

MCPModPack package: Technical Manual (Version 0.3)

1	Introduction	1
2	MCPMod methodology	1
3	Hypothesis testing	3
4	Dose-response modeling	6

1 Introduction

This technical manual provides a detailed description of statistical methods implemented in the MCPModPack package. This package implements the Multiple Comparison-Modeling (MCPMod) methodology in dose-finding trials with continuous, binary and count endpoints.

To learn more about the MCPModPack package, visit its page on the CRAN web site

<https://CRAN.R-project.org/package=MCPModPack>

The most current version of the package is available at

<https://github.com/medianainc/MCPModPack>

2 MCPMod methodology

The MCPMod method is an efficient dose-finding method, which combines frequentist tests of dose-response signals (using multiple-contrast tests) with parametric dose-response modeling approaches.

This method was introduced in Bretz, Pinheiro and Branson (2005) and was later expanded in multiple directions. For a detailed discussion of the general MCPMod methodology and related approaches, see Pinheiro, Bretz and Branson (2006), Bretz, Tamhane and Pinheiro (2009), Bornkamp, Bezlyak and Bretz (2015) and Nandakumar, Dmitrienko and Lipkovich (2017). Extensions of the original method can be found in Klingenberg (2009), Pinheiro et al. (2013) and other publications.

2.1 Hypothesis testing and dose-response modeling

The objective of the MCPMod method is to improve dose-finding in drug development by applying efficient statistical methodology that addresses model uncertainty. At the end of a dose-finding Phase II trial, the trial's sponsor is interested in establishing a dose-response relationship, while controlling the probability of continuing

to the next development stage with a non-effective treatment. While parametric dose-response modeling methods provide an efficient approach to the detection of drug-related effects, they may not control the Type I error rate, e.g., when the model is misspecified.

The key idea behind the MCPMod method is that the sponsor of a dose-finding Phase II trial needs to simultaneously account for the uncertainty around the dose-response model and protect the probability of an incorrect conclusion about the dose-response (Type I error rate). The shape of the true dose-response function is unknown at the beginning of a Phase II trial and thus model uncertainty needs to be accounted for in dose-finding trials. Instead of considering a single dose-response contrast that might not be representative of the true dose-response model, MCPMod defines a set of candidate models (and corresponding dose-response contrasts) and focuses on the goal of demonstrating that at least one contrast-based test produces a significant result. The evaluation of relevant contrasts is performed as part of the first component that can be referred to as the multiple comparison or hypothesis testing component.

It is important to note that, within the contrast-based framework, a hypothesis testing approach is emphasized with appropriate multiplicity adjustments, but no estimation methods are supported, e.g., the functional form of the true dose-response model is not explicitly estimated. The dose is treated as an ordinal variable at best, rather than a continuous variable. After the global hypothesis of no treatment effect is rejected and the experimental treatment is shown to be effective, it is not clear how to approach the problem of estimating target doses. To address this limitation of the contrast-based testing approach, the MCPMod method includes the dose-response modeling component, which expands the hypothesis testing framework to obtain estimates of the underlying dose-response function and target doses.

In general, the MCPMod method relies on the following algorithm:

- Step 1: Derive the optimal contrasts from the candidate models.
- Step 2: Carry out the dose-response tests based on the optimal contrasts.
- Step 3: Select the dose-response models corresponding to the significant dose-response tests.
- Step 4: Estimate the target doses from the selected dose-response models.

Steps 1 through 3 are related to the hypothesis testing component and will be defined in Section 3 and the last step, which is included in the dose-response modeling component, will be described in Section 4.

2.2 General setting and dose-response models

Consider a dose-finding trial with m arms ($m-1$ doses of an experimental treatment versus placebo). The dose levels are denoted by $d_1 = 0, \dots, d_m$. In general, the trial's design does not have to be balanced and n_i denotes the number of patients in the i th arm, $i = 1, \dots, m$, with the total sample size in the trial denoted by n . The primary endpoint is a continuous, binary or count endpoint and the response of the j th patient in the i th arm is denoted by y_{ij} . Let

$$\bar{y}_i = \frac{1}{n_i} \sum_{j=1}^m y_{ij},$$

denote the treatment effect estimate in the i th trial arm that corresponds to the dose d_i , $i = 1, \dots, m$. This quantity serves as an estimate of the mean response in

trials with continuous endpoints, the response rate in trials with binary endpoints and the mean number of events in trials with count endpoints.

