of this prior is given by

$$f(\boldsymbol{p}) \propto 1_\Omega(\boldsymbol{p}) \prod_{i,j} b_{ij}(p_{ij}). \tag{7}$$

In this paper, we take $b_{ij}(\cdot)$ to be Beta densities:

$$b_{ij}(p_{ij}) \propto p_{ij}^{\alpha_{ij}-1}(1-p_{ij})^{\beta_{ij}-1}. \tag{8}$$

Note that we can write (7) as

$$f(\boldsymbol{p}) \propto \prod_{k=1}^{K} 1_{\Omega_k}(\boldsymbol{p}) \prod_{i,j} b_{ij}(p_{ij}) \tag{9}$$

which is a form suitable for Gibbs sampling since the lattice structure encodes local relations between elements of $\boldsymbol{p}$. Given the neighbors of a node in the lattice, its distribution is independent of the rest of the variables. In other words, $f(\boldsymbol{p})$ is a graphical model (Koller and Friedman, 2009) with the undirected lattice diagram serving as its independence graph. For future reference, we will make the following definition:

Definition 1. The *lattice-restricted Beta* distribution with shape parameters $\boldsymbol{\alpha} = (\alpha_{ij})$ and $\boldsymbol{\beta} = (\beta_{ij})$ is the multivariate distribution defined by (8) and (9). We refer to $\alpha_{ij}+\beta_{ij}$ as the (effective) *sample sizes* of the distribution (ESS).

It is worth noting that the marginals of $\boldsymbol{p}$ under a lattice-restricted Beta distribution are not Beta distributions themselves, due to the restrictions imposed by the lattice constraint. The notion of the sample size in Definition 1 is based on the common practice of referring to $\alpha+\beta$ as the effective sample size of Beta(α, β) distribution. The rationale behind this naming is well-known from the posterior inference in Beta-Binomial models; see also (15). In a simple Beta(α, β) the effective sample size $\alpha + \beta$ can directly control the variance of the distribution. For the multivariate *lattice-restricted Beta*, the relation between the sample sizes $\alpha_{ij} + \beta_{ij}$ and the variances of the components of $\boldsymbol{p}$ are much more complicated. In fact, the lattice constraint indirectly restricts how much $\alpha_{ij} + \beta_{ij}$ influences the variance, making it challenging to design diffuse priors. In Section 3.4, we propose a simple discounting scheme to work around this issue.

3.2 Gibbs sampler for the prior

It is easy to sample from the lattice-restricted distribution (9) using a Gibbs sampler. The updates are as follows: Let us derive the updates for the Gibbs sampler:

$$\begin{aligned} f(p_{11} \mid \boldsymbol{p}_{-11}) &\propto b_{11}(p_{11}) \cdot 1_{\Omega_{11}}(\boldsymbol{p}) \\ &\vdots \\ f(p_{ij} \mid \boldsymbol{p}_{-ij}) &\propto b_{ij}(p_{ij}) \cdot 1_{\Omega_{ij}}(\boldsymbol{p}) \\ &\vdots \\ f(p_{IJ} \mid \boldsymbol{p}_{-IJ}) &\propto b_{IJ}(p_{IJ}) \cdot 1_{\Omega_{IJ}}(\boldsymbol{p}) \end{aligned} \tag{10}$$

where $\boldsymbol{p}_{-ij}$ is the vector $\boldsymbol{p}$ with p_{ij} removed (i.e., all variables are included except p_{ij}). Note that although each p_{ij} also appears in some other constraint sets besides Ω_{ij}, we do not need to include them in the above conditional calculation since those constraints are also enforced by Ω_{ij}. In other words, there is some redundancy in the condition of $\Omega_1, \Omega_2 \ldots, \Omega_K$ that we have introduced to simplify deriving the Gibbs sampler.

Each conditional distribution in (10) is a truncated Beta distribution which is easy to sample from, where the truncated Beta density is defined as

$$T(x; \alpha, \beta, a, b) \;\propto\; x^{\alpha-1}(1-x)^{\beta-1} 1\{x \in (a, b)\}. \tag{11}$$

Using this notation, for example,

$$f(p_{11} \mid \boldsymbol{p}_{-11}) = T\big(p_{11};\, \alpha_{11}, \beta_{11}, 0, \min(p_{12}, p_{21})\big).$$

Thus, all these conditional distributions are easily derived and they are all truncated Beta distributions that can be simulated efficiently.

