A Bayesian Phase I/II Design for Oncology Clinical Trials of Combining Biological Agents

Ying Yuan

Department of Biostatistics, MD Anderson Cancer Center

March 12, 2014

Outline

- Introduction
- Probability model
- Dose finding algorithm
- Simulation study
- Conclusion

Biological agents

- The paradigm of oncology drug development is expanding from traditional cytotoxic agents to novel biological (or molecularly targeted) agents.
- Examples of biological agents:
  - Biospecimens targeting a specific tumor pathway.
  - Gene products aiming for DNA repair.
  - Immunotherapies stimulating the immune system to attack a tumor.

Biological agents versus cytotoxic agents

- Cytotoxic agents
  - Toxicity and efficacy are assumed to monotonically increase with dose.
  - The goal is to find the maximum tolerated dose (MTD).
- Biological agents
  - The toxicity is usually tolerable within the therapeutic dose range and may plateau at higher dose levels.
  - The dose-efficacy curves often follow a non-monotonic pattern.
  - The goal is to find the optimal biological dose (OBD), defined as the dose yielding the most desirable treatment effect.

Drug-combination Trials

- Treating patients with a combination of agents is becoming common in cancer clinical trials.
- Most existing drug-combination trial designs concern cytotoxic agents (e.g., Thall et al., 2003; Wang and Ivanova, 2005; Yin and Yuan, 2009), thus are not applicable to the trials combining biological agents.
- A phase I/II trial design is imperative for biological agent combination trials because of non-monotonic dose-efficacy and -toxicity relationship.

Motivating trial

- A lymphoma trial combining two novel biological agents to target two components in the PI3K/AKT/mTOR signaling pathway.
  - Agent A is a PI3K kinase inhibitor.
  - Agent B inhibits mTOR kinase downstream in the pathway.
- 4 doses of agent A combined with 4 doses of agent B.
- Goal: to find the biologically optimal dose combination (BODC), defined as the dose combination with the highest efficacy and tolerable toxicity.

Motivating trial

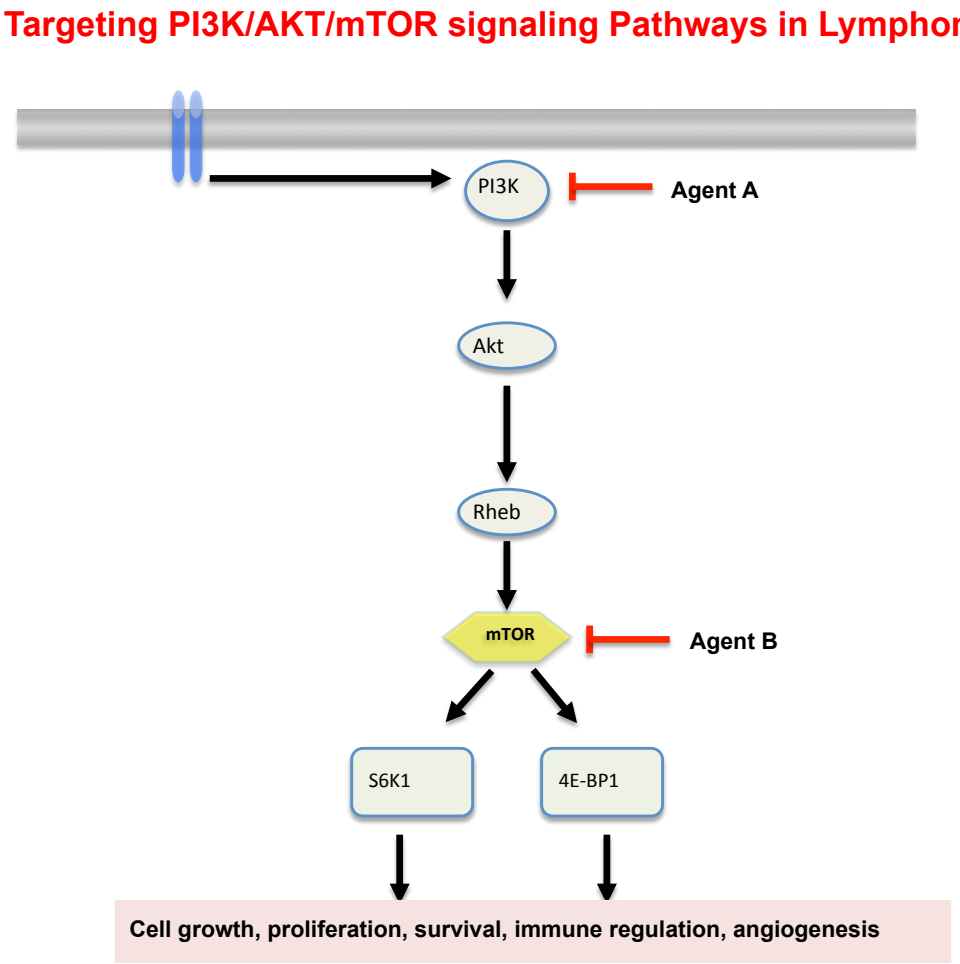

Proposed design

We propose a phase I/II trial design to identify the BODC.

