

Estimating the Design Operating Characteristics in Clinical Trials with the Ordinal Scale Disease Progression Endpoint

Shirin Golchi¹

¹Department of Epidemiology and Biostatistics, McGill University

Abstract

Bayesian adaptive designs have gained popularity in all phases of clinical trials in the recent years. The COVID-19 pandemic, however, has brought these designs to the centre stage. The need for establishing evidence for the effectiveness of vaccines, therapeutic treatments and policies that could resolve or control the crisis has resulted in development of efficient designs for clinical trials that can be concluded with smaller sample sizes in a shorter time period. Design of Bayesian adaptive trials, however, requires extensive simulation studies that is considered a disadvantage in time-sensitive settings such as the pandemic. In this paper, we propose a set of methods for efficient estimation and uncertainty quantification for the design operating characteristics of Bayesian adaptive trials. The proposed approach is tailored to address design of clinical trials with the ordinal disease progression scale endpoint but can be used generally in the clinical trials context where design operating characteristics cannot be obtained analytically.

Keywords: Bayesian test statistic; Constrained design; COVID-19; Ordinal scale outcome; Proportional odds model; Sampling distribution; Trial simulation

1 Introduction

Bayesian adaptive designs are flexible alternatives to conventional fixed size randomized clinical trial (RCT). In adaptive designs stopping or adaptation decisions may be made at interim analyses according to accumulating evidence which can result in sample size and budget savings. In addition, participants benefit from an increased chance of receiving a beneficial treatment ([Berry et al., 2010](#)).

Various stopping rules may be defined and incorporated into a Bayesian adaptive trial design with multiple interim analysis. Common decision rules include stopping the trial for efficacy or futility, or eliminating an arm due to small probability of effectiveness for the corresponding treatment. Stopping rules are most commonly defined based on a Bayesian test statistic driven from the posterior distribution of model parameters, for example the posterior probability of effectiveness (alternative hypothesis). Efficacy or futility are then defined based on high and low thresholds for the posterior probability of effectiveness (PPE), respectively.

Despite the growing use of Bayesian methods in clinical trials, proposed designs are primarily assessed by regulatory agencies based on their frequentist operating characteristics such as power and false positive rate (FPR). The Food and Drug Administration (FDA), for example, emphasizes the importance of simulation studies for evaluation of operating characteristics in adaptive trials of drugs and biologics in a guideline document prepared for industry ([FDA, 2019](#)).

Similar to the critical region of a frequentist test statistic, the probability thresholds for PPE can be specified to achieve satisfactory design operating characteristics (DOC). Power, for instance, is defined as the probability of observing a high PPE given the assumed effect size. Therefore, the sampling distribution of the PPE (or any other Bayesian test statistic) is required to obtain power. Unlike in classical hypothesis testing framework, however, the sampling distribution of a Bayesian probability statement is not available in closed form. Therefore, design operating characteristics are estimated through Monte Carlo, i.e., by simulating the trial and sampling distribution of PPE for a given set of parameters.

Simulation studies for design of Bayesian adaptive trial can be time consuming since the combination of plausible ranges of model parameters (effect size and baseline measure assumptions) together with possible design parameters (efficacy and futility thresholds, sample sizes, frequency of interim looks, etc.) can result in a large number of simulation scenarios. In many cases, a complex trial design and/or model without analytically tractable posteriors involve a significant amount of computation for a single trial simulation which is multiplied by the number of simulation scenarios as well as the number of iterations.

Therefore, despite the appeal of Bayesian adaptive designs for COVID-19 trials, time consuming simulation studies were considered a major hurdle given the urgency of the situation. A common endpoint used in this context is the ordinal scale of disease progression defined by the World Health Organization (WHO) ([WHO, 2020](#)) with the goal of evaluating the effect of therapeutics in reducing the severity of COVID-19. The Proportional Odds (PO) ordinal logistic regression can then be used for the analysis ([Harrell and Lindsay, 2020](#)). Bayesian inference for the PO model, however, requires MCMC sampling since the analytic form of the posterior is not available. Therefore, assessment of DOC for Bayesian trials with the ordinal scale outcome requires extensive simulations with significant computational burden.

In this paper, we propose a set of methods to overcome the computational hurdles for evaluation of DOC in design of Bayesian adaptive trials. We focus on the PO model for the ordinal scale outcome in COVID-19 trials and explain the specifics of the proposed approach in this context. However, the methodology can be employed in general for design of clinical trials where power analysis and DOC assessment relies on simulations rather than analytic forms for the sampling distribution of the test statistic. The proposed approach relies on estimating the sampling distribution of PPE (the test statistic) through the model parameter space which is then used to provide fast estimates of DOC for a wide range of assumptions and design parameters without the need to run additional simulations. A simple parametric density estimation approach is employed where the sampling distribution of PPE is assumed to follow a Beta distribution whose parameters are modelled as Gaussian processes (GP) with distance-based correlation across the parameter space. The GPs are trained over an initial set of simulated distributions for the test statistic at select set of parameter values.

One of the challenging aspects of the problem is to define an initial set of points over the constrained space that represents the probabilities associated with the levels of the ordinal scale outcome, i.e., the assumed risk associated with each level of disease severity. The vector of probabilities must add up to one, resulting in a simplex. Considering the evolving situation during the COVID-19 pandemic, DOC assessment for a wide range of risk assumptions is specifically important for COVID-19 trials.

