Bayesian Adaptive Designs for Early-Phase Oncology Trials

Guosheng Yin Associate Professor Department of Statistics & Actuarial Science The University of Hong Kong



- Background and motivation
- Phase I trial
- Seamless phase I/II trial
- Drug-combination trial
- Concluding remarks

Clinical Trial

- A clinical trial is a prospective study that examines the treatment effects of a new drug, therapy or any medical intervention in humans.
- Conventionally, clinical trials are classified into four sequential phases: I, II, III and IV.
- In oncology, phase I clinical trials focus on the induced toxicities of the new drug.
- The main purpose of the study is to find the maximum tolerated dose (MTD), which is the highest possible dose while still tolerable.

Phase II

- Phase II clinical trials examine the preliminary evidence of efficacy and continue to monitor the safety of the drug.
- At this stage, compounds found to be ineffective or unsafe should be dropped to avoid wasting more resources.
- Generally speaking, non-working or unsafe drugs should be "killed" as early as possible.
- Once passing through the phase II stage, the new drug will proceed into phase III clinical trials, which are more rigorous and long-term efficacy studies.



- Phase III trials involve comparison with multiple treatments (controlled), randomization, and definitive clinical endpoints.
- Phase IV clinical trials are conducted to learn more about rare side effects of the approved intervention and its interaction with other therapies after the regulatory approval of the new treatment.



- In conventional phase I trials, the primary objective is often to find the maximum tolerated dose (MTD).
- A sequence of doses is screened in order to find the target dose associated with the maximum level of tolerable toxicity.
- Many methods have been proposed for phase I trials, see Chevret (2006) and Ting (2006) for comprehensive reviews.
- Typically, we assume that toxicity monotonically increases with the dose.



7



3+3 Design

- This is an algorithm-based procedure that is conservative and typically finds the MTD as the highest dose with a toxicity probability less than 33%.
- Due to its simplicity, many phase I clinical trials are carried out using the 3 + 3 design in practice.



Figure 3: Diagram of the standard 3 + 3 design.



- The 3 + 3 design is "memoryless" because dose escalation or de-escalation is solely based on the toxicity outcomes observed at the current dose level without any borrowing information from other doses.
- Moreover, this design does not have any statistical convergence property and is only suitable for targeting a toxicity probability less than 33%.

Continual Reassessment Method

- The continual reassessment method (CRM) proposed by O'Quigley, Pepe and Fisher (1990) is a popular dose-finding design.
- The CRM is model-based and requires practitioners to prespecify the toxicity probability at each dose.
- The true toxicity probabilities are linked with the prespecified probabilities in a parametric model via a single unknown parameter.
- During the trial, the CRM continuously updates the unknown parameter as more data are collected, and eventually identifies the MTD.



Figure 4: Different dose-toxicity curves lead to different MTDs.

CRM

- Let (d_1, \ldots, d_J) denote a set of J prespecified doses for the drug under investigation, and let (p_1, \ldots, p_J) be the prespecified toxicity probabilities at those doses; $p_1 < \cdots < p_J$ are known as the skeleton of the CRM.
- The CRM assumes a working dose-toxicity model,

Pr(toxicity at
$$d_j$$
) = $\pi_j(\alpha) = p_j^{\exp(\alpha)}$

for j = 1, ..., J, where α is an unknown parameter (O'Quigley and Shen, 1996).



Figure 5: Dose-toxicity curves under the CRM power function.

Likelihood

- Suppose among n_j patients treated at dose level j, there are y_j patients experiencing the DLT.
- Let $D = \{(n_j, y_j), j = 1, ..., J\}$ denote the observed data, then the likelihood function is

$$L(D|\alpha) \propto \prod_{j=1}^{J} \{p_j^{\exp(\alpha)}\}^{y_j} \{1 - p_j^{\exp(\alpha)}\}^{n_j - y_j}$$

• Let $f(\alpha)$ denote a prior distribution for α , for example, $\alpha \sim N(0, \sigma^2)$.

Bayesian Estimation

• Using Bayes' theorem, the dose toxicity probabilities can be estimated by the **posterior means**

$$\hat{\pi}_j = \int p_j^{\exp(\alpha)} \frac{L(D|\alpha) f(\alpha)}{\int L(D|\alpha) f(\alpha) \mathbf{d}\alpha} \mathbf{d}\alpha,$$

• A new cohort of patients is assigned to dose level j^* such that

$$j^* = \operatorname{argmin}_{j \in (1,\dots,J)} |\hat{\pi}_j - \phi_T|.$$

• The trial continues until the total sample size is exhausted, and the dose with a posterior toxicity probability closest to ϕ_T is selected as the MTD.