The set of candidate dose-response models supported by the MCPModPack package is presented in Table 1. Let r denote the number of candidate models. As shown in the table, each model is defined using the dose-response function $f(d, \beta)$ which depends on the model-specific vector of parameters, i.e., $\beta = (\beta_1, \dots, \beta_k)$, where k is the number of parameters in this particular model. The dose-response function provides a link between the dose and an appropriate characteristic of the endpoint:

- Continuous endpoints: $f(d, \beta) = \mu(d)$, where $\mu(d)$ is the mean of the continuous response at the dose level d .
- Binary endpoints: $f(d, \beta) = \ln[\pi(d)/(1 - \pi(d))]$, where $\pi(d)$ is the probability of the outcome of interest at the dose level d .
- Count endpoints: $f(d, \beta) = \ln \nu(d)$, where $\nu(d)$ is the average number of events at the dose level d .

TABLE 1 Dose-response models supported by the MCPModPack package

Name	Dose-response model	Coefficients
Linear	$f(d, \beta) = \beta_1 + \beta_2 d$	$\beta_1 = E_0, \beta_2 = \delta$
Quadratic	$f(d, \beta) = \beta_1 + \beta_2 d + \beta_3 d^2$	$\beta_1 = E_0, \beta_2 = \delta_1, \beta_3 = \delta_2$
Exponential	$f(d, \beta) = \beta_1 + \beta_2(\exp(d/\beta_3) - 1)$	$\beta_1 = E_0, \beta_2 = E_1, \beta_3 = \delta$
E _{max}	$f(d, \beta) = \beta_1 + \beta_2 d / (\beta_3 + d)$	$\beta_1 = E_0, \beta_2 = E_{\max},$ $\beta_3 = ED_{50}$
Logistic	$f(d, \beta) = \beta_1 + \beta_2 / [1 + \exp((\beta_3 - d)/\beta_4)]$	$\beta_1 = E_0, \beta_2 = E_{\max},$ $\beta_3 = ED_{50}, \beta_4 = \delta$
SigE _{max}	$f(d, \beta) = \beta_1 + \beta_2 d^{\beta_4} / (\beta_3^{\beta_4} + d^{\beta_4})$	$\beta_1 = E_0, \beta_2 = E_{\max},$ $\beta_3 = ED_{50}, \beta_4 = h$

The following notation will be used in the subsequent sections: z_α denotes the upper 100α th percentile of the standard normal distribution.

3 Hypothesis testing

3.1 Step 1: Dose-response contrasts

As the first step, a set of optimal dose-response contrasts is found. An optimal contrast under a particular dose-response model is defined as the contrast that results in the highest probability of rejecting the null hypothesis of no effect. Optimal contrasts are easily derived as follows. Considering the i th model from the candidate set defined in Table 1, let u_{ij} denote the predicted effect at the j th dose based on guesstimates of the model parameters (β_0), i.e.,

$$u_{ij} = f_i(d_j, \beta_0), \quad i = 1, \dots, r, \quad j = 1, \dots, m.$$

The models are standardized in the sense that the linear parameters (β_1 and β_2) are chosen in each model to ensure that $u_{i1} = 0$ and $u_{im} = 1$, $i = 1, \dots, r$.

The optimal contrasts are found from the vector of predicted effects using the variances of the adjusted treatment effect estimates. These estimates and their variances are defined as follows:

- Continuous endpoints: The adjusted treatment effect is simply equal to the sample mean, i.e., $\tilde{y}_i = \bar{y}_i$, and its variance is proportional to $1/n_i$, i.e., $s_i^2 = 1/n_i$, $i = 1, \dots, m$.
- Binary endpoints: The adjusted treatment effect is defined as the logit of the response rate, i.e., $\tilde{y}_i = \ln(\bar{y}_i) - \ln(1 - \bar{y}_i)$, and its variance is equal to $s_i^2 = 1/(n_i \bar{y}_i (1 - \bar{y}_i))$, $i = 1, \dots, m$.
- Count endpoints: The adjusted treatment effect is defined as the logarithm of the mean number of events, i.e., $\tilde{y}_i = \ln(\bar{y}_i)$, and its variance is equal to $s_i^2 = (\theta_i + \bar{y}_i)/(n_i \theta_i \bar{y}_i)$, $i = 1, \dots, m$. Here θ_i is the assumed value of the overdispersion parameter in the i th arm.