- A change-line model is used to reflect the property that the dose-toxicity surface of the combinational agents may plateau at higher dose levels.
- A logistic model with quadratic terms is used to accommodate the possible non-monotonic pattern for the dose-efficacy relationship.
- We devise a novel adaptive dose-finding algorithm to encourage sufficient exploration of the two-dimensional dose space.

Notation

- Consider a trial of combinational biological agents
  - J doses of agent A: $a_1 < a_2 < \cdots < a_J$
  - K doses of agent B: $b_1 < b_2 < \cdots < b_K$
  - (a_j, b_k): combination of dose a_j and dose b_k
  - p_{jk} and q_{jk} denote the toxicity and efficacy probabilities of dose combination (a_j, b_k)
- Goal: identify the BODC in the $J \times K$ dose matrix.

Change-line model for toxicity

We model toxicity probability p_{jk} using a change-line model:

$$\text{logit}(p_{jk}) = (\beta_0 + \beta_1 a_j + \beta_2 b_k) I(\beta_0 + \beta_1 a_j + \beta_2 b_k \leq \omega) + \omega I(\beta_0 + \beta_1 a_j + \beta_2 b_k > \omega)$$

- $I(\cdot)$: indicator function
- $\beta_1 > 0$ and $\beta_2 > 0$ such that p_{jk} initially increases with the doses of A and B
- When it reaches a plateau, the toxicity probability: $e^{\omega}/(1 + e^{\omega})$.
- We did not include an interactive effect for the two agents because the estimation of that needs large sample

Figure : Surface of the toxicity probabilities

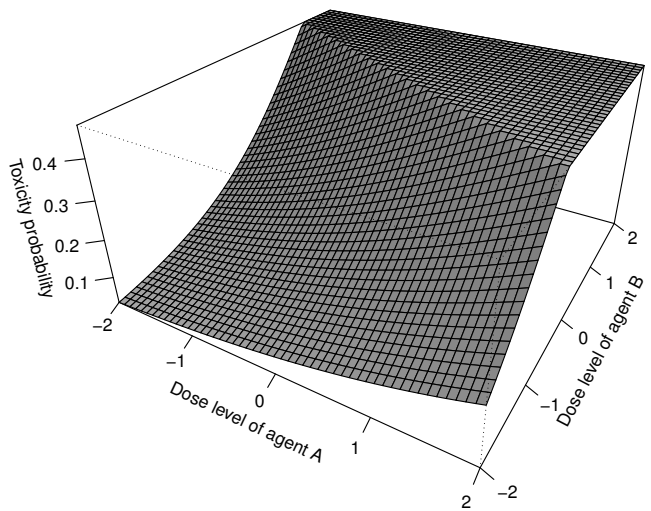

Logistic model for efficacy

Assume the efficacy probability q_{jk} follows a logistic model

$$\text{logit}(q_{jk}) = \gamma_0 + \gamma_1 a_j + \gamma_2 b_k + \gamma_3 a_j^2 + \gamma_4 b_k^2$$

- The quadratic terms render the model adequate flexibility to capture the non-monotonic pattern.
- We model the marginal distributions of toxicity and efficacy.
- Joint modeling is possible, but small sample size $\rightarrow$ cannot reliably estimate the correlation parameter.

Likelihood

Suppose that at a certain stage of the trial

- n_{jk} patients are treated at the paired dose (a_j, b_k)
- x_{jk} and y_{jk} patients have experienced toxicity and efficacy, respectively.
- The marginal likelihood for the toxicity data $\mathbf{x}$ is

$$L(\mathbf{x}|\omega, \boldsymbol{\beta}) \propto \prod_{j=1}^{J} \prod_{k=1}^{K} p_{jk}^{x_{jk}} (1 - p_{jk})^{n_{jk} - x_{jk}};$$

for the efficacy data $\mathbf{y}$ is

$$L(\mathbf{y}|\boldsymbol{\gamma}) \propto \prod_{j=1}^{J} \prod_{k=1}^{K} q_{jk}^{y_{jk}} (1 - q_{jk})^{n_{jk} - y_{jk}}.$$

- The posterior distribution is

$$f(\omega, \boldsymbol{\beta}, \boldsymbol{\gamma} | \mathbf{x}, \mathbf{y}) \propto L(\mathbf{x}|\omega, \boldsymbol{\beta}) L(\mathbf{y}|\boldsymbol{\gamma}) f(\omega) f(\boldsymbol{\beta}) f(\boldsymbol{\gamma})$$

  where $f(\omega)$, $f(\boldsymbol{\beta})$, and $f(\boldsymbol{\gamma})$ denote the prior distributions for ω, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$, respectively.

