Information Theory in Dose-Finding: Improving Safety of the CRM

Pavel Mozgunov, **Thomas Jaki**

Medical and Pharmaceutical Statistics Research Unit,
Department of Mathematics and Statistics, Lancaster University, UK

November 2, 2018

6th Early Phase Adaptive Trials Workshop

Acknowledgement: This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No 633567.

Motivation

Consider a dose-finding trial with binary responses and two doses: d_1, d_2

Goal is to find the maximum tolerated dose (MTD): $\gamma = 0.30$.

10 patients were assigned to each dose, 2 and 4 toxicities observed

Q: Which dose should be administered to the next patient?

Motivation

Consider a dose-finding trial with binary responses and two doses: d_1, d_2
Goal is to find the maximum tolerated dose (MTD): $\gamma = 0.30$.
10 patients were assigned to each dose, 2 and 4 toxicities observed

Q: Which dose should be administered to the next patient?

$$(\hat{p}_i - \gamma)^2 \tag{1}$$

The next patient can be assigned to either of doses, but one can argue that doses are not ‘equal‘ for two reasons.

Motivation

Consider a dose-finding trial with binary responses and two doses: d_1, d_2
Goal is to find the maximum tolerated dose (MTD): $\gamma = 0.30$.
10 patients were assigned to each dose, 2 and 4 toxicities observed

Q: Which dose should be administered to the next patient?

$$(\hat{p}_i - \gamma)^2 \tag{1}$$

The next patient can be assigned to either of doses, but one can argue that doses are not 'equal' for two reasons.

1. Criterion (1) ignores the randomness of the estimates.

$$\mathbb{P}\left(p_2 \in (0.25, 0.35)\right) > \mathbb{P}\left(p_1 \in (0.25, 0.35)\right).$$

Motivation

Consider a dose-finding trial with binary responses and two doses: d_1, d_2
Goal is to find the maximum tolerated dose (MTD): $\gamma = 0.30$.
10 patients were assigned to each dose, 2 and 4 toxicities observed

Q: Which dose should be administered to the next patient?

$$(\hat{p}_i - \gamma)^2 \tag{1}$$

The next patient can be assigned to either of doses, but one can argue that doses are not 'equal' for two reasons.

1. Criterion (1) ignores the randomness of the estimates.

$$\mathbb{P}\left(p_2 \in (0.25, 0.35)\right) > \mathbb{P}\left(p_1 \in (0.25, 0.35)\right).$$

2. $\hat{p}_2 = 0.4$ is an unacceptably high toxicity.

Motivation

It is usually of interest to balance both aims in a Phase I clinical trial

Current solutions

Safety:

Escalation with Overdose Control (EWOC) design (Babb et al., 1998):

$$\mathbb{E}\left(\alpha(\gamma - P_i)^+ + (1-\alpha)(P_i - \gamma)^+\right) \tag{2}$$

+ Low average number of DLTs
− Underestimation of the MTD
- Modifications: α_n by Tighiouart et al. (2010) and Wheeler et al. (2017)

Current solutions

Safety:

Escalation with Overdose Control (EWOC) design (Babb et al., 1998):

$$\mathbb{E}\left(\alpha(\gamma - P_i)^+ + (1-\alpha)(P_i - \gamma)^+\right) \tag{2}$$

+ Low average number of DLTs
− Underestimation of the MTD
  - Modifications: α_n by Tighiouart et al. (2010) and Wheeler et al. (2017)

Safety & Uncertainty

Bayesian Logistic Regression Model (BLRM, Neuenschwander et al., 2008). uses the distribution of DLT probabilies. For example, for $\gamma = 0.33$

$$L = \begin{cases} 1 \text{ if } p \in (0.00, 0.26); & 0 \text{ if } p \in (0.26, 0.41); \\ 1 \text{ if } p \in (0.41, 0.66); & 2 \text{ if } p \in (0.66, 1.00) \end{cases}$$

Goal

We propose a new criterion for selecting doses in dose-escalation trials that accounts for

1. Uncertainty in the estimates
2. Ethical constraints

and requires only **one additional parameter** to be specified.

Goal

We propose a new criterion for selecting doses in dose-escalation trials that accounts for

1. Uncertainty in the estimates
2. Ethical constraints

and requires only **one additional parameter** to be specified.

We incorporate the proposed criterion to the one-parameter Bayesian continual reassessment method (O'Quigley et al., 1990, CRM)

Novel Criterion

The main object of estimation is the probability of DLT $p_i \in (0, 1)$

Squared distance is not a reliable measure for objects on the unit interval (Aitchison, 1992).

Novel Criterion

The main object of estimation is the probability of DLT $p_i \in (0, 1)$

Squared distance is not a reliable measure for objects on the unit interval (Aitchison, 1992).

Instead, we propose a distance satisfying the desirable properties

$$\delta(p, \gamma) = \frac{(p - \gamma)^2}{p(1 - p)}. \tag{3}$$

- $\delta(\cdot) = 0$ at $p = \gamma$
- $\delta(\cdot) \to \infty$ as $p \to 0$ or $p \to 1$
- The variance in denominator (Criterion 3 is a score statistic)

Novel Criterion

The main object of estimation is the probability of DLT $p_i \in (0, 1)$
Squared distance is not a reliable measure for objects on the unit interval (Aitchison, 1992).

Instead, we propose a distance satisfying the desirable properties

$$\delta(p, \gamma) = \frac{(p - \gamma)^2}{p(1 - p)}. \tag{3}$$

- $\delta(\cdot) = 0$ at $p = \gamma$
- $\delta(\cdot) \to \infty$ as $p \to 0$ or $p \to 1$
- The variance in denominator (Criterion 3 is a score statistic)

In the illustration example above

$$\delta(\hat{p}_1 = 0.2, \gamma = 0.3) = 1/16 \quad \text{and} \quad \delta(\hat{p}_2 = 0.4, \gamma = 0.3) = 1/24$$

(!) Single point estimate summarizes the information about uncertainty.

Introducing safety compound

The target toxicity γ is always less than 0.5.

Then for estimates $\hat{p}_1 = \gamma - \theta$ and $\hat{p}_2 = \gamma + \theta$, symmetric criterion favours $\hat{p}_2$.

Introducing safety compound

The target toxicity γ is always less than 0.5.
Then for estimates $\hat{p}_1 = \gamma - \theta$ and $\hat{p}_2 = \gamma + \theta$, symmetric criterion favours $\hat{p}_2$.

We introduce an asymmetry parameter a:

$$\delta(p, \gamma) = \frac{(p - \gamma)^2}{p^a (1 - p)^{2-a}}. \tag{4}$$

$0 < a < 1$ implies more severe penalty for more toxic doses.

(!) Selection of under toxic doses remain to be undesirable as well.

Introducing safety compound

The target toxicity γ is always less than 0.5.
Then for estimates $\hat{p}_1 = \gamma - \theta$ and $\hat{p}_2 = \gamma + \theta$, symmetric criterion favours $\hat{p}_2$.

We introduce an asymmetry parameter a:

$$\delta(p, \gamma) = \frac{(p - \gamma)^2}{p^a (1 - p)^{2-a}}. \tag{4}$$

$0 < a < 1$ implies more severe penalty for more toxic doses.

(!) Selection of under toxic doses remain to be undesirable as well.

In the illustration example above, for $a = 0.5$

$$\delta(\hat{p}_1 = 0.2, \gamma = 0.3, a = 0.5) < \delta(\hat{p}_2 = 0.4, \gamma = 0.3, a = 0.5).$$

Asymmetry parameter (I)

Parameter *a* balances the trade-off between ethical concerns and uncertainty

How can *a* be choosen?

Asymmetry parameter (I)

Parameter a balances the trade-off between ethical concerns and uncertainty

How can a be choosen?

Value $a = 2\gamma$ leads to the same allocation as the squared distance $\rightarrow$

$a < 2\gamma$ leads to more conservative allocation of patients.

Asymmetry parameter (I)

Parameter a balances the trade-off between ethical concerns and uncertainty

How can a be choosen?

Value $a = 2\gamma$ leads to the same allocation as the squared distance $\rightarrow$
$a < 2\gamma$ leads to more conservative allocation of patients.

Let $(\gamma - \theta, \gamma + \theta)$ be an interval such that among two estimates standing on the same squared distance from γ, the lower estimate would be preferred

$$a = 2 \times \left(1 + \left(\log \frac{\gamma - \theta}{\gamma + \theta}\right) / \left(\log \frac{1 - \gamma - \theta}{1 - \gamma + \theta}\right)\right)^{-1}$$

Bayesian continual reassessment method

DLT probability has the functional form $\psi(d_i, \beta) = d_i^{\exp(\beta)}$.

$f_0(.)$ is prior distribution of β. After j patients have already been assigned to doses $d(1), \ldots, d(j)$ and binary responses $\mathbb{Y}_j = [y_1, \ldots, y_j]^{\mathrm{T}}$ were observed the posterior $f_j(\beta)$ is obtained.

Bayesian continual reassessment method

DLT probability has the functional form $\psi(d_i, \beta) = d_i^{\exp(\beta)}$.

$f_0(.)$ is prior distribution of β. After j patients have already been assigned to doses $d(1), \ldots, d(j)$ and binary responses $\mathbb{Y}_j = [y_1, \ldots, y_j]^{\mathrm{T}}$ were observed the posterior $f_j(\beta)$ is obtained.

Then, the dose d_k minimising

$$\mathbb{E}\left(\frac{(\psi(d_i, \beta) - \gamma)^2}{\psi(d_i, \beta)^a (1 - \psi(d_i, \beta))^{2-a}}\right) \tag{5}$$

among all $d_1, \ldots, d_m$ is recommended for the next group of patients

Bayesian continual reassessment method

DLT probability has the functional form $\psi(d_i, \beta) = d_i^{\exp(\beta)}$.

$f_0(.)$ is prior distribution of β. After j patients have already been assigned to doses $d(1), \ldots, d(j)$ and binary responses $\mathbb{Y}_j = [y_1, \ldots, y_j]^{\mathrm{T}}$ were observed the posterior $f_j(\beta)$ is obtained.

Then, the dose d_k minimising

$$\mathbb{E}\left(\frac{(\psi(d_i, \beta) - \gamma)^2}{\psi(d_i, \beta)^a (1 - \psi(d_i, \beta))^{2-a}}\right) \tag{5}$$

among all $d_1, \ldots, d_m$ is recommended for the next group of patients

Convex Infinite Bounds Penalization with parameter a as CIBP(a).

Illustration (I)

We revisit the Everolimus Trial in patients with HER2-overexpressing Metastatic Breast Cancer $\gamma = 0.3$. The study considers 3 regimens given together with Paclitaxel and Trastuzumab (PT):

1. Daily dosing of Everolimus 5mg plus PT (d_1)
2. Daily dosing of Everolimus 10mg plus PT (d_2)
3. Weekly dosing of Everolimus 30mg plus PT (d_3)

Table: Aggregated data of the Everolimus trial

Dose	d_1	d_2	d_3
Number of Patients assigned	6	17	10
Number of DLTs	3	6	7

Illustration (I)

We revisit the Everolimus Trial in patients with HER2-overexpressing Metastatic Breast Cancer $\gamma = 0.3$. The study considers 3 regimens given together with Paclitaxel and Trastuzumab (PT):

1. Daily dosing of Everolimus 5mg plus PT (d_1)
2. Daily dosing of Everolimus 10mg plus PT (d_2)
3. Weekly dosing of Everolimus 30mg plus PT (d_3)

Table: Aggregated data of the Everolimus trial

Dose	d_1	d_2	d_3
Number of Patients assigned	6	17	10
Number of DLTs	3	6	7

We compare original CRM and CIBP (0.3) using the same prior parameters.