To compute the coefficients of the optimal contrast for the i th model, let

$$u_{ij}^* = \frac{u_{ij} - \bar{u}_i}{s_j^2}, \text{ where } \bar{u}_i = \sum_{j=1}^m u_{ij} s_j^2 / \sum_{j=1}^m s_j^2, \quad i = 1, \dots, r, \quad j = 1, \dots, m$$

and

$$\bar{u}_i^* = \frac{1}{m} \sum_{j=1}^m u_{ij}^*.$$

The coefficients of the optimal contrast are standardized as follows:

$$c_{ij} = (u_{ij}^* - \bar{u}_i^*) / \sqrt{\sum_{l=1}^m (u_{il}^* - \bar{u}_i^*)^2}, \quad i = 1, \dots, r, \quad j = 1, \dots, m.$$

3.2 Step 2: Dose-response tests

Using the model-specific optimal contrasts derived in Step 1, the significance of the dose-response trend is evaluated based on the test statistics computed from these contrasts. The test statistics are denoted by t_1, \dots, t_r and are defined for each class of endpoints as shown below.

Continuous endpoints

The test statistic for assessing the significance of a dose-response trend for the i th model is given by

$$t_i = \sum_{j=1}^m c_{ij} \tilde{y}_j / \sqrt{\sum_{j=1}^m \frac{c_{ij}^2 s_j^2}{n_j}}, \quad i = 1, \dots, r,$$

where $\tilde{y}_j = \bar{y}_j$ and s is the sample pooled standard deviation.

Furthermore, the $100(1 - \alpha)\%$ confidence interval for the mean effect in the i th arm is given by

$$\left(\bar{y}_i - z_\alpha \frac{s}{\sqrt{n_i}}, \bar{y}_i + z_\alpha \frac{s}{\sqrt{n_i}} \right), \quad i = 1, \dots, m.$$

Binary endpoints

To define the test statistics, \tilde{y}_i will denote the adjusted treatment effect in the i th arm, i.e., the logit of the response rate, $i = 1, \dots, m$. The sample response rate \bar{y}_i

will be re-defined as follows in the extreme cases where there are no responders or no non-responders in the i th arm:

$$\bar{y}_i = \frac{1}{3n_i + 2} \text{ if no patient experienced the outcome of interest in the } i\text{th arm,}$$

$$\bar{y}_i = \frac{3n_i + 1}{3n_i + 2} \text{ if all patients experienced the outcome of interest in the } i\text{th arm.}$$

The test statistic for the i th model is defined as

$$t_i = \sum_{j=1}^m c_{ij} \tilde{y}_i / \sqrt{\sum_{j=1}^m c_{ij}^2 s_j^2}, \quad i = 1, \dots, r,$$

where s_j^2 is the variance of \tilde{y}_j defined in Step 1.

The $100(1 - \alpha)\%$ confidence interval for the response rate in the i th arm is given by

$$\left(\bar{y}_i - z_\alpha \sqrt{\frac{\bar{y}_i(1 - \bar{y}_i)}{n_i}}, \bar{y}_i + z_\alpha \sqrt{\frac{\bar{y}_i(1 - \bar{y}_i)}{n_i}} \right), \quad i = 1, \dots, m.$$

Count endpoints

As above, \tilde{y}_i will denote the adjusted treatment effect in the i th arm, i.e., the logarithm of the mean number of events, $i = 1, \dots, m$. When computing the adjusted treatment effect, \tilde{y}_i will re-defined if there are no events in the i th trial arm, i.e., \tilde{y}_i will be set to $\Gamma^{-1}(0.5)$. The shape and scale parameters of this gamma distribution are equal to $1/3$ and $1/n_i$.

The test statistic for the i th model is given by

$$t_i = \sum_{j=1}^m c_{ij} \tilde{y}_j / \sqrt{\sum_{j=1}^m c_{ij}^2 s_j^2}, \quad i = 1, \dots, r,$$

where s_j^2 is the variance of \tilde{y}_j defined in Step 1 based on θ_j , i.e., the assumed overdispersion parameter in the j th arm.

