Bayesian adaptive designs for early phase clinical trials

Yuan Ji

yuanji@mdanderson.org

Division of Quantitative Science

University of Texas M.D. Anderson Cancer Center

Outline

I will describe three Bayesian adaptive designs for three types of early phase clinical trials.

- Phase I dose-finding trials based on a binary toxicity response – phase I tox
- Phase I dose-finding trials based on a time-to-event toxicity response with late onset toxicity – phase I tite tox
- Phase II multiple-arm randomized trials with adaptive randomization – phase II AR

Phase I tox

In phase I dose-finding based on toxicity:

- Oncologists want to find an appropriate dose level that is effective to the disease and yet is not "too toxic"
- For example, $< 30\%$ of the patients will experience the dose-limiting toxicity (DLT) – A binary random outcome
- The highest dose of which the probability of toxicity is less than p_T, say, $p_T = 0.30$, is called the maximum tolerated dose, or the MTD.

Underlying assumption: a higher dose is more effective

Trial set up and an up/down principle

- Align the doses from the lowest to the highest, say dose $1, 2, \ldots, 8$.
- Treat the first cohort of patients (cohort size ≥ 1) at the starting dose
- Depending on the observed binary toxicity outcomes, make a decision on the dose level for treating the next cohort
  - If the observed toxicity rate is much greater than the target p_T, the decision should be to de-escalate;
  - If the observed toxicity rate is close to the target p_T, the decision should be to stay – continue to treat patients at the current dose;
  - If the observed toxicity rate is much lower than the target p_T, the decision should be to escalate.
- By adaptively changing the dose levels at which patients are treated, the goal of the trial is to find the MTD

A typical set of observation

Notation: E = Escalation; S = Stay; D = De-escalation.

Current dose	observed toxicity	Decision
1	0/3	E
2	0/3	E
3	2/3	D
2	2/6	S
2	...	...

At the end of the trial, one dose is selected as the estimated MTD.

In the above case, probably dose 2 will be selected.

A practical Bayesian deign

The proposed design 1) provides the decisions need to be made at every step of the trial and 2) selects a final dose as the estimated MTD at the end.

- The method is implemented in an Excel macro to be demonstrated next (http://odin.mdacc.tmc.edu/$\sim$yuanj)
- Suppose patients are treated at dose i
- Identify the number of patients treated at this dose and go to the corresponding column in the table;
- Identify the number patients experienced toxicity and go to the corresponding row in the table;
- the corresponding entry in the cell provides the dose-assignment decision that one needs to take.

Probability model

- Likelihood function is a product of binomial densities: $l(\boldsymbol{p}) \propto \prod_{i=1}^{d} p_i^{x_i}(1-p_i)^{n_i - x_i}$, where n_i and x_i are the numbers of patients treated at dose i and experienced DLT, respectively.
- The priors of p_i are i.i.d. Beta(α, α), where α takes a small value, e.g., $\alpha = 0.005$, resulting in a U-shaped prior.
- Posteriors are beta with known parameter values.

Note: One can assume a dose-response curve (the gain of doing this for phase I trial is not clear)