Illustration (II)

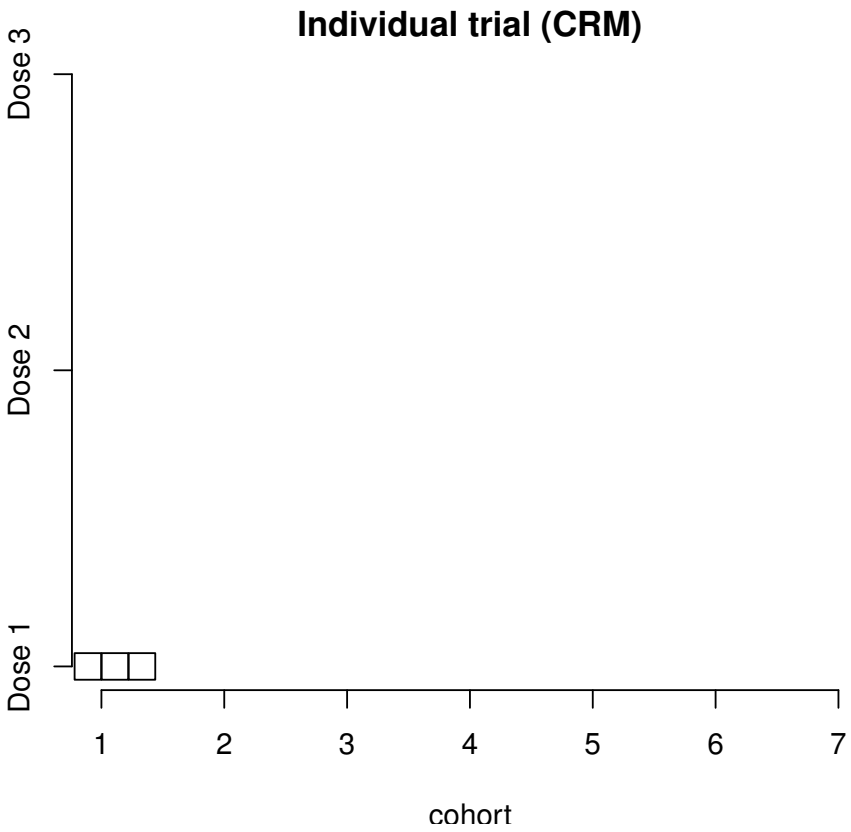

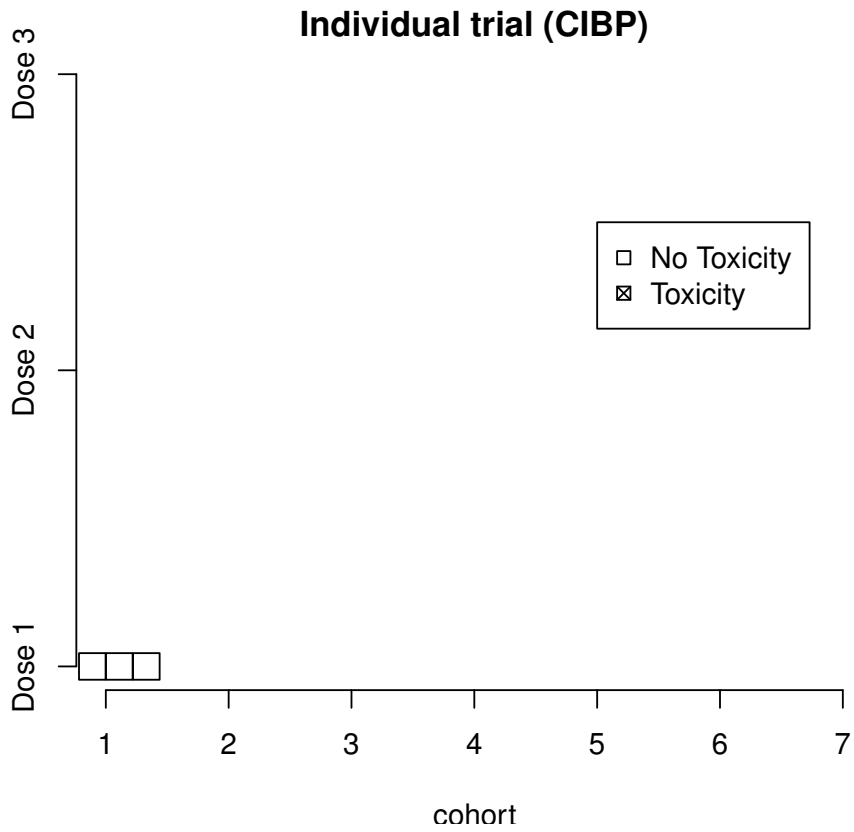

Illustration (II)

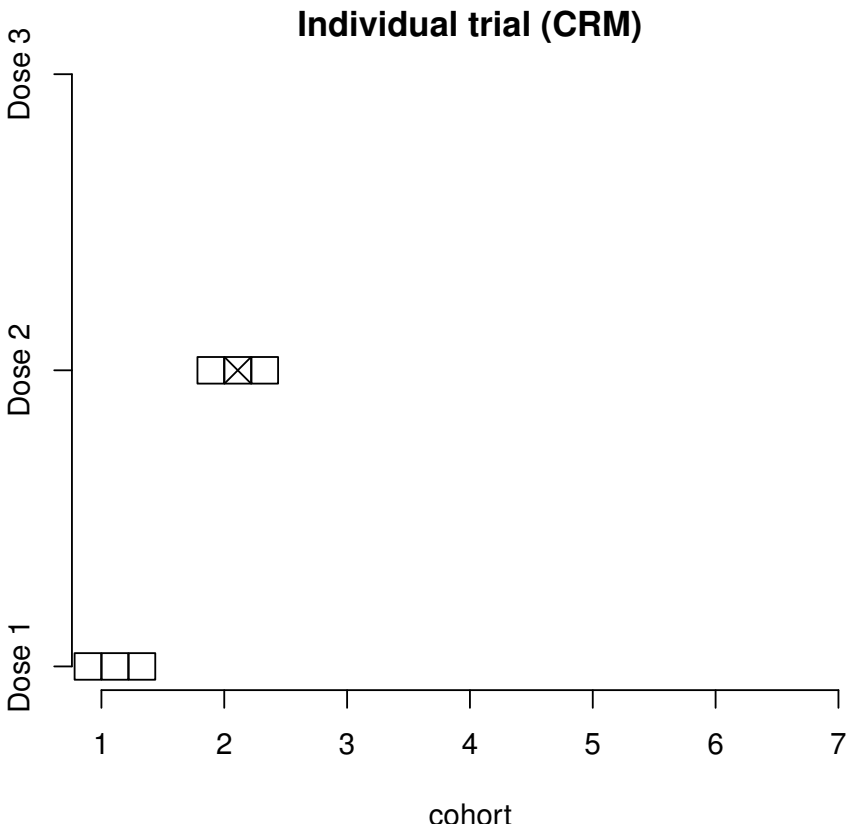

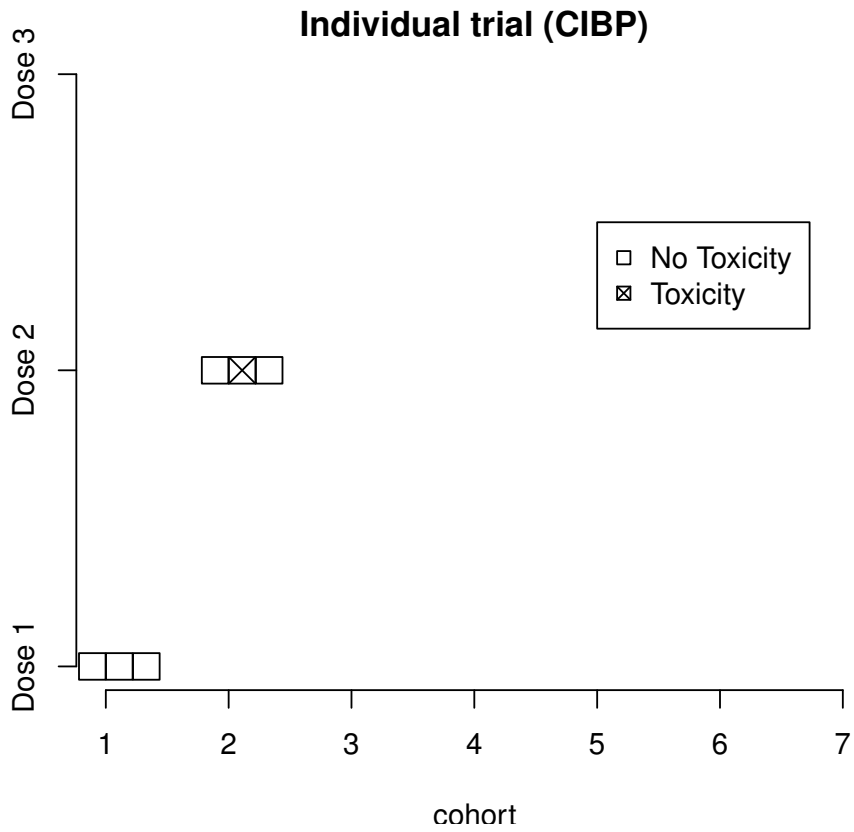

Illustration (II)

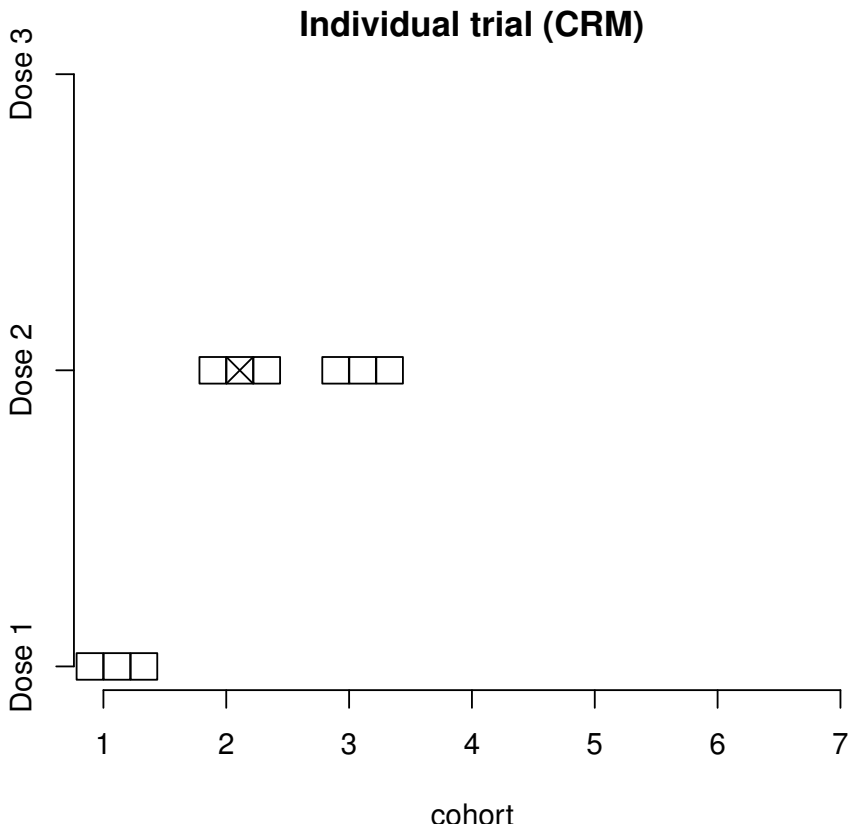

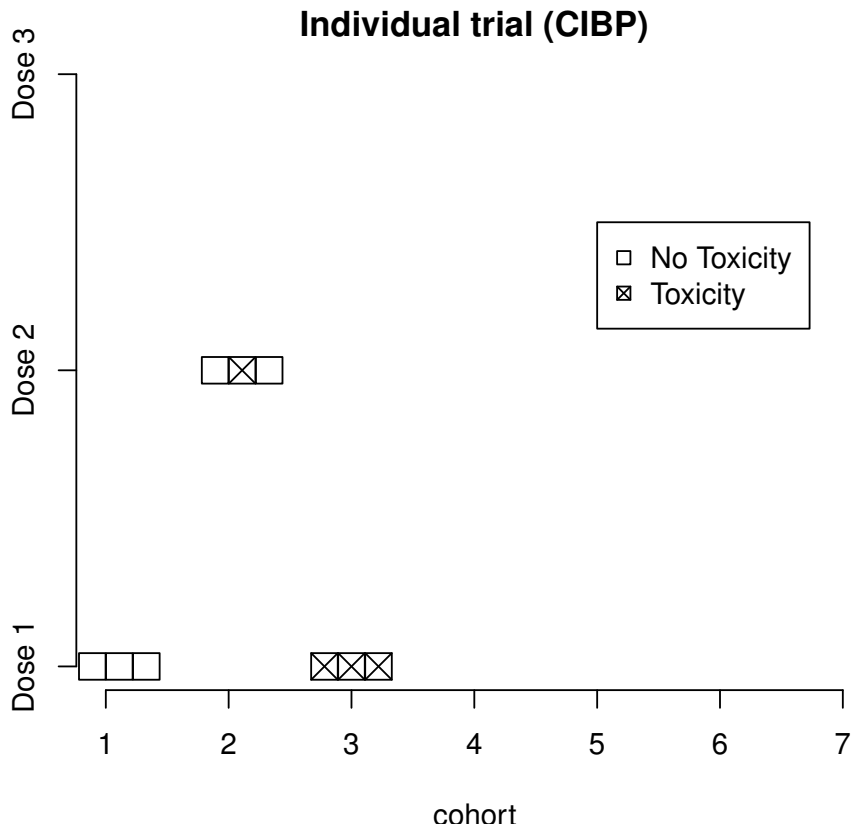

Illustration (II)