The $100(1 - \alpha)\%$ confidence interval for the mean number of events in the i th arm is

$$\left(\exp \left[\ln \bar{y}_i - z_\alpha \sqrt{\frac{(\theta_i + \bar{y}_i)}{n_i \theta_i \bar{y}_i}} \right], \exp \left[\ln \bar{y}_i + z_\alpha \sqrt{\frac{(\theta_i + \bar{y}_i)}{n_i \theta_i \bar{y}_i}} \right] \right), \quad i = 1, \dots, m.$$

Joint distribution of the dose-response test statistics

The test statistics t_1, \dots, t_r are utilized to choose the set of relevant models (models corresponding to significant dose-response contrasts) that play a key role in dose selection decisions. To identify the relevant models, an adjustment for multiplicity is applied to control the Type I error rate. The joint distribution of the r test statistics is taken into account to compute an adjusted critical value that defines the threshold for statistical significance.

Under the global null hypothesis of no treatment effect, the test statistics follow a central multivariate t distribution with ν degrees of freedom and correlation matrix

R. Here $\nu = n - m$, where n is the total sample size in the trial and m is the number of trial arms. The correlation coefficients are known at the design stage and the correlation between the test statistics t_i and t_j , $i \neq j$, is given by

$$\rho_{ij} = \frac{\sum_{l=1}^m c_{il}c_{jl}s_l^2}{\sqrt{\sum_{l=1}^m c_{il}^2s_l^2 \sum_{l=1}^m c_{jl}^2s_l^2}}.$$

The multiplicity-adjusted critical value is calculated as the value q that controls the probability of rejecting the hypotheses of no dose-response at level α :

$$P\{T_1 \geq q \text{ or } \dots \text{ or } T_r \geq q\} = \alpha.$$

Here T_1, \dots, T_r , denote the random variables that have the same joint distribution as t_1, \dots, t_r under the global null hypothesis of no treatment effect for any of the r contrasts. Using the correlation matrix defined above, numerical integration routines for computing multivariate t probabilities (Genz and Bretz, 2002) are used to compute the adjusted critical value. It is important to note that the adjusted critical value q depends on the initial values of the model parameters selected in Step 1.

The set of relevant dose-response models is defined as the set of models that correspond to significant dose-response contrasts after the multiplicity adjustment, i.e., the set of models with $t_i \geq q$, $i = 1, \dots, r$. Let s denote the number of relevant dose-response models.

3.3 Step 3: Best dose-response models

Using the set of s relevant dose-response models, the step deals with the process of selecting an appropriate dose-response model based on appropriate model-selection criteria. The best model can be chosen by finding the contrast that corresponds to the maximum test statistic. This criterion relies only on the strength of evidence against the null hypothesis of no effect and does not account for the number of parameters in the underlying dose-response model. Alternatively, a rule that incorporates information on the functional form of each dose-response model can be considered, e.g., a rule is based on the popular information criterion known as the Akaike information criterion (AIC). In a general dose-response model with k parameters, the AIC is given by:

$$\text{AIC} = 2(L(\hat{\beta}) + k),$$

where $L(\hat{\beta})$ is the negative log-likelihood function evaluated at the maximum likelihood estimate of the parameter vector β . For more information on log-likelihood functions for the candidate dose-response models defined in Table 1, see Section 4. Under this criterion, the best model corresponds to the lowest value of the AIC. In general, in addition to identifying the single best model, the trial's sponsor can consider examining a set of two or three most promising dose-response models based on the same criterion.

4 Dose-response modeling

To implement Step 4, i.e., to identify the target doses based on the selected dose-response models, the models defined in Table 1 are fitted using the method of

maximum likelihood. The details of the model fitting algorithms are presented in Sections 4.1 through 4.3 and identification of the target doses is discussed in Section 4.4.

To simplify notation, the patients will be numbered sequentially and x_i and y_i will denote the dose level and the endpoint's value for the i th patient, $i = 1, \dots, n$. Along the same line, $f_i = f(x_i, \boldsymbol{\beta})$, i.e., the dependence on the dose level (x_i) and vector of dose-response parameters ($\boldsymbol{\beta}$) will be suppressed.