Prior-posterior

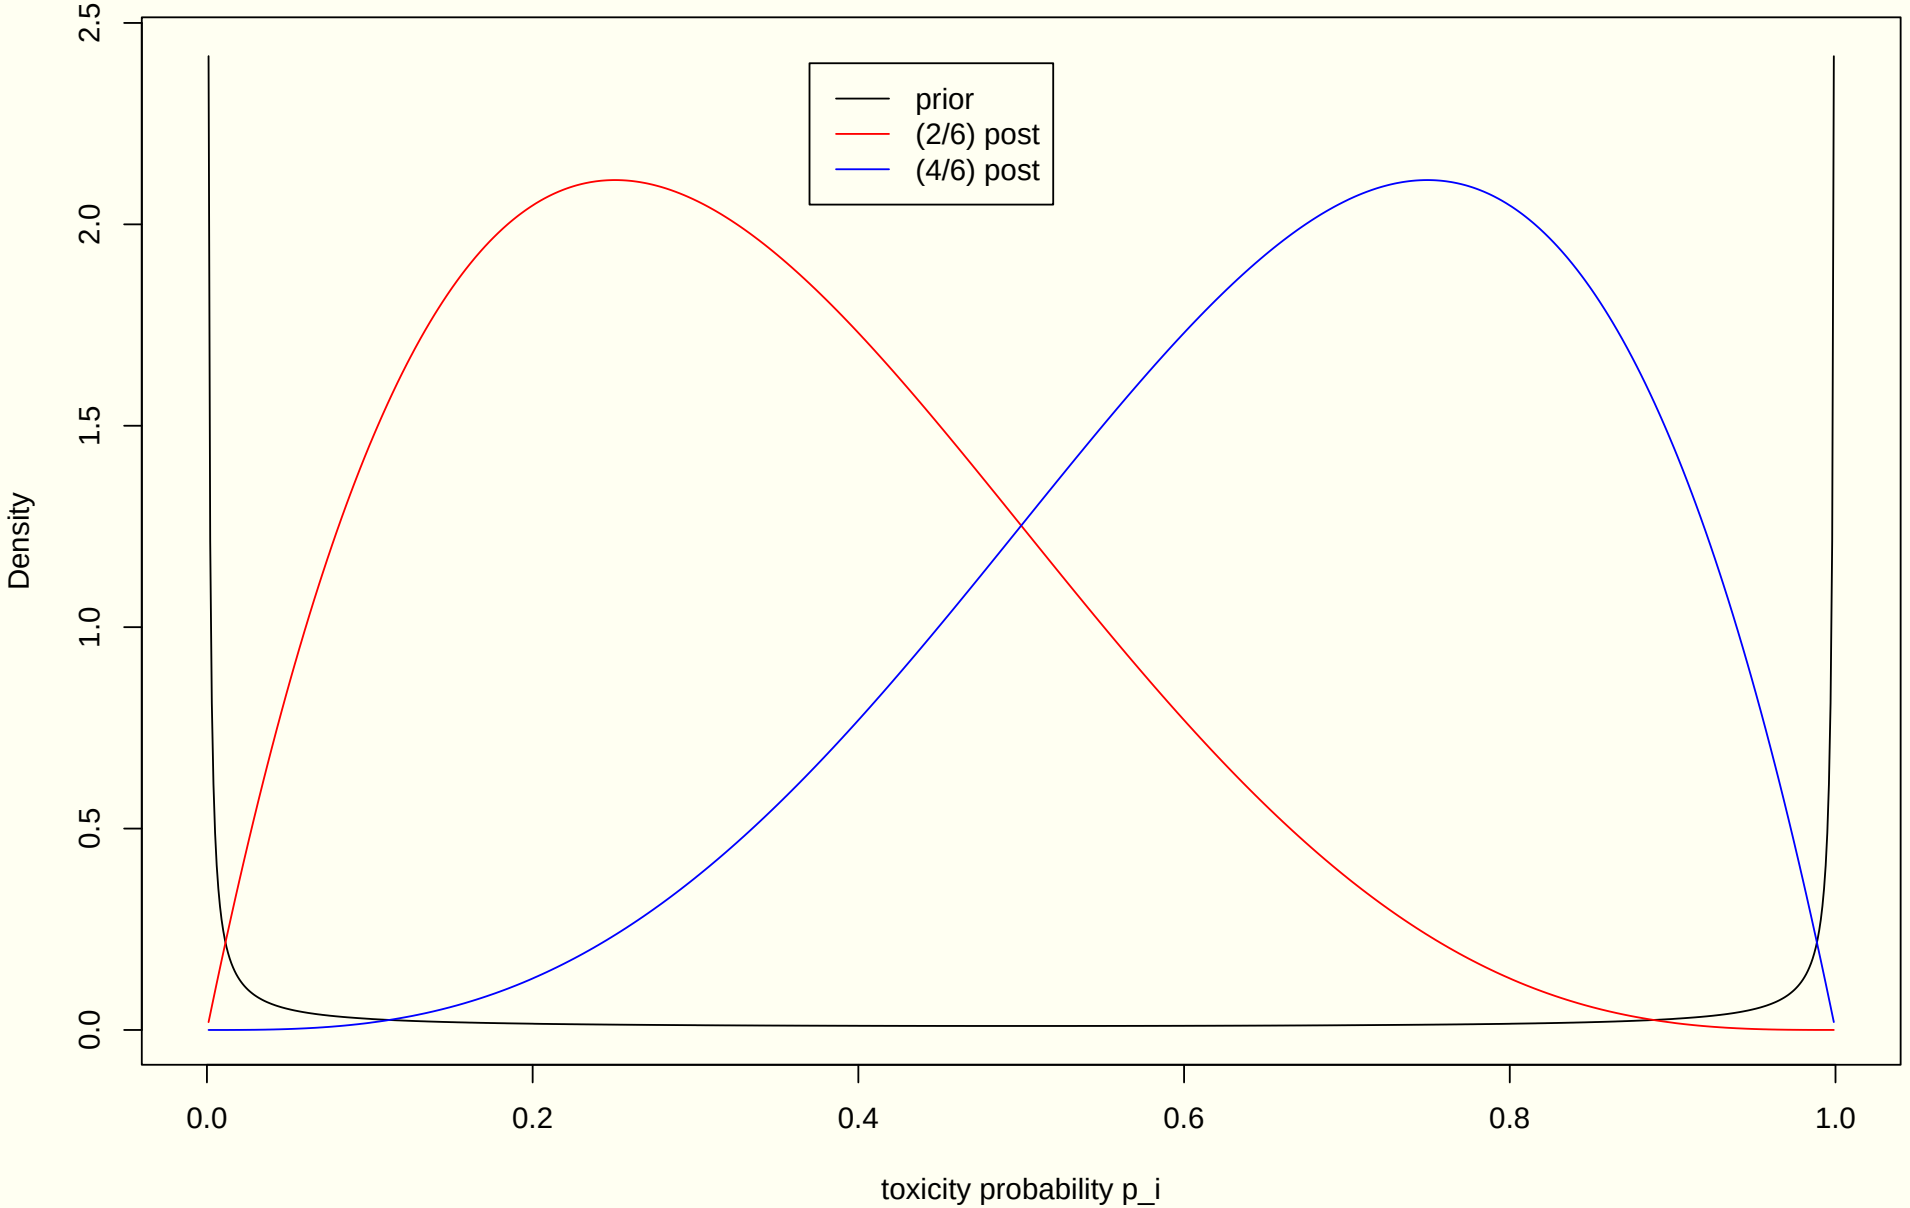

Decision rules

Let D, S, E denote the decision to de-escalate to dose $(i-1)$, stay at dose i, and escalate to dose $(i+1)$, respectively. Define the posterior probabilities for the three intervals:

$$q(D,i) = P(p_i - p_T > K_1\sigma_i|\text{data}),$$

$$q(S,i) = P(-K_2\sigma_i \leq p_i - p_T \leq K_1\sigma_i|\text{data}),$$

$$q(E,i) = P(p_i - p_T < -K_2\sigma_i|\text{data}).$$

The dose-assignment rule

$$\mathcal{B}_i = \arg \max_{m\in\{D,S,E\}} q(m,i),$$

i.e. take the decision that has the maximum posterior probability.

- A decision-theoretic interpretation of these rules can be found in Ji et al. (2007, Stat Sinica)

Two issues

- What if the first dose is very toxic?
- What if dose $i-1$ is safe, but dose i is very toxic? For example, $p_{i-1} = 0.05$ and $p_i = 0.6$ (while the $p_T = 0.3$).

Exclusion rule

Define

$$\mathcal{T}_i = 1\left\{P(p_i > p_T|\text{data}) > \xi\right\},$$

where $1\{\}$ is the indicator function and $\xi \in (0,1)$ is a cutoff value (e.g., $\xi = 0.95$). For a large value of ξ, $\mathcal{T}_i = 1$ implies that dose i is very likely to be highly toxic, and escalation to this dose should be permanently prohibited. To incorporate this rule, modified decision rule is given by

$$\mathcal{B}_i^{(e)} = \arg \max_{m \in \{D,S,\tilde{E}\}} q(m,i),$$

where $q(\tilde{E},i) = q(E,i)(1-\mathcal{T}_{i+1})$.Therefore, if $\mathcal{T}_{i+1} = 1$, the probability $q_{\tilde{E},i}$ equals zero and the assignment rule $\mathcal{B}_i^{(e)}$ can be only D, to de-escalate, or S, to stay.

Dose-finding algorithm

- Suppose that the current tried dose is i for $i = 1, \ldots, d$. After the toxicity outcomes of the last cohort are observed, select the dose for treating the next cohort among $\{(i-1), i, (i+1)\}$ based on the assignment rule $\mathcal{B}_i^{(e)}$. There are two exceptions: if $i = 1$, the next available doses are $\{i, (i+1)\}$; if $i = d$, the next available doses are $\{(i-1), i\}$.

Dose-finding algorithm con't

- Suppose that dose 1 is a dose that has been tried previously. If $\mathcal{T}_1 = 1$, terminate the trial due to excessive toxicity. Otherwise, terminate the trial when the maximum sample size is reached.
- In the special case of cohort of size 1, by convention, do not apply the exclusion rule $\mathcal{T}_i$ until two or more patients have been evaluated at a dose.

Dose-finding algorithm con't

- At the end of the trial, select the dose as the estimated MTD with the smallest difference $|\hat{p}_i^* - p_T|$ among all the tried doses i for which $\mathcal{T}_i = 0$.
- Quantity $\hat{p}_i^*$ is the isotonic transformation estimator of the posterior mean $\hat{p}_i$ so that $\hat{p}_j^* \leq \hat{p}_i^*$ for $j > i$.
- If two or more doses tie for the smallest difference, perform the following rule. Let p^* denote the transformed posterior mean of the tied doses.
  - If $p^* < p_T$, choose the highest dose among the tied doses.
  - If $p^* > p_T$, choose the lowest dose among the tied doses.

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$								Tox %	Ave. n
	Dose	1	2	3	4	5	6	7	8		
Scenario 1		5	25	50	60	70	80	90	95	none		
Bayes	% MTD	13	**79**	8	0	0	0	0	0	0	25	30
	# Pts	7.7	16.1	5.8	0.5	0	0	0	0			
3+3	% MTD	24	**58**	16	2	0	0	0	0	0	25	12
	# Pts	4.0	5.0	2.6	0.4	0	0	0	0			
BCD	% MTD	10	**78**	11	1	0	0	0	0	0	24	30
	# Pts	11.4	11.5	5.2	1.4	0.3	0.1	0	0			
CFM	% MTD	6	**80**	14	0	0	0	0	0	0	29	30
	# Pts	5.2	16.3	7.5	0.9	0	0	0	0			
CRM	% MTD	6	**83**	11	0	0	0	0	0	0	27	30
	# Pts	5.7	18.6	4.9	1.0	0	0	0	0			

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$									Tox %	Ave. n
	Dose	1	2	3	4	5	6	7	8			
Scenario 2		1	2	3	4	5	6	25	50	none		
Bayes	% MTD	0	0	0	0	2	22	**62**	14	0	12	30
	# pt	3.2	3.2	3.4	3.5	3.7	4.5	5.9	2.6			
3+3	% MTD	0	0	0	2	3	21	**46**	8	0	11	27
	# pt	3.1	3.2	3.3	3.4	3.3	3.7	4.5	2.2			
BCD	% MTD	0	0	1	2	7	24	**56**	10	0	10	30
	# pt	3.2	3.6	3.6	3.5	3.8	5.4	4.8	2.1			
CFM	% MTD	0	0	0	0	1	22	**61**	16	0	12	30
	# pt	3.1	3.0	3.1	3.5	3.7	5.1	6.3	2.1			
CRM	% MTD	0	0	1	1	5	22	**50**	21	0	13	30
	# pt	3.1	3.4	3.3	3.7	3.6	4.4	5.1	3.4			

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$									Tox	Ave.
											%	n
	Dose	1	2	3	4	5	6	7	8			
Scenario 3		1	5	50	60	70	80	90	95	none		
Bayes	% MTD	0	**79**	21	0	0	0	0	0	0	22	30
	# pt	5.5	13.2	10.2	1.0	0	0	0	0			
3+3	% MTD	0	**70**	28	2	0	0	0	0	0	22	13
	# pt	3.1	5.2	4.4	0.7	0.1	0	0	0			
BCD	% MTD	0	**60**	39	1	0	0	0	0	0	22	30
	# pt	4.9	14.3	8.2	2.2	0.4	0	0	0			
CFM	% MTD	0	**56**	44	0	0	0	0	0	0	28	30
	# pt	3.1	11.7	13.1	2.0	0.1	0	0	0			
CRM	% MTD	0	**49**	51	0	0	0	0	0	0	26	30
	# pt	3.1	13.0	12.0	1.8	0	0	0	0			

