

Bayesian Adaptive Designs for Phase I Oncology Trials

Joseph S. Koopmeiners

Division of Biostatistics
University of Minnesota

PubH 7440
April 14, 2011



- ① Phase I Oncology Trials
- ② Rule Based Designs
- ③ CRM
- ④ EffTox
- ⑤ Combination Therapy
- ⑥ Summary

1 Phase I Oncology Trials

2 Rule Based Designs

3 CRM

4 EffTox

5 Combination Therapy

6 Summary

Phases of Drug Development

- **Phase 1: Safety trials**
- Phase 2: Efficacy trials
- Phase 3: Confirmatory trials
- Phase 4: Post-marketing surveillance

Goals of Phase 1

- First application of a new drug in humans
- Establish safety profile of a new drug
- pharmacokinetics and pharmacodynamics
- Determine appropriate dosing for future studies



Why are Oncology trials different?

- Healthy volunteers are used when a new drug is expected to be relatively nontoxic
- Cytotoxic agents are known to be toxic
- Phase 1 oncology trials include patients for whom standard treatments have failed
- Goals:
 - Limit toxicities
 - Maximize the number of patients that receive a therapeutic dose

The dose-toxicity curve

- We are interested in the probability of dose limiting toxicity (DLT) for each dose
- For cytotoxic agents, we assume that as dose increases, the probability of DLT and the probability of efficacy will also increase
- Goal is to find the maximum tolerated dose (MTD): Highest dose with probability of DLT less than some pre-specified cut-off (usually 0.2 or 0.33)

The dose-toxicity curve



① Phase I Oncology Trials

② Rule Based Designs

③ CRM

④ EffTox

⑤ Combination Therapy

⑥ Summary

Rule-based designs: 3+3 design

- **Step 1:** Enter first cohort of 3 patients at the lowest dose
- **Step 2:** Observe toxicity outcomes
 - if 0/3 DLT, treat next cohort of 3 patients at next highest dose
 - if 1/3 DLT, treat next cohort of 3 patients at same dose
 - if 1/3 + 0/3 DLT, treat next cohort of 3 patients at next highest dose
 - if 1/3 + 1/3 DLT, define dose as MTD
 - if 1/3 + 2/3 or 3/3 DLT, dose exceeds MTD
 - if 2/3 or 3/3 DLT, dose exceeds MTD
- **Step 3:** Repeat until MTD is reached.
- **Step 4:** MTD is defined as a dose with $\leq 2/6$ DLT

3+3 example

- Cohort 1: dose level 1, 0/3 DLT
- Cohort 2: dose level 2, 0/3 DLT
- Cohort 3: dose level 3, 1/3 DLT
- Cohort 4: dose level 3, 0/3 DLT
- Cohort 5: dose level 4, 2/3 DLT

Conclude that dose level 3 is the MTD

Pros and Cons of the 3+3 design

- Pros
 - Simple!
- Cons
 - ad hoc
 - imprecise
 - only allows for a target DLT probability of .33
 - treats a maximum of 6 patients at the MTD

① Phase I Oncology Trials

② Rule Based Designs

③ CRM

④ EffTox

⑤ Combination Therapy

⑥ Summary

Continual Reassessment Method

- The continual reassessment method (CRM) was the first Bayesian adaptive phase 1 design (O'Quigley, Pepe and Fisher, 1990)
- The CRM uses a one-parameter parametric model for the probability of DLT at each dose

One-parameter models for the probability of DLT

Let $p(d) = P(DLT | dose = d)$

- Power Model

$$p(d) = d^{\exp(a)}$$

- Logistic Model

$$p(d) = \frac{\exp(3 + ad)}{1 + \exp(3 + ad)}$$

- Hyperbolic Tangent Model

$$p(d) = \left(\frac{\exp(d)}{\exp(d) + \exp(-d)} \right)^a$$

Estimating the probability of DLT for each dose

As a model based method, we are able to estimate the probability of DLT for each dose using standard Bayesian methods. Assuming prior distribution, $\pi(\mathbf{a})$,

$$\pi(\mathbf{a}|\vec{y}) \propto L(\vec{y}|\mathbf{a}) \pi(\mathbf{a}),$$

where,

$$L(\vec{y}|\mathbf{a}) = \prod_{i=1}^n p(d_i)^{y_i} (1 - p(d_i))^{1-y_i}.$$

The posterior is generally not available in closed form and must be estimated using numerical methods (i.e. MCMC)

CRM algorithm

- **Step 1:** Assume a vague or fully non-informative prior for α
- **Step 2:** Treat 1 patient at the level closest to the current estimate of MTD
- **Step 3:** Observe toxicity outcome
- **Step 4:** Update posterior for α
- **Step 5:** Treat next patient at the level closest to the current estimate of MTD based on posterior for α
- **Step 6:** Repeat 1-5 until maximum sample size is reached

