Improving the BLRM in Oncology Dose-Finding

Hongtao Zhang

Principal Scientist

Merck & Co., Inc.

BAYES 2022

10/13/2022

Public

Main Reference and Coauthors

- Coauthors
  - Alan Y Chiang, Lyell Immunopharma
  - Jixian Wang, Bristol Myers Squibb

- Zhang, H., Chiang, A. Y., & Wang, J. (2022). Improving the performance of Bayesian logistic regression model with overdose control in oncology dose-finding studies. *Statistics in Medicine*.

Public

Dose-Finding Designs

- An accurately identified maximum tolerable dose (MTD) serves as the cornerstone of successful subsequent phases in oncology drug development
- "Uncertainly around the optimal dose" has been identified as one of top reasons for phase 3 failures and delays/denials in regulatory filings
- Dose-finding designs can be categorized into 3 classes
  - Rule-based designs: 3+3, i3+3
  - Model-based designs: CRM, **BLRM**
  - Model-assisted designs: BOIN, mTPI(-2), keyboard
- Almost all designs have mechanisms to prevent assigning patients to overly toxic doses, a.k.a., **overdose control**
- Underdosing is also an unfavorable outcome

Public

Bayesian Logistic Regression Model (Neuenschwander et al., 2008)

- At dose level d, denote the number of DLTs r_d, out of n_d subjects

$$r_d|\pi_d \sim Bin(n_d, \pi_d)$$

- Two parameter logistic model for dose-DLT relationship

$$\log\left(\frac{\pi_d}{1-\pi_d}\right) = \log(\alpha) + \beta \log(d/d_R)$$

- Bivariate normal prior for the parameter vector $\theta \overset{\text{def}}{=} (\log(\alpha), \log(\beta))'$

$$\theta \sim N_2(\mu, \Sigma), \text{ where } \mu = (\mu_1, \mu_2)' \text{ and } \Sigma = \begin{bmatrix} \tau_1^2 & \rho\tau_1\tau_2 \\ \rho\tau_1\tau_2 & \tau_2^2 \end{bmatrix}$$

Public

Dose Recommendation by BLRM

- Define three toxicity intervals for each dose level d
  - Underdosing: $0 \leq \pi_d < 0.16$; **target toxicity**: $0.16 \leq \pi_d < 0.33$; overdosing: $\pi_d \geq 0.33$
- With cumulative data, compute the posterior probabilities that π_d falls into each toxicity interval
- Recommended dose for the next cohort
  - Has the highest probability of falling into the target toxicity interval
  - Binding overdose control (EWOC): probability of overdosing must be less than a prespecified threshold, e.g., 0.25
  - Additional rules can be applied, e.g., cannot escalate more than one dose level

Public

A Hypothetical Data Scenario

- EWOC cutoff is 0.25
- Escalating to 200 mg is prohibited due to EWOC requirement, although 100 mg is clearly underdosing
- This frustrates clinicians and sometimes results in overriding the recommendation made by BLRM

TABLE 1 A Hypothetical Data Scenario

Dose (mg)	#DLT	#Patient	P(Under)	P(Target)	P(Over)
10	0	3	0.998	0.002	0.000
25	0	3	0.992	0.008	0.000
50	0	3	0.972	0.027	0.001
100	0	3	0.777	0.186	0.037
200	0	0	0.510	0.200	**0.290**
400	0	0	0.344	0.190	0.467
800	0	0	0.239	0.167	0.595

Public

BLRM + EWOC is Overly Conservative

Statistics in CCR

Accuracy, Safety, and Reliability of Novel Phase I Trial Designs

Heng Zhou[1], Ying Yuan[1], and Lei Nie[2]

with higher accuracy of identifying the MTD. For the CRM, skipping a dose is not recommended, as it substantially increases the chance of overdosing patients while providing limited gain for identifying the MTD. EWOC and BLRM appear excessively conservative. They are safe but have relatively poor accuracy of finding the MTD. The BOIN and keyboard (equivalently mTPI-2) designs have

reduced safety (i.e., treating a high percentage of patients above the MTD) and reliability (i.e., high risk of overdosing a large percentage of patients). The overdose control rule commonly used in BLRM seems to be too conservative, and a more appropriate overdose control rule may be needed to make BLRM work

Public

Could EWOC Be the Problem?

