MAP-curvature: a model-free approach for analysis of dose-finding trials

Linxi Han[a], Qiqi Deng[b], Zhangyi He[c,1,*], Feng Yu[a,*]

[a] *School of Mathematics, University of Bristol, Bristol, UK*
[b] *Biostatistics, ModernaTX Inc, Cambridge, Connecticut, USA*
[c] *Department of Epidemiology and Biostatistics, School of Public Health, Imperial College London, London W2 1PG, UK*

Abstract

MCP-Mod (Multiple Comparison Procedure-Modelling) is an efficient statistical method for the analysis of Phase II dose-finding trials, although it requires specialised expertise to pre-specify plausible candidate models along with model parameters. This can be problematic given limited knowledge of the agent/compound being studied. Misspecification of candidate models and model parameters can severely degrade its performance. To circumvent this challenge, in this work, we introduce MAP-curvature, a Bayesian model-free approach for the detection of the dose-response signal in Phase II dose-finding trials. MAP-curvature is built upon a Bayesian hierarchical method incorporating information about the total curvature of the dose-response curve. Through extensive simulations, we show that MAP-curvature has comparable performance to MCP-Mod if the true underlying dose-response model is included in the candidate model set of MCP-Mod. Otherwise, MAP-curvature can achieve performance superior to that of MCP-Mod, especially when the true dose-response model drastically deviates from candidate models in MCP-Mod.

Keywords: Dose finding, Dose-response signal detection, Minimum effective dose estimation, Model-free approach, Bayesian hierarchical model, Curvature prior

1. Introduction

Characterising the dose-response relationship and finding the right dose are important but challenging in the pharmaceutical drug development process. Nearly half of all failures in Phase III trials result in part from a lack of understanding of the dose-response relationship in Phase II

*Corresponding author.
Email addresses: z.he@beatson.gla.ac.uk (Zhangyi He), feng.yu@bristol.ac.uk (Feng Yu)
[1] Present address: Cancer Research UK Beatson Institute, Glasgow G61 1BD, UK

trials (Sacks et al., 2014). Over the last decade, the Multiple Comparison Procedure-Modelling (MCP-Mod) method, developed by Bretz et al. (2005), has been increasingly popular for Phase II trials as it can provide superior statistical evidence for dose selection.

MCP-Mod is a two-step approach that combines MCP principles and modelling techniques. In the first MCP step, it establishes a dose-response signal (Proof of Concept, PoC). In the second Mod step, it estimates the dose-response curve and target doses of interest. It overcomes shortcomings of traditional approaches for dose-finding studies (see *e.g.*, Ting, 2006, for an excellent introduction). MCP-Mod requires the pre-specification of plausible candidate models and model parameters to capture model uncertainty, which however is mainly based on the limited knowledge of the agent/compound being studied, if it is available at all (Chen & Liu, 2020). Misspecification of candidate models and model parameters in MCP-Mod may cause a loss in power and unreliable model selection (see Saha & Brannath, 2019, and references therein).

Motivated by these shortcomings, in this work, we develop a model-free Bayesian approach for the detection of the dose-response trend and estimation of the dose-response relationship in Phase II trials. We introduce a novel Bayesian hierarchical method incorporating the total (in the L^2 sense) curvature of the dose-response curve as a prior parameter. Our approach avoids the requirement of a set of pre-specified candidate models. The responses at the given set of doses are estimated through maximum *a posteriori* (MAP), with which we construct a test statistic to establish PoC through simulations. We can then estimate the dose-response relationship using simple interpolation.

The remainder of this work is organised as follows. In Section 2, we describe in detail our MAP approach with a curvature prior, abbreviated MAP-curvature. In Section 3, we assess the operating characteristics of MAP-curvature through simulations, and compare its performance to that of MCP-Mod. We present concluding remarks and future directions in Section 4.