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$									Tox %	Ave. n
	Dose	1	2	3	4	5	6	7	8			
Scenario 4**		40	50	60	70	80	90	95	99	none		
Bayes	% MTD	31	2	0	0	0	0	0	0	**67**	41	19
	# pt	16.8	1.8	0.2	0	0	0	0	0			
3+3	% MTD	38	9	1	0	0	0	0	0	**52**	43	6
	# pt	4.7	0.5	0.6	0.7	0	0	0	0			
BCD	% MTD	38	2	0	0	0	0	0	0	**60**	45	18
	# pt	12.6	4.6	1.2	0.2	0	0	0	0			
CFM	% MTD	38	3	1	0	0	0	0	0	**58**	42	14
	# pt	11.7	1.9	0.5	0.1	0	0	0	0			
CRM	% MTD	47	2	0	0	0	0	0	0	**51**	42	23
	# pt	20.2	2.5	0.2	0	0	0	0	0			

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$								Tox	Ave.
	Dose	1	2	3	4	5	6	7	8	%	n
Scenario 5		15	25	35	45	55	65	75	85	none	
Bayes	% MTD	31	**41**	21	7	0	0	0	0	24	30
	# pt	12.4	9.5	5.5	1.9	0.3	0	0	0		
3+3	% MTD	29	**37**	20	7	1	0	0	0	26	12
	# pt	4.4	3.9	2.4	0.9	0.2	0	0	0		
BCD	% MTD	21	**46**	22	6	1	0	0	0	26	29
	# pt	10.6	9.2	5.7	2.5	0.8	0.1	0	0		
CFM	% MTD	15	**44**	32	7	0	0	0	0	27	30
	# pt	8.0	10.6	8.0	2.6	0.4	0	0	0		
CRM	% MTD	36	**47**	14	2	0	0	0	0	23	30
	# pt	13.8	11.4	3.6	0.9	0.2	0	0	0		

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$								Tox %	Ave. n
	Dose	1	2	3	4	5	6	7	8		
Scenario 6		5	15	25	35	45	55	65	75	none		
Bayes	% MTD	2	24	**42**	24	7	0	0	0	0	22	30
	# pt	5.1	8.2	9.2	5.7	1.6	0.3	0	0			
3+3	% MTD	9	28	**34**	22	5	0	0	0	0	21	15
	# pt	3.6	4.3	3.8	2.3	0.8	0.2	0	0			
BCD	% MTD	1	29	**44**	19	6	1	0	0	0	21	30
	# pt	6.8	8.7	7.5	4.4	1.9	0.6	0.1	0			
CFM	% MTD	0	14	**49**	29	6	0	0	0	0	24	30
	# pt	3.9	6.2	10.7	7.1	1.8	0.3	0	0			
CRM	% MTD	4	37	**45**	12	2	0	0	0	0	20	30
	# pt	5.5	11.5	8.9	3.4	0.7	0.1	0	0			

Useful links

- Paper and software http://odin.mdacc.tmc.edu/∼yuanj

So far

we talked about

- phase I dose-finding trials
- with binary toxicity
- assume toxicity can be observed in a short time (the first cycle of treatment usually)

Phase I tite tox

- Most phase I trials use a binary variable indicating that a DLT has occurred within a time interval of fixed length t^*, which is usually called the assessment window
- Late onset toxicities refer to the toxicities occurs late toward the end of the assessment period
- Statistical methods that do not specifically address late onset toxicities may treat an undesirably large number of patients at toxic doses before any toxicities are observed

The TITE-CRM, by Cheung and Chappell (2000), attempts to address the late onset toxicity by modeling the time-to-toxicity.

Proposed methodology

- We propose a Bayesian method that possesses two new features:
  - The method contains a set of decision rules that temporarily suspend accrual if the risk of toxicity at prospective doses for future patients is unacceptably high
  - The method allows for restarting accrual if the risk of toxicity reduces to an acceptable level after additional followup data are observed.

Probability model

Notations

- Suppose we have K doses and let $d_1 < \cdots < d_K$ denote the K dose levels.
- First fix a sequence of times $0 = t_0 < t_1 < \cdots < t_{C-1} < t_C = \infty$, where $[t_0,\ t_{C-1}]$ = $[0,\ t^*]$ is the assessment window
- Denote T_i to be the time-to-toxicity random variable for patient i
- Define $Y_i = j$ if $t_{j-1} \leq T_i < t_j$, for $j = 1, \cdots, C$. ($Y_i = C$ means no toxicity within the assessment window)
- Let T_i^o be the observed time to toxicity or right censoring and Y_i^o be the observed index so that $Y_i^o = j$ if $t_{j-1} \leq T_i^o \leq t_j$.
- Finally, let $\delta_i = 1$ if $T_i^o = T_i$ (i.e., toxicity has occurred) and $\delta = 0$ if $T_i^0 < T_i$ (i.e.,toxicity has not occurred).

Probability model

Discrete time hazard model

- Define discrete time hazard

$$\Phi(\beta_{j,k}) = \Pr(Y_i = j \mid Y_i \geq j, d_k) = \frac{\Pr(Y_i = j \mid d_k)}{\Pr(Y_i \geq j \mid d_k)}$$

- The probability of toxicity during the jth interval is

$$\Pr(Y_i = j \mid d_k) = \Phi(\beta_{j,k}) \prod_{h=1}^{j-1} \{1 - \Phi(\beta_{h,k})\},$$

- The probability of toxicity not occurring by t_j is

$$\Pr(Y_i > j \mid d_k) = \prod_{h=1}^{j} \{1 - \Phi(\beta_{h,k})\} \text{ for } j \leq C - 1.$$

Probability model

Likelihood function

- Denote $k(i)$ be the index (“level") of the dose administered to the i^{th} patient.
- At any point in the trial, the discretized data from the current n patients take the form $D_n = \{(Y_i^o, k(i), \delta_i), i = 1, ..., n\}$.
- Denoting $\boldsymbol{\beta} = (\beta_{1,1}, \cdots, \beta_{C-1,K})$, the likelihood is

$$L(\boldsymbol{\beta}|D_n) = \prod_{i=1}^{n} \Pr(Y_i = Y_i^0)^{\delta_i} \Pr(Y_i > Y_i^0 - 1)^{1-\delta_i}$$

$$= \prod_{i=1}^{n} \Phi\left(\beta_{Y_i^o,k(i)}\right)^{\delta_i} \prod_{h=1}^{Y_i^o-1} \left\{1 - \Phi(\beta_{h,k(i)})\right\}.$$

Latent variables

To facilitate computation of posterior quantities, following Albert and Chib (1993, 2001), and Chib and Greenberg (1998), we express the likelihood using a latent variable formulation.

- For patient i, define the vector of latent variables $\mathbf{Z}_i$ = $(Z_{i,1}, \cdots, Z_{i,Y_i^o})$ if $Y_i^o < C$ and $\mathbf{Z}_i$ = $(Z_{i,1}, \cdots, Z_{i,C-1})$ if $Y_i^o = C$.
- Assume that $Z_{i,j} \sim N(\beta_{j,k}, 1)$ if the i^{th} patient received dose d_k.
- The likelihood may be augmented with the latent variables and re-expressed as

$$L(\boldsymbol{\beta}|D_n, \mathbf{Z}) = \prod_{i=1}^{n} \left\{\phi(Z_{i,Y_i^o}; \beta_{Y_i^o,k(i)}, 1) I(Z_{i,Y_i^o} < 0)\right\}^{\delta_i} \prod_{j=1}^{Y_i^o - 1} \phi(Z_{i,j}; \beta_{j,k(i)}, 1)\ I(Z_{i,j} > 0)$$

- The augmented likelihood only involves the normal pdf while the likelihood involves normal CDF.

Prior

For each dose d_k, we assume a state space model for the prior of $\boldsymbol{\beta}_j = (\beta_{j,1}, \cdots, \beta_{j,K})$, defined by the recursive relationship $\beta_{j,k} \mid \beta_{j,k-1} \sim N(\beta_{j,k-1}, \sigma_\beta^2)$ for $j = 2, ..., C-1$, with $\beta_{j1} \sim N(\beta_{j0}, \sigma_\beta^2)$ and β_{0k} fixed to ensure identifiability.

This gives the joint prior

$$f(\boldsymbol{\beta}_j) \propto \prod_{k=1}^{K} \phi(\beta_{j,k}; \beta_{j,k-1}, \sigma_\beta^2)$$

for each $k = 1, \cdots, K$.