3.3 Posterior

Given prior (9) and the likelihood in (4), we can readily obtain the posterior,

$$\begin{aligned}
\pi(\boldsymbol{p} \mid \boldsymbol{z}) &\propto L(\boldsymbol{z} \mid \boldsymbol{p})\; f(\boldsymbol{p}) \\
&\propto \prod_{i,j} p_{ij}^{z_{ij}} (1-p_{ij})^{n_{ij}-z_{ij}} \prod_{k=1}^{K} 1_{\Omega_k}(\boldsymbol{p}) \prod_{i,j} b_{ij}(p_{ij}) \\
&\propto \prod_{i,j} p_{ij}^{\alpha_{ij}+z_{ij}-1} (1-p_{ij})^{\beta_{ij}+n_{ij}-z_{ij}-1} \prod_{k=1}^{K} 1_{\Omega_k}(\boldsymbol{p})
\end{aligned} \tag{12}$$

using (8). We note that the posterior is of the form

$$\pi(\boldsymbol{p} \mid \boldsymbol{z}) \propto \prod_{i,j} b'_{ij}(p_{ij}) \prod_{k=1}^{K} 1_{\Omega_k}(\boldsymbol{p}) \tag{13}$$

where $b'_{ij}(\cdot)$ is the density of Beta distribution with parameters $\alpha_{ij}+z_{ij}$ and $\beta_{ij} + n_{ij} - z_{ij}$. That is, posterior (13) is of the exact same form as (9), that is, a lattice-restricted Beta distribution, with updated parameters. Thus, the Gibbs sampler derived earlier for the prior works for the posterior as well, using the new Beta parameters.

3.4 Discounting the prior

The relatively high-dimensional prior in (9) will have reduced variances for components of $\boldsymbol{p}$, relative to those, one would expect when the components are independent. This is due to the restrictions imposed by the lattice constraints and is a challenging aspect of specifying priors in high dimensions under many constraints on the coordinates. We believe the difficulty is present as long as one insists on the coordinates satisfying strict order constraints and is not an artifact of the particular choice of the Beta densities.

At the early stages of the trial, due to the data having a small sample size and the prior having diminished variances (hence high concentration), the posterior inference will be dominated by the prior. This can be mitigated by controlling the sample size of the data relative to the (effective) sample size of the prior. To do so, we evaluate a pseudo-posterior by raising the likelihood to power $\omega > 1$ as follows:

$$\begin{aligned}
\pi(\boldsymbol{p} \mid \boldsymbol{z}) &\propto \Big(\prod_{ij} p_{ij}^{z_{ij}} (1-p_{ij})^{n_{ij}-z_{ij}} \Big)^{\omega} \prod_{ij} p_{ij}^{\alpha_{ij}-1} (1-p_{ij})^{\beta_{ij}-1} 1_{\Omega}(\boldsymbol{p}) \\
&= \prod_{ij} p_{i}^{\omega z_{ij}+\alpha_{ij}-1} (1-p_{ij})^{\omega(n_{ij}-z_{ij})+\beta_{ij}-1} 1_{\Omega}(\boldsymbol{p}).
\end{aligned} \tag{14}$$

The resulting pseudo-posterior is again an instance of a latticed-restricted Beta distribution, as in Definition 1, with shape parameters $\omega \boldsymbol{z} + \boldsymbol{\alpha}$ and $\omega(\boldsymbol{n} - \boldsymbol{z}) + \boldsymbol{\beta}$.

The idea of raising the likelihood to a power has been explored in the literature to address model misspecification (Royall and Tsou, 2003; Grünwald and Van Ommen, 2017; Bissiri, Holmes, and Walker, 2016) and to incorporate historical data in a Bayesian analysis (Ibrahim and Chen, 2000). Using this idea to simulate diffuse priors from high-dimensional concentrated priors, as we intend here, is new to the best of our knowledge. It is a natural approach for tuning the relative effects of the data and the prior on the posterior as the following simple example illustrates.

Consider the simple univariate model where $z \mid p \sim \text{Bin}(n, p)$ and $p \sim \text{Beta}(\alpha, \beta)$. The *$\omega$-reweighed pseudo-posterior* is a Beta distribution with parameters $\omega z + \alpha$ and $\omega(n - z) + \beta$, whose mean is given by

$$\mathbb{E}[p \mid z] = \frac{\omega z + \alpha}{\omega n + \alpha + \beta} =: \lambda \frac{z}{n} + (1 - \lambda) \frac{\alpha}{\alpha + \beta} \tag{15}$$

where $\lambda = \omega n / (\omega n + \alpha + \beta) \in (0, 1)$. (Note that in (15), $\alpha + \beta$ plays the same role in the prior term as does the sample size n in the data-driven term.) That is, the posterior mean is a weighted (in fact, convex) combination of the maximum likelihood estimate z/n, which is solely based on data, and the prior mean, with weights that are controlled by λ. Parameter ω allows us a degree of freedom beyond the sample size n to control the effect of the prior, effectively tuning its overall variance. In particular, the weight of the data relative to the prior is given by $\rho := \frac{\lambda}{1-\lambda} = \frac{n\omega}{\alpha+\beta}$. For a desired level of ρ, which can be thought of as a user-specified level of confidence in the prior, we can solve for the appropriate ω as

$$\omega = \frac{\rho(\alpha + \beta)}{n}. \tag{16}$$

By choosing ρ one can calibrates the relative influence of the prior and data on the posterior. When $\rho = 1$, the relative influence of the prior and the data are as given by the traditional Bayesian approaches. For a more outcome-adaptive inference, one sets ρ to be greater than 1.