In GP-based models, satisfactory prediction performance throughout the input (parameter) space, is achieved by spreading the initial set of target function evaluations uniformly using a “space-filling design” (O’Hagan, 1978; McKay et al., 1979; Johnson et al., 1990; Morris and Mitchell, 1995; Sacks et al., 1989). Space-filling designs over non-rectangular and constrained spaces has recently received considerable attention (Lin et al., 2010; Lekivetz and Jones, 2014; Mak and Joseph, 2018). Generating a space-filling design on a simplex is, however, not straightforward. We use the method of Lekivetz and Jones (2014) which is to first generate a covering sample of the target space, cluster the sample points and select the cluster centroids as the design points. We propose a sampling technique for generating the initial sample on the simplex as an efficient and effective alternative to simple Monte Carlo.

The novelty of the present manuscript is in proposing a set of methodology for the design of Bayesian adaptive trials beyond simple Monte Carlo simulations for the assessment of DOC. Currently, estimation of operating characteristics for a given set of parameters relies solely on simulations and is not accompanied with adequate uncertainty quantification. The present work sets the foundation for development of methods that facilitate the use of Bayesian measures for decision-making in clinical trials while satisfying traditional requirements.

The remaining of the article is organized as follows – the problem setting including the PO model and a generic Bayesian adaptive trial that employs the PO model is described in Section 2. In Section 3, the emulation and prediction of the sampling distribution of the Bayesian test statistic using Gaussian processes including the design of the initial simulation set is described. A simulation study is provided in Section 4 to assess the performance of the proposed approach in compare to trial simulation in a simple sitting. We present the results

Table 1: Ordinal scale disease severity endpoint levels and descriptions

Patient State	Descriptor	Level
Uninfected	Uninfected; no viral RNA detected	0
Ambulatory mild disease	Asymptomatic; viral RNA detected Symptomatic; independent Symptomatic; assistance needed	1
Hospitalised: moderate disease	Hospitalised; no oxygen therapy Hospitalised; oxygen by mask or nasal prongs	2
Hospitalised: severe disease	Hospitalised: oxygen by NIV or high flow Intubation and mechanical ventilation	3
Dead	Dead	4

of applying the proposed methodology to the PO model example in Section 5. Section 6 follows with a discussion of the limitations of the proposed approach together with potential extensions.

2 The ordinal scale endpoint and the proportional odds model

The clinical progression scale proposed by the WHO is a recommended endpoint for measuring the therapeutic response to viral infection (WHO, 2020). The proposed ordinal scale has ten disease progression level from uninfected to dead. The challenge in use of a 10-level endpoint in clinical trials is lack of granular data to inform the base rate assumptions as well as the dimensionality of the parameter space in statistical inference. Therefore, a simplification of the 10-level ordinal scale is one that merges levels to reduce the definition of the endpoint to include uninfected, mild disease, moderate disease, severe disease and dead (Table 1). Considering that being infected is an inclusion criteria in most clinical trials of therapeutics, we focus on a 4-level ordinal scale endpoint.

An infected patient’s outcome can then be assumed to follow a multinomial distribution,

$$Y \sim \text{Multinom}(1, \mathbf{p} = (p_1, \dots, p_4))$$

where $p_k, k = 1, \dots, 4$ is the risk associated with the k^{th} level of disease severity and $\sum_{k=1}^4 p_k = 1$. The proportional odds model is defined as a logistic regression of the tail probability,

$$P(Y \geq k | A) = \frac{1}{1 + \exp[-(\alpha_k + \beta A)]}$$

where

$$\alpha_k = -\log \frac{\sum_{i=1}^k p_i}{1 - \sum_{i=1}^k p_i}$$

and A is the treatment assignment indicator, i.e., $A = 1$ indicates that the patient has received the treatment and $A = 0$ indicates assignment to the control arm. The ratio of the odds of disease severity is then used to represent the effect of the treatment,

$$OR = \frac{P(Y \geq k | A = 1)}{P(Y \geq k | A = 0)} = \exp(\beta).$$

We will be focusing on \mathbf{p} and OR and will refer to them as the *model parameters* in the remaining of the manuscript.

Under the Bayesian framework prior distributions are specified for the α_k and β , and the inference is based on the posterior distribution of these parameters given the observed data \mathbf{y} . This posterior distribution is analytically intractable and needs to be approximated or estimated by Monte Carlo sampling. [James \(2020\)](#) developed an R package that samples from the posterior distribution of the PO model parameters using Hamiltonian Monte Carlo implemented in Stan. The posterior distribution of \mathbf{p} and OR is obtained respectively as transformations of posterior samples of α_k and β .

Consider a Bayesian adaptive trial where adaptation and stopping decisions as well as inference rely on the posterior PPE given the accumulated data.

$$\pi(\mathbf{y}) = P(OR < 1 | \mathbf{y}). \tag{1}$$

We will be denoting the PPE as π hereafter for the sake of brevity. At any interim analysis, the trial may be stopped for efficacy if $\pi > U$ or for futility if $\pi < \ell$ where U and ℓ are pre-specified upper and lower probability thresholds. The design operating characteristics

are therefore defined as probability of a superiority or futility conclusion under the null and alternative hypotheses. For example, false positive rate is

$$P_{\mathbf{y}}(\pi > U \mid OR_T \geq 1),$$

where OR_T is the “true” odds ratio. The DOC, therefore, depend on the sampling distribution of the test statistic π which is not available in closed form and is estimated using simulations in design of Bayesian adaptive trials.

When the evaluation of π for every simulated set of \mathbf{y} requires MCMC sampling, as in the PO model, these simulation studies are computationally intensive and time consuming. Especially for estimation of type I error rates a large number of simulations is required to obtain reasonably precise estimates of the probability of a “rare event”. See [Harrell and Lindsell \(2020\)](#) for a comprehensive source on design of Bayesian adaptive trials with the ordinal scale outcome. In the following section we propose a set of methods to predict the sampling distribution of π based on an initial number of simulations across the space spanned by the model parameters.