Posterior computation

Denote $A_0 = (-\infty, 0]$, $A_1 = (0, \infty)$, and $\bar{A}_{i,j} = A_{1-I(Y_i^o=j,\ \delta_i=1)}$. The following process is initialized using the prior mean of $\boldsymbol{\beta}$, steps 1 and 2 are iterated until convergence.

Step 1. Generation of the latent variables. Generate each $Z_{i,j}$ independently from the full conditional which follows a truncated normal distribution $\phi(z; \beta_{j,k(i)}, 1) I(z \in \bar{A}_{i,j})$.

Posterior computation

Step 2. Generation of $\boldsymbol{\beta}$. Denote

$$S_{j,k} = \sum_{i=1}^{n} \sum_{h=1}^{Y_i^o} I\{k[i] = k, h = j\}\sigma_\beta^2 \text{ and } Z_{j,k}^{+} = \sum_{i=1}^{n} \sum_{h=1}^{Y_i^o} Z_{i,h} I(k[i] = k, h = j).$$

Given $\mathbf{Z}$ and the current data, generate $\boldsymbol{\beta}$ from its full conditional distribution under which, for $k = 1, ..., K$, $\beta_{j,k}$ is normal with mean $\tilde{\beta}_{j,k} = \{\sigma_\beta^2 Z_{j,k}^{+} + \beta_{j,k-1} + I(k < K)\beta_{j,k+1}\}/\{1 + I(k < K) + S_{j,k}\,\sigma_\beta^2\}$ and variance $\tilde{\sigma}_{\beta,k}^2 = \sigma_\beta^2/\{1 + I(k < K) + S_{j,k}\,\sigma_\beta^2\}$.

Posterior inference

Our posterior inference will be based on the conditional probabilities

$$\pi(\boldsymbol{\beta}, d_k, j) = \Pr(Y \leq C-1 \mid Y \geq j, \boldsymbol{\beta}, d_k)$$

for $j = 1, \cdots, C-1$ and $k = 1, \cdots, K$.

- $\pi(\boldsymbol{\beta}, d_k, Y^o)$ is the probability that a patient who has survived $Y^o - 1$ intervals without toxicity will experience toxicity by t^* = t_{C-1} at dose d_k.
- Since Pr($Y \geq 1$) = 1, the unconditional probability of toxicity within the window $[0,\ t^*]$ is $\pi(\boldsymbol{\beta}, d_k, 1)$ (denoted as $\pi(\boldsymbol{\beta}, d_k)$)
- It follows from that $\pi(\boldsymbol{\beta}, d_k, j) = 1 - \prod_{h=j}^{C-1} \{1 - \Phi(\beta_{h,k})\}$, and in particular $\pi(\boldsymbol{\beta}, d_k) = 1 - \prod_{h=1}^{C-1} \{1 - \Phi(\beta_{h,k})\}$.

Bayesian isotonic regression

From the previous two-step computational algorithm, we obtain a posterior sample of $\boldsymbol{\beta}$, which leads to a posterior sample of $\pi(\boldsymbol{\beta}, d_k, j)$.

Since it is often assumed in phase I trials that toxicity increases with dose, we apply the Bayesian isotonic regression transformation (Dunson and Neelon, 2003) to the posterior sample of $\pi(\boldsymbol{\beta}, d_k, j)$.

The order-transformed posterior is denoted as $\tilde{\pi}(\boldsymbol{\beta}, d_k, j)$.

Bayesian isotonic regression

After applying **Step 1** and **Step 2** until convergence, we apply to the following step to the resulting posterior samples of $\pi(\boldsymbol{\beta}, d_k, j)$.

Step 3. Apply the Dunson-Neelon algorithm (2003) to $\boldsymbol{\pi}(\boldsymbol{\beta}, j) = (\pi(\boldsymbol{\beta}, d_1, j), ..., \pi(\boldsymbol{\beta}, d_K, j))$ as follows.

The vector obtained by the Dunson-Neelon Bayesian isotonic regression transformation is

$$\tilde{\pi}(\boldsymbol{\beta}, d_k, j) = \min_{k_2 \in U_k} \max_{k_1 \in L_k} \left(\frac{\mathbf{1}'_{k_2-k_1+1} \mathbf{V}^{-1}_{j,[k_1:k_2]} \boldsymbol{\pi}(\boldsymbol{\beta}, j)_{[k_1:k_2]}}{\mathbf{1}'_{k_2-k_1+1} \mathbf{V}^{-1}_{j,[k_1:k_2]} \mathbf{1}_{k_2-k_1+1}} \right),$$

where $L_k = \{s : s \leq k\}$, $U_k = \{s : s \geq k\}$ and $\mathbf{1}_k$ is the k-vector with all entries 1.

This transformation ensures that
$\tilde{\pi}(\boldsymbol{\beta}, d_1, j) \leq \tilde{\pi}(\boldsymbol{\beta}, d_2, j) \leq ... \leq \tilde{\pi}(\boldsymbol{\beta}, d_K, j)$ for all j.

Posterior quantities

The decision rules for our method rely on two different types posterior quantities:

- the posterior probabilities $\xi_k(D_n)$ = Pr$\{\tilde{\pi}(\boldsymbol{\beta}, d_k, 1) > \pi^* | \ D_n\}$ for $k = 1, \cdots, K,$
- and the predictive probabilities based on approximate values of the ξ_k's that involve both D_n and future outcomes.

Posterior predictive

Let m_k be the number of patients treated at dose d_k who have not been fully evaluated, indexed by $i_1, \ldots, i_{m_k}$.

Define the indicator $W_{i_r} = I(T_{i_r} \leq t^*)$ that patient i_r will eventually have toxicity within the assessment window, $r = 1, \ldots, m_k$.

Let $S(\boldsymbol{W}_k) = \sum_{r=1}^{m_k} W_{i_r}$ be the number of patients among the m_k who will have toxicity by time t^*.

Denote $p_k(\boldsymbol{w}, m_k)$ the posterior probability of toxicity for dose d_k after $S(\boldsymbol{W}_k)$ additional patients experience toxicity.

Posterior predictive

Suppose π^* is target toxicity rate of the MTD (e.g, 0.3), $.05 \leq \underline{\xi} \leq .30$ and $.70 \leq \overline{\xi} \leq .95$ are predetermined cutoffs. Define two predicted risks of toxicity (PRT) as

$$PN_k(D_n) = \sum_{\mathbf{w}} I[\Pr\{p_k(\mathbf{w}, m_k) > \pi^*\} \leq \underline{\xi}] \Pr(\mathbf{W}_k = \mathbf{w} \mid D_n),$$

and

$$PE_k(D_n) = \sum_{\mathbf{w}} I[\Pr\{p_k(\mathbf{w}, m_k) > \pi^*\} \geq \overline{\xi}] \Pr(\mathbf{W}_k = \mathbf{w} \mid D_n),$$

where $\Pr(\mathbf{W}_k = \boldsymbol{w}|D_n) = \int \prod_{r=1}^{m_k} \Pr(W_{i_r} = w_{i_r}|\tilde{\pi}(\boldsymbol{\beta}, d_k, Y_{i_r}^o)) f(\boldsymbol{\beta}|D_n) d\boldsymbol{\beta}$

is the posterior predictive probability.

PRT

The predicted risks of toxicity

- $PN_k(D_N)$ and $PE_k(D_n)$ are approximately predictive probabilities that d_k has negligible or excessive toxicity, respectively.
- $PA_k(D_n) = 1 - PN_k(D_N) - PE_k(D_n)$ is the predictive probability that d_k has acceptable toxicity.

The key idea is that the accrual will be suspended if the posterior probabilities $\xi_k(D_n)$ suggests a different dose for future patients from that suggested by the PRT

Decision rules

Let d_k denote the current dose. Recall that d_k has negligible toxicity if $\xi_k(D_n) < \underline{\xi}$, and it is excessively toxic if $\xi_k(D_n) > \overline{\xi}$. The trial is conducted as follows:

1) The first cohort of patients are treated at a starting dose chosen by the physicians.

2) No untried dose may be skipped when escalating.

3) At any point in the trial, if $\xi_1(D_n) > \overline{\xi}$, then stop the trial and conclude that none of the doses are acceptably safe.

4) If $\xi_k(D_n) > \overline{\xi}$ and $k > 1$ then de-escalate to the highest dose $k' < k$ such that $\xi_{k'}(D_n) \leq \overline{\xi}$.

Decision rules

5) For lower probability cut-off ϵ, if $\underline{\xi} \leq \xi_k(D_n) \leq \overline{\xi}$ and

5.1) $PE_k(D_n) \leq \epsilon$ then treat next cohort at d_k,

5.2) $PE_k(D_n) > \epsilon$ then suspend accrual and reconsider enrolling the patient when the data for the patients currently enrolled in the trial have been updated, which occurs when a toxicity is observed or a patient advances from one interval to the next.