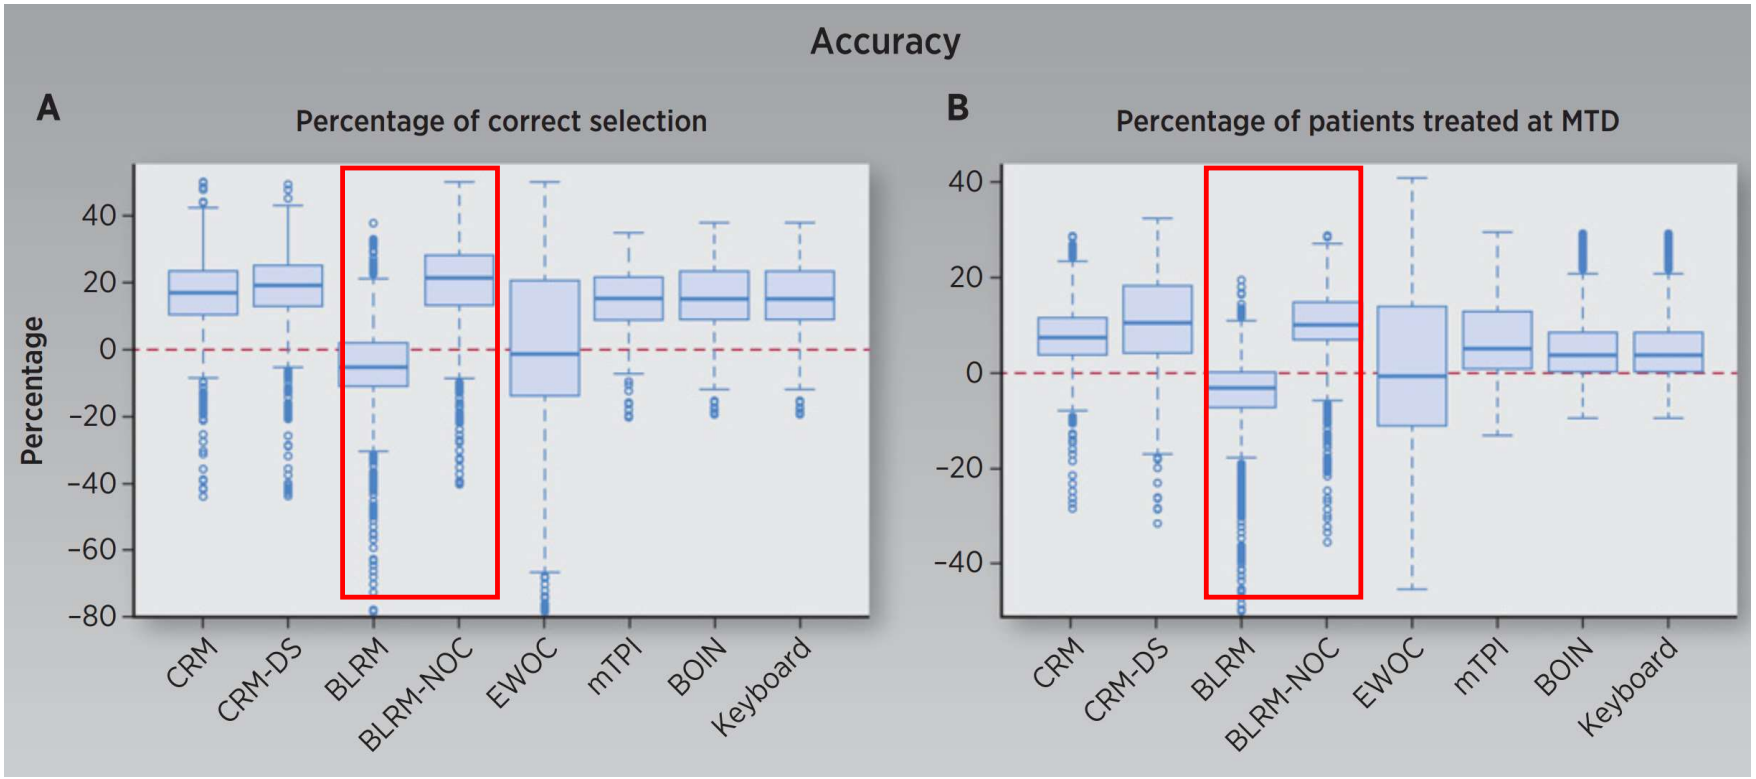

Public

A Review of Overdose Control Mechanisms

- EWOC (Babb et al., 1998)
- BOIN (Liu and Yuan, 2015)
- mTPI (Ji et al., 2010), mTPI-2 (Guo et al., 2017), keyboard (Yan et al., 2017)