3 Estimating the data distribution of PPE

As described above, the sampling distribution of π across the parameter space is the key to assessment of DOC. Let us denote the parameter space by $\Theta = \mathcal{P} \times \mathcal{O}$, where \mathcal{P} is the 4-dimensional simplex of all plausible values for the vector of probabilities \mathbf{p} and \mathcal{O} is the interval of OR values that can be realistically assumed as the effect of intervention.

Denoting the sampling distribution of π for a given $\theta \in \Theta$ by $f(\pi \mid \theta)$ we assign

$$f(\pi \mid \theta) := \text{Beta}(a(\theta), b(\theta)) \tag{2}$$

and

$$a(\theta) \sim \mathcal{GP}(\mu_a, \rho_a(\theta, \theta')), \quad b(\theta) \sim \mathcal{GP}(\mu_b, \rho_b(\theta, \theta')) \tag{3}$$

where $\mathcal{GP}(\mu, \rho(\theta, \theta'))$ denotes a Gaussian process (GP) with a constant mean μ and a covariance function $\rho(\theta, \theta')$. The subscripts show that the mean and covariance parameters of the GP prior are trained independently for the shape and scale parameters of the beta distribution in (2).

Note that parameters of the Beta distribution are supposed to be positive and specifying a GP prior that assigns non-zero probabilities to negative values is problematic. This issue is addressed by rejection sampling from the posterior predictive distribution. A variety of methods have been proposed for fitting Gaussian process models to observations from constrained functions (Riihimäki and Vehtari, 2010; Lin and Dunson, 2014; Golchi et al., 2015; Wang and Berger, 2016). A common approach, referred to as “warped Gaussian process” (Snelson et al., 2004) is to transform the observations using a monotonic function mapping them onto the real line. However, transformations result in a non-Gaussian process prior over the original function which can exhibit undesirable properties. For example, the log transformation resulting in a log-Gaussian process prior over the original function whose prediction variance is significantly larger than that obtained by a GP (Diggle and Ribeiro Jr., 2007).

The next step is to simulate $f(\pi | \theta)$ at a set of training points $(\theta_1, \dots, \theta_n)$, where n is the size of the training set. A common recommendation in the GP literature is the use of a space-filling design to select the training set uniformly spread across the input space (O’Hagan, 1978; McKay et al., 1979; Johnson et al., 1990; Morris and Mitchell, 1995; Sacks et al., 1989). The challenge here is that the parameter space is a simplex since \mathcal{P} is defined by the following constraints,

$$\sum_{k=1}^4 p_k = 1,$$

and

$$l_k < p_k < u_k$$

where l_k and u_k determine the realistic range of values for the risk associated with each category.

Most space-filling design algorithms either assume a rectangular input space or rely on an

initial sample covering the target space which is not straightforward to achieve with simple Monte Carlo. In the following section we describe the sampling method used to generate an initial set of points covering the simplex that defines \mathcal{P} . Then, we will employ the method of [Lekivetz and Jones \(2014\)](#) that groups the covering sample over the target space by K-mean clustering and selects the cluster centroids as the design points. The design over \mathcal{P} is then combined with a set of equally distanced OR values to provide the training and test sets over the parameter space Θ .

3.1 Space-filling design over \mathcal{P}

Consider the constrained space $\mathcal{P} \in [0, 1]^4$. Generating a uniform sample over \mathcal{P} is equivalent to sampling from the following distribution,

$$\pi^T(\mathbf{p}) = \frac{\mathcal{U}(\mathbf{p})\mathbf{1}_{\mathcal{P}}(\mathbf{p})}{\int_{\mathcal{P}} \mathcal{U}(\mathbf{p})d\mathbf{p}}, \quad (4)$$

where $\mathcal{U}(\cdot)$ is the density function of a uniform distribution with the domain $[0, 1]^4$ and $\mathbf{1}_{\mathcal{P}}(\mathbf{p})$ is an indicator function that is equal to 1 if $\mathbf{p} \in \mathcal{P}$ and is zero otherwise.

We use the sequentially constrained Monte Carlo (SCMC) sampling approach proposed by [Golchi and Campbell \(2016\)](#) to sample from (4). Specifically, we define the deviation of a given point \mathbf{p} from the constraints that define the design region \mathcal{P} by

$$C_{\mathcal{P}}(\mathbf{p}) = (|\sum_{k=1}^4 p_k - 1|, \{p_k - u_k\}_{k=1}^4, \{l_k - p_k\}_{k=1}^4).$$

For a point $\mathbf{p} \in \mathcal{P}$ the first component of the deviation vector is zero and the rest are negative. For a point $\mathbf{p} \notin \mathcal{P}$, $C_{\mathcal{P}}(\mathbf{p})$ measures the deviation from each of the constraints.