- Vague priors are used:

  $\gamma_0 \sim$ Cauchy(0,10), $\gamma_1, \cdots, \gamma_4 \sim$ Cauchy(0, 2.5). $\beta_0 \sim$ Cauchy(0, 10), $\beta_1, \beta_2 \sim$ Gamma(0.5, 0.5) $\omega \sim N(0, 4)$

Trial design

Our design is conducted in two stages:

- Stage I (run in): We escalate doses along the diagonal to explore the dose-combination space quickly and collect some preliminary data.
- Stage II (dose finding): Based on observed efficacy and toxicity data, we continuously update the posterior estimates of toxicity and posterior means of efficacy and assign patients to the most appropriate dose.

Def: dose (a_j, b_k) is deemed safe if $\Pr(p_{jk} < \phi | \mathcal{D}) > \delta$; otherwise toxic.

- ϕ is the target toxicity upper limit and δ is a prespecified safety cutoff.

Stage I: Run-in period

The trial starts with the treatment of the first cohort of patients at the lowest dose (a_1, b_1).

- I1 If current dose is safe, escalate the dose along the diagonal. If (a_1, b_1) is deemed toxic, terminate the trial.
- I2 Stage I completes when either current dose is deemed toxic or the highest dose combination is reached. Stage II starts.

g-degree admissible dose set

Assume that the current dose combination is (a_j, b_k),

- Define g-degree neighbors of (a_j, b_k), denoted by $\mathcal{N}_g$, as dose combinations $\{(a_{j'}, b_{k'})\}$ whose dose levels are different from (a_j, b_k) no more than g levels, i.e.,
  $\mathcal{N}_g = \{(a_{j'}, b_{k'}) : |j' - j| \leq g \text{ and } |k' - k| \leq g\}$.
- Further define a g-degree admissible dose set $\mathcal{A}_g$, which is a subset of the g-degree neighbors $\mathcal{N}_g$ satisfying the pre-specified safety requirement $Pr(p_{j'k'} < \phi_T | \mathcal{D}) > \delta$.

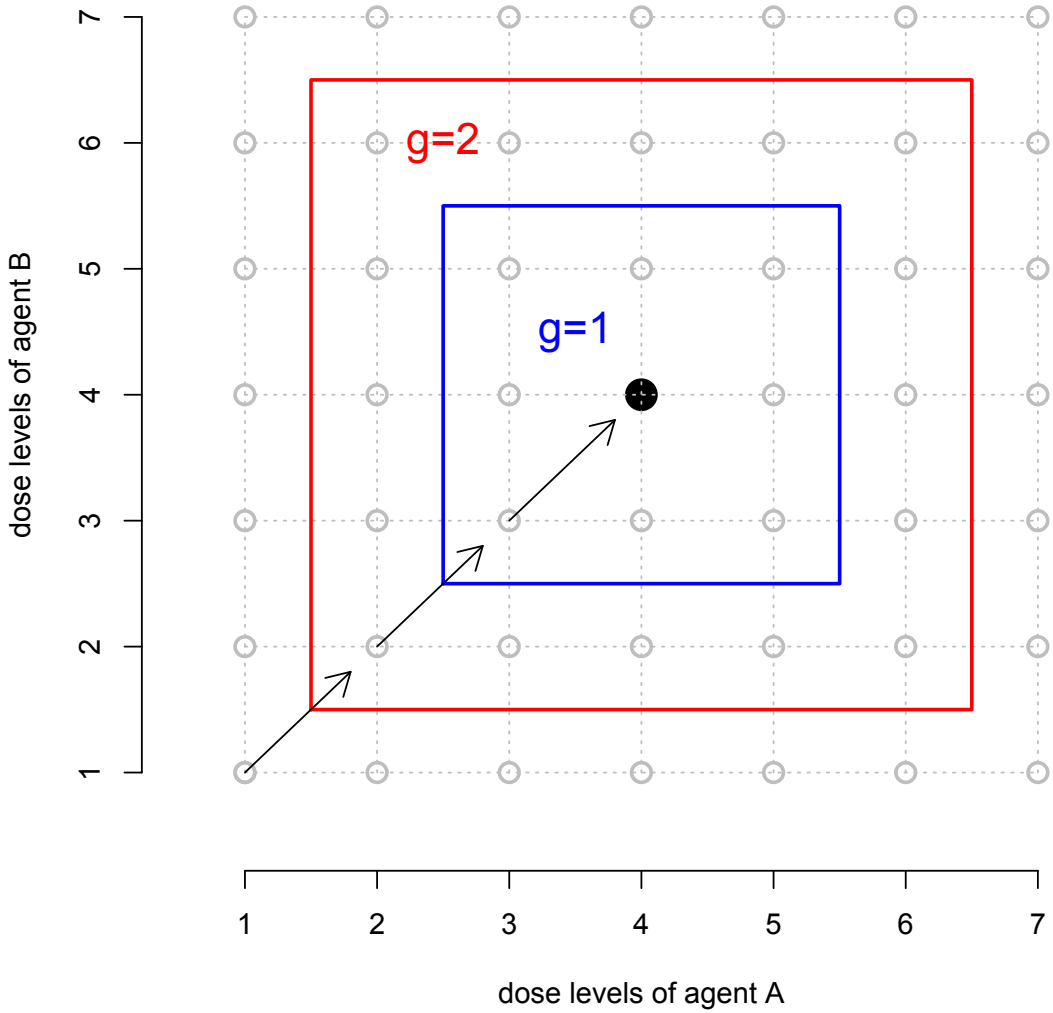
g=2
g=1
dose levels of agent B
dose levels of agent A
1
2
3
4
5
6
7