Decision rules

6) If $\xi_k(D_n) < \underline{\xi}$ and

6.1) $k = K$ then treat the next patient at d_K.

6.2) $k < K$, $PN_k(D_n) \geq 1 - \epsilon$ and $PE_{k+1}(D_n) \leq \epsilon$ then treat next cohort at d_{k+1},

6.3) $k < K$, $PN_k(D_n) \geq 1 - \epsilon$ and $PE_{k+1}(D_n) > \epsilon$ then suspend accrual as in 5.2,

6.4) for $k < K$ and $PN_k(D_n) < 1 - \epsilon$ then suspend accrual as in 5.2.

7) At the end of the trial, among set of acceptable doses $\{j : \xi_j(D_{N_{max}}) \leq \overline{\xi},\ j = 1, \cdots, K\}$, select the dose minimizing $\mid \mathrm{E}\{\pi(\boldsymbol{\beta}, d_j) \mid D_n\} - \pi^* \mid$.

Decision rules

- Rules (5) and (6) utilize the PRT criteria when the risk of toxicity at the current dose based on the current data is either acceptable or negligible.
- These rules exploit the fact that predictive probabilities provide information about the risk of future toxicities that cannot be obtained from posterior probabilities alone.

Trial example

We illustrate our method with a clinical trial of a single agent that was conducted in patients with advanced leukemia.

- The compound were shown to be safe at 1 and 1.5 units.
- A higher dose may be needed to achieved to durable remission for advanced cancer.
- A trial was begun at a dose of 3 units, which was higher than previously tried doses shown to be safe.
- During a period of 6 weeks, a total of 7 patients were enrolled.
- Within 6 weeks thereafter, 6 of the 7 patients experienced severe irreversible DLTs.

Trial example

In contrast, using our proposed method we would design this trial to evaluate six dose levels 1, 1.5, 2, 2.5, 3, 3.5 units for three days (denoted doses 1 through 6, respectively) with starting dose 1.5 units.

We assumed that

- accrual rate 4/month
- assessment window $t^* = 3$ month
- Max. sample size 36 patients
- Target toxicity rate π^* = .30
- ϵ = .05, $\underline{\xi}$ = .30, $\overline{\xi}$ = .90
- Cohort size 3

Event chart

A representative example of how such a phase I trial would proceed is given in the Event Chart displayed in the following figure.

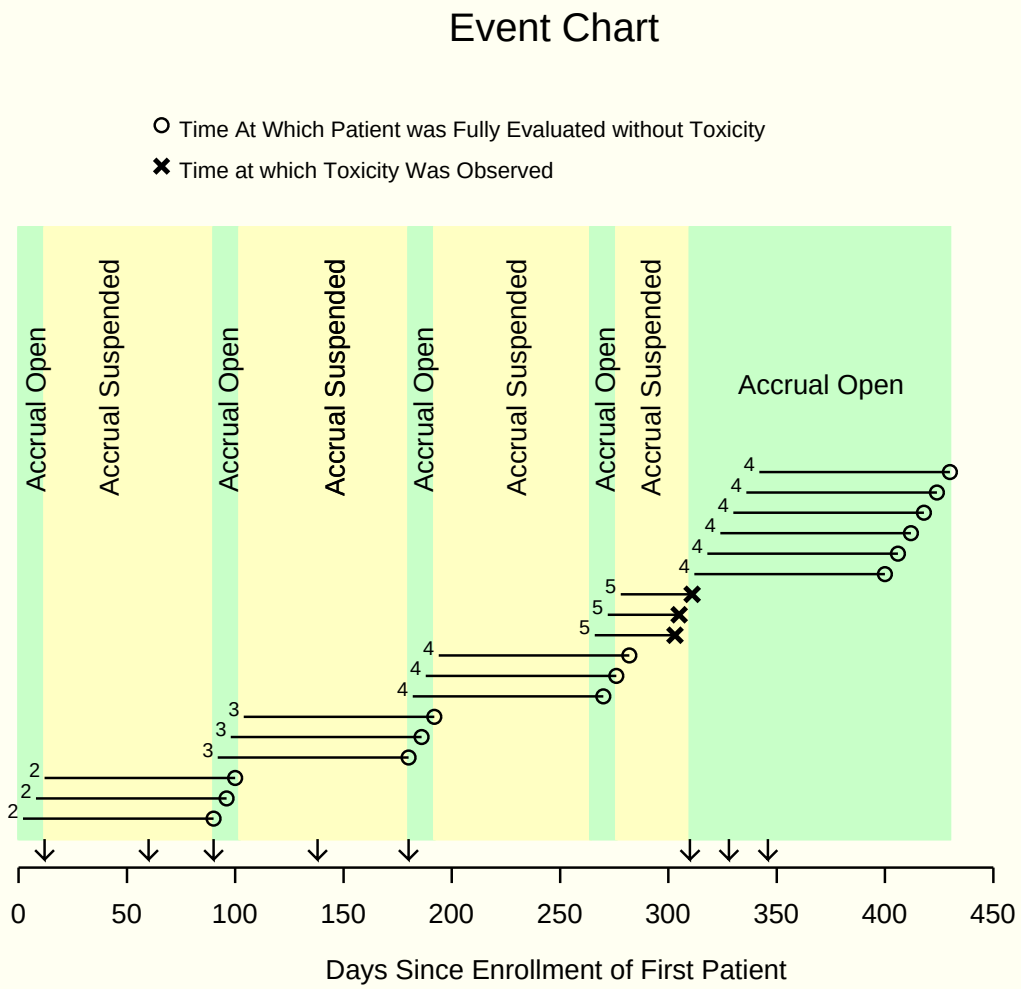

Event chart

- The PRT method would have suspended accrual long enough to observed at least 2 of the dose limiting toxicities associated with this dose.
- It ultimately would have saved at least 3 lives because subsequent patients would not have been treated at this excessively toxic dose.