- Denote
  - ϕ the target toxicity level
  - (a, b) the target toxicity interval

TABLE 2 Correct and Incorrect Actions

Scenario	Correct Action	Incorrect Action
$p_i \in [0, a]$	$\mathcal{E}$	$\bar{\mathcal{E}} = \{\mathcal{S}, \mathcal{D}\}$
$p_i \in (a, b)$	$\mathcal{S}$	$\bar{\mathcal{S}} = \{\mathcal{E}, \mathcal{D}\}$
$p_i \in [b, 1]$	$\mathcal{D}$	$\bar{\mathcal{D}} = \{\mathcal{E}, \mathcal{S}\}$

Public

Original EWOC

- DLT rate at dose d is $F(\alpha + \beta d)$, where $F(\cdot)$ is a distribution function
- With cumulative data, the posterior marginal distribution of MTD can be derived, $\pi(\cdot)$
- CRM tends to assign the next patient to the median of the distribution, i.e., $\pi^{-1}(0.5)$
- With EWOC, the assigned dose is $\pi^{-1}(\alpha_f)$, where $\alpha_f < 0.5$ is the **feasibility bound**
  - α_f may increase gradually to 0.5 or depends on # of DLTs
  - Underdosing is accounted for, although the nomenclature focuses on overdosing
- In an equivalent loss function formulation, the loss of the recommended dose level d'

$$l_{\alpha_f}(d', d_\phi) = \begin{cases} \alpha_f(d_\phi - d') & , \text{if } d' \le d_\phi, \\ (1-\alpha_f)(d' - d_\phi) & , \text{if } d' > d_\phi. \end{cases}$$

Public

BOIN

- Three point hypotheses

$$H_{0i} : p_i = \phi, \quad H_{1i} : p_i = a, \quad H_{2i} : p_i = b,$$

- The recommended action by BOIN minimizes the decision error rate

$$\alpha(\lambda_{1i}, \lambda_{2i}) = pr(H_{0i})pr(\bar{S}|H_{0i}) + pr(H_{1i})pr(\bar{\mathcal{E}}|H_{1i}) + pr(H_{2i})pr(\bar{\mathcal{D}}|H_{2i}),$$

- Underdosing, or failure to escalate, is penalized

Public

mTPI, mTPI-2 and Keyboard

- The recommended action by mTPI minimizes an expected loss function that depends on three components:
  - $L(\mathcal{S}, p_i)$ is the loss when a stay action is taken

$$L(\mathcal{S}, p_i) = \begin{cases} K_{\mathcal{S}}, p_i \in [0, a) \\ 0, p_i \in [a, b] \\ M_{\mathcal{S}}, p_i \in (b, 1] \end{cases}$$

  - $L(\mathcal{E}, p_i)$ and $L(\mathcal{D}, p_i)$ are defined likewise
- mTPI-2 and keyboard designs partition the domain of p_i to blunt the Occam's razor observed in mTPI design
  - The optimizations are also based on an expected loss, with 0-1 loss function: 0 for correct decision, 1 for incorrect decisions
- Underdosing is penalized in the loss function in all three designs

Public

Only BLRM+EWOC Lacks Underdose Control

Method	Overdosing penalized	Underdosing penalized	Asymmetric penalty
Original EWOC	Yes	Yes	Yes
BOIN	Yes	Yes	Flexible
mTPI	Yes	Yes	Flexible
Keyboard/ mTPI-2	Yes	Yes	No
BLRM + EWOC	Yes	No	-