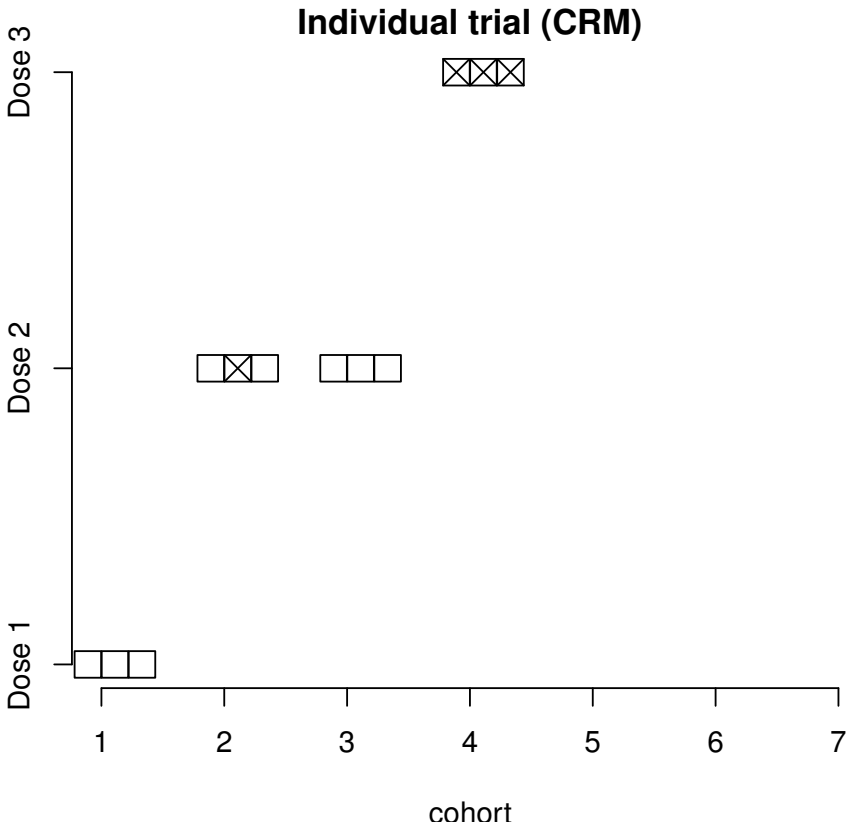

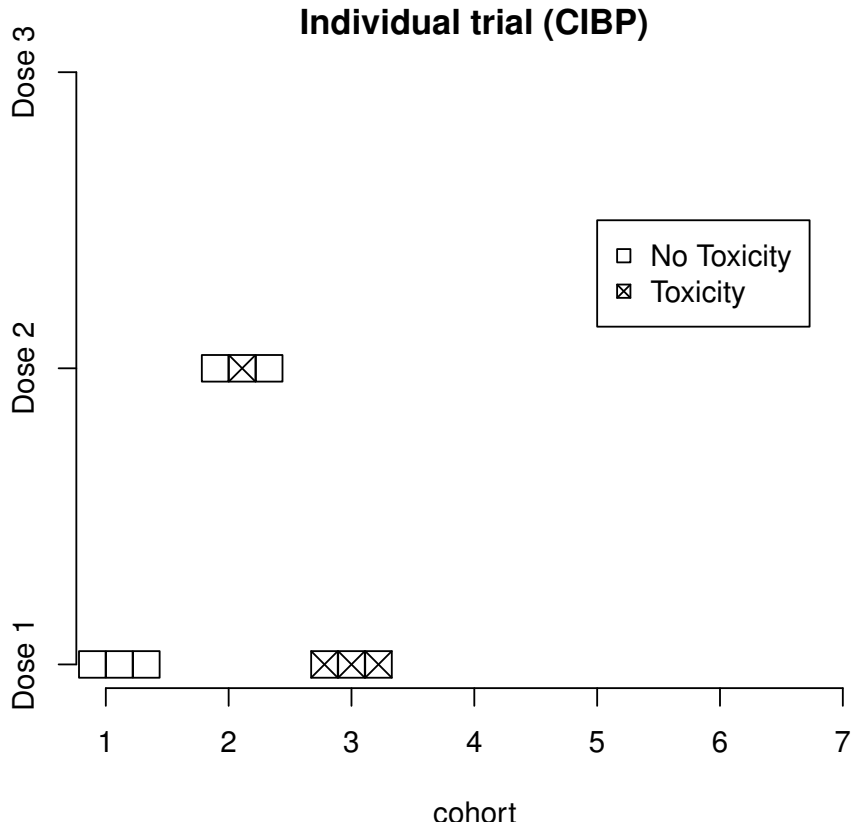

Illustration (II)

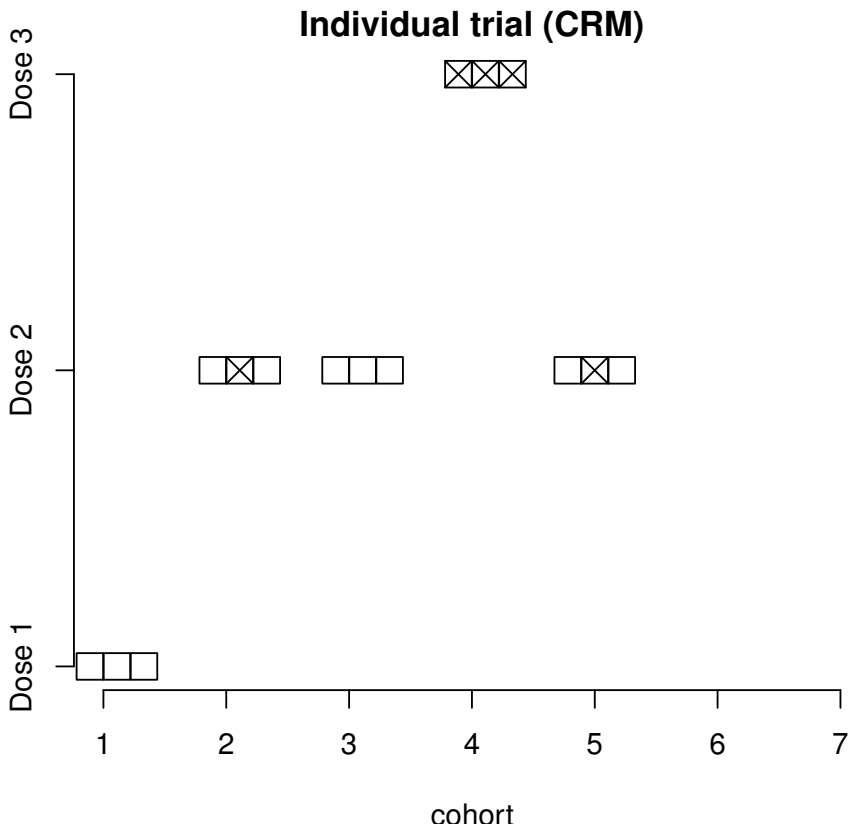

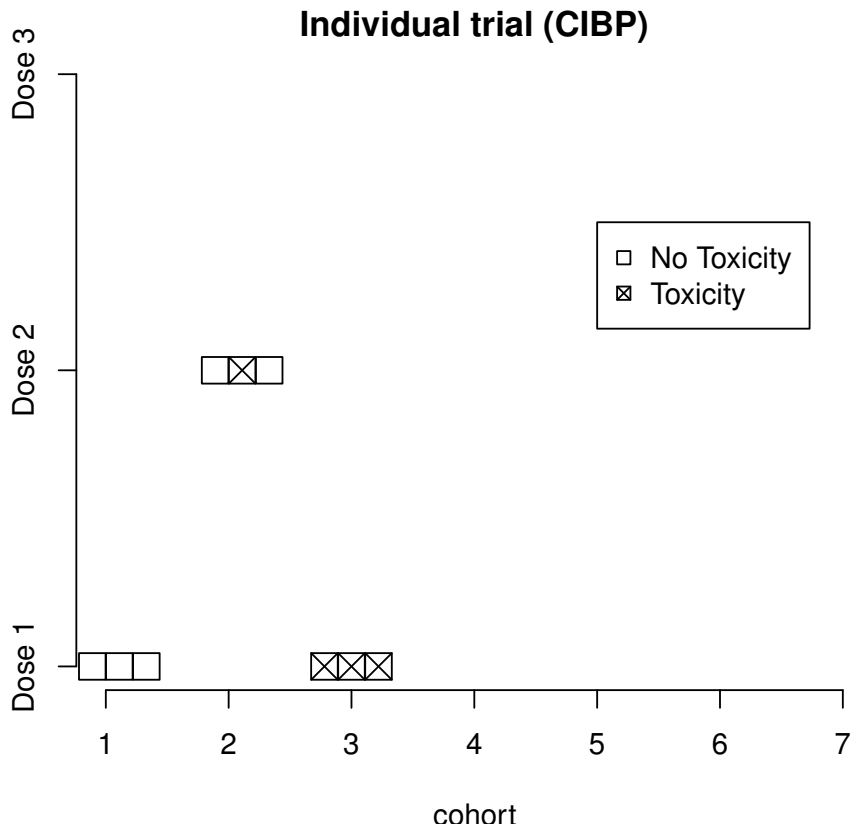

Illustration (II)

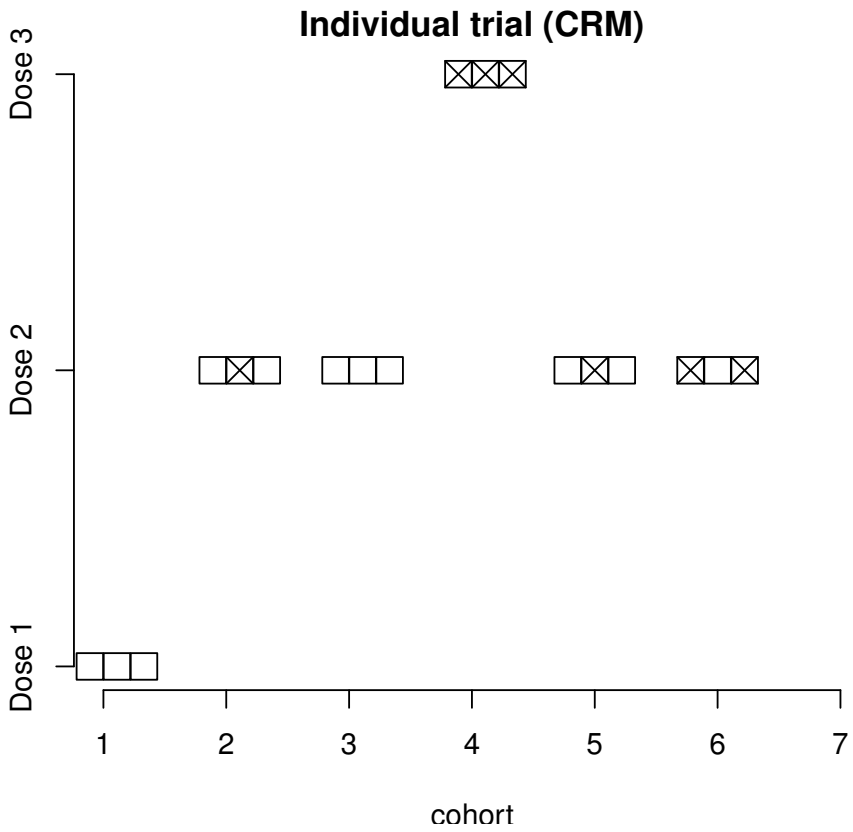

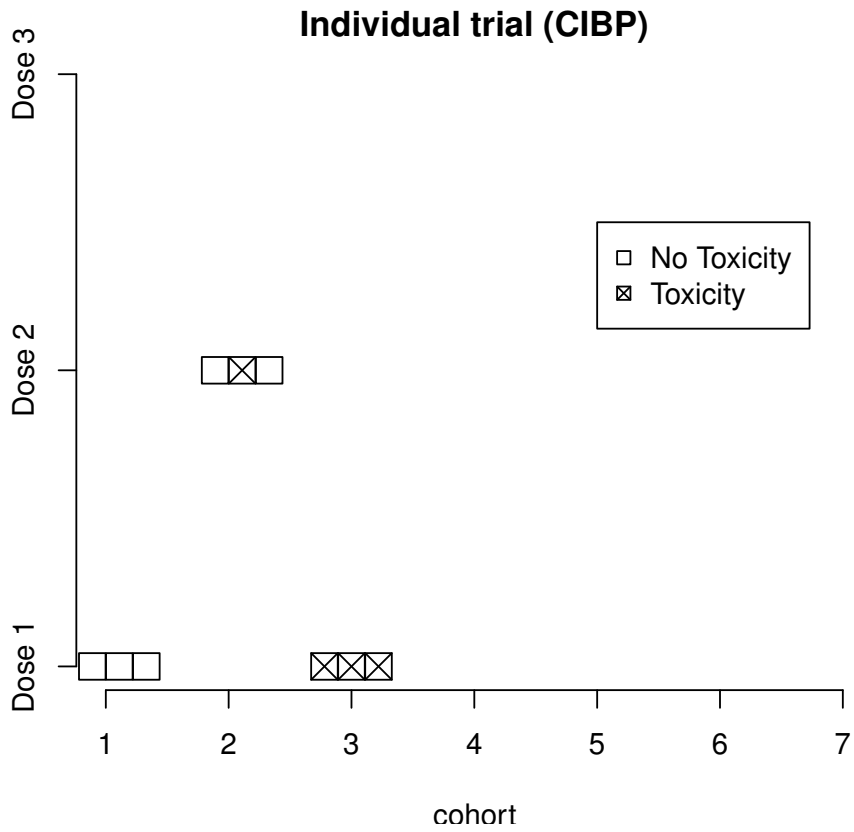

Illustration (II)