Simulation Comparison

- We considered five increasing dose levels with monotonically increasing toxicity.
- The target toxic probability was $\phi_T = 30\%$, and the initial guesses of the toxicity probabilities were $(p_1, \ldots, p_5) = (0.1, 0.2, 0.3, 0.4, 0.5).$
- The first cohort of patients was treated at the lowest dose level, and the maximum sample size was 30.
- We simulated 10,000 trials, and implemented the 3+3 design for comparison.

Table 1: Simulations using the CRM and 3 + 3 designs with a target toxicity probability $\phi_T = 30\%$.

	Select	tion pe	$\mathbf{Ave}~\#$	$\mathbf{Ave}~\#$			
Design	1	2	3	4	5	\mathbf{tox}	pats
Scenario 1	0.30	0.40	0.55	0.60	0.65		
\mathbf{CRM}	70.6	27.7	1.7	0	0	10.1	30
# patients	20.6	8.2	1.2	0.1	0		
3 + 3	44.2	20.9	2.9	0.2	0	2.8	7.9
# patients	4.8	2.4	0.6	0.1	0		

_

	Selec	tion pe	$\mathbf{Ave}~\#$	$\mathbf{Ave}\ \#$			
Design	1	2	3	4	5	tox	pats
Scenario 2	0.10	0.20	0.30	0.40	0.50		
\mathbf{CRM}	2.9	37.2	44.4	14.0	1.6	7.1	30
# patients	5.6	11.9	9.1	3.0	0.5		
3 + 3	18.9	33.8	28.5	12.7	1.4	3.0	13.3
# patients	4.0	4.2	3.1	1.5	0.4		
Scenario 3	0.02	0.06	0.10	0.20	0.30		
\mathbf{CRM}	0	0.5	9.3	39.3	50.9	5.0	30
# patients	3.2	4.2	6.1	8.6	7.8		
3 + 3	1.6	5.8	18.2	32.2	11.8	2.2	17.3
# patients	3.2	3.5	3.9	4.0	2.7		

Bayesian Model Averaging CRM

- Despite the popularity, a major criticism of the CRM is the arbitrariness and subjectivity in the prespecification of (p_1, \ldots, p_J) , which is known as the "skeleton" of the CRM.
- Model misspecification may lead to poor operating characteristics, incorrectly select the MTD, and may even result in treating a substantial number of patients at excessively toxic doses.
- In practice, we have no information to justify whether a specific skeleton is reasonable.

SKELETON

- CRM model: $\pi_j(\alpha) = p_j^{\exp(\alpha)}$ for $j = 1, \dots, J$.
- **Prefixed** p_1, p_2, \ldots, p_J .
- For example, eight dose levels with a target $\phi = 30\%$:

$$(p_1, \dots, p_8) = \begin{cases} (.02, .06, .08, .12, .20, .30, .40, .50), \text{ Skeleton 1} \\ (.01, .05, .09, .14, .18, .22, .26, .30), \text{ Skeleton 2} \\ (.10, .20, .30, .40, .50, .60, .70, .80), \text{ Skeleton 3} \\ (.20, .30, .40, .50, .60, .65, .70, .75), \text{ Skeleton 4.} \end{cases}$$



23



Selection percentage

Toxicity Probabilities of Doses 1–8



Toxicity Probabilities of Doses 1–8

BMA-CRM

- To overcome the arbitrariness and further enhance the robustness of the design, we can use multiple parallel CRM models, each with a different skeleton (Yin and Yuan, 2009).
- Instead of using a single CRM for the trial conduct, we carry out multiple parallel CRMs and rely upon the BMA approach for decision making.
- BMA is known to provide a better predictive performance than any single model (Raftery, Madigan and Hoeting, 1997; and Hoeting et al., 1999).

CRM Model

• Let model M_k be associated with the *k*th skeleton (p_{k1}, \ldots, p_{kJ}) , for $k = 1, \ldots, K$, and the corresponding CRM model is given by

$$\pi_{kj}(\alpha_k) = p_{kj}^{\exp(\alpha_k)}, \qquad j = 1, \dots, J.$$

• Let $P(M_k)$ be the prior probability that model M_k is the true model, and $P(M_k) = 1/K$ if there is no preference a priori for any single CRM model.