A value of $\rho > 1$ is what we suggest for the high-dimensional prior we are using (Definition 1). As discussed earlier, the lattice constraint causes any prior distribution to have diminished variances. A choice of $\rho > 1$ deflates the effect of the prior, in effect simulating a more diffuse overall prior (i.e., having larger variance). Empirically, we have found that setting $\rho = 2$ significantly improves the performance. Thus, we choose $\omega = 1 + 2 \sum_{ij} (\alpha_{ij} + \beta_{ij}) / \sum_{ij} n_{ij}$ as suggested by (16) and since we need $\omega > 1$.

3.5 Choosing the hyperparameters

Let m_{ij} be the effective sample size of the Beta prior p_{ij} (Definition 1). To choose the hyperparameters, we do a grid search over different choices of values for m_{11} and m_{IJ} with the the rest of m_{ij}s being equal to a value that is less than $\min(m_{11}, m_{IJ})$. If the dose space is too large, one can limit the search space even further and assume that $m_{11} = m_{IJ}$. The grid search is done by running our gibbs sampler algorithm (10) many

times with all these different combination of hyperparameters. At the end, we choose the hyperparameters that match our prior guess of the toxicity probabilities. The details for hyperparameter selection can be found in the Appendix A.

4. Trial Design

To limit the exposure of patients to toxic combinations and provide better posterior estimation, we enroll more patients to the first two cohorts. For better exploration of the dose space, following Tighiouart, Li, and Rogatko, 2017, we enroll patients to different dose combinations in each cohort $c > 1$. However, rather than alternating each time between the vertical and horizontal direction, we choose the direction randomly.

Thus, the design of a Phase I trial for two agents using the proposed NBCD proceeds as follows:

1. The first 4 patients in the first cohort receive the minimum dose combination (d_1^A, d_1^B)
2. In the second cohort, patients 5 and 6 receive $(d_1^A, d_{j^*}^B)$, where

$$j^* = \underset{j}{\operatorname{argmin}} |\hat{p}_{1j} - \theta| \tag{17}$$

   Similarly, patient 7 and 8 receive $(d_{i^*}^A, d_1^B)$, where

$$i^* = \underset{i}{\operatorname{argmin}} |\hat{p}_{i1} - \theta| \tag{18}$$

3. In the c-th cohort ($c \geq 3$) of two patients, from each of the two dose combinations in cohort $c-1$, choose between horizontal and vertical direction randomly to fix one drug level and vary the other drug level and find the dose combination with posterior median probability of DLT closest to θ similar to (17) and (18). If the posterior median DLT probability of minimum dose of one direction or both directions is greater than $1.5 \times \theta$, choose the direction with the lowest minimum one.
4. Repeat step 3 and terminate the trial when all the patients are enrolled, or the following stopping rule holds.

Stopping rule: Stop the trial after n patients are accrued if

$$\mathbb{P}(\hat{p}_{11}^{(n)} > \theta + \gamma) > \epsilon \tag{19}$$

where, γ and ϵ are the stopping rules parameters.

4.1 Recommended phase II doses

At the end of the trial, we recommend one or more dose combinations to be used in future phase II studies. To achieve this, we first set the margins δ_l and δ_u and consider an asymmetric neighborhood $\mathcal{N}$ around θ, that is $\mathcal{N} = [\theta - l, \theta - u] \subseteq [\theta - \delta_l, \theta - \delta_u]$. We start with small l and u and gradually increase them until for some $(i, j), \hat{p}_{ij} \in \mathcal{N}$. Among these dose levels, we recommend the ones that were experimented more than once (If no dose levels are available as such, we recommend the ones that are experimented once). If no dose levels belong to $[\theta - \delta_l, \theta - \delta_u]$, then we do not recommend any doses. Algorithm 1 summarizes the steps for recommending doses for phase II.

Algorithm 1 Dose recommendation for phase II

Set δ_l, δ_u, the step sizes γ_u, γ_l, η_u and pick initial values l_0, u_0.
Set $\mathcal{I} \leftarrow \varnothing$; $l \leftarrow l_0$; $u \leftarrow u_0$
if $\left\{\frac{\sum_{ij} 1\{\hat{p}_{ij} > \theta\}}{K} \geq \frac{1}{2}\right\}$ **then**
    The scenario is toxic; Set toxic $\leftarrow 1$.
end if
while $\mathcal{I} = \varnothing$ AND $(l \leq \delta_l$ OR $u \leq \delta_u)$ **do**
    Update $\mathcal{I} \leftarrow \left\{(i, j) : -l \leq \hat{p}_{ij} - \theta \leq u\right\}$
    Update $l \leftarrow l + 1\{l \leq \delta_l\} \times \gamma_l$
    Update $u \leftarrow u + 1\{u \leq \delta_u\}\Big(\eta_u \times \text{toxic} + \gamma_u \times (1 - \text{toxic})\Big)$
end while
Return $\widehat{\text{MTD}}(\delta) = \left\{d_{ij} : (i, j) \in \mathcal{I}, n_{ij} > 1\right\}$