Public

Proposed BLRM+ Diagram

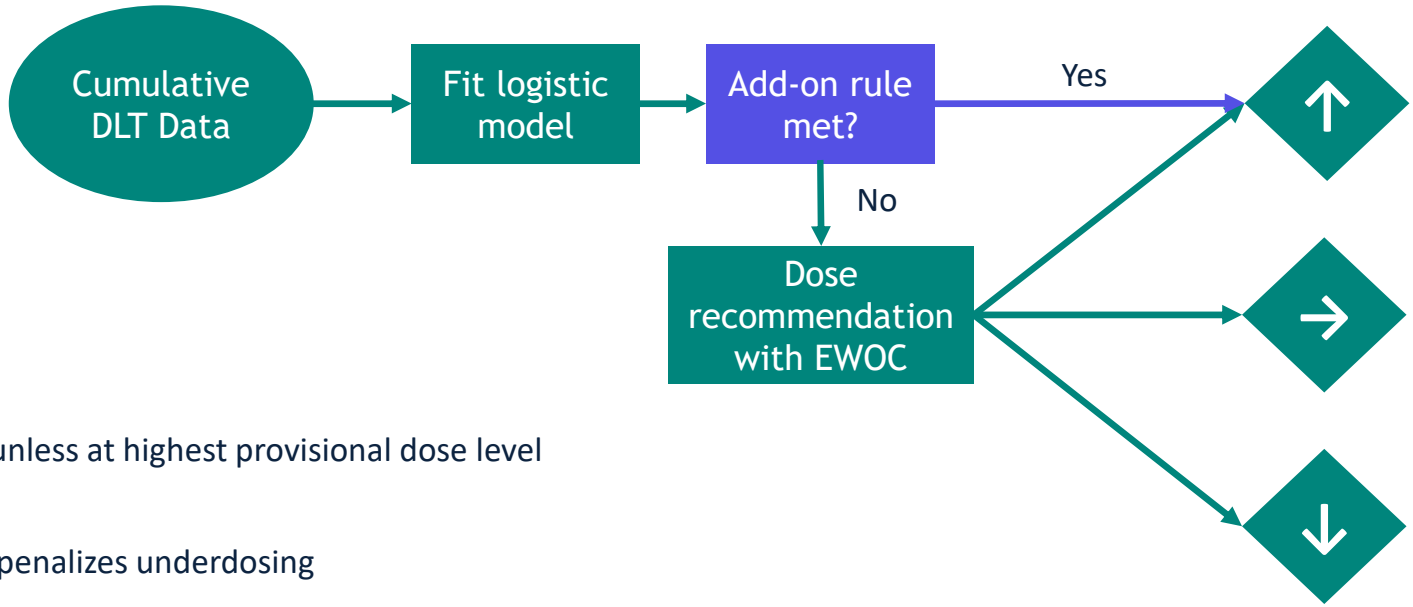

The add-on rule:

- Is assessed unless at highest provisional dose level
- Reasonably penalizes underdosing
- Leads to an **escalation** decision if met, overriding EWOC requirement

Public

Designs 1 and 2

Design 1

- Reminiscent of the EWOC by Babb et al. (1998)
- Introduces the feasibility bound $\alpha_f < 0.5$
- The add-on rule is met if

$$\alpha_f P_i(Under) > (1 - \alpha_f) P_i(Over),$$

Design 2

- Fact: dose recommendation by BLRM depends on the dose increment
- Denote relative dose strength $r_i = d_{i+1}/d_i$
- A non-decreasing function $g(r)$
- The add-on rule is met if

$$P_i(Under) > g(r_i) \cdot P_{i+1}(Over),$$

Public

Designs 3 and 4

- Designs 1 and 2 are based on the probability of DLT in underdosing/target/overdosing interval
- Like mTPI, we replace them with the corresponding unit probability mass, which account for the interval width ("model size")
  - For example, $U_i(target) = P_i(target)/(b-a)$

Design 3: This design is similar to design 1, except that the add-on rule is met if

$$\alpha_f \mathcal{U}_i(Under) > (1-\alpha_f)\mathcal{U}_i(Over).$$

Design 4: This design is similar to design 2, except that the add-on rule is met if

$$\mathcal{U}_i(Under) > g(r_i) \cdot \mathcal{U}_{i+1}(Over).$$

Public

Hypothetical Data Scenario Revisited

- Design 1 with $\alpha_f = 0.25$ is met since

  0.25 * 0.777 > (1- 0.25) * 0.037

- Design 2 with $g(r) = r$ is met since $r_i = 2$ and

  0.777 > 2 * 0.290

- UPM-based designs 3 and 4 are also met

- Recommendation would be escalation, overriding the EWOC requirement

Dose (mg)	#DLT	#Patient	P(Under)	P(Target)	P(Over)
10	0	3	0.998	0.002	0.000
25	0	3	0.992	0.008	0.000
50	0	3	0.972	0.027	0.001
100	0	3	0.777	0.186	0.037
200	0	0	0.510	0.200	**0.290**
400	0	0	0.344	0.190	0.467
800	0	0	0.239	0.167	0.595