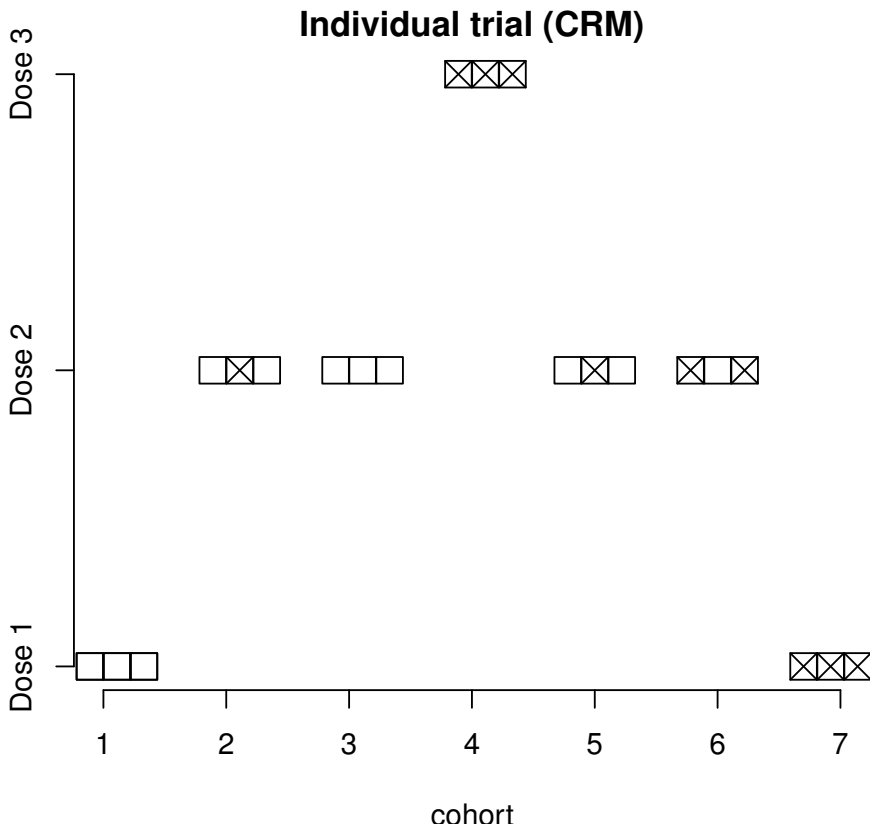

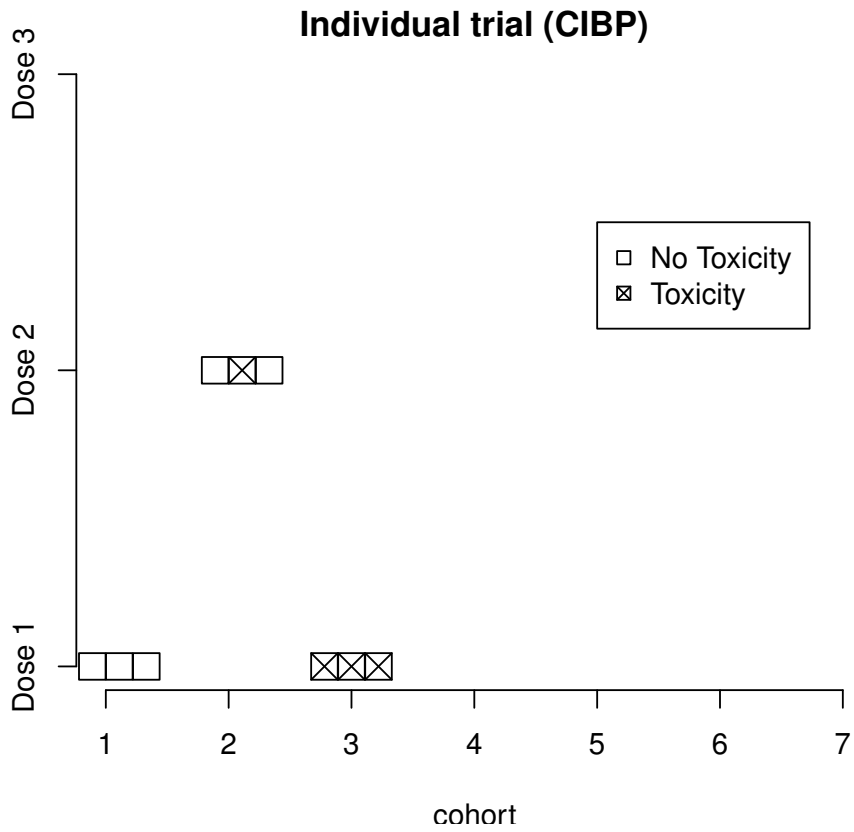

Numerical Study

Setting by Wheeler et al. (2017).

- $n = 40$ patients; $m = 6$ doses; $c = 1$ cohort size; target $\gamma = 0.33$
- $\beta \sim \mathcal{N}(0, 1.34)$
- $a = \{0.5, 0.25, 0.10\}$.

Numerical Study

Setting by Wheeler et al. (2017).

- $n = 40$ patients; $m = 6$ doses; $c = 1$ cohort size; target $\gamma = 0.33$
- $\beta \sim \mathcal{N}(0, 1.34)$
- $a = \{0.5, 0.25, 0.10\}$.

We study the performance of designs in terms of
(i) *Accuracy*

$$\mathcal{A} = 1 - m\frac{\sum_{i=1}^{m}(p_i - \gamma)^2 \pi_i}{\sum_{i=1}^{m}(p_i - \gamma)^2}$$

(*ii*) mean number of toxic responses (DLTs) and focus on the mean performance.

Scenarios

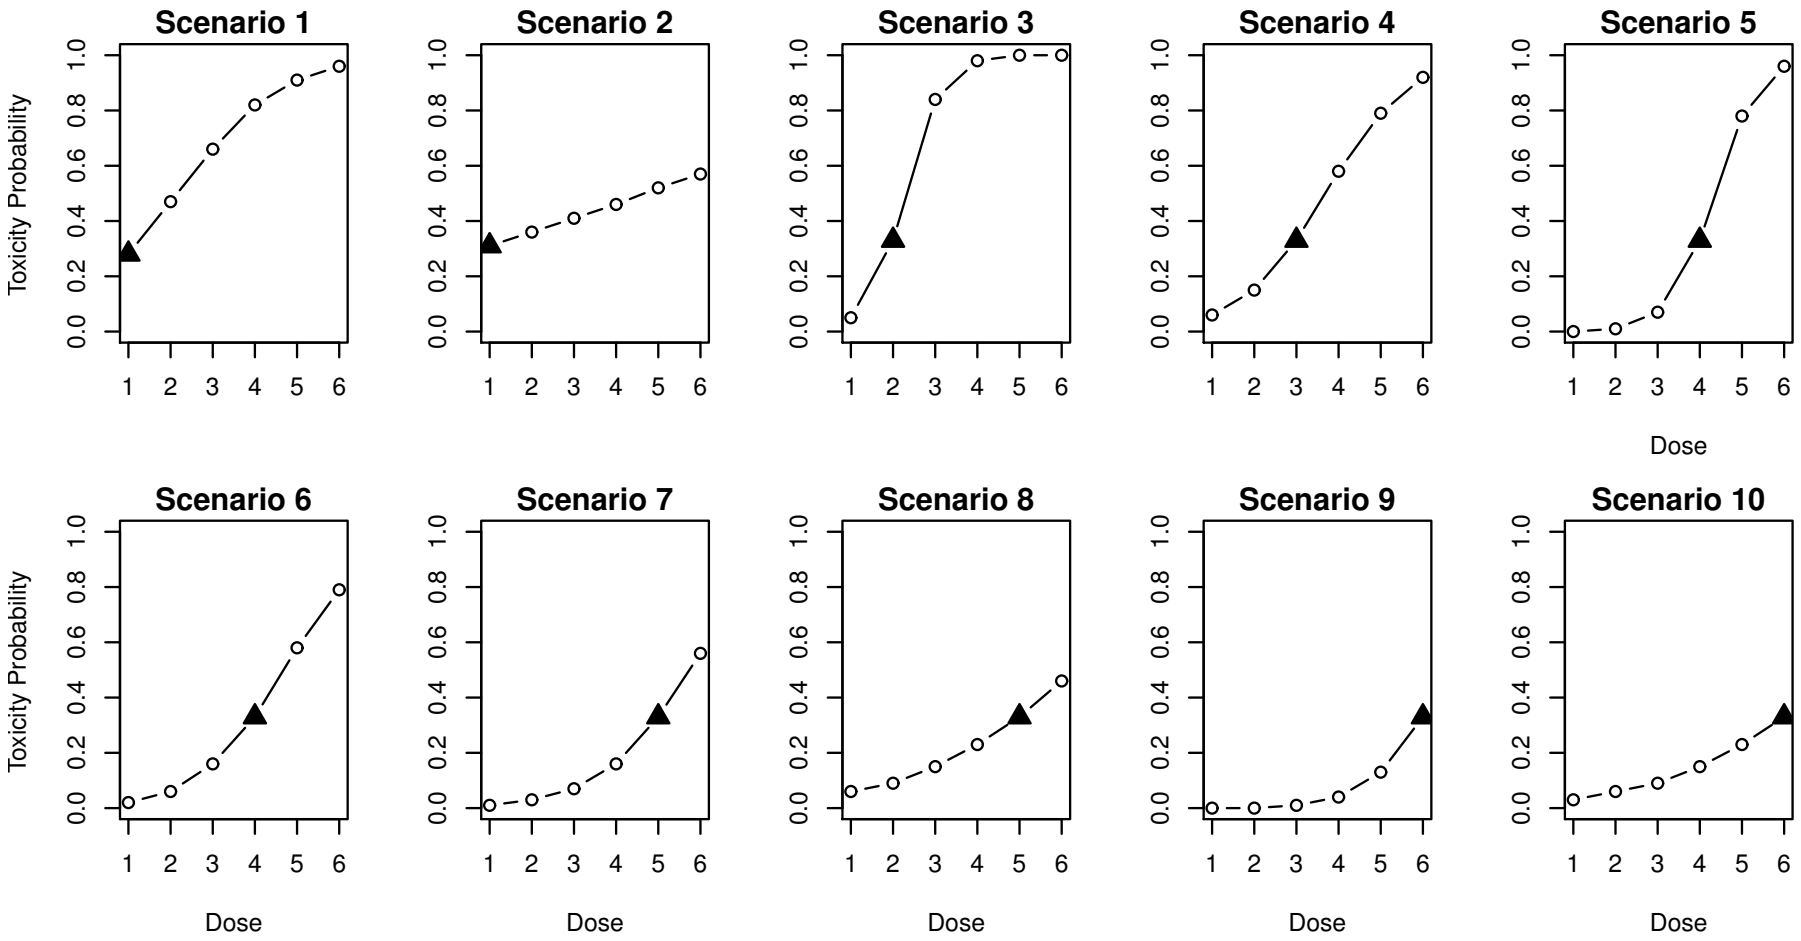
Scenario 1
Scenario 2
Scenario 3
Scenario 4
Scenario 5
Scenario 6
Scenario 7
Scenario 8
Scenario 9
Scenario 10
Toxicity Probability
Dose
0.0 0.2 0.4 0.6 0.8 1.0
1 2 3 4 5 6

Comparators

We compare the performance of the proposed approach to

- **EWOC**
- **TR** design by Tighiouart et al. (2010)
- Toxicity-dependent feasibility bound (**TDFB**) by Wheeler et al. (2017)
- **BLRM** by Neuenschwander et al. (2008)

We use the same prior distribution as Neuenschwander et al. (2008).