4.1 Continuous endpoints

Assume that the endpoint follows a normal distribution, i.e., y_i is normally distributed with the mean μ_i and common standard deviation σ . The dose-response parameters are estimated by minimizing the negative log-likelihood function. Since $\mu_i = f_i(\boldsymbol{\beta})$, this function is equal to

$$L(\boldsymbol{\beta}, \sigma) = n \ln \sqrt{2\pi} + n \ln \sigma + \frac{1}{2\sigma^2} \sum_{i=1}^n (f_i(\boldsymbol{\beta}) - y_i)^2.$$

To apply the BFGS algorithm, the gradient functions need to be evaluated. The gradients are given by

$$\frac{dL}{d\sigma} = \frac{n}{\sigma} - \frac{1}{\sigma^3} \sum_{i=1}^n (f_i(\boldsymbol{\beta}) - y_i)^2, \quad \frac{dL}{d\beta_j} = \frac{1}{\sigma^2} \sum_{i=1}^n (f_i(\boldsymbol{\beta}) - y_i) \frac{df_i(\boldsymbol{\beta})}{d\beta_j}, \quad j = 1, \dots, k.$$

The model-specific derivatives ($df_i(\boldsymbol{\beta})/d\beta_j$) can be found in the Appendix.

4.2 Binary endpoints

Assume that the endpoint is binary, i.e., $y_i = 0$ or 1 , and π_i denotes the probability of the outcome of interest (response rate) for the i th patient. In this case, $f_i(\boldsymbol{\beta})$ is defined as the logit of the response rate for the i th patient, i.e.,

$$f_i(\boldsymbol{\beta}) = \text{logit } \pi_i = \ln \frac{\pi_i}{1 - \pi_i}.$$

The negative log-likelihood as a function of the parameter vector $\boldsymbol{\beta}$ is given by

$$L(\boldsymbol{\beta}) = \sum_{i=1}^n (\ln(1 + \exp f_i(\boldsymbol{\beta})) - y_i f_i(\boldsymbol{\beta}))$$

and it is easy to verify that the gradients are given by

$$\frac{dL}{d\beta_j} = \sum_{i=1}^n (f_i(\boldsymbol{\beta}) - y_i) \frac{df_i(\boldsymbol{\beta})}{d\beta_j}, \quad j = 1, \dots, k.$$

The model-specific derivatives ($df_i(\boldsymbol{\beta})/d\beta_j$) are given in the Appendix.

4.3 Count endpoints

The endpoint is defined in terms of the number of events experienced by a patient, i.e., y_i is the number of events for the i th patient. It is assumed that y_i follows a Poisson distribution with the parameter λ_i , where λ_i is gamma-distributed with

the shape parameter θ_i and rate parameter θ_i/ν_i . It is easy to show that the mean number of events experienced by the i th patient is ν_i and it is assumed that $\nu_i = \exp f_i(\boldsymbol{\beta})$.

Under the assumption that the overdispersion parameters $\theta_1, \dots, \theta_n$ are fixed by the user, the negative log-likelihood as a function of the parameter vector $\boldsymbol{\beta}$ is equal to

$$L(\boldsymbol{\beta}) = \sum_{i=1}^n \ln \frac{\Gamma(\theta_i) y_i!}{\Gamma(\theta_i + y_i)} - \sum_{i=1}^n (y_i \ln \pi_i + \theta_i \ln(1 - \pi_i)),$$

where

$$\pi_i = \frac{\exp f_i(\boldsymbol{\beta})}{\theta_i + \exp f_i(\boldsymbol{\beta})},$$

and therefore

$$L(\boldsymbol{\beta}) = \sum_{i=1}^n \ln \frac{\Gamma(\theta_i) y_i!}{\Gamma(\theta_i + y_i)} + \sum_{i=1}^n [(\theta_i + y_i) \ln(\theta_i + \exp f_i(\boldsymbol{\beta})) - y_i f_i(\boldsymbol{\beta}) - \theta_i \ln \theta_i].$$

The gradients are given by

$$\frac{dL}{d\beta_j} = \sum_{i=1}^n \frac{\theta_i (\exp f_i(\boldsymbol{\beta}) - y_i)}{\theta_i + \exp f_i(\boldsymbol{\beta})} \frac{df_i(\boldsymbol{\beta})}{d\beta_j}, \quad j = 1, \dots, k.$$

As before, the model-specific derivatives $(df_i(\boldsymbol{\beta})/d\beta_j)$ are defined in the Appendix.