5. Simulations

In this section, we show the effectiveness of our proposed method in comparison with the existing methods through various simulation studies. For all the following simulations, we define the MTD as any dose combination that is within $\delta = 0.1$ of the the target probability and the MTD is estimated using Algorithm 1. We also use $\delta_l = 0.1$ and $\delta_u = 0.05$, $l_0 = 0.05$, $u_0 = 0$, $\gamma_l = \frac{\delta_l}{2}$, $\gamma_u = \frac{\delta_u}{2}$ and $\eta_u = \frac{\delta_u}{5}$, throughout. All the trials start from the lowest dose level $d_{11} = (d_1^A, d_1^B)$. We use a cohort size of 4 for the first two cohorts and a cohort size of 2 for the rest. The toxicity outcome is generated as a Bernoulli random variable that takes a value of 1 with probability p_{ij} and 0 from the corresponding scenario. For finding the median posterior, we took 11000 posterior samples and discarded 1000 burn-in iterations in Gibbs sampling procedure. To select the hyperparameters, we used a grid search as explained in section 3.5. For each scenario, 2000 simulated trials are replicated to evaluate the operating characteristics of NBCD and other methods. Specifically, we calculated the mean percentage that each dose combination was selected as the MTD(δ) at the end of the trial (rec-

ommendation percentage) and the mean percentage of patients assigned to each dose combination (experimentation percentage). For PIPE design, the (weak) prior sample size of $\frac{1}{I*J}$ was used as suggested by Mander and Sweeting, 2015. The dose escalation is done by a neighborhood constraint, with admissible doses chosen from those closest to the estimated MTD contour.

5.1 Simulation Study I

For the first simulation study, we compare the performance of NBCD with the results from Mander and Sweeting, 2015 (previously examined by Braun and Jia, 2013) under scenarios A-G that are reproduced in Table 1. These methods include the generalized CRM (gCRM) Braun and Jia, 2013, the latent contingency method by Yin and Yuan, 2009a and the coupla model of Yin and Yuan, 2009b and the product of independent beta probabilities (PIPE) Mander and Sweeting, 2015. Among all these methods, PIPE and NBCD are nonparametric and the rest are parametric models. The target toxicity probability θ is set at 0.2 and the total sample size is 50. The median prior for the probability of DLT at the smallest and largest dose combinations is set to 0.04 and 0.34, respectively to match that of scenario A. The stopping rules parameters $\gamma = 0.1$ and $\epsilon = 0.8$ are used.

The operating characteristics of the NBCD method and all the other methods are shown in Table 2, where the results from the parametric models and the PIPE method were produced from Table IV of Mander and Sweeting, 2015. In scenario D where all doses are toxic, all methods perform well in the sense that they do not recommend an MTD. Our method outperforms all the other methods in scenarios A,E,F in dosing at the target and in scenarios A,F,G in dose recommendation within 10% of the target. In particular, the percent recommendation within 10% of the target for NBCD exceeds that of PIPE by an absolute 13% and 23% under scenarios A and G, respectively. The percent of patients allocated to doses within 10% of the target is higher for PIPE relative to NBCD under scenarios A, B, and E and they are fairly close for scenarios C, F, and G. Given that the primary goal of phase I trials is estimation of the MTD, we conclude that our method is competitive with the other approaches under the scenarios studied by Braun and Jia, 2013; Mander and Sweeting, 2015.

5.2 Simulation Study II

For our second simulation study, we investigate the performance of our method under an asymmetric dose-combination space with seven 4 by 5 scenarios, see Table 3. These scenarios cover a wide range of dose–response relationships and include cases where the MTD is achieved at the highest dose combination (scenario 2), lowest dose combination (scenario 3), and more complex structures (scenarios 5 and 7). We compare our method with PIPE as it is a nonparametric model. The target toxicity probability θ is set at 0.3 and the total sample size is 50 with a cohort size of 2 for PIPE. For a fair comparison, we do not impose any early termination for NBCD and PIPE. The median prior for the probability of DLTs at the smallest and largest dose combinations was set to 0.05 and