Stage II: Systematic dose finding

II1 Based on the observed data, identify $\mathcal{A}_{g^*}$ as the nonempty set of safe neighbors of (a_j, b_k) with minimum degree g^*. If $\mathcal{A}_{g^*}$ does not exist (i.e., all experimental doses are deemed toxic), terminate the trial.

II2 Among the doses in $\mathcal{A}_{g^*}$, identify the dose (a_{j^*}, b_{k^*}) with the highest posterior mean of efficacy $\hat{q}_{j^*k^*}$.

First-degree neighbors of current dose combination, $\mathcal{N}_1$

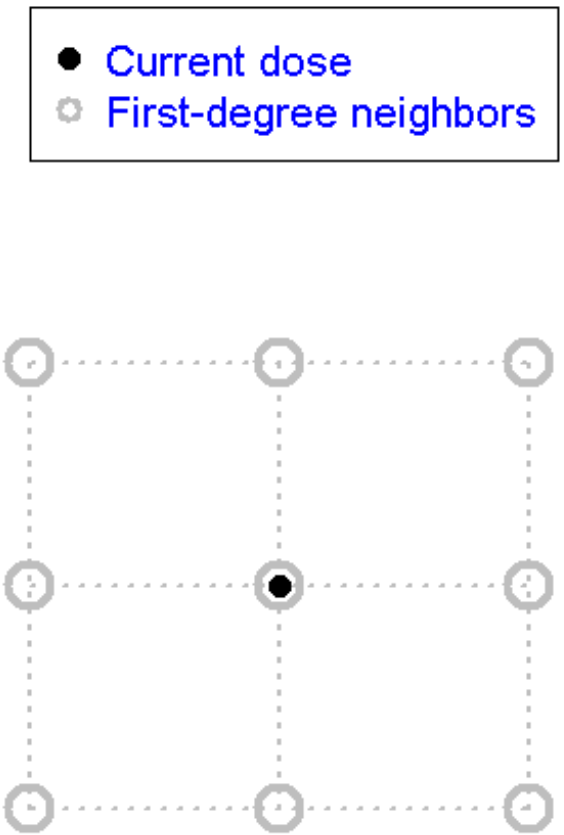

First-degree admissible dose set of current dose combination, $\mathcal{A}_1$

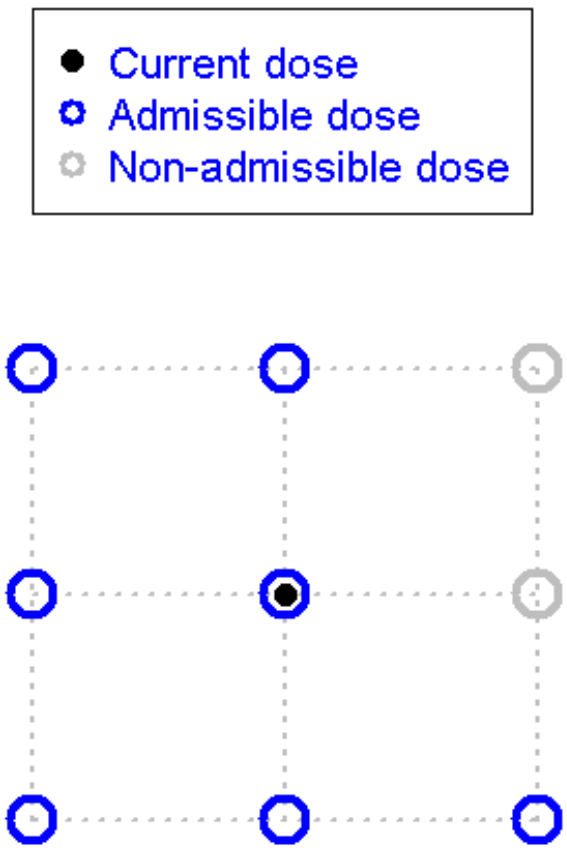

The dose (a_{j^*}, b_{k^*}) with the highest posterior mean of efficacy $\hat{q}_{j^*k^*}$