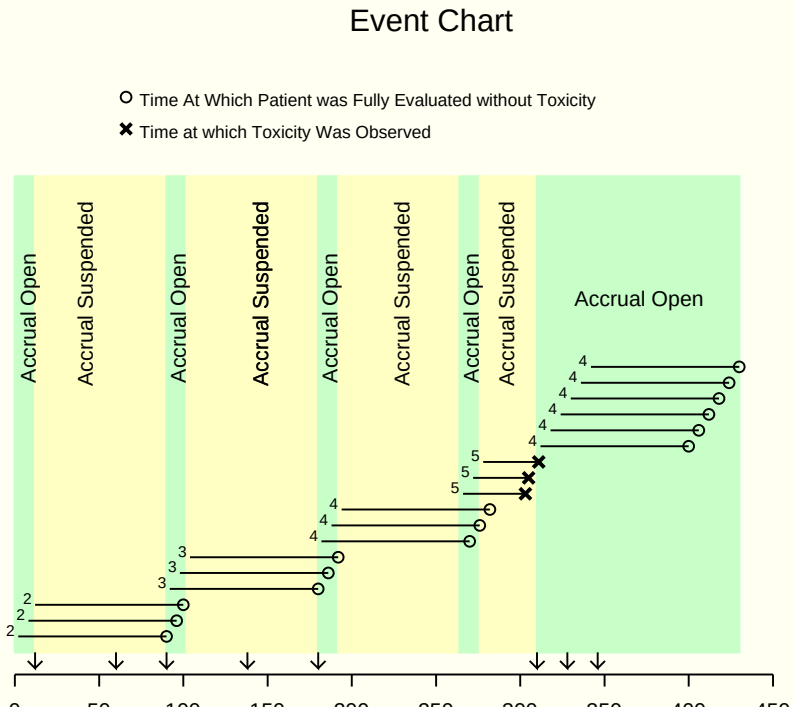

Hazard function

Late Onset Hazard Function

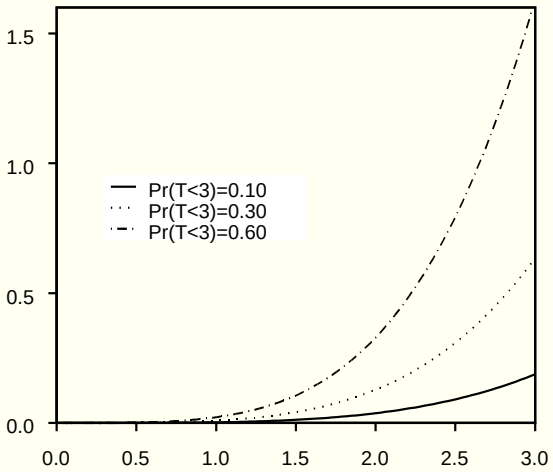

Cum. Prob. of Tox- Late Onset

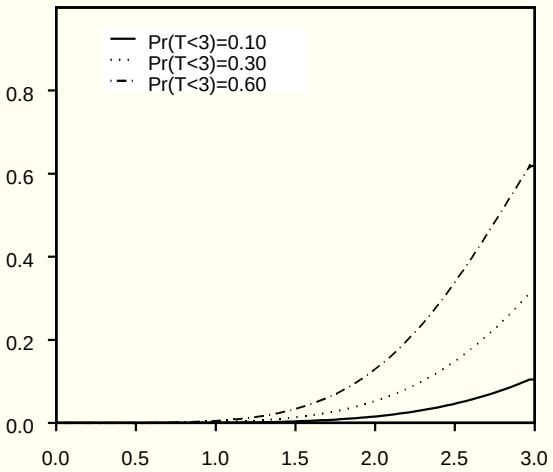

Early Onset Hazard Function

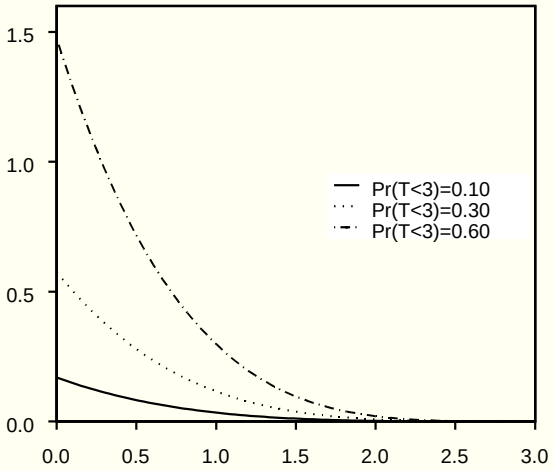

Cum. Prob. of Tox- Early Onset

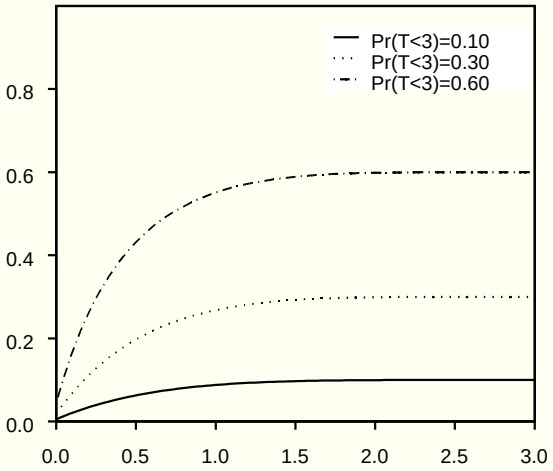

Simulation results

		True Prob($T < 3\ months \mid d_k$)								
		d_1	d_2	d_3	d_4	d_5	d_6			
Scenario 1 (Late Onset)		0.03	0.05	0.10	0.30	0.50	0.60	None	Total	Duration
PRT	*% Selected*	0	1	26	64	9	0	0	–	1.8
	# Patients	3.5	4.8	9.8	12.5	4.6	0.8		36.0	–
	# Toxicities	0.1	0.3	1.0	3.7	2.2	0.5		7.8	–
Tite CRM	*% Selected*	0	0	12	70	19	0	0	–	1.0
	# Patients	3.0	3.2	4.5	11.2	9.4	4.7		36.0	–
	# Toxicities	0.1	0.2	0.4	3.4	4.7	2.8		11.6	–

Simulation results

		True Prob($T < 3\ months \mid d_k$)								
		d_1	d_2	d_3	d_4	d_5	d_6			
Scenario 2 (Late Onset)		0.50	0.60	0.68	0.73	0.76	0.78	None	Total	Duration
PRT	*% Selected*	11	0	0	0	0	0	89	–	0.8
	# Patients	13.0	5.1	0.5	0.0	0.0	0.0		18.6	–
	# Toxicities	5.6	2.6	0.3	0.0	0.0	0.0		8.6	–
Tite CRM	*% Selected*	18	0	0	0	0	0	82	–	0.7
	# Patients	15.1	5.6	3.7	1.8	0.7	0.3		27.2	–
	# Toxicities	5.7	3.1	2.2	1.2	0.5	0.2		12.8	–

Simulation results

		True Prob($T < 3\ months \mid d_k$)								
		d_1	d_2	d_3	d_4	d_5	d_6			
Scenario 3 (Early Onset)		0.250	0.350	0.500	0.600	0.680	0.730	None	Total	Duration
PRT	*% Selected*	44	38	3	0	0	0	14	–	1.0
	# Patients	16.6	12.6	2.7	0.3	0.0			32.1	–
	# Toxicities	4.4	4.7	1.6	0.2	0.0			10.9	–
Tite CRM	*% Selected*	40	41	6	1	0	0	12	–	0.9
	# Patients	17.0	11.4	3.8	0.6	0.0	0.0		32.8	–
	# Toxicities	4.2	4.0	1.9	0.3	0.0	0.0		10.5	–