Public

Simulation Studies

- Compare BLRM+EWOC and four new designs
- Aligned specs across methods
  - Provisional doses: 10, 25, 50, 100, 200, 400, 800 (mg)
  - Target toxicity interval (0.16, 0.33), EWOC cutoff is 0.3
  - Weakly informative prior
  - Maximum of 45 subjects
  - Skipping dose not allowed
  - The MTD can be declared if (1) it is recommended to stay at current dose level, AND (2) at least 6 patients at current dose level, AND (3) the probability of target toxicity is at least 0.5
  - All doses are deemed overly toxic if $P_i(over) > 0.3$ at all doses
- For designs 1 and 3: $\alpha_f = 0.25$
- For designs 2 and 4: $g(r) = r$
- Performance metrics
  - MTD accuracy
  - Number of patients treated at MTD
  - Overall sample size
  - Overall DLT rate

Public

Dose-Toxicity Scenarios

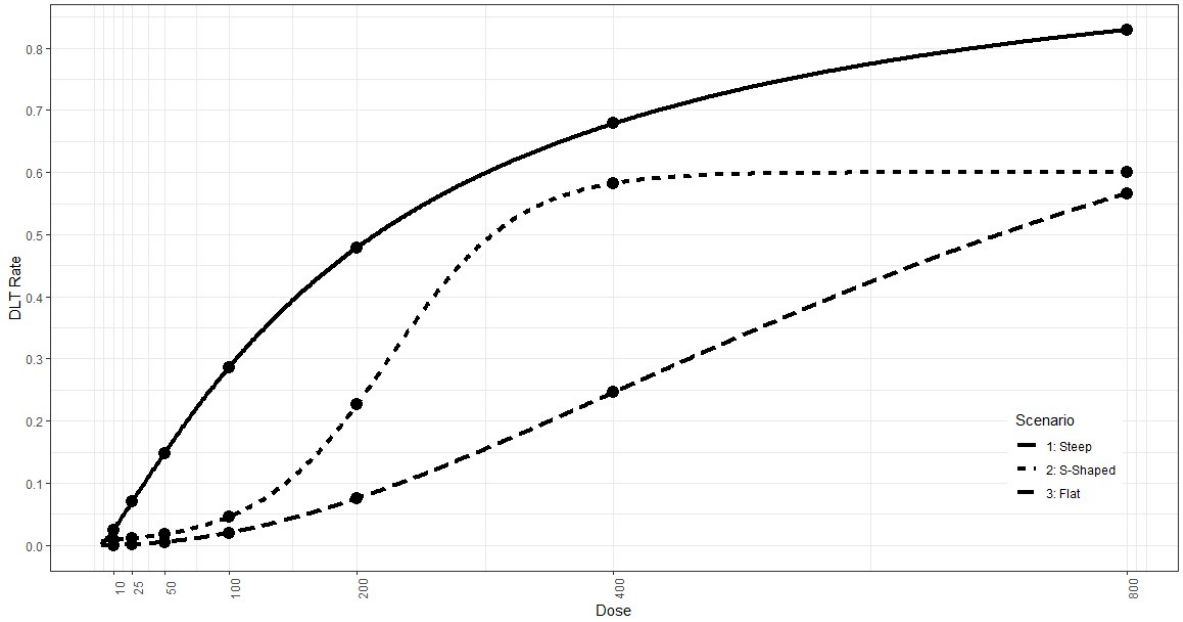

- Steep, S-shaped and flat
- Two additional scenarios considered
  - All doses are safe (DLT rate < 0.16)
  - All doses are toxic (DLT rate > 0.33)

Public

S-shaped Scenario

S-shaped Scenario										
	Original BLRM		Design 1		Design 2		Design 3		Design 4	
MTD	Frequency	N	Frequency	N	Frequency	N	Frequency	N	Frequency	N
AllToxic	0.017	-	0.017	-	0.017	-	0.017	-	0.017	-
10	0.000	3.01	0.000	3.01	0.000	3.01	0.000	3.00	0.000	3.00
25	0.000	3.07	0.000	3.03	0.000	3.07	0.000	2.97	0.000	2.97
50	0.000	3.74	0.000	3.11	0.000	3.70	0.001	3.02	0.001	3.03
100	0.047	14.30	0.051	7.91	0.045	12.08	0.060	6.62	0.057	6.70
200	**0.561**	**9.03**	**0.749**	**10.58**	**0.621**	**9.56**	**0.737**	**9.63**	**0.742**	**9.67**
400	0.028	2.36	0.044	3.73	0.026	2.53	0.052	3.99	0.053	3.88
800	0.002	0.22	0.002	0.30	0.001	0.21	0.004	0.71	0.003	0.69
NotFound	0.345	-	0.137	-	0.290	-	0.129	-	0.127	-
Overall		35.74		31.67		34.16		29.94		29.93
%DLT	12.9		16.4		13.6		17.6		17.4	