Results. Accuracy

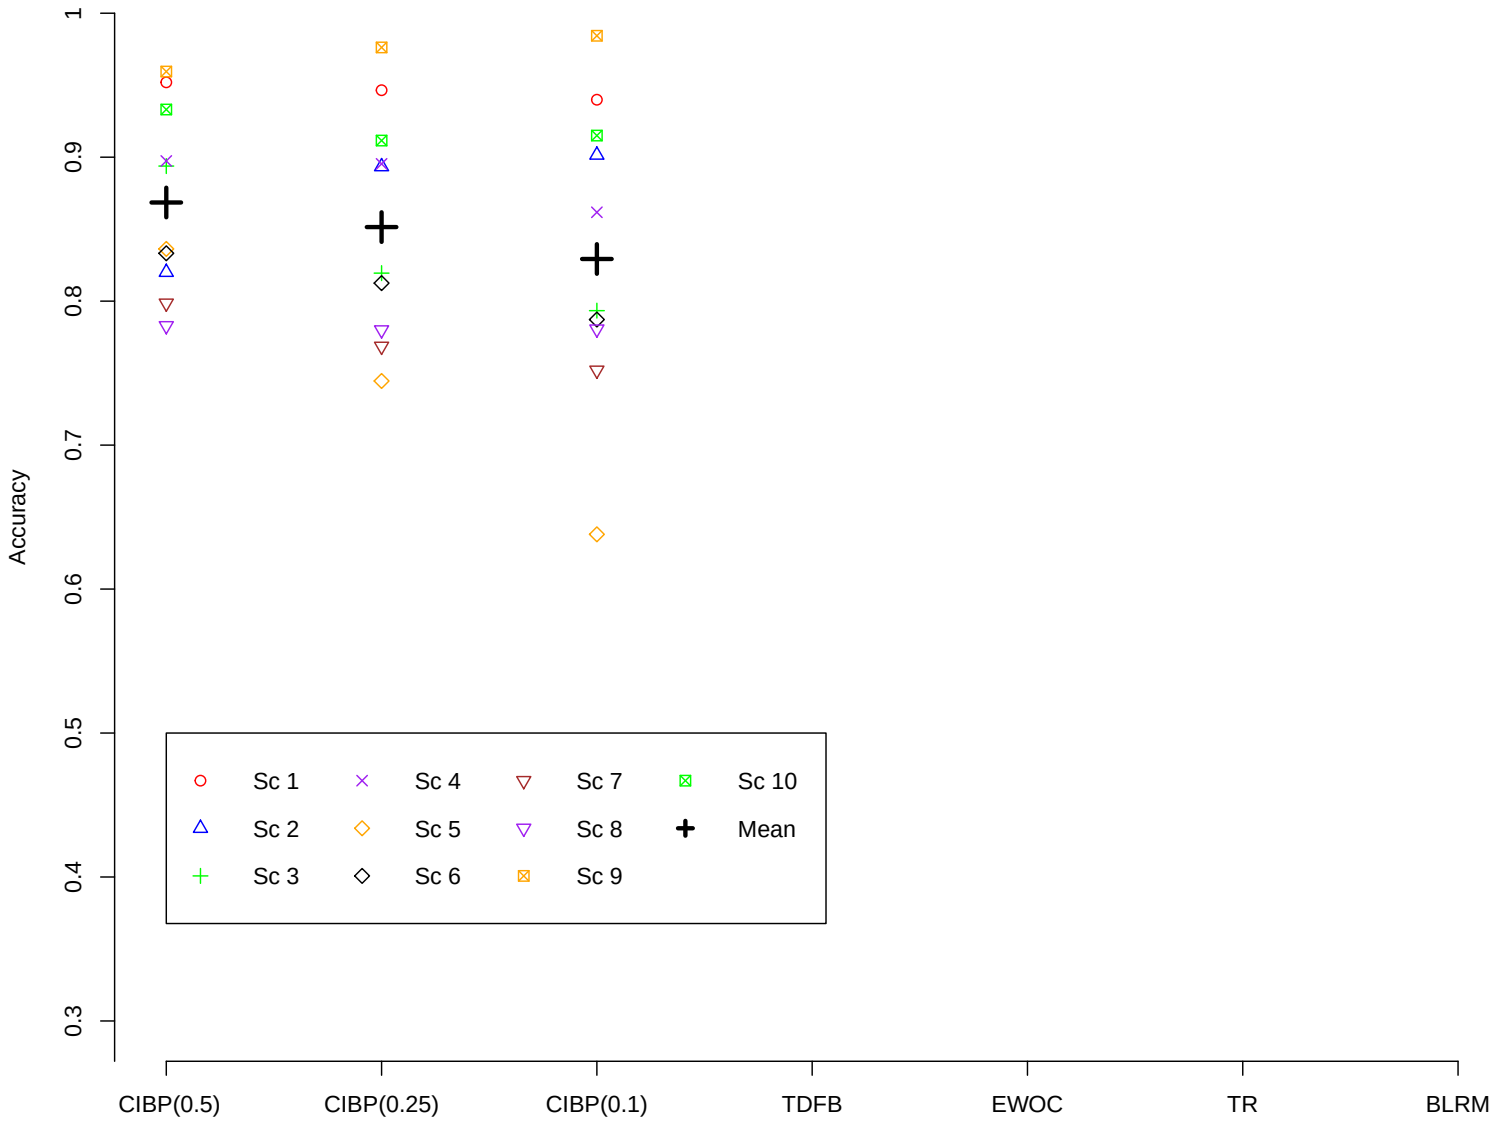

Results. Accuracy

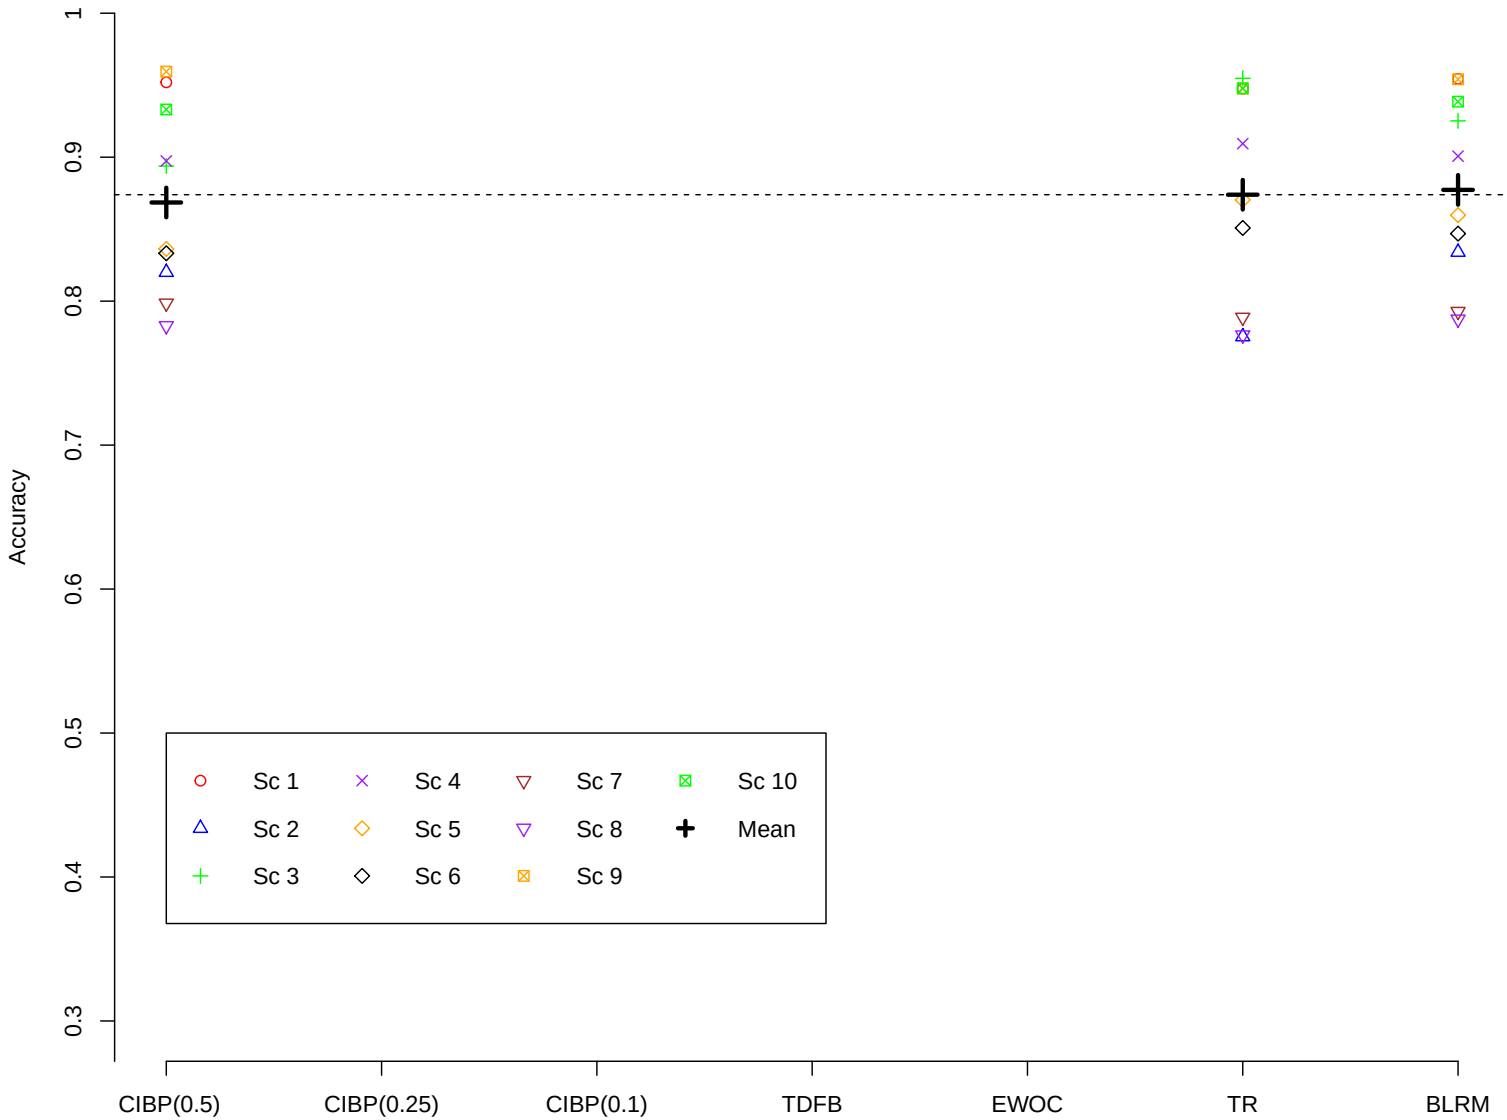

Results. Accuracy

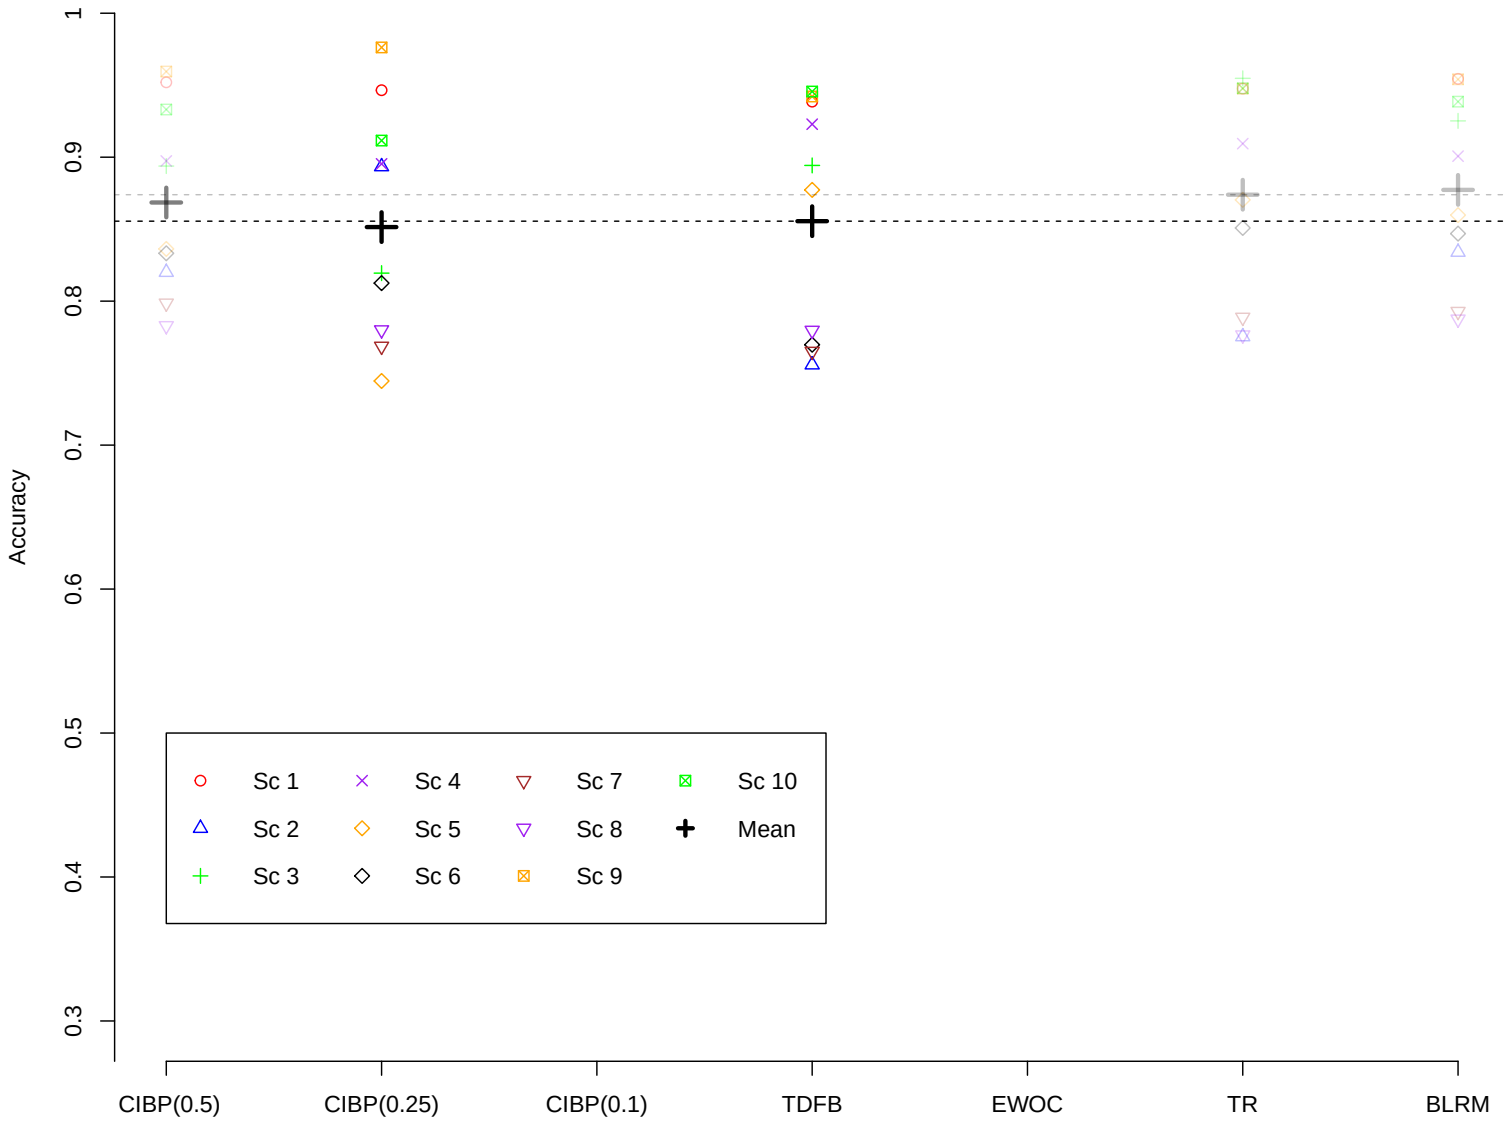

Results. Accuracy

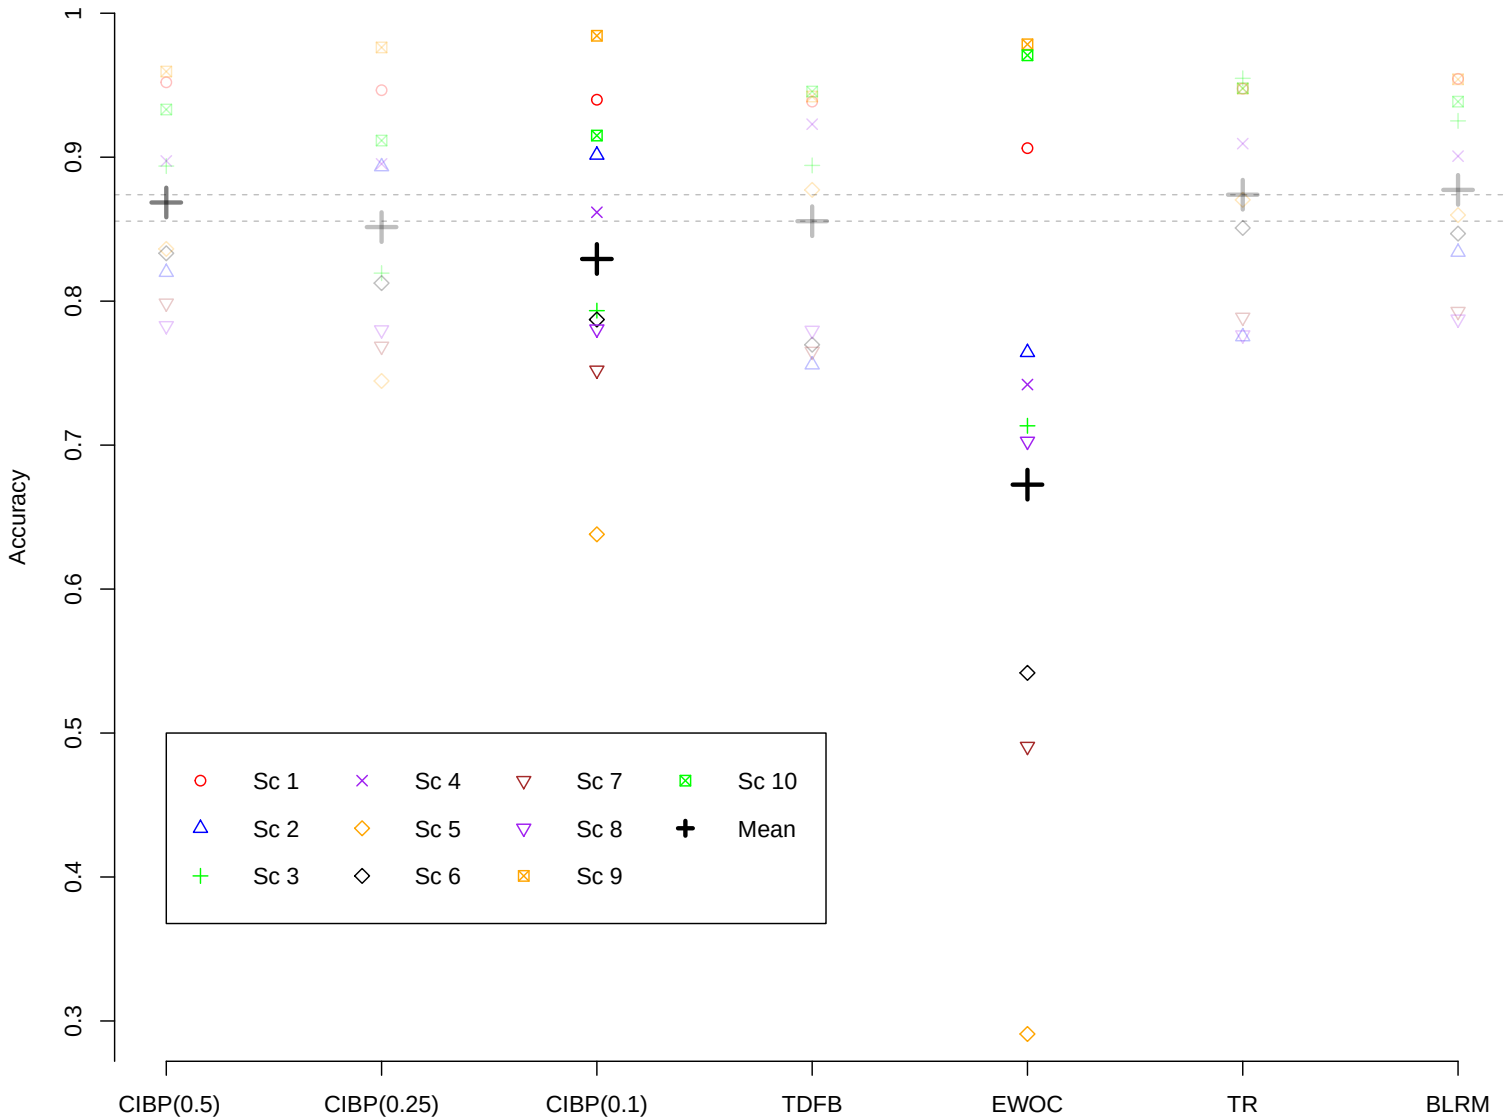

Results.DLTs

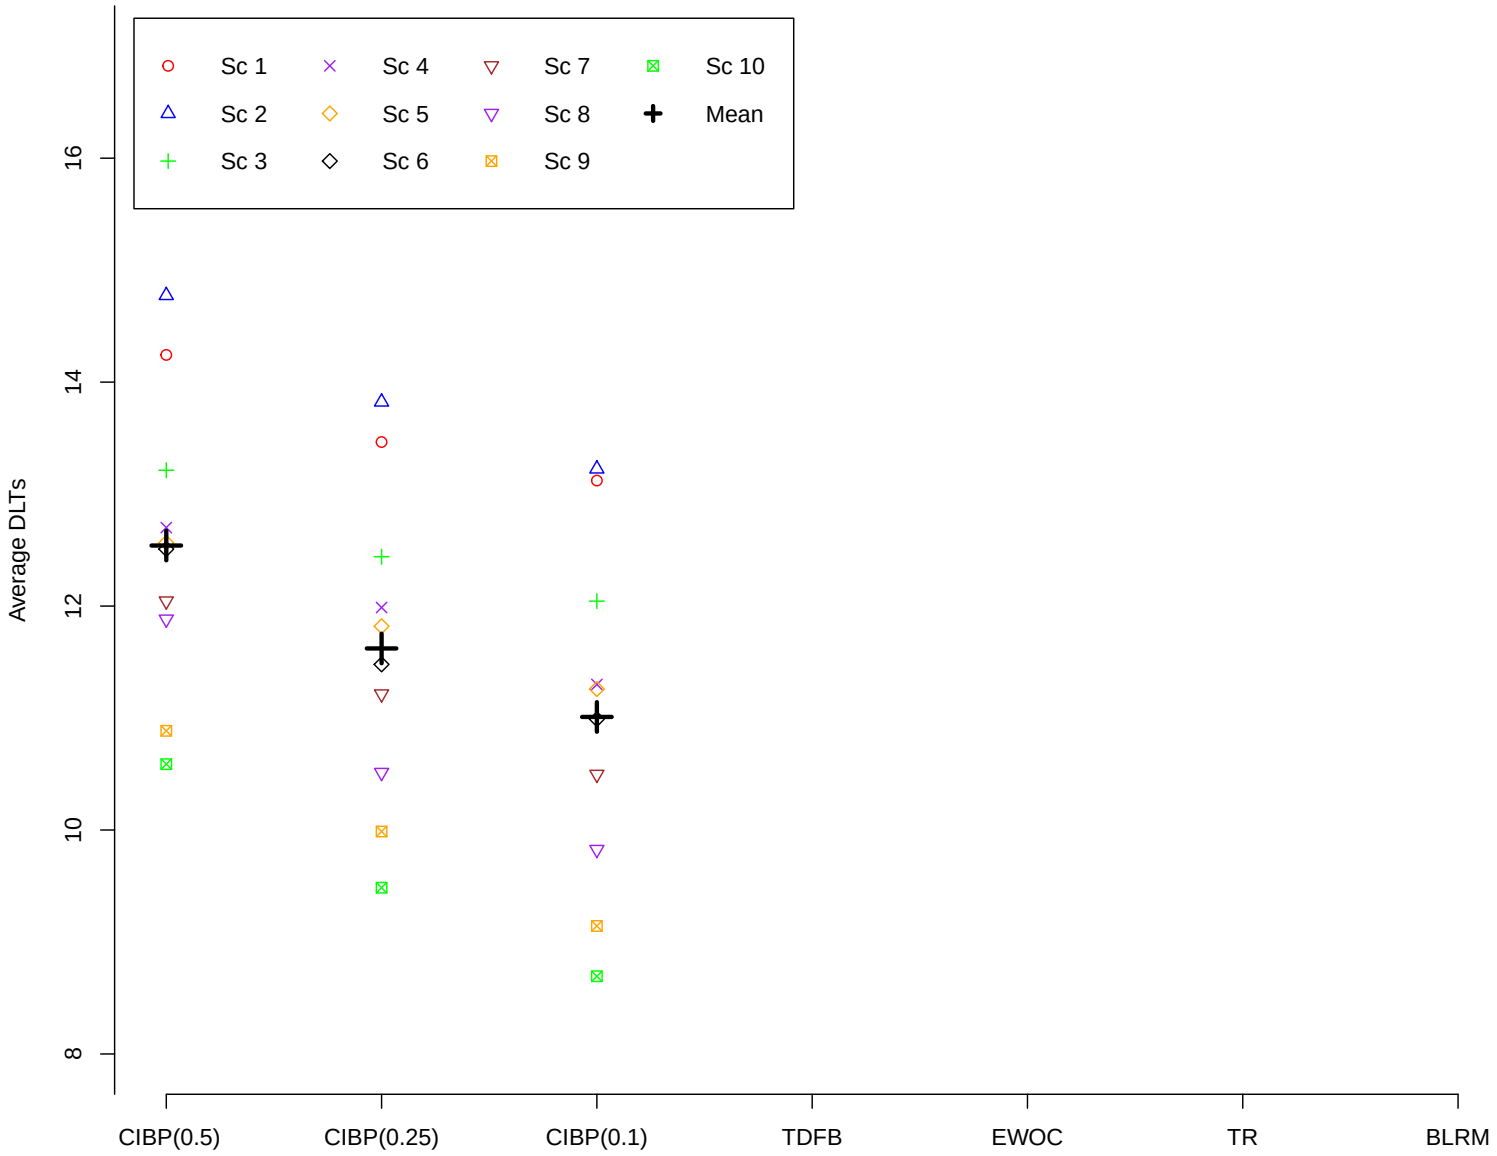

Results.DLTs

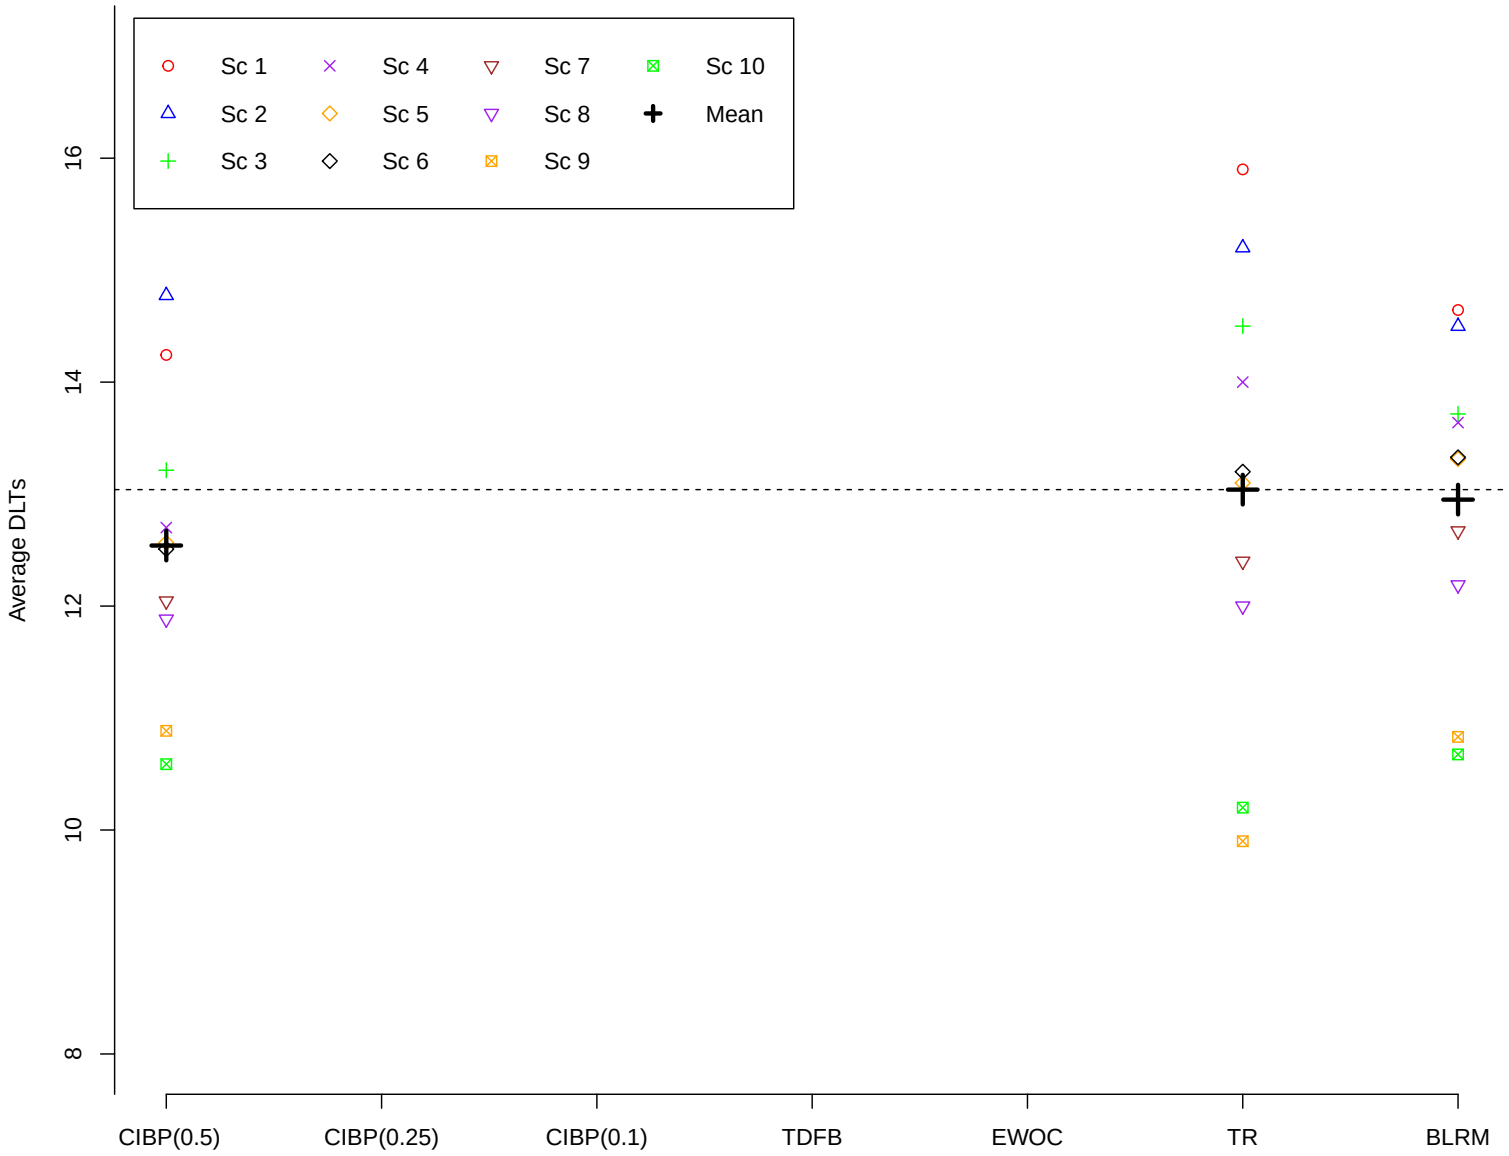

Results.DLTs

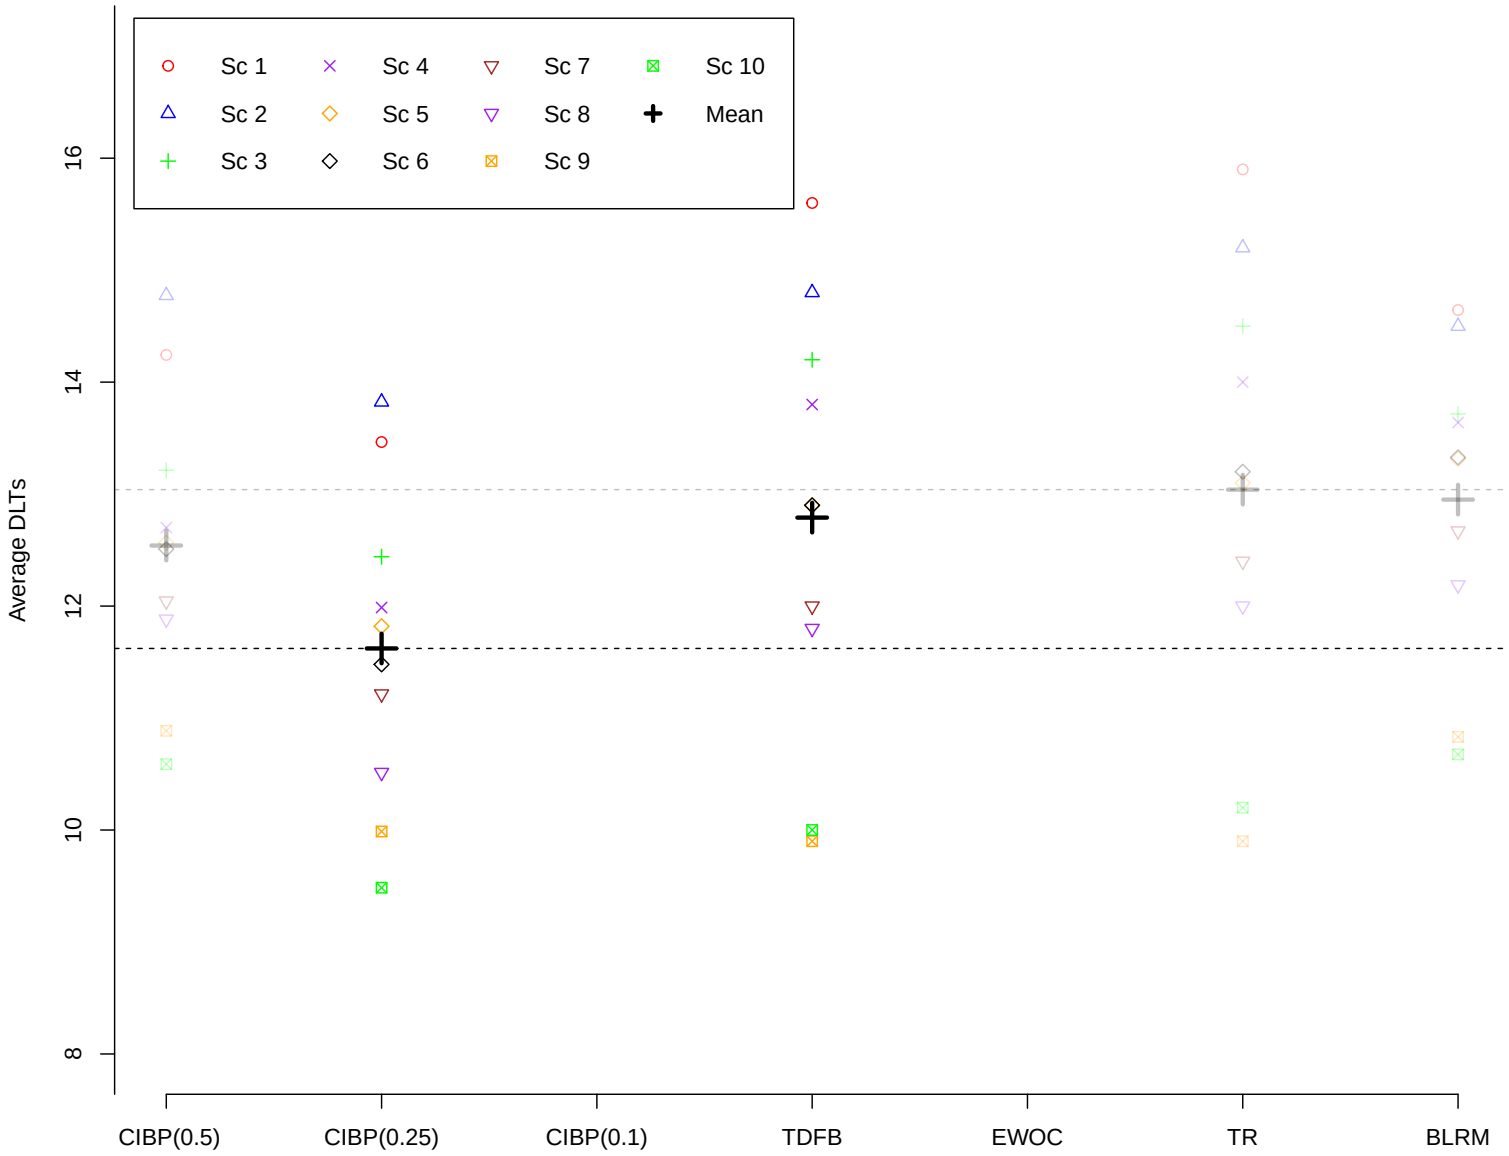

Results.DLTs

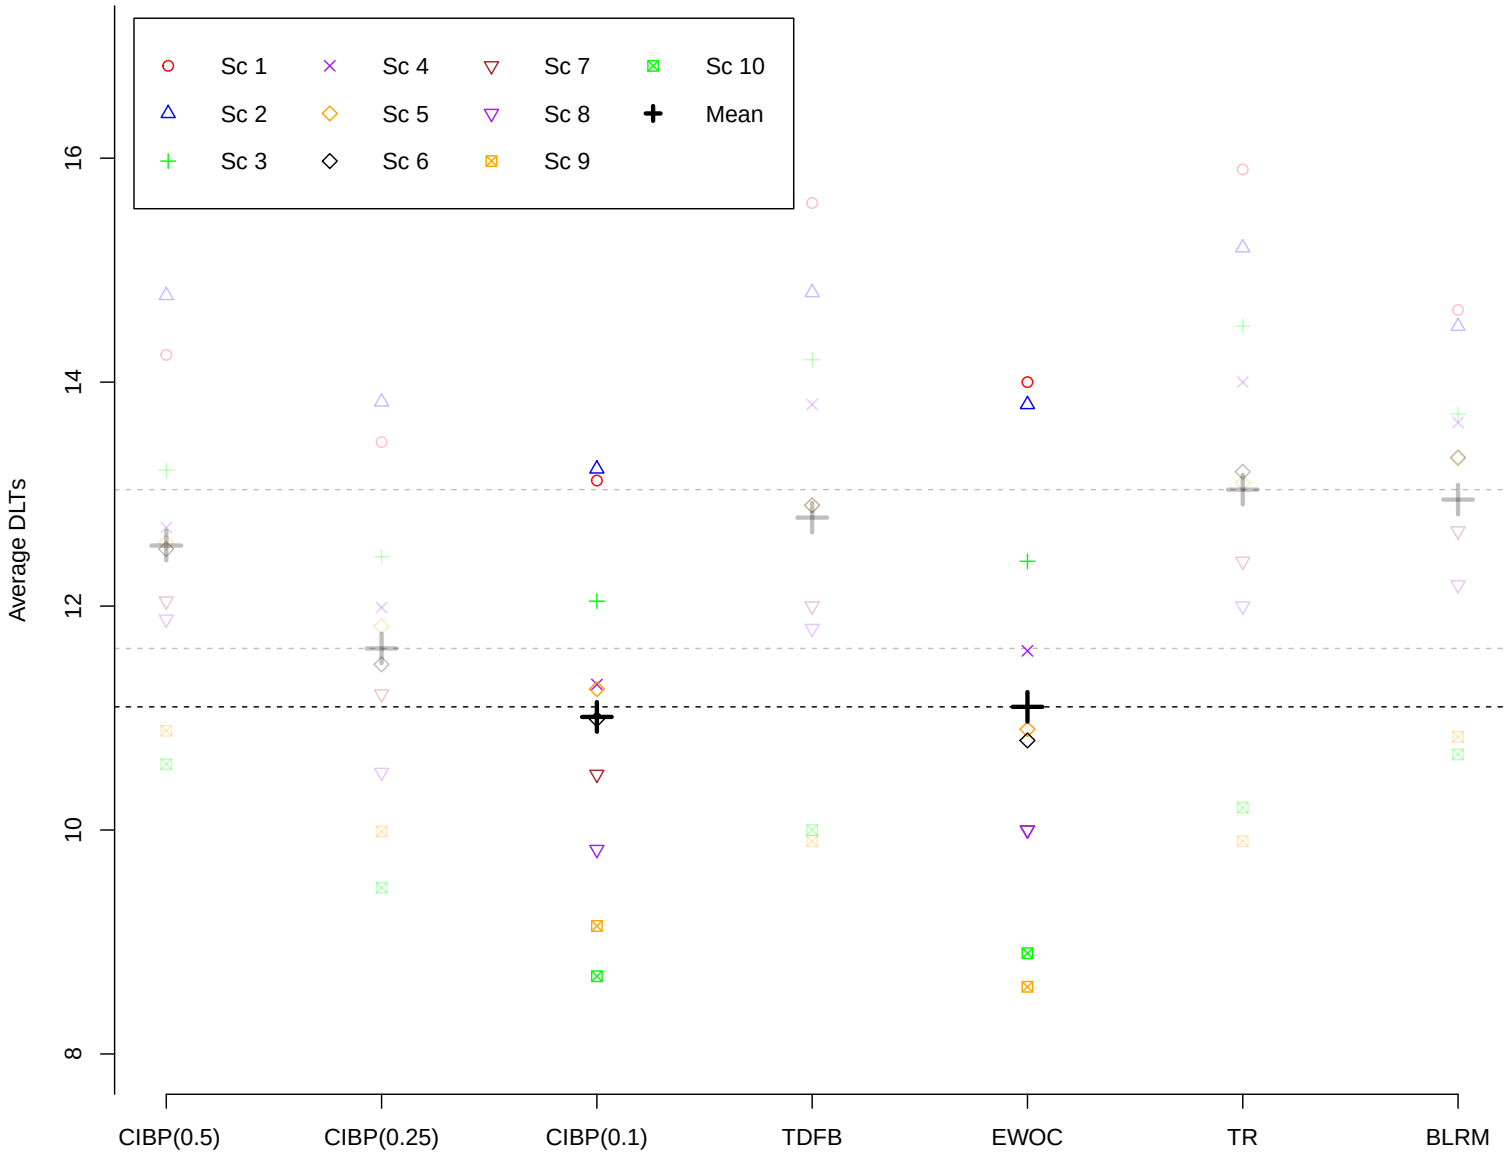

Conclusions

- The novel criterion requires **one additional parameter only**.
- The criterion incorporated into the one-parameter CRM method is found to result in
  1. **Similar** accuracy, but **fewer** mean number of DLTS.
  2. **Greater** accuracy, but **similar** mean number of DLTs.
- The new criterion allows to make model-based design **more ethical** as it does not lead to any decrease in accuracy.
- Criterion can be motivated by information theory and used by itself (Mozgunov and Jaki, 2018)

Aitchison, J. (1992) On criteria for measures of compositional difference. *Mathematical Geology*, **24**, 365–379.