2. Methods

We consider a trial with a total of $M+1$ distinct doses $x_0, x_1, \ldots, x_M$, where x_0 represents placebo. Let N_i be the number of patients in dose group i. We assume the true dose response function is $f(x)$ at dose x. We assume $f(x)$ is defined on the interval $[0, 1]$ in what follows

unless otherwise specified. For $i = 0, 1, \ldots, M$ and $j = 1, 2, \ldots, N_i$, we let Y_{ij} be the response observed for patient j allocated to dose x_i. We assume

$$Y_{ij} = \mu_i + \epsilon_{ij}, \tag{1}$$

where $\mu_i = f(x_i)$ denotes the mean response at dose x_i, and $\epsilon_{ij} \overset{iid}{\sim} N(0, \sigma^2)$ denotes the error term for patient j in dose group i. As in MCP-Mod (Bretz et al., 2005), we assume σ is known. Detailed discussion of the assumption on the standard deviation σ can be found in Fleischer et al. (2022).

In MCP-Mod, a set of plausible candidate model curves are specified. This constrains the possible set of dose-response curves. However, with limited knowledge of the agent/compounds in the trial, it is possible to mis-specify this candidate set of models. The procedure for specifying models is also somewhat cumbersome. We would like to avoid the pre-specification of possible models beforehand, but still want to impose a degree of smoothness to the dose-response curve. To this end, we introduce the L^2-total curvature

$$Sf = \left(\int_{x_0}^{x_M} f''(x)^2 \, dx \right)^{1/2}$$

to measure how far the dose-response curve $f(x)$ is from being a straight line. We will impose a half-normal $HN(\gamma^2)$ prior on Sf to give low prior probabilities to the dose-response that is very curved, where the standard deviation γ controls the trade-off between the L^2-total curvature of $f(x)$ and fidelity to data $\boldsymbol{Y}$. Given the dose-response function $f(x)$ being available at $M + 1$ distinct doses, for $i = 1, 2, \ldots, M - 1$, we have

$$f''(x_i) \approx 2 \left(\frac{\mu_{i+1} - \mu_i}{(x_{i+1} - x_i)(x_{i+1} - x_{i-1})} - \frac{\mu_i - \mu_{i-1}}{(x_i - x_{i-1})(x_{i+1} - x_{i-1})} \right)$$

through the second-order central difference scheme (Burden et al., 2015), therefore the L^2-total curvature Sf being approximated through numerical integration with

$$S_{\boldsymbol{\mu}} = 2 \left(\sum_{i=1}^{M-1} \left(\frac{\mu_{i+1} - \mu_i}{(x_{i+1} - x_i)(x_{i+1} - x_{i-1})} - \frac{\mu_i - \mu_{i-1}}{(x_i - x_{i-1})(x_{i+1} - x_{i-1})} \right)^2 \Delta x_i \right)^{1/2},$$

where $\Delta x_i = (x_{i+1} - x_{i-1})/2$ for $i = 2, 3, \ldots, M - 2$ with $\Delta x_1 = (x_2 + x_1)/2 - x_0$ and $\Delta x_{M-1} =$

$x_M - (x_{M-1} + x_{M-2})/2$.

We let

$$\boldsymbol{Y} = \{Y_{ij}, i = 0, 1, \ldots, M, j = 1, 2, \ldots, N_i\}$$

and define our Bayesian hierarchical model to be

$$p(\boldsymbol{\mu}, \gamma \mid \boldsymbol{Y}) \propto p(\boldsymbol{\mu}, \gamma) p(\boldsymbol{Y} \mid \boldsymbol{\mu}), \tag{2}$$

where the prior

$$p(\boldsymbol{\mu}, \gamma) = p(\gamma) \prod_{i=0}^{M} p(\mu_i) p(S_{\boldsymbol{\mu}} \mid \gamma)$$

with

$$S_{\boldsymbol{\mu}} \mid \gamma \sim HN(\gamma^2),\ \gamma \sim HN(\tau^2),\ \mu_i \overset{iid}{\sim} U(0,1) \text{ for } i = 0, 1, \ldots, M,$$

and the likelihood

$$p(\boldsymbol{Y} \mid \boldsymbol{\mu}) = \prod_{i=0}^{M} \prod_{j=1}^{N_i} p(Y_{ij} \mid \mu_i).$$