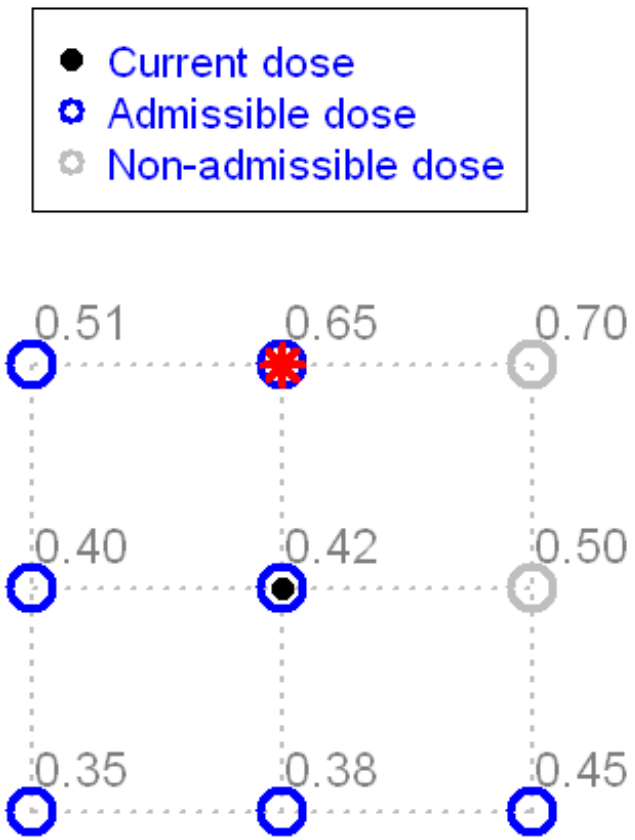

- The commonly used algorithm is to assign the next cohort of patients to (a_{j^*}, b_{k^*}).
- Problem: this greedy algorithm is easily trapped in locally optimal doses due to
  - small sample size
  - model misspecification
- Solution: a novel dose-finding algorithm to adaptively encourage the exploration of untried doses

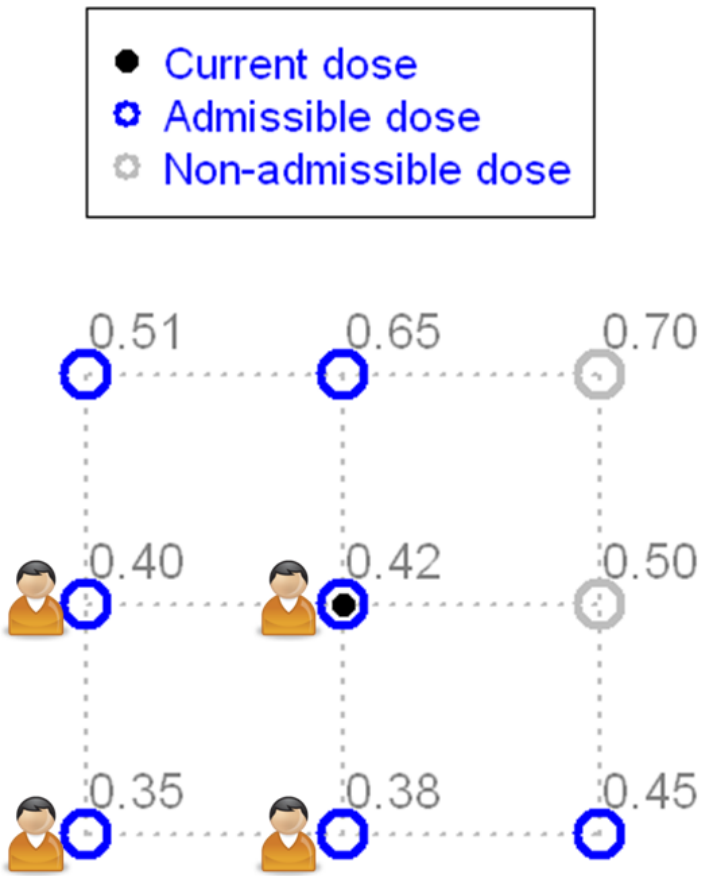
Current dose
Admissible dose
Non-admissible dose
0.51
0.65
0.70
0.40
0.42
0.50
0.35
0.38
0.45