Babb, J., Rogatko, A. and Zacks, S. (1998) Cancer phase I clinical trials: efficient dose escalation with overdose control. *Statistics in Medicine*, **17**, 1103–1120.

Mozgunov, P. and Jaki, T. (2018) An information-theoretic approach for selecting arms in clinical trials. *Preprint, arXiv:1708.02426*.

Neuenschwander, B., Branson, M. and Gsponer, T. (2008) Critical aspects of the Bayesian approach to phase I cancer trials. *Statistics in Medicine*, **27**, 2420–2439.

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

Tighiouart, M., Rogatko, A. et al. (2010) Dose finding with escalation with overdose control (ewoc) in cancer clinical trials. *Statistical Science*, **25**, 217–226.

Wheeler, G. M., Sweeting, M. J. and Mander, A. P. (2017) Toxicity-dependent feasibility bounds for the escalation with overdose control approach in phase I cancer trials. *Statistics in Medicine*.

Information theory

1) A statistical experiment of estimation of a toxicity probability.

The Shannon differential entropy (DE) $h(f_n)$ of the PDF f_n is defined as

$$h(f_n) = -\int_0^1 f_n(p)\log f_n(p)\mathrm{d}p \tag{6}$$

with the convention $0\log 0 = 0$.

Information theory

1) A statistical experiment of estimation of a toxicity probability.
The Shannon differential entropy (DE) $h(f_n)$ of the PDF f_n is defined as

$$h(f_n) = -\int_0^1 f_n(p)\log f_n(p)\mathrm{d}p \tag{6}$$

with the convention $0\log 0 = 0$.

2) A statistical experiment of a sensitive estimation.
The weighted Shannon differential entropy (WDE) , $h^{\phi_n}(f_n)$, of the RV $Z^{(n)}$ with positive weight function $\phi_n(p) \equiv \phi_n(p, \alpha, \gamma)$ is defined as

$$h^{\phi_n}(f_n) = -\int_0^1 \phi_n(p) f_n(p)\log f_n(p)\mathrm{d}p. \tag{7}$$

Weight Function

The Beta-form weight function

$$\phi_n(p) = \Lambda(\gamma, x, n) p^{\gamma\sqrt{n}} (1-p)^{(1-\gamma)\sqrt{n}}. \tag{8}$$

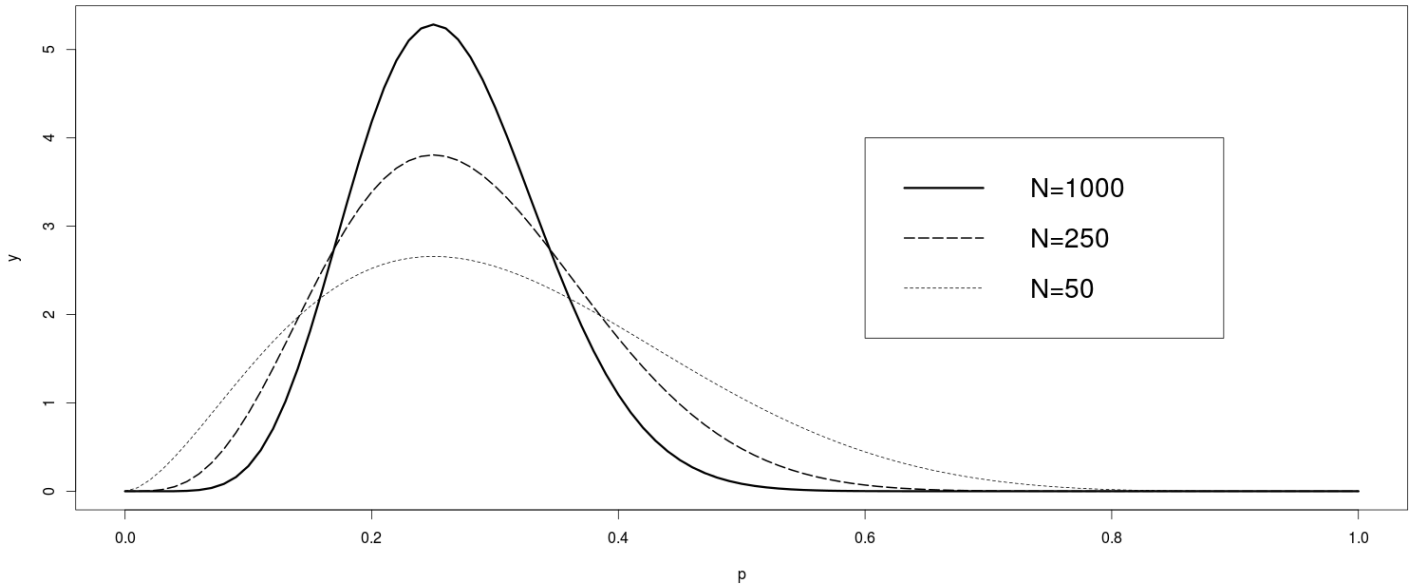

Additional information for sensitive estimation

$$h^{\phi_n}(f_n) - h(f_n) \tilde{=} \frac{(\alpha - \gamma)^2}{\alpha(1-\alpha)}$$

Additional information for sensitive estimation

$$h^{\phi_n}(f_n) - h(f_n) \tilde{=} \frac{(\alpha - \gamma)^2}{\alpha(1-\alpha)}$$

Can be estimated for each regimen j