Suppose the standard deviation τ in the hyperprior for γ is pre-specified. The MAP estimates of all parameters in the model defined in Eq. (2) can be obtained by maximising the corresponding log-likelihood

$$\hat{\boldsymbol{\mu}}, \hat{\gamma} = \arg\max_{\boldsymbol{\mu}, \gamma} \left\{ -\sum_{i=0}^{M} \sum_{j=1}^{N_i} \left(\frac{Y_{ij} - \mu_i}{\sigma} \right)^2 - \log \gamma^2 - \left(\frac{S_{\boldsymbol{\mu}}}{\gamma} \right)^2 - \left(\frac{\gamma}{\tau} \right)^2 \right\}$$

through a numerical optimisation algorithm like the Broyden–Fletcher–Goldfarb–Shanno (BFGS) method and its variants (see, *e.g.*, Nocedal & Wright, 1999, for more details).

To establish PoC, we propose a test statistic

$$T = \max\{\hat{\mu}_1 - \hat{\mu}_0, \hat{\mu}_2 - \hat{\mu}_0, \ldots, \hat{\mu}_M - \hat{\mu}_0\}$$

with hypotheses

$$H_0 : \hat{\mu}_0 = \hat{\mu}_1 = \ldots = \hat{\mu}_M$$

$$H_1 : \max\{\hat{\mu}_1, \hat{\mu}_2, \ldots, \hat{\mu}_M\} > \hat{\mu}_0.$$

We define a significance level α for a dose-response signal, such that the corresponding critical value c satisfies

$$\mathbb{P}(T > c \mid H_0) < \alpha.$$

For a given α, we compute the critical value c via simulation

$$\mathbb{P}(T > c \mid H_0) \approx \frac{1}{R} \sum_{r=1}^{R} \mathbb{1}_{\{T^{(r)} > c\}},$$

where $T^{(r)}$ is the test statistic computed from the data $\boldsymbol{Y}^{(r)}$ simulated under the null hypothesis H_0, R is the total number of replicates, and $\mathbb{1}_A$ is the indicator function equal to 1 if condition A holds and 0 otherwise.

Finally, the mean response estimates $\hat{\boldsymbol{\mu}}$ can be linearly interpolated (or using a more sophisticated interpolation scheme) to obtain an estimate of the dose-response curve $f(x)$, which will also yield an estimate of the target dose of interest. In this work, we only perform simple linear interpolation as we shall focus on the PoC stage of dose-finding trials.

3. Simulations

In this section, we assess the performance of MAP-curvature, in terms of power to detect dose-response signals as well as estimates of the dose-response curve and target doses of interest. We compare MAP-curvature to MCP-Mod.

3.1. Simulation settings

We simulate randomised, double-blind, parallel-group trials with patients being equally allocated to placebo (0) or one of four active doses (0.15, 0.50, 0.80 and 1). We take the placebo response rate to be 0% and the maximum treatment effect to be 50%, respectively. We vary the sample size per dose group in $\{10, 20, 30, 40, 50, 60\}$. We choose one of 12 common dose-response shapes to be the true dose-response model. These models are plotted in Figure 1, with corresponding parameters summarised in Supplemental Material, Table S1. We simulate each patient's response according to Eq. (1) with a standard deviation of $\sigma = 1$. For each of 72 combinations of parameters, consisting of sample size and dose-response shape, we run 10,000 simulated trials.