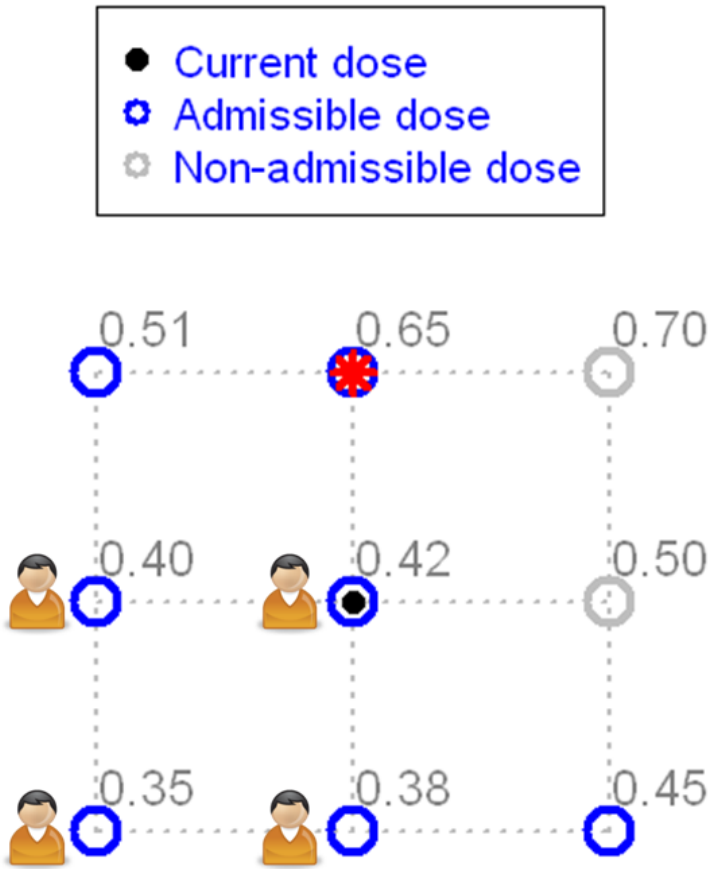
Current dose
Admissible dose
Non-admissible dose
0.51
0.65
0.70
0.40
0.42
0.50
0.35
0.38
0.45

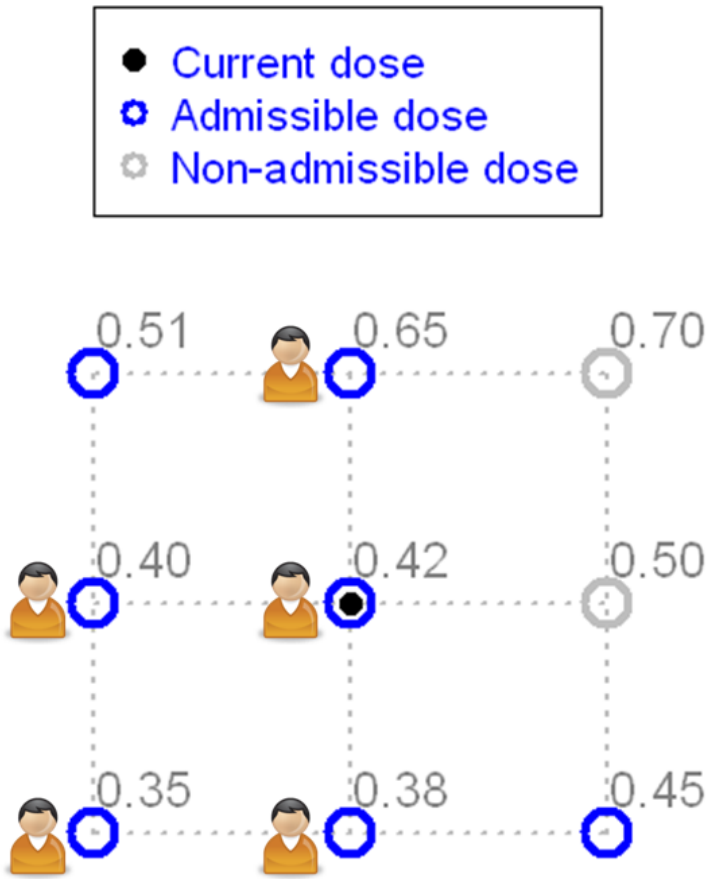
Current dose
Admissible dose
Non-admissible dose
0.51
0.65
0.70
0.40
0.42
0.50
0.35
0.38
0.45