Posterior Model Probability

• The likelihood function under model M_k is

$$L(D|\alpha_k, M_k) \propto \prod_{j=1}^{J} \{p_{kj}^{\exp(\alpha_k)}\}^{y_j} \{1 - p_{kj}^{\exp(\alpha_k)}\}^{n_j - y_j}$$

• The posterior model probability for M_k is given by

$$P(M_k|D) = \frac{L(D|M_k)P(M_k)}{\sum_{i=1}^{K} L(D|M_i)P(M_i)},$$

where the marginal likelihood of model M_k ,

$$L(D|M_k) = \int L(D|\alpha_k, M_k) f(\alpha_k|M_k) \mathbf{d}\alpha_k.$$

BMA Estimation

• The BMA estimate for the toxicity probability at each dose level is given by

$$\bar{\pi}_j = \sum_{k=1}^K \hat{\pi}_{kj} P(M_k | D), \quad j = 1, \dots, J,$$

where $\hat{\pi}_{kj}$ is the posterior mean of the toxicity probability at dose level j under model M_k ,

$$\hat{\pi}_{kj} = \int p_{kj}^{\exp(\alpha_k)} \frac{L(D|\alpha_k, M_k) f(\alpha_k | M_k)}{\int L(D|\alpha_k, M_k) f(\alpha_k | M_k) \mathbf{d}\alpha_k} \mathbf{d}\alpha_k.$$

• By assigning $\hat{\pi}_{kj}$ a weight of $P(M_k|D)$, the BMA method automatically favors the best model.



- The target toxicity probability was $\phi_T = 30\%$. We took the prior distribution $\alpha \sim N(0, 4)$, and a discrete uniform prior model probability $P(M_k) = 1/4$ for $k = 1, \dots, 4$.
- We used a cohort size of 3, and treated the first cohort of patients at the lowest dose level. The maximum sample size was 30, and for each scenario we carried out 10,000 simulated trials.

Selection percentage Scenario 1

Toxicity Probabilities of Doses 1–8



Accumulating Cohorts under Scenario 1



Toxicity Probabilities of Doses 1–8



Toxicity Probabilities of Doses 1–8

34



- Higher doses are assumed to induce more severe toxicities, while efficacy is typically not considered in phase I studies.
- In phase II trials, patients would be treated at the MTD to examine the potential efficacy of the drug.
- Efficacy is often modeled as a short-term and binary endpoint.
- In conventional settings, phase I and phase II trials are conducted separately without any kind of formal borrowing of information or strength across them.

Combining Phase I and II

- There has been a growing trend to seamlessly combine phase I and phase II clinical trials to
 - (1) speed up the drug development process,
 - (2) improve the dose-finding procedure by maximizing the drug's efficacy as well as controlling its toxicity, and
 - (3) enlarge the sample size by pooling patients in phase I and phase II trials to produce more reliable estimates of toxicity and efficacy than would be achieved in each separate trial.

Notation

- Let p_j and q_j be the probabilities of toxicity and efficacy at dose d_j .
- For toxicity, a monotonic order is assumed, $p_1 < \cdots < p_J$, while no constraint is imposed for q_j .
- Let X_{ij} denote the toxicity outcome for subject i at dose level j,

 $X_{ij} = \begin{cases} 1 & \text{with probability } p_j \\ 0 & \text{with probability } 1 - p_j, \end{cases}$

and the efficacy outcome $Y_{ij} \sim Ber(q_j)$.

• We measure the association between the bivariate

outcomes using the global cross ratio (Dale, 1986).

• Define $\pi_{j(xy)} = P(X_{ij} = x, Y_{ij} = y)$ where x, y = 0, 1, and at dose level j,

$$\gamma_j = \frac{\pi_{j(00)}\pi_{j(11)}}{\pi_{j(01)}\pi_{j(10)}},$$

quantifies the association between toxicity and efficacy.

• The joint probabilities $\pi_{j(xy)}$ can be derived from γ_j

and the marginal probabilities p_j and q_j ,

$$\pi_{j(11)} = \begin{cases} \frac{a_j - (a_j^2 + b_j)^{1/2}}{2(\gamma_j - 1)} & \gamma_j \neq 1 \\ p_j q_j & \gamma_j = 1, \end{cases}$$

$$\pi_{j(10)} = p_j - \pi_{j(11)}, \\ \pi_{j(01)} = q_j - \pi_{j(11)}, \\ \pi_{j(00)} = 1 - p_j - q_j + \pi_{j(11)}, \end{cases}$$

where $a_j = 1 + (p_j + q_j)(\gamma_j - 1)$ and $b_j = -4\gamma_j(\gamma_j - 1)p_jq_j$.