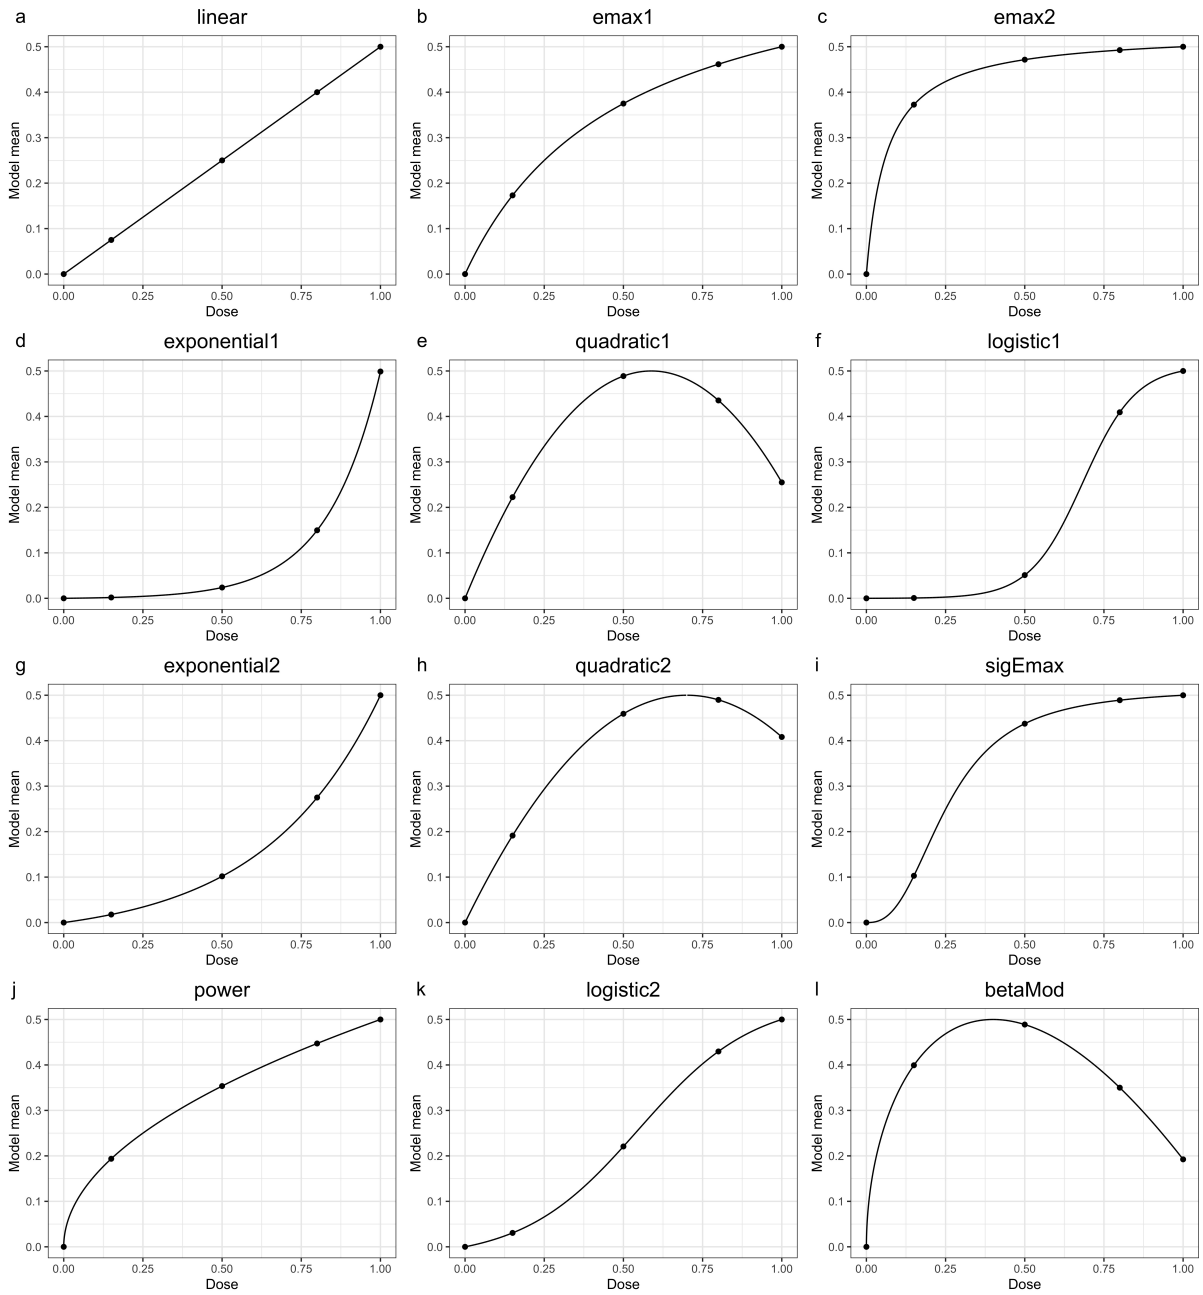

Figure 1: Dose-response shapes selected for the true dose-response model.