Table 1: Dose limiting toxicity for simulation study I

Drug B	Drug A: 1	2	3	4		Drug A: 1	2	3	4
			Scenario A					Scenario E	
1	0.04	0.08	0.12	0.16		0.08	0.18	0.28	0.29
2	0.10	0.14	0.18	0.22		0.09	0.19	0.29	0.30
3	0.16	0.20	0.24	0.28		0.10	0.20	0.30	0.31
4	0.22	0.26	0.30	0.34		0.11	0.21	0.31	0.41
			Scenario B					Scenario F	
1	0.02	0.04	0.06	0.08		0.12	0.13	0.14	0.15
2	0.05	0.07	0.09	0.11		0.16	0.18	0.20	0.22
3	0.08	0.10	0.12	0.14		0.44	0.45	0.46	0.47
4	0.11	0.13	0.15	0.17		0.50	0.52	0.54	0.55
			Scenario C					Scenario G	
1	0.10	0.20	0.30	0.40		0.01	0.02	0.03	0.04
2	0.25	0.35	0.45	0.55		0.04	0.10	0.15	0.20
3	0.40	0.50	0.60	0.70		0.06	0.15	0.30	0.45
4	0.55	0.65	0.75	0.85		0.10	0.30	0.50	0.80
			Scenario D						
1	0.44	0.48	0.52	0.56					
2	0.50	0.54	0.58	0.62					
3	0.56	0.60	0.64	0.68					
4	0.62	0.66	0.70	0.74					

Table 2: Experimentation and recommendation percentages for simulation study I

		Recommendation percentages				Experimentation percentages			
Scenario	Model	At θ	1-10% of θ	$>10\%$ of θ	None	At θ	1-10% of θ	$>10\%$ of θ	None
A	gCRM	10	82	3	5	6	72	17	5
	YY09a	13	82	5	0	13	72	15	0
	YY09b	11	81	6	2	10	70	20	0
	PIPE	10	88	3	0	8	87	5	0
	NBCD	16	83	1	0	10	76	14	0
B	gCRM	0	94	3	3	0	87	13	0
	YY09a	0	99	1	0	0	86	14	0
	YY09b	0	96	4	0	0	71	29	0
	PIPE	0	83	17	0	0	82	18	0
	NBCD	0	96	4	0	0	72	28	0
C	gCRM	45	39	5	11	30	41	18	11
	YY09a	41	50	5	4	27	54	16	3
	YY09b	42	47	5	6	29	55	11	5
	PIPE	29	59	7	5	19	46	34	2
	NBCD	32	54	14	0	20	48	31	1
D	gCRM	0	0	4	96	0	0	22	78
	YY09a	0	0	1	99	0	0	20	80
	YY09b	0	0	1	99	0	0	16	84
	PIPE	0	0	1	99	0	0	37	63
	NBCD	0	0	4	96	0	0	41	59
E	gCRM	9	70	14	7	5	56	32	7
	YY09a	6	65	27	2	9	55	34	2
	YY09b	7	67	25	1	6	54	38	2
	PIPE	11	84	4	1	9	77	13	1
	NBCD	15	78	7	0	10	70	20	0
F	gCRM	13	70	6	11	10	64	16	10
	YY09a	14	76	6	4	7	75	14	4
	YY09b	12	74	7	7	7	77	9	7
	PIPE	12	75	11	2	12	69	18	2
	NBCD	16	72	11	1	14	65	20	1
G	gCRM	25	68	5	2	18	57	24	1
	YY09a	12	76	12	0	3	71	26	0
	YY09b	15	72	13	0	7	61	32	0
	PIPE	9	62	29	0	14	54	31	0
	NBCD	20	74	6	0	15	56	29	0

0.50, respectively. We then chose the hyperparameters that match our prior guess of the toxicity probabilities as discussed in section 3.5.

We used scenario 1 as the prior for both NBCD and PIPE. Table 4 shows that NBCD outperforms PIPE in all scenarios with respect to percent recommendation within 10% of the target probability of DLT θ with the highest percent equal to 20% achieved in scenarios 2 and 3. The percent of patients allocated to doses within 10% of the target are fairly similar between the two methods except for scenario 3 where NBCD allocates 23% more patients than PIPE. The safety profiles of the two methods are presented in Table 5. In general, the average percent of DLTs across all simulated trials are fairly close in all seven scenarios but PIPE tends to allocate more patients to overtoxic doses and is more likely to recommend overtoxic doses under scenarios 3 and 6. Finally, PIPE is more likely to result in a trial with an excessive rate of DLTs relative to NBCD as assessed by the percent of trials with a DLT rate more than $\theta + 0.1$. In particular, the probability that a prospective trial using PIPE will result in an excessive rate of DLTs under scenario 3 exceeds that of NBCD by 32.4%. Based on these scenarios, we conclude that NBCD approach is safer than PIPE and more efficient in recommending the MTD.

A. Details for hyperparameter selection

We choose $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ vectors to have the following forms:

$$\begin{aligned} \boldsymbol{\alpha} &= \begin{bmatrix} m(1-t) & l & \dots & l & Ms \end{bmatrix}, \\ \boldsymbol{\beta} &= \begin{bmatrix} mt & u & \dots & u & M(1-s) \end{bmatrix}. \end{aligned}$$

That is, all the elements of $\boldsymbol{\alpha}$, except the first and last, are taken to be equal, and similarly for $\boldsymbol{\beta}$. The values m, M, s, t, l and u are determined by a grid search as follows.