A probabilistic version of the constraint indicator, $\mathbf{1}_{\mathcal{P}}(\mathbf{p})$, can then be defined as,

$$\prod \Phi(-\tau C_{\mathcal{P}}(\mathbf{p})),$$

where Φ is the standard normal cumulative distribution function and the parameter τ con-

trols the slope of the probit function. The number of terms in the above product is equal to the number of constraints defined by $C_{\mathcal{P}}(\mathbf{p})$ which would be nine in the present example. We have,

$$\lim_{\tau \rightarrow \infty} \prod \Phi(-\tau C_{\mathcal{P}}(\mathbf{p})) = \mathbf{1}_{\mathcal{P}}(\mathbf{p}).$$

Starting from a uniform sample of size N on the unit hypercube, $[0, 1]^4$, the goal is to filter this sample through a sequence of increasingly constrained densities towards the simplex, \mathcal{P} . The SMC sequence of densities is defined as,

$$\pi^t(\mathbf{p}) \propto \mathcal{U}(\mathbf{p}) \prod \Phi(-\tau_t C_{\mathcal{P}}(\mathbf{p})), \quad t = 0, \dots, T$$

$$0 = \tau_0 < \tau_1 < \dots < \tau_T \rightarrow \infty,$$

An effective sequence of constraint parameters τ_t is achieved adaptively (Jasra et al., 2011). At each step the next value in the sequence is determined such that the effective sample size (ESS) does not fall below a certain threshold, e.g. half the sample size, $\frac{N}{2}$. This is done by numerically solving the following equation for τ_t ,

$$\text{ESS} = \frac{\left(\sum_{n=1}^N w_n^t(\tau_t)\right)^2}{\sum_{n=1}^N (w_n^t(\tau_t))^2} = \frac{N}{2},$$

where

$$w_n^t(\tau_t) = \frac{\prod \Phi(-\tau_t C_{\mathcal{P}}(\mathbf{p}_n^{t-1}))}{\prod \Phi(-\tau_{t-1} C_{\mathcal{P}}(\mathbf{p}_n^{t-1}))}.$$

The target value for the constraint parameter is 10^6 .

Figure 1 shows the two-dimensional marginals of the sample generated over \mathcal{P} using the SMC sampling scheme. The lower and upper limits for \mathbf{p} are arbitrarily specified as $\mathbf{l} = (0.5, 0.05, 0.01, 0.005)$ and $\mathbf{u} = (0.9, 0.30, 0.05, 0.025)$, respectively. Note that \mathbf{l} and \mathbf{u} do not have to belong to \mathcal{P} ; i.e., the limit vectors do not need to satisfy the constraint, $\sum_{k=1}^4 p_k = 1$.

Following the generating of a covering sample over the target space we use the method proposed by Lekivetz and Jones (2014) which relies on K-mean clustering of the sample




Figure 1: Two-dimensional marginals of the sample generated over the simplex, \mathcal{P} .

and selecting the cluster centroids as the design points. This sampling-based approach aims to generate a design in which no two points are too close to each other while not relying on a distance measure such as the Euclidean distance that does not realistically represent distances on a subspace created by the constraints. It is worth noting that there have been various other methods developed within the field of design of experiments that can be employed here, e.g. [Welch et al. \(1996\)](#); [Draguljić et al. \(2012\)](#); [Joseph and Gul \(2015\)](#). The reason behind our choice is simplicity of implementation.

3.2 Gaussian process fit and prediction

The GP prior in (3) is based on the assumption that parameters of the target distribution $f(\pi | \theta)$ are smooth in Θ . This, in turn, will result in power predictions that are also smooth in Θ . The covariance function of the GP processes in (3) are selected as the squared exponential covariance functions that are one of the most common choices for the covariance function

in GP modelling [Rasmussen and Williams \(2006\)](#). This choice of the covariance function assumes infinite differentiability with respect to θ . This may prove an unrealistic assumption in some applications resulting in convergence issues in estimation of GP parameters. We have not encountered such issues in this example.

To illustrate the proposed method and assess the quality of DOC estimates that are obtained from the proposed approach for estimating the sampling distribution of the Bayesian test statistic, we consider a simple but common scenario within the clinical trials framework where the goal is to evaluate the effect of a treatment in reducing (or increasing) the odds of an event. In such a setting, a Beta-binomial model may be used where the event risk is assigned beta priors under each arm. Together with a binomial likelihood the posterior distribution for the probability of event is obtained as a Beta distribution. The posterior distribution for OR is respectively obtained by drawing samples from the risk posterior distributions under each arm.

In the conjugate framework, obtaining the test statistic that is derived from the posterior does not require MCMC sampling for each simulated trial. The analytic posterior results in much more efficient simulations of the sampling distribution. Therefore, we can simulate the trial and estimate the power over a wide range of parameter values. We define a fine grid of size 100 over the parameter space given by the Cartesian product of $(0.25, 0.7)$ (for the base event risk p_0) and $(0.65, 1)$ (for the OR). We then simulate the sampling distribution of the PPE by simple Monte Carlo for each of the (p_0, OR) pairs over this grid. The power is obtained for each set of parameter values as the upper 5% tail of the sampling distribution, estimated as the 95% sample quantile.

The above-mentioned grid is then treated as the test set. A training set is generated similarly as a coarse grid of size 20. Beta distributions are fit to the simulated sampling distribution of the probability of effectiveness over the training set and their corresponding parameter values are used to train GP models that predict the parameters of the sampling distribution over the test set. Power estimates together with 95% credible intervals are obtained respectively and compared with the power estimates that were previously given through simple Monte Carlo.