$$\hat{\Delta}_j = \frac{(\hat{p}_j - \gamma)^2}{\hat{p}_j(1-\hat{p}_j)}$$

Escalation design

NMA (Mozgunov and Jaki, 2018)

Let $d_j(i)$ be a regimen d_j recommended for cohort i.

- The procedure starts from $\hat{\Delta}_j^{(0)}$
- l cohorts were already assigned

The $(l+1)^{th}$ cohort of patients will be assigned to regimen k such that

$$d_j(l+1): \ \hat{\Delta}_k^{(l)} = \inf_{i=1,\ldots,m} \hat{\Delta}_i^{(l)}, \ l = 0, 1, 2, \ldots, C.$$

We adopt regimen $d_j(C+1)$ as the final recommended regimen.

Asymmetry parameter (II)

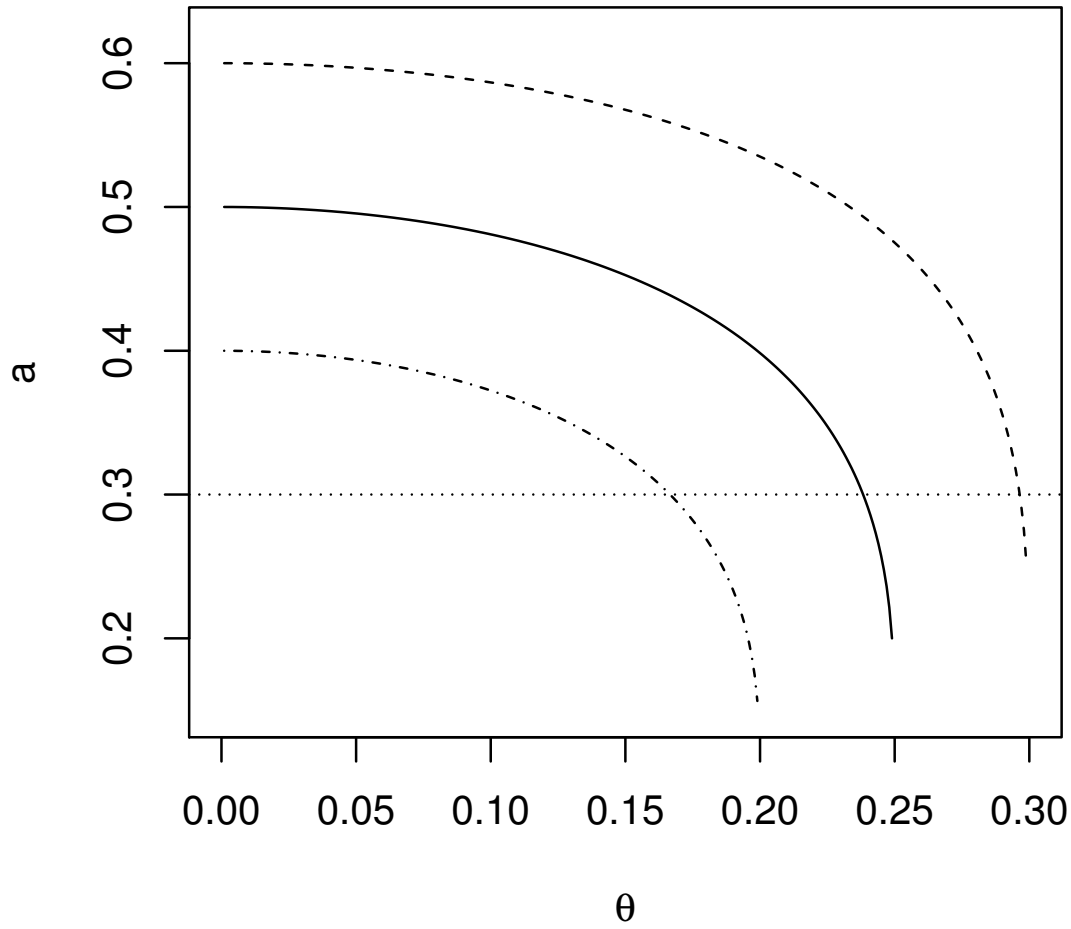

Comparators

We compare the performance of the proposed approach to

- **EWOC** design using fixed $\alpha = 0.25$
- **TR** design by Tighiouart et al. (2010) using $\alpha_2 = \ldots = \alpha_9 = 0.25$, $\alpha_n = \min(\alpha_{n-1} + 0.05, 0.50)$.
- Toxicity-dependent feasibility bound (**TDFB**) by Wheeler et al. (2017)

$$\alpha_{n+1} = \min\left(0.50, 0.25 + (0.50 - 0.25\frac{n - 1 - \sum_{i=1}^{n} y_i}{12\frac{2}{3}}\right)$$

- **BLRM** by Neuenschwander et al. (2008)
  We use the same prior distribution as Neuenschwander et al. (2008).