We vary t and s in a subinterval of $(0, 1)$ away from the boundaries, e.g., $(0.2, 0.5)$. We choose u and l such that $u + l \leq \min(m, M)$. To reduce the size of the grid search, we can fix one of u or l. For example, we can set $u = \min(m, M)/2 - l$ and let l vary in the interval $[0.2, 0.4] \times \min(m, M)$. The two parameters m and M are also varied, independently, over an interval $[\min(\alpha_0, \beta_0), \alpha_0 + \beta_0]$ where α_0 and β_0 are some heuristic prespecified values. For example, $\alpha_0 + \beta_0$ roughly specify the overall sample size of the prior. Below we will discuss a heuristic for choosing α_0 and β_0 that we have found effective in practice.

The goal of the gird search is to find a combination of hyperparameters such that the resulting prior satisfies some specified criteria for the median dose combinations. Let $\tilde{p}_{ij}$ be the prior median for the (i, j)-th does combination. Often the median for the smallest and largest dose combinations, i.e., $\tilde{p}_{11}$ and $\tilde{p}_{IJ}$, are required to match certain values and we have a range for the intermediate dose combinations. For example, the following is a possible set of criteria for a 4×4 lattice:

$$\{\tilde{p} \mid \ \tilde{p}_{11} \approx 0.04,\ \tilde{p}_{44} \approx 0.34,\ \tilde{p}_{12} < 0.1\}.$$

This set is often rewritten by setting a tolerance, say $\delta = 0.01$,

$$\{\tilde{p} \mid \ |\tilde{p}_{11} - 0.04| < \delta,\ |\tilde{p}_{44} - 0.34| < \delta,\ \tilde{p}_{12} < 0.1\}. \tag{20}$$

Table 3: Dose limiting toxicity scenarios in simulation study II

		Drug A									
		1	2	3	4	5	1	2	3	4	5
		Scenario 1					Scenario 2				
Drug B	1	0.05	0.07	0.11	0.16	0.23	0.01	0.03	0.07	0.09	0.11
	2	0.07	0.12	0.17	0.24	0.33	0.04	0.06	0.08	0.10	0.22
	3	0.12	0.18	0.25	0.33	0.43	0.09	0.13	0.22	0.25	0.27
	4	0.18	0.27	0.35	0.43	0.50	0.12	0.16	0.23	0.28	0.30
		Scenario 3					Scenario 4				
	1	0.30	0.35	0.40	0.50	0.55	0.01	0.03	0.08	0.12	0.15
	2	0.40	0.55	0.65	0.75	0.85	0.02	0.05	0.10	0.16	0.30
	3	0.50	0.60	0.70	0.80	0.90	0.07	0.09	0.15	0.25	0.35
	4	0.55	0.70	0.75	0.85	0.95	0.10	0.26	0.30	0.33	0.50
		Scenario 5					Scenario 6				
	1	0.07	0.12	0.20	0.25	0.30	0.10	0.15	0.20	0.30	0.45
	2	0.10	0.18	0.23	0.30	0.35	0.11	0.20	0.30	0.40	0.50
	3	0.30	0.48	0.56	0.65	0.68	0.15	0.30	0.35	0.50	0.60
	4	0.40	0.55	0.60	0.66	0.70	0.30	0.40	0.50	0.60	0.65
		Scenario 7									
	1	0.11	0.12	0.13	0.14	0.15					
	2	0.14	0.20	0.25	0.30	0.35					
	3	0.16	0.25	0.40	0.55	0.60					
	4	0.20	0.40	0.60	0.90	0.95					

Table 4: Experimentation and recommendation percentages for simulation study II

		Recommendation %				Experimentation %		
Scenario	Model	At θ	1-10% of θ	$>10\%$ of θ	None	At θ	1-10% of θ	$>10\%$ of θ
1	NBCD	0.0	84.0	15.5	0.0	0.0	58.1	41.9
	PIPE	0.0	75.0	25.0	0.0	0.0	60.9	39.1
2	NBCD	6.8	81.3	11.5	0.4	6.3	53.8	39.9
	PIPE	4.2	64.3	31.5	0.0	5.5	56.8	37.6
3	NBCD	24.6	60.7	10.5	4.2	26.4	49.6	24.0
	PIPE	24.0	40.5	35.4	0.1	19.4	33.3	47.3
4	NBCD	29.7	52.5	17.5	0.3	20.4	36.6	43.1
	PIPE	25.4	38.4	36.2	0.0	19.4	32.4	48.2
5	NBCD	36.7	41.4	20.5	1.4	25.3	29.5	45.2
	PIPE	36.5	39.6	23.6	0.3	22.3	29.3	48.4
6	NBCD	52.3	38.8	8.6	0.3	35.3	34.1	30.6
	PIPE	40.7	35.9	23.0	0.4	30.7	36.9	32.4
7	NBCD	16.8	72.9	9.8	0.5	9.8	56.4	33.7
	PIPE	10.0	60.6	29.2	0.0	10.2	52.9	36.9

For each combination of the hyperparameters (m, M, t, s, l), we run the Gibbs sampler and estimate the corresponding median and variance for each variable p_{ij}. We then choose the combination for which the estimated prior medians satisfy the criteria (20). If there are multiple solutions, we choose the one that maximizes the total prior variance: $\sum_{i,j} \text{var}(p_{ij})$.