Simulation results

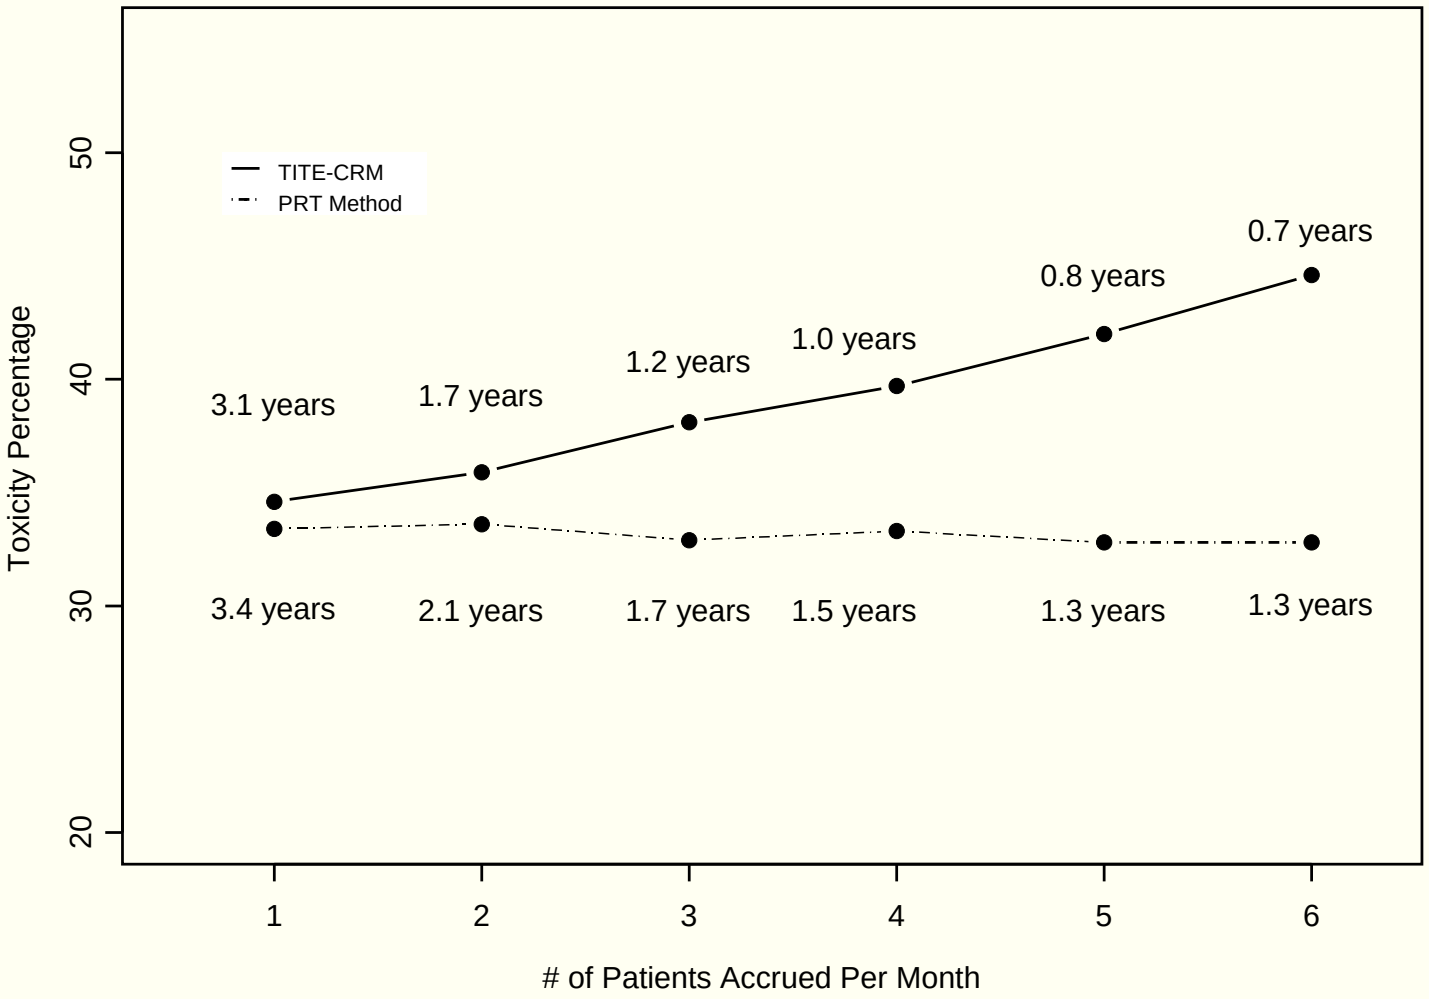

Summary

If patient accrual is rapid and toxicities occur at the targeted rate in the toxicity evaluation window $[0,\ t^*]$ but are likely to occur late in this interval then, on average,

- the PRT method gives a trial with fewer toxicities but a longer duration compared to the TITE-CRM
- the percentage of toxicities allowed by the PRT method does not increase with accrual rate,
- in some cases (scenarios 2 and 4 of our study) the TITE-CRM is more likely than the PRT method to select a final dose having excessive toxicity probability (≥ 0.50).

Phase II AR

Now consider the third type of trials

- there are multiple treatment arms under comparison
- each patient is to be randomized to one of the arms
- a treatment arm is superior if it is more effective with similar toxicity or if it is less toxic with similar efficacy
- the assessment window of efficacy/toxicity is quite long, e.g., 52 weeks

Challenges for AR

Adaptive randomization (AR) is desirable since more patients are usually assigned to the better treatment arms. The challenges in the AR are

- How to use the information contained in the patients that have not completed followup (e.g., 52 weeks) for randomization?
- How to incorporate the joint efficacy/toxicity responses from patients into the AR probability?

Main idea

- Use a failure time regression to estimate the joint efficacy/toxicity probability at the end of followup
- Use a latent modeling approach to achieve simple posterior computation
- Elicit a probability measure to incorporate joint efficacy/toxicity responses into the AR probabilities

Notation

Let us focus on models for one treatment arm first.

- Let t_i be the followup time for patient i; let $t_{\max}$ be the followup duration; apparently, $t_i \leq t_{\max}$.
- Let $Z_{kli}(t_i)$ be the joint efficacy/toxicity indicator at time t_i for patient i, $k, l = 0, 1$.
- Let $\pi_{kli} = \Pr(Z_{kli}(t_i) = 1)$ be the probability of the joint efficacy/toxicity at time t_i.
- For example, $\pi_{01i} = \Pr(Z_{01i}(t_i) = 1)$ is the probability of no-efficacy/toxicity at time t_i.
- Last, let $p_{kl} = \Pr(Z_{kli}(t_{\max} = 1)$, which are the parameters of interests.

Failure time model

Outcomes by time $t_{\max}$	Outcomes by time t_i				Row total
	No-Eff & No-Tox	*No-Eff & Tox*	*Eff & No-Tox*	*Eff & Tox*	
No-Eff & No-Tox	p_{00}	0	0	0	p_{00}
No-Eff & Tox	$p_{01}(1-w_i)$	$p_{01}w_i$	0	0	p_{01}
Eff & No-Tox	$p_{10}(1-w_i)$	0	$p_{10}w_i$	0	p_{10}
Eff & Tox	$p_{11}(1-w_i)^2$	$p_{11}w_i(1-w_i)$	$p_{11}w_i(1-w_i)$	$p_{11}w_i^2$	p_{11}
Column total	π_{00}	π_{01}	π_{10}	π_{11}	1

The weight $w_i = t_i/t_{\max}$ (see Cheung and Chappell, 2000, Biometrics).

- A general Bayesian approach for estimating w_i based on interim data is given in Ji and Bekele (2008, Biometrics).

Failure time model

Summarizing the table, we have

$$
\begin{aligned}
\pi_{00i} &= p_{00} + p_{01}(1-w_i) + p_{10}(1-w_i) + p_{11}(1-w_i)^2, \\
\pi_{01i} &= p_{01}w_i + p_{11}w_i(1-w_i), \\
\pi_{10i} &= p_{10}w_i + p_{11}w_i(1-w_i), \\
\pi_{11i} &= p_{11}w_i^2.
\end{aligned}
$$

Likelihood

Given the failure time model,

- The likelihood function for patient i is

$$L_i(p_{00}, p_{01}, p_{10}, p_{11}) \propto \prod_{k=0}^{1} \prod_{l=0}^{1} \pi_{kli}^{Z_{kli}(t_i)}$$

- The full likelihood function is then $L = \prod_{i=1}^{n} L_i$.

Latent modeling

- If we plug in the π_{kli}'s as functions of p_{kl}'s, the likelihood function involves quardrinomial expansions.
- A latent modeling approach is proposed to simply the computation.

Latent variables

Z's are observed. Given the column total Z, the cell entries in that column are the latent variables which follow a multinomial distribution with parameters Z and the corresponding probabilities in the previous table.

Outcomes by time $t_{\max}$	Outcomes by time t_i			
	No-Eff & No-Tox	*No-Eff & Tox*	*Eff & No-Tox*	*Eff & Tox*
No-Eff & No-Tox	y_{001i}	0	0	0
No-Eff & Tox	y_{002i}	y_{012i}	0	0
Eff & No-Tox	y_{003i}	0	$y_{103i}w_i$	0
Eff & Tox	$Z_{00i}(t_i) - \sum_{h=1}^{3} y_{00hi}$	$Z_{01i}(t_i) - y_{012i}$	$Z_{10i}(t_i) - y_{103i}$	$Z_{11i}(t_i)$
Column total	$Z_{00i}(t_i)$	$Z_{01i}(t_i)$	$Z_{10i}(t_i)$	$Z_{11i}(t_i)$

Computation

- Assume $(p_{00}, p_{01}, p_{10}, p_{11})$ follows a Dirichlet prior.
- With the latent variables, one can write down an augmented likelihood involving products of multinomials.
- By assuming multinomial distributions for the latent variables, the posterior computation can be carried out in a standard fashion.
- Using the results in Tanner and Wang, 1987, JASA, we show that the random samples of $(p_{00}, p_{01}, p_{10}, p_{11})$ computed using the augmented likelihood and multinomial latent distributions are posterior samples under the original likelihood function and the Dirichlet prior.

AR probability

- Extend the notation so that $(p_{00j}, p_{01j}, p_{10j}, p_{11j})$ denote the joint probabilities of efficacy/toxicity for treatment arm j
- Using proposed models, we can compute posterior samples of $(p_{00j}, p_{01j}, p_{10j}, p_{11j})$ for each arm
- The AR probability

$$q_j = \frac{\xi_j}{\sum_{j=1}^{T} \xi_j},$$

in which ξ_j is a measure of desirability of arm j.

Measure ξ_j

- Let $\pi_j^E = p_{10j} + p_{11j}$ be the marginal probability of efficacy for arm j
- Let $\pi_j^T = p_{01j} + p_{11j}$ be the marginal probability of toxcity for arm j
- We define

$$\xi_j = \text{Pr}((\pi_j^E, \pi_j^T) \in \mathcal{A} | \text{data})$$

as the measure of desirability, where $\mathcal{A}$ is an acceptable region.