Modifications to the CRM

The standard CRM dose-escalation can be too aggressive. The following modifications have been suggested as a remedy:

- Start at the lowest dose and do not skip untried doses
- Treat patients in cohorts greater than 1 (2 or 3)
- No dose escalation until all treatment in patients from previous doses are completed
- More aggressive stopping rules for safety

Advantages and disadvantages to the CRM

- Advantages
 - In general, the CRM is more likely to find the true MTD than 3+3
 - Can be adapted for target probabilities other than .33
 - Model-based methods uses all data to model the dose-toxicity curve
 - Treats more patients with doses close to the MTD (i.e. patients more likely to receive efficacious dose)
- Disadvantages
 - More complicated (for clinicians, at least)
 - Model misspecification is a potential concern

① Phase I Oncology Trials

② Rule Based Designs

③ CRM

④ **EffTox**

⑤ Combination Therapy

⑥ Summary

Simultaneously considering Efficacy and Toxicity

- Standard Phase I designs only consider toxicity
- This is not ideal in some situations
 - MTD has unacceptable efficacy
 - Efficacy does not continue to increase as dose increases

EffTox (Thall and Cook, 2004) is a Bayesian adaptive phase 1 design that simultaneously considers toxicity and efficacy

- EffTox considers two binary outcomes
 - Toxicity: Y_T
 - Efficacy: Y_E
- We must model the joint probability of efficacy and toxicity

$$\pi_{a,b}(x, \theta) = Pr(Y_E = a, Y_T = b | x, \theta)$$

Modeling the marginal probabilities

Thall and Cook (2004) consider logistic regression models for the marginal probabilities of efficacy and toxicity

- The marginal probabilities for efficacy and toxicity are:

$$\pi_E(x, \theta) = \pi_{1,0}(x, \theta) + \pi_{1,1}(x, \theta)$$

and

$$\pi_T(x, \theta) = \pi_{0,1}(x, \theta) + \pi_{1,1}(x, \theta)$$

- We fit the following logistic regression models for the marginal probabilities:

$$\text{logit}(\pi_E(x, \theta)) = \beta_{0,E} + x\beta_{1,E} + x^2\beta_{2,E},$$

and

$$\text{logit}(\pi_T(x, \theta)) = \beta_{0,T} + x\beta_{1,T},$$

Modeling the joint probabilities

The joint probability of efficacy and toxicity is modeled using a copula model

$$\pi_{a,b} = \pi_E^a (1 - \pi_E)^{1-a} \pi_T^b (1 - \pi_T)^{1-b} + (-1)^{a+b} \pi_E (1 - \pi_E) \pi_T (1 - \pi_T) \frac{e^\psi - 1}{e^\psi + 1}$$

- ψ characterizes the correlation between efficacy and toxicity
- $\psi = 0$ indicates that efficacy and toxicity are independent

Estimating the parameters

The likelihood, $L(Y_E, Y_T | x, \theta)$, for a patient at dose x is

$$\pi_{1,1}(x, \theta)^{Y_E Y_T} \pi_{1,0}(x, \theta)^{Y_E(1-Y_T)} \pi_{0,1}(x, \theta)^{(1-Y_E)Y_T} \pi_{0,0}(x, \theta)^{(1-Y_E)(1-Y_T)}$$

- Specify priors for all parameters
- Proceed with standard Bayesian inference using MCMC

Efficacy-Toxicity trade-off

Clinicians are asked to specify three points

- $(\pi_E^*, 0)$ - minimum efficacy with no toxicity
- $(1, \pi_T^*)$ - maximum toxicity with perfect efficacy
- (π'_E, π'_T) - a point that is equally desirable to the first two

Acceptable Doses

Dose levels are considered acceptable if

$$Pr(\pi_E(x, \theta) > \pi_E^* | Y_E, Y_T) > p_E$$

and

$$Pr(\pi_T(x, \theta) < \pi_T^* | Y_E, Y_T) > p_T$$

Desirability Index

A dose with probability of efficacy, π_E , and probability toxicity, π_T , has desirability $\delta = 1 - r$ where

$$r = \left(\left(\frac{1 - \pi_E}{1 - \pi_E^*} \right)^p + \left(\frac{\pi_T}{\pi_T^*} \right)^p \right)^{1/p}$$

where p is found by plugging (π'_E, π'_T) in for (π_E, π_T) , setting $r = 1$ and solving for p .