4.4 Step 4: Target doses

For each relevant dose-response model selected in Step 3, the next step is to estimate the target dose or doses of interest, e.g., the minimum effective dose (MED). The MED is defined as the lowest dose that results in a clinically meaningful improvement over placebo. Let Δ denote a threshold that corresponds to the clinically relevant effect over placebo and, secondly, let $f_i(d)$ and $f_i(0)$ define the values of the i th dose-response function at the dose d and in the placebo arm, $i = 1, \dots, r$.

Assume first a positive dose-response relationship, i.e., a larger value of the primary endpoint corresponds to a beneficial treatment effect. Using the i th model, the MED, denoted by d_i^* , is the lowest dose d for which the treatment difference is greater than Δ , i.e.,

$$d_i^* = \min\{d : f_i(d) \geq f_i(0) + \Delta\}.$$

Similarly, under a negative dose-response relationship,

$$d_i^* = \min\{d : f_i(d) \leq f_i(0) - \Delta\}.$$

Under a positive dose-response relationship, the MED is computed for each of the dose-response models defined in Table 1 as shown below. First, consider the i th dose-response model and, if the primary endpoint is continuous, let $\Delta_i^* = \Delta$. Furthermore, if the primary endpoint is binary,

$$\Delta_i^* = \text{logit}(f_i(0, \boldsymbol{\beta}) + \Delta) - \text{logit } f_i(0, \boldsymbol{\beta}).$$

and, if the primary endpoint is a count endpoint,

$$\Delta_i^* = \ln(f_i(0, \boldsymbol{\beta}) + \Delta) - \ln f_i(0, \boldsymbol{\beta}).$$

The MED is found as follows

- Linear model: $d_i^* = \Delta_i^*/\beta_2$ if $\beta_2 \neq 0$, and d_i^* is undefined otherwise.
- Quadratic model:

$$d_i^* = \min(d_1, d_2),$$

where

$$d_1 = (-\sqrt{d_3} - \beta_2)/2\beta_3, \quad d_2 = (\sqrt{d_3} - \beta_2)/2\beta_3, \quad d_3 = \beta_2^2 + 4\beta_3\Delta_i^*,$$

if $d_3 \geq 0$, and d_i^* is undefined otherwise.

- Exponential model:

$$d_i^* = \beta_3 \ln(\Delta_i^*/\beta_2 + 1)$$

if $\beta_2 \neq 0$ and $\Delta_i^*/\beta_2 \geq 0$, and d_i^* is undefined otherwise.

- Emax model:

$$d_i^* = \Delta_i^* \beta_2 / (\beta_2 - \Delta_i^*)$$

if $\Delta_i^* < \beta_2$, and d_i^* is undefined otherwise.

- Logistic model:

$$d_i^* = \beta_3 - \beta_4 \ln(d_2/d_3),$$

where

$$d_1 = \exp(\beta_3/\beta_4), \quad d_2 = \beta_2 d_1 - \Delta_i^* (d_1 + 1), \quad d_3 = \beta_2 + \Delta_i^* (d_1 + 1),$$

if $\Delta_i^* < \beta_2(1 - d_1)$, and d_i^* is undefined otherwise.

- SigEmax model:

$$d_i^* = d_1^{1/\beta_4},$$

where

$$d_1 = \Delta_i^* \beta_3^{\beta_4} / (\beta_2 - \Delta_i^*),$$

if $\Delta_i^* < \beta_2$, and d_i^* is undefined otherwise.

Finally, the MED estimate is computed by replacing the parameters of the dose-response models with their estimates.

The discussion presented above assumed that a single model is identified at the end of Step 3 to describe the dose-response relationship in the trial. If several dose-response models are selected, model averaging techniques can be used to define a weighted MED. Model weights are often computed using an information criterion such as the AIC (Pinheiro et al., 2013). Suppose that s dose-response models were chosen in Step 3, and let p_1, \dots, p_s denote the prior model probabilities (these probabilities quantify how well these models are believed to approximate the true dose-response function). If a non-informative prior is assumed, $p_i = 1/s$, $i = 1, \dots, s$.