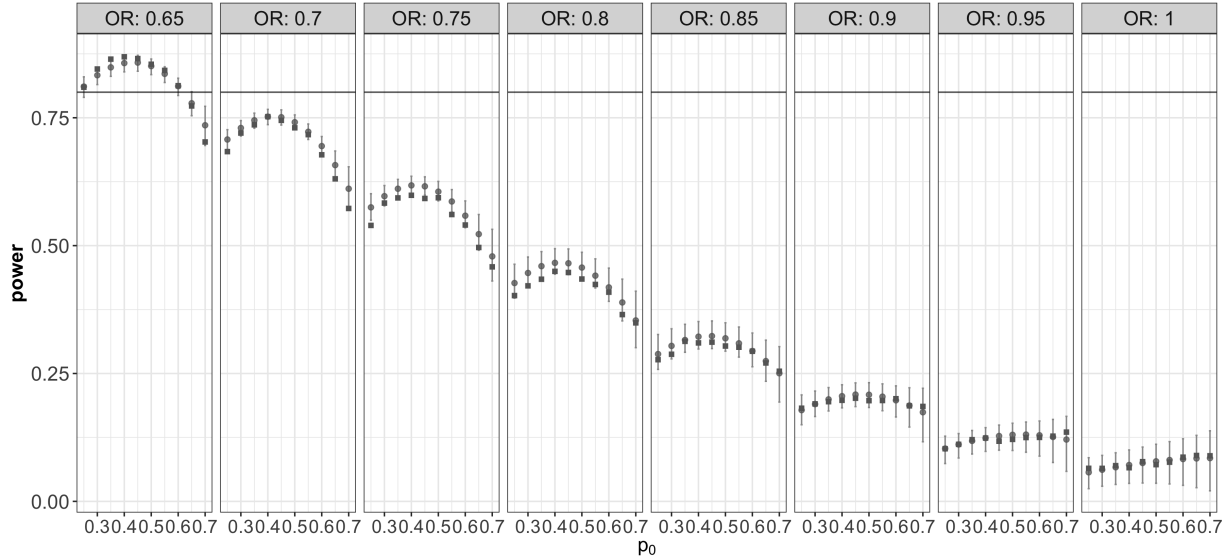


Figure 2: Point estimates for power (probability of concluding superiority) and 95% credible intervals (round dots with error bars). The square dots are estimates of power obtained by trial simulation. The horizontal line is drawn at 80% power.

Figure 2 shows the results for this example. The square dots show the power obtained from 1000 trial simulations at every set of parameter values. The round dots show the power estimated using the proposed approach based on the simulated sampling distribution within the training set. And the vertical error bars show the 95% credible intervals. Note that “power” when $OR = 1$ represents the false positive rate. The horizontal line is drawn at 80% power that is most often targeted in trial design.

The results presented in Figure 2 are reassuring with near perfect coverage and point estimates obtained from the surrogate model closely following the simulated power estimates. However, the performance of GP-based surrogate models heavily rely on the training set design which, as discussed above, cannot be defined as a grid over the input space in the PO model example. For a more realistic assessment of the proposed approach a simulation study that captures the design uncertainty is described in the following section.

4 Simulation study

We continue working with the binary endpoint example of the previous section in the simulation study. One hundred training sets are generated using the methods described in Section 3.1. Specifically, at every iteration, 100 points are generated over the two dimensional parameter space $\Theta = (0.25, 0.7) \times (0.6, 1)$ and a training set of size 20 is constructed using the clustering method of [Lekivetz and Jones \(2014\)](#). The trial is simulated over these 20 points and GP models are trained over the outputs obtained for the 20 sets of parameter values.

The test set for the simulation study is defined as the same grid used in Figure 2. The simulation outputs are therefore defined to capture the performance of the proposed method in predicting power over the test set. We compute root mean squared error (RMSE), bias and posterior standard deviation (PSD) for any point θ in the test set and for each of the 100 training sets, as follows:

$$\text{RMSE} = \sqrt{\frac{1}{K} \sum_{k=1}^K (\phi_k - \phi_t)^2}$$

where $\phi_k = P(\pi > 0.95; a_k(\theta), b_k(\theta))$ is the estimate of power obtained as the upper tail of a beta distribution with parameters given as the k^{th} posterior samples $a_k(\theta)$ and $b_k(\theta)$; and ϕ_t is the “true” power obtained by trial simulation at θ ;

$$\text{bias} = \hat{\phi} - \phi_t,$$

where the posterior mean is used as the point estimate of power, $\hat{\phi} = \frac{1}{K} \sum_{k=1}^K (\phi_k)$;

$$\text{PSD} = \sqrt{\frac{1}{K} \sum_{k=1}^K (\phi_k - \hat{\phi})^2}.$$

Figures 3-5 show the above measures averaged over the 100 training sets throughout the parameter space. The RMSE is on average below 4.5% and only increases for extreme values




Figure 3: Root mean squared error averaged over 100 training sets over the input space.

of the base risk, p_0 (small effect size and low power). Similarly, averaged over the design, absolute bias does not exceed 0.04, the largest bias occurs at the two ends of the range of p_0 values. As for posterior standard error, the highest precision is achieved where the effect size is relatively large $OR < 0.8$ with the posterior variance increasing under the null hypotheses and for small/large values of p_0 .

5 Application to the ordinal scale endpoint and the PO model

In this section we implement the proposed methods to assess the DOC for a clinical trial with the ordinal scale outcome where the PO model is used for analysis as described in Section 2. We consider a simple trial design with sample size 1000 and a range of assumptions for the baseline risk probabilities \mathbf{P} and OR to be explored. Estimation of DOC for more complex designs with multiple interim analyses and adaptive features is performed similarly. The design complexities will only need to be implemented into the initial simulations and do not have any implications on the prediction methods thereafter.




Figure 4: Bias averaged over 100 training sets over the input space.




Figure 5: posterior standard deviation averaged over 100 training sets over the input space.

The goal is to predict the probability that

$$P(OR < 1 \mid \mathbf{y}) > U$$

over the input space $\Theta = \mathcal{P} \times \mathcal{O}$ where

$$\mathcal{P} = \left\{ \mathbf{p} : \sum_{k=1}^4 p_k = 1, 0.5 < p_1 < 0.9, 0.05 < p_2 < 0.3, 0.01 < p_3 < 0.05, \right. \\ \left. 0.005 < p_4 < 0.025 \right\}$$

For the training set, a space-filling design of size 20 is generated over \mathcal{P} using the methods described in Section 3.1. The Cartesian product of this set and a grid of size 4 over $\mathcal{O} = (0.7, 1)$ is used as the final design of size 80 over Θ . Trial data is then simulated from the PO model described in Section 2 for each of the (\mathbf{P}, OR) pairs in the training set. For every simulated trial the probability of superiority of the treatment π (given in (1)) is obtained to estimate the sampling distribution of π . A Beta density function is fit to the Monte Carlo samples of π at every $\theta = (\mathbf{P}, OR)$ to obtain estimates of $a(\theta)$ and $b(\theta)$ and GPs are fit to these simulated values.