EffTox Algorithm

- **Step 1:** Treat the first cohort (2 or 3 patients) at the lowest dose level
- **Step 2:** Observe the efficacy and toxicity outcomes
- **Step 3:** Update posterior and calculate posterior means, $E(\pi_E(x, \theta) | Y_E, Y_t)$ and $E(\pi_T(x, \theta) | Y_E, Y_t)$ for each dose level
- **Step 4:** Determine the acceptable doses
- **Step 5:** Compute desirability index for all acceptable doses
- **Step 7:** Terminate the trial if no doses are acceptable, otherwise, treat the next cohort at the most desirable dose under the restriction that no untried doses may be skipped
- **Step 8:** Continue until the maximum sample size is reached. The dose with the highest desirability is selected for further evaluation

Advantages and Disadvantages of EffTox

- Advantages
 - Efficacy and toxicity considered simultaneously
 - Considers MTD undesirable if the MTD has unacceptable efficacy
 - Does not escalate unless there is an efficacy benefit
- Disadvantages
 - Complicated!
 - Requires a larger sample size than standard phase 1 designs

① Phase I Oncology Trials

② Rule Based Designs

③ CRM

④ EffTox

⑤ Combination Therapy

⑥ Summary

Combination Therapy

- It is increasingly common that our interest lies in the safety profile for a combination of two drugs
- This requires dose finding with two agents
- Yin and Yuan (2009) describe a Bayesian adaptive design for two-agent dose finding

Combination Therapy

- Let A_1, \dots, A_J be the dose levels for drug A and B_1, \dots, B_K be the dose levels for drug B
- Let X_{jk} and Y_{jk} be binary indicators for whether or not a patient experienced a toxicity from drugs A and B, respectively, from combination A_j and B_k
- Yin and Yuan use a similar approach to the EffTox design in that they directly model the marginal probabilities and combine them using a copula model

Modeling the marginal probabilities

- Yin and Yuan use the one-parameter power model for the marginal probabilities

$$P(\text{DLT from Drug A} | \text{dose} = j) = p_j^\alpha$$

and

$$P(\text{DLT from Drug B} | \text{dose} = j) = q_k^\beta$$

Modeling the joint probabilities

The joint probability of efficacy and toxicity is modeled using a copula model

$$\pi_{j,k}^{(x,y)} = p_j^{\alpha x} (1 - p_j^{\alpha})^{1-x} q_k^{\beta y} (1 - q_k)^{1-y} + (-1)^{x+y} p_j^{\alpha} (1 - p_j^{\alpha}) q_k^{\beta} (1 - q_k) \frac{e^{\psi} - 1}{e^{\psi} + 1}$$

- ψ characterizes the correlation between drug A and drug B
- $\psi = 0$ indicates that drug A and drug B are independent

Estimating the parameters

The Likelihood for Yin and Yuan's model is

$$\prod_{j=1}^J \prod_{k=1}^K \left(\pi_{j,k}^{(0,0)} \right)^{n_{j,k}^{(0,0)}} \left(\pi_{j,k}^{(0,1)} \right)^{n_{j,k}^{(0,1)}} \left(\pi_{j,k}^{(1,0)} \right)^{n_{j,k}^{(1,0)}} \left(\pi_{j,k}^{(1,1)} \right)^{n_{j,k}^{(1,1)}}$$

- Specify priors for α , β and ψ
- Proceed with standard Bayesian inference using MCMC

Problem

- A problem arises if we are unable to assign DLTs to a specific drug
- In this case, we only observe whether or not a DLT occurs
- To accommodate, we alter the likelihood to

$$\prod_{j=1}^J \prod_{k=1}^K \left(\pi_{j,k}^{(0,0)} \right)^{n_{j,k}^{(0,0)}} \left(1 - \pi_{j,k}^{(0,0)} \right)^{n_{j,k} - n_{j,k}^{(0,0)}}$$

Combination dose-finding algorithm

- **Step 1:** Treat patients in the first cohort at the lowest dose combination
- **Step 2:** If for the current dose (j, k) we have

$$P(\pi_{j,k} > \pi_* | Data) < c_e$$

escalate to an adjacent dose level with probability of toxicity greater than that for the current dose level.

- **Step 3:** If for the current dose (j, k) we have

$$P(\pi_{j,k} > \pi_* | Data) > c_d$$

de-escalate to an adjacent dose level with probability of toxicity lower than that for the current dose level.

- **Step 5:** The MTD is the dose level with probability of DLT closest to the target when the maximum sample size is reached.

① Phase I Oncology Trials

② Rule Based Designs

③ CRM

④ EffTox

⑤ Combination Therapy

⑥ **Summary**

Summary

- The goal of a phase 1 clinical trial is to assess the safety profile of a new drug and determine the optimal dose for further investigation
- Phase 1 oncology trials are unique for enrolling cancer patients who's previous treatments have failed
- We would like to treat as many patients as possible at therapeutic dose levels while identifying the MTD
- Bayesian adaptive designs allow us to treat more patients at therapeutic dose levels while identifying the true MTD at a higher rate than standard designs
- Bayesian adaptive designs provide the flexibility to accommodate deviations from standard phase 1 designs (toxicity and efficacy, two drugs, etc.)