Let AIC_i denote the AIC value corresponding to the i th model in the set of relevant dose-response models and let \hat{d}_i^* denote the MED's estimate computed from the i th model, $i = 1, \dots, s$. The weight of the i th model is defined as

$$w_i = \frac{p_i \exp(-\text{AIC}_i/2)}{\sum_{j=1}^s p_j \exp(-\text{AIC}_j/2)}.$$

The estimated MED is defined as a weighted sum of the model-specific minimum effective doses, i.e.,

$$\hat{d}^* = \sum_{i=1}^s w_i \hat{d}_i^*.$$

Note that the best dose-response model is not selected if model averaging is performed.

References

- [1] Bornkamp, B., Bezlyak, V., Bretz, F. (2015). Implementing the MCP-Mod procedure for dose-response testing and estimation. *Modern Approaches to Clinical Trials Using SAS*. Menon, S., Zink, R. (editors). SAS Press: Cary, NC.
- [2] Bretz, F., Pinheiro, J.C., Branson, M. (2005). Combining multiple comparisons and modeling techniques in dose response studies. *Biometrics*. 61, 738-748.
- [3] Bretz, F., Tamhane, A.C., Pinheiro, J. (2009). Multiple testing in dose response problems. *Multiple Testing Problems in Pharmaceutical Statistics*. Dmitrienko, A., Tamhane, A.C., Bretz, F. (editors). New York: Chapman and Hall/CRC Press.
- [4] Genz, A., Bretz, F. (2002). Methods for the computation of multivariate t -probabilities. *Journal of Computational and Graphical Statistics*. 11, 950-971.
- [5] Nandakumar, S., Dmitrienko, A., Lipkovich, I. (2017). Dose-finding methods. *Analysis of Clinical Trials Using SAS: A Practical Guide* (Second Edition). Dmitrienko, A., Koch, G.G. (editors). SAS Press: Cary, NC.
- [6] Pinheiro, J.C., Bornkamp, B., Bretz, F. (2006). Design and analysis of dose finding studies combining multiple comparisons and modeling procedures. *Journal of Biopharmaceutical Statistics*. 16, 639-656.
- [7] Pinheiro J., Bornkamp B., Glimm E., Bretz F. (2013). Model-based dose finding under model uncertainty using general parametric models. *Statistics in Medicine*. 33, 1646-1661.

Appendix

Dose-response modeling: Calculation of gradients

The gradients used in the derivation of maximum likelihood estimates for the dose-response models listed in Table 1 are defined below.

Linear model

$$\frac{df_i}{d\beta_1} = 1, \quad \frac{df_i}{d\beta_2} = x_i.$$

Quadratic model

$$\frac{df_i}{d\beta_1} = 1, \quad \frac{df_i}{d\beta_2} = x_i, \quad \frac{df_i}{d\beta_3} = x_i^2.$$

Exponential model

$$\frac{df_i}{d\beta_1} = 1, \quad \frac{df_i}{d\beta_2} = \exp(x_i/\beta_3) - 1, \quad \frac{df_i}{d\beta_3} = -\frac{\beta_2 \exp(x_i/\beta_3)}{\beta_3^2}.$$

Emax model

$$\frac{df_i}{d\beta_1} = 1, \quad \frac{df_i}{d\beta_2} = \frac{x_i}{\beta_3 + x_i}, \quad \frac{df_i}{d\beta_3} = -\frac{\beta_2 x_i}{(\beta_3 + x_i)^2}.$$

Logistic model

$$\frac{df_i}{d\beta_1} = 1, \quad \frac{df_i}{d\beta_2} = \frac{1}{z}, \quad \frac{df_i}{d\beta_3} = -\frac{(z-1)\beta_2}{z^2\beta_4}, \quad \frac{df_i}{d\beta_4} = \frac{(z-1)(\beta_4 - x_i)\beta_2}{z^2\beta_4^2},$$

where

$$z = 1 + \exp((\beta_4 - x_i)/\beta_4).$$

SigEmax model

$$\frac{df_i}{d\beta_1} = 1, \quad \frac{df_i}{d\beta_2} = \frac{x_i^{\beta_4}}{z}, \quad \frac{df_i}{d\beta_3} = -\frac{x_i^{\beta_4} \beta_2 \beta_3^{\beta_4-1} \beta_4}{z^2}, \quad \frac{df_i}{d\beta_4} = \frac{x_i^{\beta_4} \beta_2 \beta_3^{\beta_4} (\ln x_i - \ln \beta_3)}{z^2},$$

where

$$z = \beta_3^{\beta_4} + x_i^{\beta_4}.$$