The sampling distribution of π is then predicted over a test set of size 400 generated over Θ in the same fashion as the training set. Figure 6 shows the predicted power derived from the distribution of π together with 95% credible intervals over a two-dimensional subspace of Θ , i.e., $p_1 \times OR$. Note that the 95% credible intervals represent two layers of uncertainty, i.e., the observation (Monte Carlo) error of the initial power evaluations and the uncertainty associated with the density estimation step to obtain the sampling distribution of the test statistic.

As mentioned earlier the strength of the proposed approach is in obtaining the sampling distribution of the test statistic rather than focusing on a single measure. This means that once the model is trained, one can explore a variety of decision thresholds, U , for concluding superiority in

$$P(\pi > U \mid \theta_T = \theta^*)$$

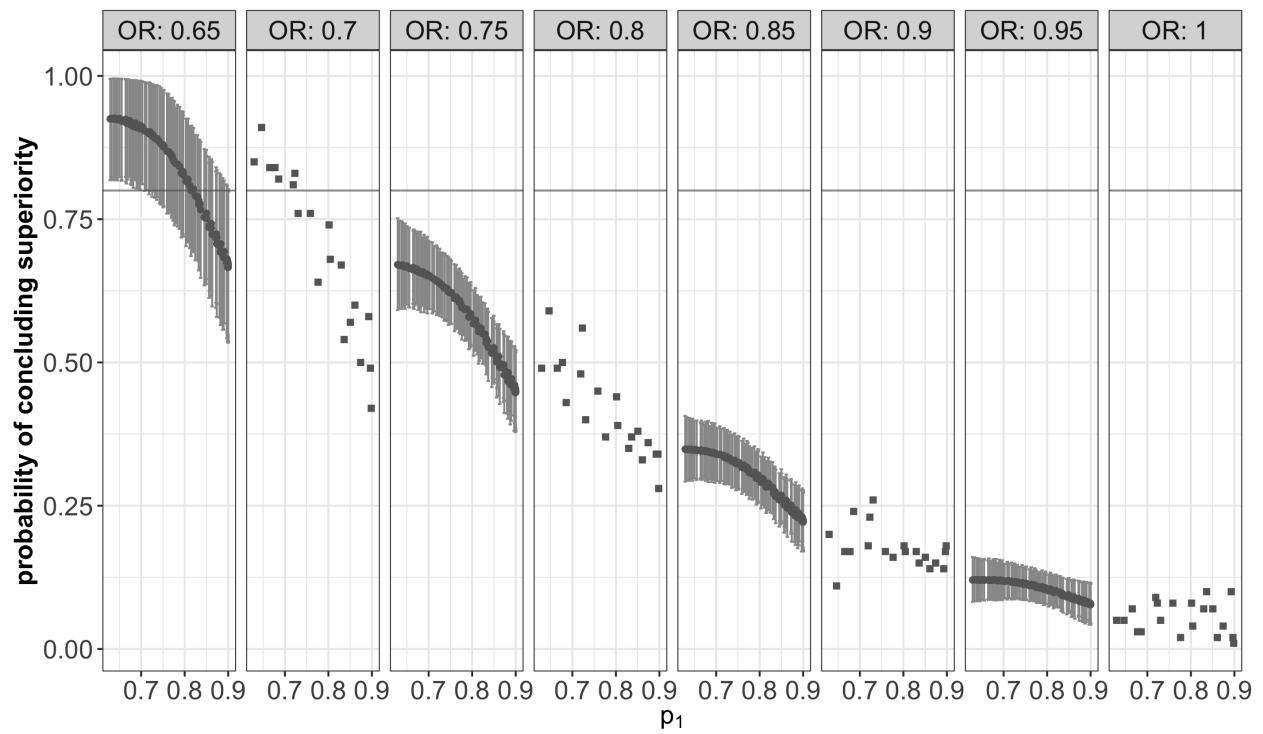


Figure 6: Point estimates of probability of concluding superiority (grey round dots) with 95% credible intervals. The square dots show probability of concluding superiority (power) obtained from trial simulation for an initial set of parameter values.

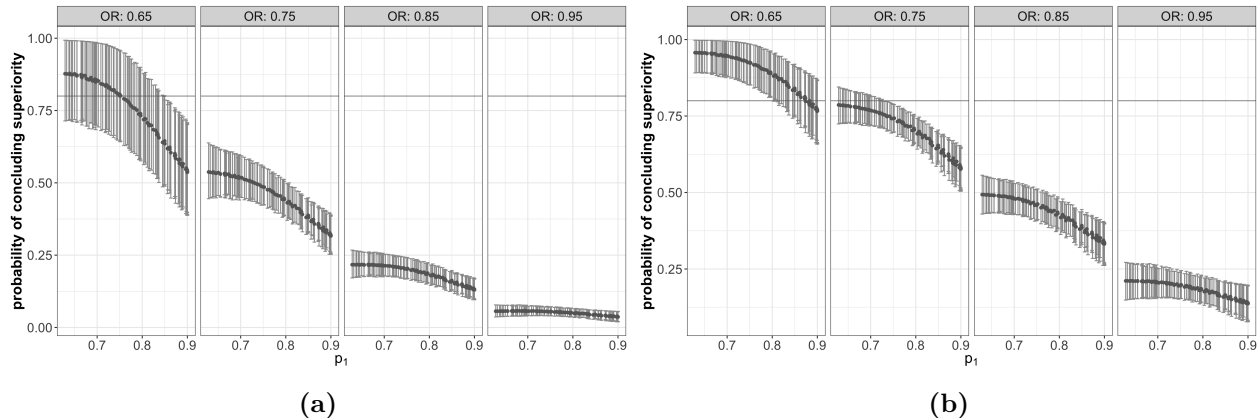


Figure 7: Point estimates and 95% credible intervals for the probability of concluding superiority with decision thresholds (a) 0.98 and (b) 0.9.

where θ^* represents a given set of parameter values.