The heuristic we follow for setting the range of m and M above is as follows: We consider a sequence of i.i.d. random variables $p'_{ij} \sim \text{Beta}(\alpha_0, \beta_0)$ for $i = 1, \ldots, I$ and $j = 1, \ldots, J$. We then consider the extreme order statistics of this sequence, i.e., the minimum and maximum of $\{p'_{ij}\}$. It is easy to analytically solve for α_0 and β_0 such that the median of the distributions of these two order statistics have specific values. In particular, we solve these equations to match the desired median values $\tilde{p}_{11}$ and $\tilde{p}_{IJ}$. The values α_0 and β_0 thus obtained provide a good heuristic to set the range $[\min(\alpha_0, \beta_0), \alpha_0 + \beta_0]$ for the grid search on m and M. (In fact, we conjecture that the distribution of the extreme order statistics of $\{p'_{ij}\}$ match those of the prior (9) when $\alpha_{ij} = \alpha_0$ and $\beta_{ij} = \beta_0$ for all i and j.)

As an example, the set of criteria for in the simulation study I is:

$$\{\tilde{p} \mid \ |\tilde{p}_{11} - 0.04| < 0.01, \ |\tilde{p}_{44} - 0.34| < 0.01\}.$$

We used $n_m = 15$ grid points for each of m and M, $n_t = 10$ points for each of t and s, and $n_l = 3$ points for l. Te grid search resulted in 625 solutions that satisfied the

Table 5: Trial safety evaluation for simulation study II

Design	1	2	3	4	5	6	7	Average
Recommendation % of overtoxic doses								
NBCD	4.4	0.0*	9.7	0.3	12.4	5.3	1.3	5.6
PIPE	5.6	0.0*	22.4	0.0	10.5	13.7	2.8	9.2
Allocation % of patients to overtoxic doses								
NBCD	6.0	0.0*	24.0	1.4	22.2	9.9	4.9	11.4
PIPE	8.7	0.0*	40.4	2.8	29.2	19.4	11.5	18.7
Average rate of DLTs								
NBCD	23.3	18.4	39.6	20.8	30.2	29.3	26.7	26.9
PIPE	25.8	20.2	43.7	22.8	34.2	32.6	29.5	29.8
% of trials with DLT rate $> \theta + \delta$								
NBCD	0.0	0.0	34.7	0.0	1.6	0.1	0.0	5.2
PIPE	0.1	0.0	67.1	0.0	9.7	4.4	1.1	11.8

* The average excludes the items with asterisk

criteria. Out of these, we chose the one that maximizes the total variance (i.e., the trace of the covariance matrix of $\boldsymbol{p}$) leading to the following choice of hyperparameters:

$$\begin{aligned} \boldsymbol{\alpha} &= \begin{bmatrix} 4.52 & 0.4 & \dots & 0.4 & 0.2 \end{bmatrix}, \\ \boldsymbol{\beta} &= \begin{bmatrix} 0.74 & 2.23 & \dots & 2.23 & 13.77 \end{bmatrix}. \end{aligned}$$

References

Storer, B. E. (1989). "Design and analysis of phase I clinical trials". In: *Biometrics*, pp. 925–937.

O'Quigley, J., M. Pepe, and L. Fisher (1990). "Continual reassessment method: a practical design for phase 1 clinical trials in cancer". In: *Biometrics*, pp. 33–48.

Goodman, S. N., M. L. Zahurak, and S. Piantadosi (1995). "Some practical improvements in the continual reassessment method for phase I studies". In: *Statistics in medicine* 14.11, pp. 1149–1161.

Korn, E. L. and R. Simon (1991). "Selecting dose-intense drug combinations: Metastatic breast cancer". In: *Breast cancer research and treatment* 20.3, pp. 155–166.

Møller, S. (1995). "An extension of the continual reassessment methods using a preliminary up-and-down design in a dose finding study in cancer patients, in order to investigate a greater range of doses". In: *Statistics in medicine* 14.9, pp. 911–922.

O'Quigley, J. and L. Z. Shen (1996). "Continual reassessment method: a likelihood approach". In: *Biometrics*, pp. 673–684.

Leung, D. H.-Y. and Y.-G. Wang (2002). "An extension of the continual reassessment method using decision theory". In: *Statistics in medicine* 21.1, pp. 51–63.

O'Quigley, J. and X. Paoletti (2003). "Continual reassessment method for ordered groups". In: *Biometrics* 59.2, pp. 430–440.

Iasonos, A. and J. O'Quigley (2011). "Continual reassessment and related designs in dose-finding studies". In: *Statistics in medicine* 30.17, pp. 2057–2061.

Daimon, T, S Zohar, and J O'Quigley (2011). "Posterior maximization and averaging for Bayesian working model choice in the continual reassessment method". In: *Statistics in Medicine* 30.13, pp. 1563–1573.