For each simulated trial, we run MAP-curvature to establish PoC and find the target doses of interest with a standard deviation of $\tau \in \{1, 3, 5\}$ for the hyperprior $HN(\tau^2)$. To benchmark MAP-curvature against MCP-Mod, we run MCP-Mod for each simulated trial using the **DoseFinding** package (version 1.0-2) in R, where we specify a fixed set of candidate models made up of linear, emax1, emax2, exponential1, quadratic and logistic1 (top six figures of Figure 1).

3.2. Simulation results

To compare the performance of MAP-curvature and MCP-Mod, in Figure 2, we plot receiver operating characteristic (ROC) curves for all true dose-response models for sample size 40. The ROC curve is produced by plotting the true positive rate against the false positive rate across a range of critical values. The closer the ROC curve approaches the top left corner, the better the method performs overall. ROC curves for all other sample sizes are illustrated in Supplemental Material, Figures S1-S5. We also summarise their powers at 5% and 10% type I error rates in Supplemental Material, Table S2 and S3, respectively.

We see from Figure 2 that MAP-curvature has better performance when the true dose-response relationship is not dramatically curved, e.g. linear, emax1, logistic1, logistic2 and sigEmax, achieving over 80% power at 5% type I error rate. Choosing an appropriate τ can further improve the performance of MAP-curvature. More specifically, MAP-curvature achieves better performance with a larger τ for the true dose-response curve that is more curved, e.g. quadratic1 and beta models. Otherwise, e.g. in logistic1 and power models, a smaller τ performs better. How to select an appropriate τ in practice will be discussed in Section 4.

To compare the performance of MAP-curvature and MCP-Mod, we divide into two cases:

1. the true dose-response curve is one of the candidate models in MCP-Mod,
2. the true dose-response curve is not one of the candidate models in MCP-Mod.

Figures 2(a-f) compare the performance of MAP-curvature and MCP-Mod for the first case, which is a fairly rare situation in practice. The resulting ROC curves are in favour of MCP-Mod as expected, but MAP-curvature with $\tau = 3$ achieves comparable performance. We note that for the true dose-response curve that is not dramatically curved, e.g. linear, emax1 and logistic1, $\tau = 1$ yields better performance than when $\tau = 3$, giving a power gain of around 2-9% for MAP-curvature with respect to MCP-Mod at 5% type I error rate.

Figures 2(g-l) compare the performance of MAP-curvature and MCP-Mod for the second case. This is the situation we expect to encounter in practice. We see that with $\tau = 3$, MAP-curvature uniformly outperforms MCP-Mod, especially when the true dose-response model drastically deviates from the candidate model set in MCP-Mod such as sigEmax and beta models. In the latter case, MAP-curvature has a power gain of approximately 5-10% compared to MCP-Mod at a type I error rate of 5%. Some true dose-response models such as exponential2,