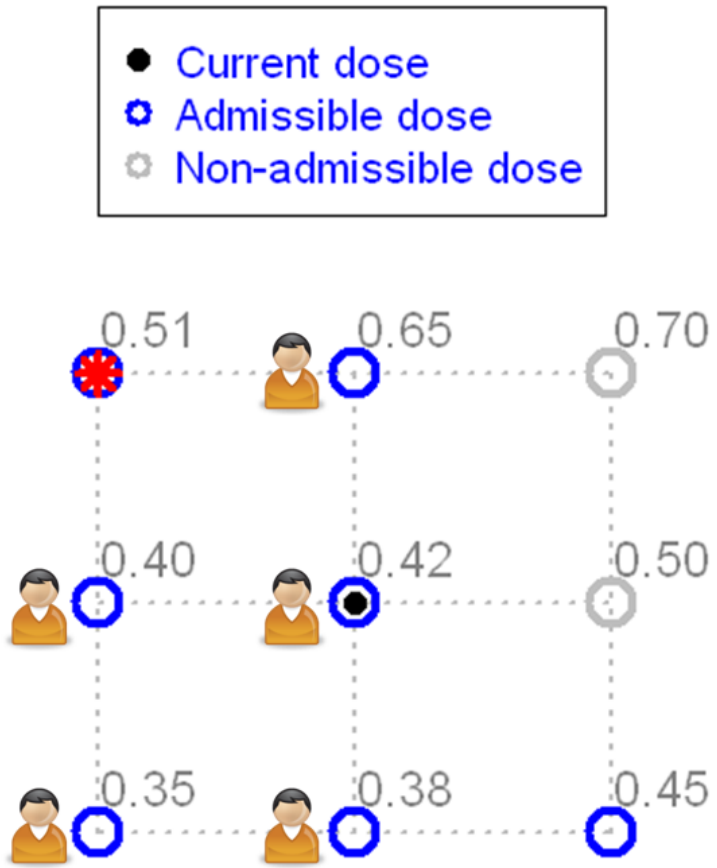
Current dose
Admissible dose
Non-admissible dose
0.51
0.65
0.70
0.40
0.42
0.50
0.35
0.38
0.45

Stage II: Systematic dose finding

II3

1. If $n_{j^*k^*} = 0$ or $n_{rs} \neq 0$ for all $(a_r, b_s) \in \mathcal{A}_{g^*}$, treat the next cohort at dose (a_{j^*}, b_{k^*}).
2. Otherwise,

$$\begin{cases} \text{If } \hat{q}_{j^*k^*} > \left(\frac{N-n}{N}\right)^{\alpha} & \text{treat the next cohort at } (a_{j^*}, b_{k^*}), \\ \text{If } \hat{q}_{j^*k^*} \leq \left(\frac{N-n}{N}\right)^{\alpha} & \text{remove dose } (a_{j^*}, b_{k^*}) \text{ from } \mathcal{A}_{g^*} \\ & \text{and go to step II2.} \end{cases}$$

- N: prespecified maximum sample size
- $n = \sum_{j,k} n_{jk}$: the total number of patients treated in the trial
- α is a known tuning parameter.

II4 Repeat steps II2-4 until exhaustion of the sample size. Select as the BODC the dose combination with the highest $\hat{q}_{jk}$ among all safe doses.

Simulation setup

- Consider 4 dose levels for each agent:
  - Dose levels of A and B are (0.075, 0.15, 0.225, 0.3) and (0.08, 0.16, 0.24, 0.32), respectively.
- The maximum sample size was 15 cohorts of size 3.
- Set the target toxicity upper limit $\phi = 0.3$ and the safety cutoff $\delta = 0.4$.
- Set the tuning parameter $\alpha = 2$.

Simulation setup

- We compared the proposed design with a greedy design that is otherwise identical except that it uses the greedy dose-assignment rule (i.e., always assign the next cohort to the dose with the highest estimate of efficacy).
- 2000 simulated trials under each scenario.

Simulation results

Table : Scenario 1

	Agent A							
Agent	Toxicity probability				Efficacy probability			
B	1	2	3	4	1	2	3	4
4	.25	.25	.25	.25	.42	.60	.38	.32
3	.15	.25	.25	.25	.19	.44	.20	.18
2	.10	.25	.25	.25	.12	.29	.15	.10
1	.05	.10	.15	.25	.05	.22	.10	.08

The blue dose is the target BODC.

Simulation results

Table : The selection percentage and the percentage of patients treated at each dose combination (shown as the subscripts) for scenario 1.