To showcase this feature, Figure 7 shows the same results as in Figure 6 - excluding the initial set of observations - for different decision thresholds $U = 0.98$ and $U = 0.9$; these values correspond to more and less conservative decision rules in compare to $U = 0.95$, respectively.

Likewise, we may explore other operating characteristics such as probability of concluding futility according to the following decision rule for various decision thresholds ℓ ,

$$P(\pi < \ell \mid OR).$$

Figure 8 shows the probabilities of concluding futility for a restricted range of parameter values since this probability is negligible for larger effect sizes with the specified design settings.

For example, if these results correspond to considering an interim/final analysis with $N = 1000$ patients data, a superiority decision threshold of $U = 0.98$ and futility threshold of $\ell = 0.05$ there would be 42-55% chance of stopping the trial for superiority assuming a base risk vector of $\mathbf{p} = (0.75, 0.22, 0.01, 0.02)$ and $OR = 0.75$; such a design controls the false positive rate (at this specific point in the trial) at 4-6% and allows for up to 20% chance of stopping the trial for futility over negligible and non-existent effect sizes.




Figure 8: Point estimates and 95% credible intervals for the probability of concluding futility with decision thresholds (a) 0.01 and (b) 0.05.

6 Discussion

In this manuscript we have proposed a set of methods for estimating the sampling distribution of a Bayesian test statistic across the model parameter space. The goal is to estimate a variety of design operating characteristics in an efficient manner in the context of adaptive clinical trials. We take advantage of the spatial correlation throughout the model parameter space to “interpolate” the parameters of the sampling distribution. The PPE distribution parameters are modelled as independent Gaussian processes trained over a set of observations obtained by simulating the trial design over an initial set of model parameter values.

The proposed methods are specifically tailored to the proportional odds model used for the ordinal scale disease progression endpoint that has been commonly used in COVID-19 clinical trials. The vector of base risks associated with the levels of the ordinal outcome is defined as a simplex contained within the unit hypercube. This results in unique challenges in the design of the initial set of simulations used for training the GPs. We have therefore described a specialized design algorithm for this problem.

To our knowledge, the proposed approach is the first method ever proposed for estimating the DOC in Bayesian adaptive designs beyond simple Monte Carlo simulations. Main advantages are enabling exploration of a variety of operating characteristics together with adequate uncertainty quantification. This manuscript sets the foundation for development

of methods that facilitate the adoption of Bayesian measures for decision-making in clinical trials. As a starting point for this line of work, we acknowledge that it has limitations and there remains much room for improvements and further developments. Example directions for future work are laid out as follows.

One of the challenging features of the problem is the fast changing form of the sampling distribution of the test statistic in certain parts of the parameter space. The Gaussian process is therefore not the most appropriate model for the parameters of this distribution as they assume stationarity throughout the input space. Increased bias in certain parts of the model parameter space is a result of this issue. A variety of methods have been proposed to model a non-stationary response surface (Schmidt and O’Hagan, 2003; Gramacy and Lee, 2008; Gramacy and Apley, 2013; Heinonen et al., 2016). In most cases, however, the solution comes at the cost of losing the analytic form of the predictions or uncertainty estimates as well as increased computational burden. These increased complexities may therefore defeat the purpose in the current application. In the examples explored in the present manuscript the resulting bias is small enough that we consider the conventional GP fit satisfactory.

Another aspect of the proposed methodology that may be improved is incorporating the constraints into the GP fit rather than correcting for them as a post processing step in sampling from the posterior predictive distribution. Note that both the shape and scale parameters of the beta distribution used as the sampling distribution of the test statistic are bounded to be positive and are monotonic with respect to at least one of the model parameters (Figures 9a and 9b). As mentioned earlier, various techniques exist for incorporating constraints and boundary information into GP models (Tan, 2018; Solin and Kok, 2019; Ding et al., 2019). However, similar to the argument made above, when employing these methods one must assure that the additional computational complexity and potential loss of generality is justified by the gain in estimation accuracy and precision.

Another restriction of the proposed approach is the parametric form (beta distribution) assumed for the Bayesian test statistic. The beta distribution (or any other parametric distribution) alone does not realistically capture the actual sampling distribution throughout the model parameter space. Taking a non-parametric modelling approach is therefore a

potential future direction for this work.

Code

The code for the simulation study as well as the implementation of the methods for the ordinal scale endpoint and the PO model are provided in a public repository on GitHub, <https://github.com/sgolchi/DOCest>.