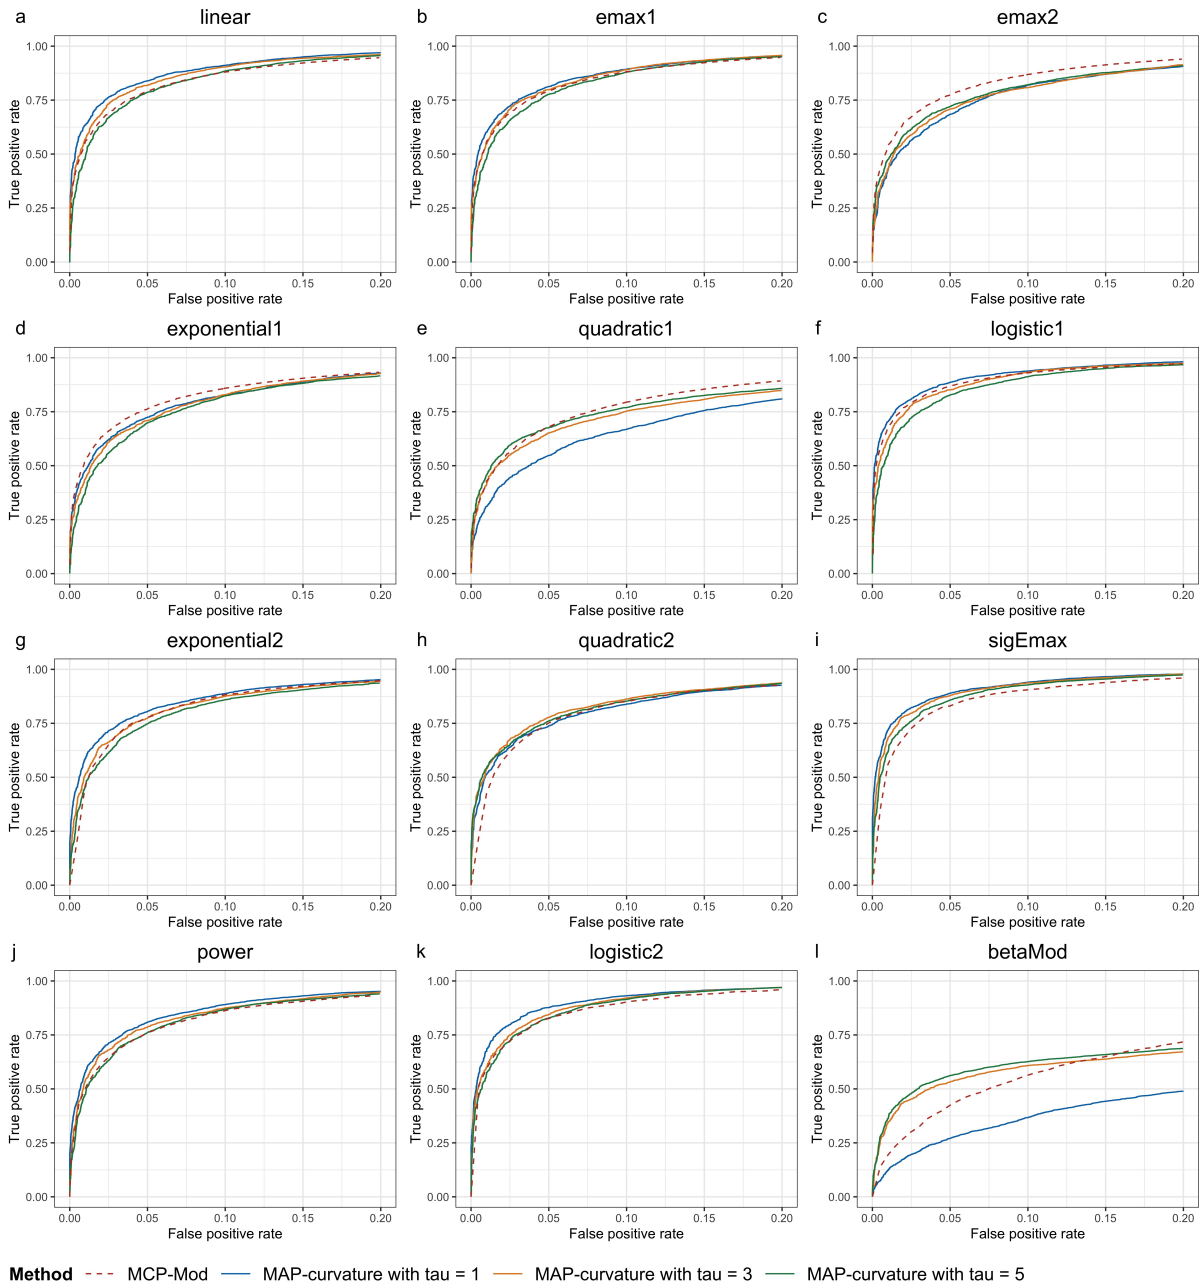

Figure 2: ROC curves of MAP-curvature and MCP-Mod across different true dose-response models with the sample size of 40 patients per arm. The ROC curves of MCP-Mod in (a)-(f) are produced with the true dose-response model not included in the candidate model set, and the ROC curves of MCP-Mod in (g)-(l) are produced with the true dose-response model not included in the candidate model set.

quadratic2 and power models, are captured by the candidate model set of MCP-Mod sufficiently well that the performance of MAP-curvature and MCP-Mod is similar. With an appropriate τ, however, MAP-curvature is still able to achieve significantly better performance, e.g. a power gained by around 2-3% for MAP-curvature over MCP-Mod at 5% type I error rate.