• If n_j subjects are treated at dose d_j , the likelihood is

$$L(D|\mathbf{p},\mathbf{q},\boldsymbol{\gamma}) \propto \prod_{j=1}^{J} \prod_{i=1}^{n_j} \prod_{x=0}^{1} \prod_{y=0}^{1} \{\pi_{j(xy)}\}^{I(X_{ij}=x,Y_{ij}=y)}$$



Figure 6: Two-dimensional toxicity-efficacy odds ratio trade-off contours with point A (q_j, p_j) . Guosheng Yin



• The odds ratio ω_j between the toxicity and efficacy at dose level j,

$$\omega_j = \frac{p_j/(1-p_j)}{q_j/(1-q_j)} = \frac{p_j(1-q_j)}{q_j(1-p_j)},$$

is exactly the ratio of the area in the lower-right rectangle to that in the upper-left.

• A smaller value of ω_j indicates a more desirable dose.

Drug-Combination Trials

- Compared with single-agent treatments, combination therapy in cancer treatment may
 - (1) lead to synergistic treatment effects,
 - (2) target tumor cells with differing drug susceptibilities and through different disease pathways, and
 - (3) achieve higher dose intensities with non-overlapping toxicities.



Drug A

Figure 7: Dose pairs in a two-drug combination study.



- A recent phase I dose-finding trial was conducted for the combination of bortezomib with gemcitabine/doxorubicin in the treatment of metastatic urothelial cancer.
- Bortezomib interferes with a substance inside the cancer cell that is responsible for cell division.
- Gemcitabine/doxorubicin are chemotherapeutic agents that disrupt the growth of cancer cells.
- Combining these two drugs is expected to show substantial synergy to enhance efficacy.

Broader Application

- The problem of two- or high-dimensional dose finding can be cast in a more general framework as one of the following choices:
 - (1) a trial evaluating several different drugs, each administered at different doses,
 - (2) a study of a single agent at a set of dose levels, adding a change to different dose schedules, or
 - (3) a single-agent dose-finding trial involving ordered patient groups.



- The toxicity order of dose combinations is only partially known.
- The dimension of the dose-searching space expands multiplicatively.
- As multiple MTD combinations may exist in the two-dimensional grid, dose searching may be more easily trapped in a local area around certain MTD combinations, and thus lack the opportunity to explore the entire searching space.

Clayton Copula-Type Regression

• Borrowing the structure of Clayton's copula, the joint toxicity probability can be written as

$$\pi_{jk} = 1 - \{ (1 - p_j^{\alpha})^{-\gamma} + (1 - q_k^{\beta})^{-\gamma} - 1 \}^{-1/\gamma},$$

where $\gamma > 0$ characterizes the drug-drug interactive effect.

• Strictly speaking, it is not a copula because it does not characterize a bivariate distribution, we in fact only observe one single DLT outcome for combined agents.

Model conditions

(1) If
$$p_j^{\alpha} = 0$$
 and $q_k^{\beta} = 0$, then $\pi_{jk} = 0$.
(2) If $p_j^{\alpha} = 0$, then $\pi_{jk} = q_k^{\beta}$; and if $q_k^{\beta} = 0$, then $\pi_{jk} = p_j^{\alpha}$.
(3) If either $p_j^{\alpha} = 1$ or $q_k^{\beta} = 1$, then $\pi_{jk} = 1$.



Drug A Toxicity Prob.

Figure 8: Toxicity probability surface and MTD equivalence contour under a Clayton copula-type model.

Likelihood

- The likelihood function can be constructed based on the binomial distribution with the probabilities of π_{jk} .
- Suppose y_{jk} out of n_{jk} patients treated at the dose combination (A_j, B_k) have experienced toxicity, then the likelihood function is given by

$$L(\alpha, \beta, \gamma | D) \propto \prod_{j=1}^{J} \prod_{k=1}^{K} \pi_{jk}^{y_{jk}} (1 - \pi_{jk})^{n_{jk} - y_{jk}},$$

where D represents the observed data.



Figure 9: Dose escalation/de-escalation in a matrix of 5×4 dose combinations.



• **BMA-CRM** user-friendly software

http://biostatistics.mdanderson.org/SoftwareDownload/

- Seamless phase I/II trial with the efficacy and toxicity odds ratio.
- A copula-type model approach for drug combination trials.