Acceptable region

The acceptable region $\mathcal{A}$ is given by

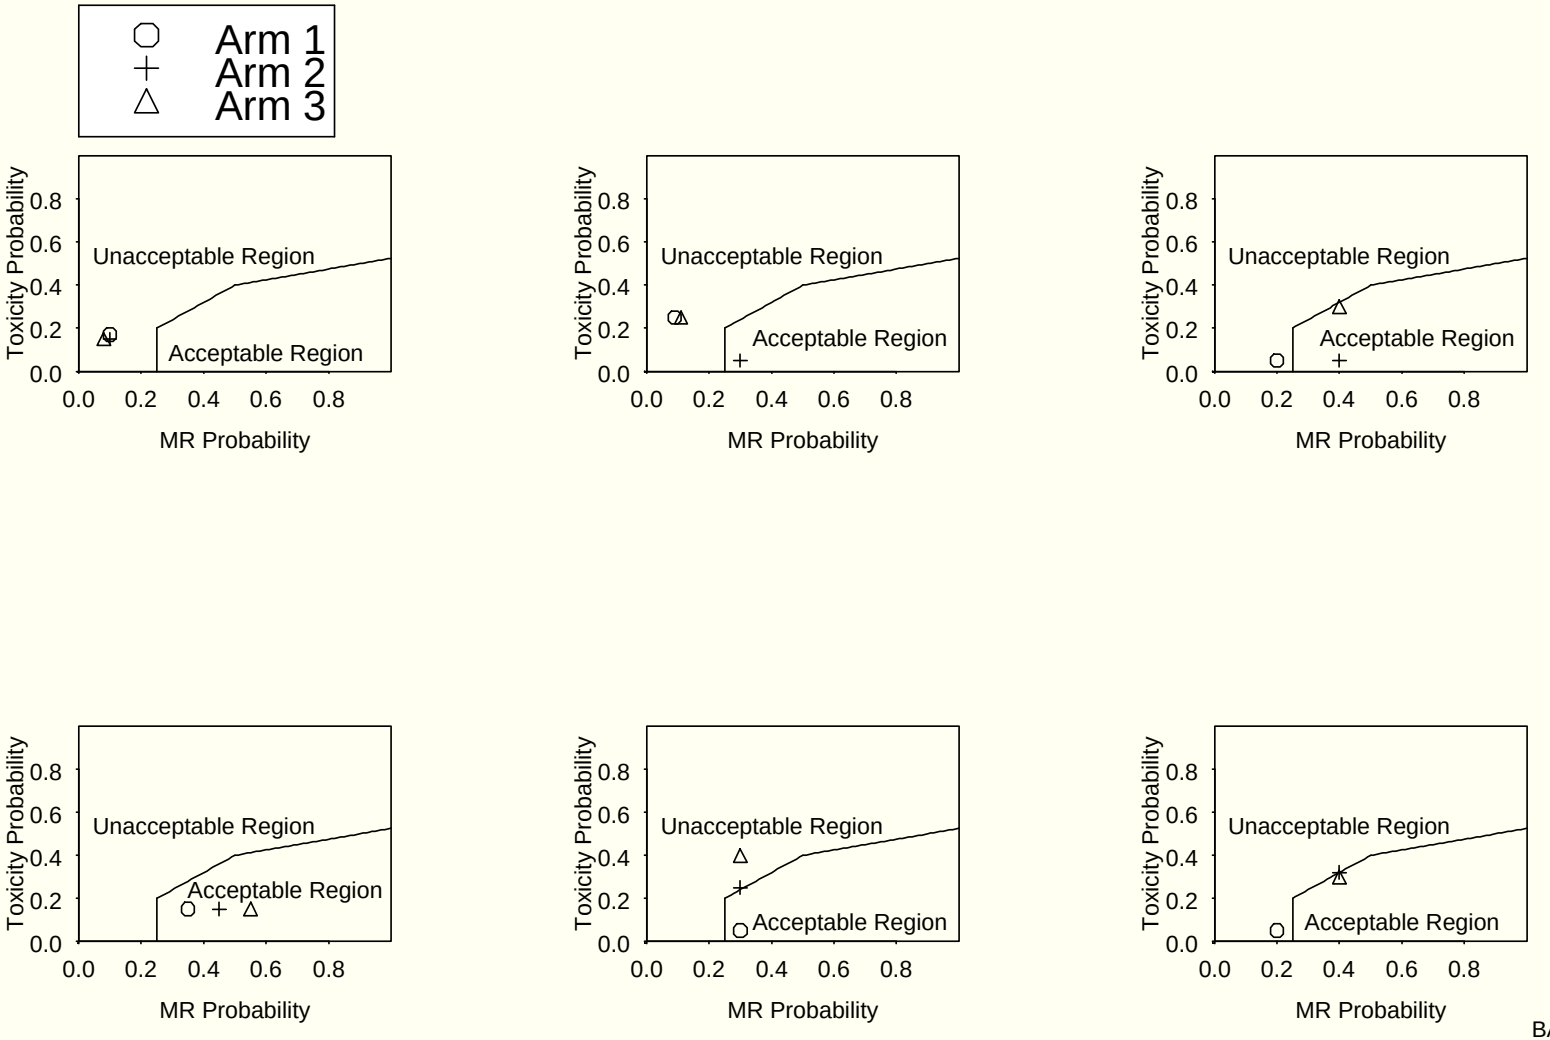

AR scheme

First, need a run-in stage when a small number (say 10 per arm) of patients are equally randomized.
Then, when a patient is enrolled,

- Compute the posterior of $(p_{00j}, p_{01j}, p_{10j}, p_{11j})$ for each arm j based on the proposed models
- Compute the measure of desirability ξ_j using the posterior samples of $(p_{00j}, p_{01j}, p_{10j}, p_{11j})$
- Randomize the patient to arm j with AR probability q_j

A simulated trial

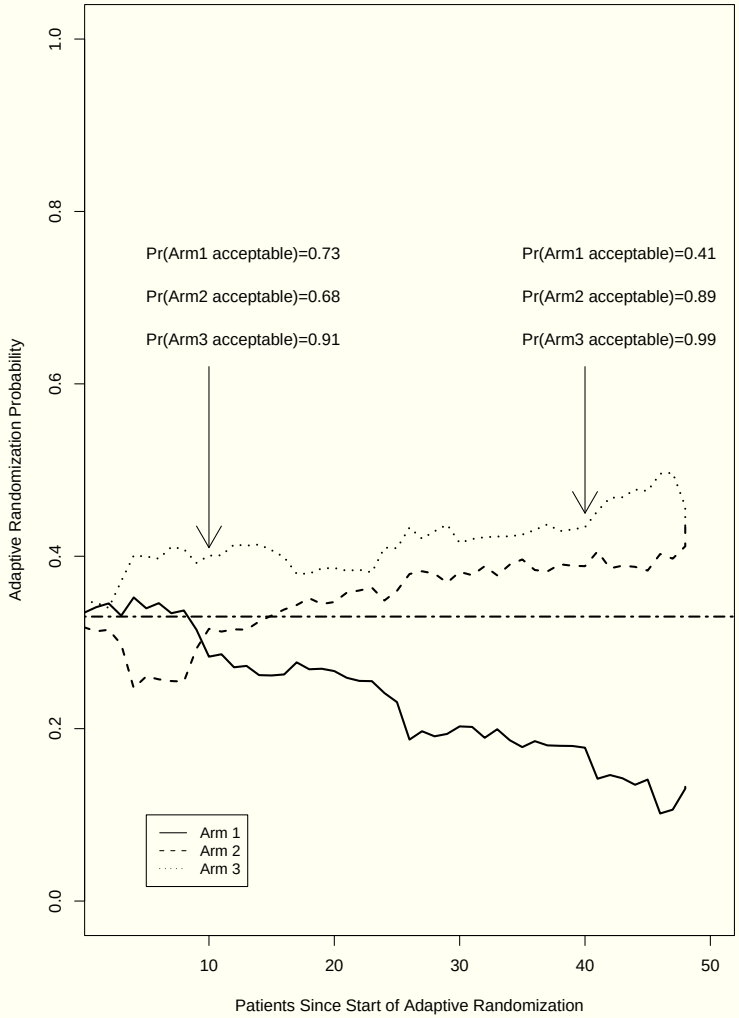

References and Collaborators

1. Ji, Y., Li, Y. and Bekele, N. (2007) Dose-finding in phase I clinical trials based on toxicity probability intervals. *Clinical Trials* **3** 235-245.
2. Ji, Y., Li, Y. and Yin, G. (2007) Bayesian dose finding in phase I clinical trials based on a new statistical framework. *Statistica Sinica* **17** 531-547.
3. Bekele, N., Ji, Y., Shen, Y. and Thall, P. (2008) Monitoring Late Onset Toxicities in Phase I Trials Using Predicted Risks. *Biostatistics* **9** 442-457.
4. Ji, Y. and Bekele, N. Bayesian adaptive randomization for multi-arm comparative clinical trials based on joint efficacy/toxicity responses. *Biometrics* In press.