We also assess and compare the performance in estimating the dose-response curve and tar-

get doses of interest, known as the minimum effective dose (MED). See Supplemental Material, Figures S6-S11 for the dose-response curves estimated using MAP-curvature and MCP-Mod for various true dose-response models and sample sizes. The corresponding results for MED estimation are summarised in Supplemental Material, Table S4.

4. Discussion

In this work, we have introduced MAP-curvature, a novel Bayesian approach for establishing PoC and estimating the dose-response curve alongside target doses of interest in Phase II trials. MAP-curvature is "model-free", in the sense that it does not require pre-specification of candidate model curves, which can influence the performance of MCP-Mod. It is built upon a Bayesian hierarchical model incorporating prior information on the L^2-total curvature of the dose-response curve. Through extensive simulations, we have shown that MAP-curvature has performance comparable to that of MCP-Mod in establishing PoC and estimating MED when the candidate model set of MCP-Mod includes the true dose-response model. When the true dose-response model deviates from the candidate model set of MCP-Mod, we have shown that MAP-curvature outperforms MCP-Mod.

To achieve optimal performance, MCP-Mod requires specialised expertise to pre-specify plausible candidate models and model parameters, but the knowledge of the agent/compounds being studied is commonly limited. Compared to MCP-Mod, the only requirement for pre-specification in MAP-curvature is the standard deviation τ for the hyperprior $\gamma \sim HN(\tau^2)$, which encodes prior knowledge of how far the dose-response curve is from a straight line. Based on additional simulations with varying values of the standard deviation τ (see Supplemental Material, Figure S12 and Table S5), we recommend choosing

1. $\tau \in [2, 4]$ if our prior knowledge is poor,
2. $\tau < 2$ if we are confident that the curvature of the dose-response curve is weak,
3. $\tau > 4$ if we are confident that the curvature of the dose-response curve is strong.

A number of relevant issues for MAP-curvature deserve further research. The current version of MAP-curvature is limited to analysing Phase II dose-finding trials with continuous endpoints, *i.e.*, a single normally distributed homoscedastic response measured at the end of the trial for each patient. To expand the applicability of MAP-curvature, extensions to other common

types of endpoints (*e.g.*, binary, counts and survival endpoints) and trials (*e.g.*, longitudinal dose-finding trials) require further investigation. Other directions of future research include the investigation of trial designs tailored to MAP-curvature and development of statistical software implementing MAP-curvature.

Acknowledgements

This work was carried out using the computational facilities of the Advanced Computing Research Centre, University of Bristol - `http://www.bristol.ac.uk/acrc/`.

Disclosure Statement

The authors report there are no competing interests to declare.

Disclaimer

The opinions expressed in this paper are solely those of the authors and not those of their affiliations. The authors' affiliations do not guarantee the accuracy or reliability of the information provided herein.

References

Bretz, F., Pinheiro, J. C., & Branson, M. (2005). Combining multiple comparisons and modeling techniques in dose-response studies. *Biometrics*, *61*, 738–748.

Burden, R. L., Faires, J. D., & Burden, A. M. (2015). *Numerical Analysis*. Boston: Cengage Learning.

Chen, J., & Liu, T. (2020). Statistical considerations on implementing the MCP-Mod method for binary endpoints in clinical trials. *Contemporary Clinical Trials Communications*, *19*, 100641.

Fleischer, F., Bossert, S., Deng, Q., Loley, C., & Gierse, J. (2022). Bayesian MCPMod. *Pharmaceutical Statistics*, *21*, 654–670.

Nocedal, J., & Wright, S. J. (1999). *Numerical Optimization*. New York: Springer.

Sacks, L. V., Shamsuddin, H. H., Yasinskaya, Y. I., Bouri, K., Lanthier, M. L. et al. (2014). Scientific and regulatory reasons for delay and denial of FDA approval of initial applications for new drugs, 2000-2012. *JAMA*, *311*, 378–384.

Saha, S., & Brannath, W. (2019). Comparison of different approaches for dose response analysis. *Biometrical Journal*, *61*, 83–100.

Ting, N. (2006). *Dose Finding in Drug Development*. New York: Springer.