References

- Berry, S. M., Carlin, B. P., Lee, J. J., and Müller, P. (2010). *Bayesian Adaptive Methods for Clinical Trials*. CRC Press.
- Diggle, P. J. and Ribeiro Jr., P. J. (2007). *Model-based Geostatistics*. Springer.
- Ding, L., Mak, S., and Wu, C. (2019). Bdrygp: a new Gaussian process model for incorporating boundary information. *arXiv:1908.08868 [stat.ME]*.
- Draguljić, D., Santner, T. J., and Dean, A. M. (2012). Non-collapsing space-filling designs for bounded non-rectangular regions. *Technometrics*, 54:169–178.
- FDA (2019). *Adaptive Design Clinical Trials for Drugs and Biologics Guidance for Industry*. U.S. Department of Health and Human Services, Food and Drug Administration, Center for Drug Evaluation and Research (CDER), Center for Biologics Evaluation and Research (CBER), <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/adaptive-design-clinical-trials-drugs-and-biologics>.
- Golchi, S., Bingham, D. R., Chipman, H., and Campbell, D. A. (2015). Monotone emulation of computer experiments. *SIAM/ASA Journal on Uncertainty Quantification*, 3(1):370–392.
- Golchi, S. and Campbell, D. A. (2016). Sequentially constrained Monte Carlo. *Computational Statistics and Data Analysis*, 97:98 – 113.


- Gramacy, R. B. and Apley, D. W. (2013). Local Gaussian process approximation for large computer experiments. *Journal of Computational and Graphical Statistics*, 24(2):561–578.
- Gramacy, R. B. and Lee, H. K. H. (2008). Bayesian treed Gaussian process models with an application to computer modeling. *Journal of the American Statistical Association*, 103(483):1119–1130.
- Harrell, F. and Lindsell, C. (2020). Statistical design and analysis plan for sequential parallel-group RCT for COVID-19. <https://hbiostat.org/proj/covid19/bayesplan.html>.
- Heinonen, M., Mannerström, H., Rousu, J., Kaski, S., and Lähdesmäki, H. (2016). Non-stationary Gaussian process regression with Hamiltonian Monte Carlo. *in Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS) 2016, Cadiz, Spain*.
- James, N. (2020). Bayesian cumulative probability models - bayescpm. *GitHub repository*, <https://github.com/ntjames/bayesCPM/tree/master/pkg>.
- Jasra, A., Stephens, D. A., and Doucet, A. (2011). Inference for Lévy-driven stochastic volatility models via adaptive sequential Monte Carlo. *Scandinavian Journal of Statistics*, 38(1-22).
- Johnson, M. E., Moore, L. M., and Ylvisaker, D. (1990). Minimax and maximin distance designs. *Journal of Statistical Planning and Inference*, 26:131–148.
- Joseph, V. R. and Gul, E. (2015). Maximum projection designs for computer experiments. *Biometrika*, pages 1–10.
- Lekivetz, R. and Jones, B. (2014). Fast flexible space-filling designs for non-rectangular regions. *Quality and Reliability Engineering*, 31.
- Lin, D. K. J., Sharpe, C., and Winker, P. (2010). Optimized U-type designs on flexible regions. *computational Statistics and Data Analysis*, 54(6):1505–1515.

- Lin, L. and Dunson, D. B. (2014). Bayesian monotone regression using Gaussian process projection. *Biometrika*, 101:303–317.
- Maceachern, S. (1999). Dependent nonparametric processes. *Proceedings of the Section on Bayesian Statistical Science, American Statistical Association*, pages 50–55.
- Maceachern, S. (2000). Dependent Dirichlet processes. Technical report, Technical Report, Department of Statistics, The Ohio State University.
- Mak, S. and Joseph, V. R. (2018). Minimax and minimax projection designs using clustering. *Journal of Computational and Graphical Statistics*, 27(1):166–178.
- McKay, M. D., Beckman, R. J., and Conover, W. J. (1979). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21:239–245.
- Morris, M. D. and Mitchell, T. (1995). Exploratory designs for computational experiments. *Journal of Statistical Planning and Inference*, 43:381–402.
- O’Hagan, A. (1978). Curve fitting and optimal design for predictions. *Journal of the Royal Statistical Society Series B*, 40:1–42.
- Quintana, F. A., Müller, P., Jara, A., and MacEachern, S. N. (2020). The dependent Dirichlet process and related models. *arXiv:2007.06129 [stat.ME]*.
- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. MIT Press.
- Riihimäki, J. and Vehtari, A. (2010). Gaussian processes with monotonicity information. In *in Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 645–652.
- Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P. (1989). Design and analysis of computer experiments (with discussion). *Statistical Science*, 4:409–435.

- Schmidt, A. and O’Hagan, A. (2003). Bayesian inference for nonstationary spatial covariance structures via spatial deformations. *Journal of the Royal Statistical Society Series B*, 65:743–758.
- Snelson, E., Rasmussen, C. E., and Ghahramani, Z. (2004). Warped Gaussian processes. In Thrun, S., Saul, L., and Schölkopf, B., editors, *Advances in Neural Information Processing System*, volume 16. MIT Press.
- Solin, A. and Kok, M. (2019). Know your boundaries: Constraining Gaussian processes by variational harmonic features. *In Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89.*, 89.
- Tan, M. H. Y. (2018). Gaussian process modelling with boundary information. *Statistica Sinica*, 28:621–648.
- Wang, X. and Berger, J. O. (2016). Estimating shape constrained functions using Gaussian processes. *SIAM/ASA Journal on Uncertainty Quantification*, 4:1–25.
- Welch, W. J., Buck, R. J., Sacks, J., Wynn, H. P., Morris, M. D., and Schonlau, M. (1996). Response to James M. Lucas. *Technometrics*, 38(199-203).
- WHO (2020). A minimal common outcome measure set for COVID-19 clinical research. *The Lancet*, 20(8):E192 – E197.



(a)



(b)

Figure 9: Point estimates and 95% credible intervals for the (a) shape and (b) scale parameters of the sampling distribution of the posterior probability of superiority