Liu, S., G. Yin, and Y. Yuan (2013). "Bayesian data augmentation dose finding with continual reassessment method and delayed toxicity". In: *The annals of applied statistics* 7.4, p. 1837.

Babb, J., A. Rogatko, and S. Zacks (1998). "Cancer phase I clinical trials: efficient dose escalation with overdose control". In: *Statistics in medicine* 17.10, pp. 1103–1120.

Tighiouart, M., S. Piantadosi, and A. Rogatko (2014). "Dose finding with drug combinations in cancer phase I clinical trials using conditional escalation with overdose control". In: *Statistics in medicine* 33.22, pp. 3815–3829.

Tighiouart, M. and A. Rogatko (2010). "Dose finding with escalation with overdose control (EWOC) in cancer clinical trials". In: *Statistical Science*, pp. 217–226.

Wheeler, G. M., M. J. Sweeting, and A. P. Mander (2017). "Toxicity-dependent feasibility bounds for the escalation with overdose control approach in phase I cancer trials". In: *Statistics in medicine* 36.16, pp. 2499–2513.

Ji, Y. and S.-J. Wang (2013). "Modified toxicity probability interval design: a safer and more reliable method than the 3+ 3 design for practical phase I trials". In: *Journal of Clinical Oncology* 31.14, p. 1785.

Yuan, Y. et al. (2016). "Bayesian optimal interval design: a simple and well-performing design for phase I oncology trials". In: *Clinical Cancer Research*, clincanres–0592.

Lin, R. and G. Yin (2017). "Nonparametric overdose control with late-onset toxicity in phase I clinical trials". In: *Biostatistics* 18.1, pp. 180–194.

Clertant, M. and J. O'Quigley (2017). "Semiparametric dose finding methods". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79.5, pp. 1487–1508.

Tighiouart, M., G. Cook-Wiens, and A. Rogatko (2018). "A bayesian adaptive design for cancer phase i trials using a flexible range of doses". In: *Journal of biopharmaceutical statistics* 28.3, pp. 562–574.

Thall, P. F. et al. (2003). "Dose-finding with two agents in phase I oncology trials". In: *Biometrics* 59.3, pp. 487–496.

Wang, K. and A. Ivanova (2005). "Two-dimensional dose finding in discrete dose space". In: *Biometrics* 61.1, pp. 217–222.

Yin, G. and Y. Yuan (2009a). "A latent contingency table approach to dose finding for combinations of two agents". In: *Biometrics* 65.3, pp. 866–875.

— (2009b). "Bayesian dose finding in oncology for drug combinations by copula regression". In: *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 58.2, pp. 211–224.

Wages, N. A., M. R. Conaway, and J. O'Quigley (2011). "Continual reassessment method for partial ordering". In: *Biometrics* 67.4, pp. 1555–1563.

Tighiouart, M., Q. Li, and A. Rogatko (2017). "A Bayesian adaptive design for estimating the maximum tolerated dose curve using drug combinations in cancer phase I clinical trials". In: *Statistics in medicine* 36.2, pp. 280–290.

Diniz, M. A., S. Kim, and M. Tighiouart (2018). "A Bayesian adaptive design in cancer phase I trials using dose combinations in the presence of a baseline covariate". In: *Journal of probability and statistics* 2018.

Jimenez, J. L., M. Tighiouart, and M. Gasparini (2019). "Cancer phase I trial design using drug combinations when a fraction of dose limiting toxicities is attributable to one or more agents". In: *Biometrical Journal* 61.2, pp. 319–332.

Lin, R., G. Yin, et al. (2016). "Bootstrap aggregating continual reassessment method for dose finding in drug-combination trials". In: *The Annals of Applied Statistics* 10.4, pp. 2349–2376.

Mander, A. P. and M. J. Sweeting (2015). "A product of independent beta probabilities dose escalation design for dual-agent phase I trials". In: *Statistics in medicine* 34.8, pp. 1261–1276.

Koller, D. and N. Friedman (2009). *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning*. The MIT Press.

Royall, R. and T.-S. Tsou (2003). "Interpreting statistical evidence by using imperfect models: robust adjusted likelihood functions". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 65.2, pp. 391–404.

Grünwald, P., T. Van Ommen, et al. (2017). "Inconsistency of Bayesian inference for misspecified linear models, and a proposal for repairing it". In: *Bayesian Analysis* 12.4, pp. 1069–1103.

Bissiri, P. G., C. C. Holmes, and S. G. Walker (2016). "A general framework for updating belief distributions". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 78.5, pp. 1103–1130.

Ibrahim, J. G., M.-H. Chen, et al. (2000). "Power prior distributions for regression models". In: *Statistical Science* 15.1, pp. 46–60.

Braun, T. M. and N. Jia (2013). "A generalized continual reassessment method for two-agent phase I trials". In: *Statistics in biopharmaceutical research* 5.2, pp. 105–115.