	Agent A							
	Proposed design				Greedy design			
B	1	2	3	4	1	2	3	4
4	$23.8_{14.1}$	$31.0_{15.9}$	$10.8_{9.4}$	$8.9_{8.5}$	$18.2_{9.5}$	$21.5_{10.0}$	$7.8_{5.3}$	$21.8_{26.5}$
3	$3.5_{3.9}$	$5.5_{6.0}$	$1.2_{6.9}$	$1.1_{4.6}$	$4.5_{3.0}$	$4.3_{3.0}$	$1.1_{9.5}$	$2.2_{3.2}$
2	$0.9_{2.3}$	$2.7_{8.1}$	$0.8_{3.7}$	$0.5_{2.3}$	$1.2_{1.6}$	$4.2_{11.4}$	$0.9_{1.6}$	$0.6_{1.9}$
1	$0.7_{7.6}$	$2.1_{2.8}$	$1.0_{2.1}$	$0.9_{1.8}$	$0.5_{8.4}$	$2.2_{1.9}$	$1.4_{2.1}$	$2.1_{1.2}$

The blue dose is the target BODC.

Simulation results

Table : Scenario 2

	Agent A							
Agent	Toxicity probability				Efficacy probability			
B	1	2	3	4	1	2	3	4
4	.25	.25	.25	.25	.10	.29	.29	.42
3	.15	.25	.25	.25	.25	.35	.43	.60
2	.10	.25	.25	.25	.12	.24	.32	.39
1	.05	.10	.15	.25	.05	.14	.28	.32

The blue dose is the target BODC.

Simulation results

Table : The selection percentage and the percentage of patients treated at each dose combination (shown as the subscripts) for scenario 2.

	Agent A							
Agent	Proposed design				Greedy design			
B	1	2	3	4	1	2	3	4
4	$1.6_{2.1}$	$3.2_{3.2}$	$4.1_{6.4}$	$17.0_{13.7}$	$2.5_{1.6}$	$3.1_{2.3}$	$3.9_{3.7}$	$30.1_{32.0}$
3	$2.5_{2.1}$	$2.8_{4.3}$	$7.1_{9.2}$	$33.1_{18.5}$	$2.4_{2.3}$	$3.1_{2.3}$	$9.0_{13.9}$	$17.9_{9.3}$
2	$0.7_{1.6}$	$1.5_{7.8}$	$3.4_{5.3}$	$9.6_{8.5}$	$0.8_{0.9}$	$1.1_{9.0}$	$3.0_{2.6}$	$8.2_{5.1}$
1	$0.3_{7.3}$	$0.8_{1.6}$	$2.5_{2.7}$	$6.0_{5.7}$	$0.1_{7.7}$	$0.6_{0.9}$	$2.2_{2.3}$	$7.1_{3.9}$

The blue dose is the target BODC.

Simulation results

Table : Scenario 3

	Agent A							
Agent	Toxicity probability				Efficacy probability			
B	1	2	3	4	1	2	3	4
4	.17	.25	.45	.55	.60	.35	.32	.28
3	.12	.16	.25	.43	.42	.30	.28	.25
2	.08	.10	.19	.22	.35	.28	.22	.20
1	.05	.08	.12	.18	.25	.23	.19	.16

The blue dose is the target BODC.

Simulation results

Table : The selection percentage and the percentage of patients treated at each dose combination (shown as the subscripts) for scenario 3.

	Agent A							
	Proposed design				Greedy design			
B	1	2	3	4	1	2	3	4
4	$46.3_{18.9}$	$6.8_{5.5}$	$3.4_{5.2}$	$1.3_{6.1}$	$39.1_{13.8}$	$7.1_{5.2}$	$3.3_{3.6}$	$0.9_{9.8}$
3	$7.8_{5.5}$	$2.7_{5.0}$	$3.1_{8.6}$	$2.2_{4.5}$	$7.3_{3.9}$	$2.6_{2.9}$	$3.5_{13.2}$	$2.9_{3.9}$
2	$5.3_{5.0}$	$1.9_{8.2}$	$1.5_{4.5}$	$3.1_{3.4}$	$3.9_{2.7}$	$3.0_{12.0}$	$1.8_{2.5}$	$3.9_{3.6}$
1	$5.5_{10.2}$	$2.3_{3.6}$	$1.7_{2.7}$	$2.9_{3.0}$	$8.6_{16.1}$	$2.5_{2.0}$	$2.5_{1.8}$	$4.9_{2.9}$

The blue dose is the target BODC.

Conclusions

- Our proposed design explicitly accounts for the unique features of the biological agents, i.e., dose-efficacy and -toxicity relationships may take non-monotonic patterns.
- The proposed design adaptively encourages dose exploration in the two-dimensional dose space.
- Our design identifies the BODC with substantially higher selection percentage and allocates more patients to the target dose combination than the greedy design.
- In the case that efficacy plateaus, a similar change-line model can be used.

Reference

- Cai, C., Yuan, Y. and Ji, Y. (2014) A Bayesian Phase I/II Design for Oncology Clinical Trials of Combining Biological Agents. *Journal of the Royal Statistical Society: Series C*, **63**, 159-173.

Thank you !