Public

Steep Scenario

Steep Scenario										
	Original BLRM		Design 1		Design 2		Design 3		Design 4	
MTD	Frequency†	N‡	Frequency	N	Frequency	N	Frequency	N	Frequency	N
AllToxic[1]	0.051	-	0.051	-	0.051	-	0.051	-	0.051	-
10	0.001	3.22	0.000	3.11	0.001	3.20	0.001	3.07	0.000	3.08
25	0.037	5.35	0.031	4.08	0.034	5.01	0.027	3.70	0.028	3.70
50	0.392	12.63	0.303	7.57	0.395	11.69	0.314	7.02	0.316	7.13
100	**0.318**	**6.54**	**0.470**	**8.80**	**0.354**	**6.54**	**0.462**	**7.91**	**0.456**	**8.07**
200	0.012	0.76	0.036	2.52	0.019	1.09	0.049	3.55	0.049	3.38
400	0.000	0.07	0.001	0.19	0.000	0.10	0.000	0.56	0.000	0.53
800	0.000	0.00	0.000	0.00	0.000	0.00	0.000	0.05	0.000	0.04
NotFound[2]	0.188	-	0.107	-	0.145	-	0.096	-	0.100	-
Overall[3]		28.60		26.30		27.65		25.85		25.93
%DLT[4]	17.7		20.7		18.2		22.5		22.2	

Public

Flat Scenario

Flat Scenario

	Original BLRM		Design 1		Design 2		Design 3		Design 4	
MTD	Frequency	N	Frequency	N	Frequency	N	Frequency	N	Frequency	N
AllToxic	0.000	-	0.000	-	0.000	-	0.000	-	0.000	-
10	0.000	3.00	0.000	3.00	0.000	3.00	0.000	3.00	0.000	3.00
25	0.000	3.02	0.000	3.01	0.000	3.02	0.000	3.00	0.000	3.00
50	0.000	3.17	0.000	3.03	0.000	3.17	0.000	3.00	0.000	3.00
100	0.003	6.67	0.005	3.93	0.002	4.71	0.003	3.18	0.003	3.25
200	0.167	10.12	0.196	7.01	0.197	9.50	0.246	6.35	0.251	6.41
400	**0.524**	**7.76**	**0.685**	**8.87**	**0.595**	**8.51**	**0.624**	**7.93**	**0.615**	**7.92**
800	0.039	2.48	0.042	3.24	0.030	2.44	0.040	3.77	0.043	3.73
NotFound	0.267	-	0.072	-	0.176	-	0.087	-	0.088	-
Overall		36.20		32.09		34.34		30.23		30.31
%DLT	12.0		14.3		12.9		15.2		15.2	

Public

Observations

- Remarkable improvement over original BLRM+EWOC
  - Higher accuracy to identify the MTD
  - More patients treated at MTD
  - Smaller total patients required
  - Similar overall DLT rate

Public

All-Doses-Safe Scenario

All-Safe Scenario

	Original BLRM		Design 1		Design 2		Design 3		Design 4	
MTD	Frequency†	N‡	Frequency	N	Frequency	N	Frequency	N	Frequency	N
AllToxic[1]	0.017	-	0.017	-	0.017	-	0.017	-	0.017	-
10	0.000	3.01	0.000	3.01	0.000	3.01	0.000	3.00	0.000	3.00
25	0.000	3.03	0.000	3.01	0.000	3.03	0.000	2.96	0.000	2.96
50	0.000	3.29	0.000	3.00	0.000	3.26	0.000	2.97	0.000	2.97
100	0.001	5.38	0.001	3.37	0.001	3.71	0.000	2.97	0.000	3.03
200	0.005	3.59	0.005	3.30	0.004	3.58	0.005	3.03	0.006	3.06
400	0.021	4.01	0.017	3.64	0.018	4.02	0.020	3.26	0.016	3.33
800	0.095	20.84	0.134	23.21	0.119	22.20	0.151	23.86	0.151	23.75
NotFound[2]	0.861	-	0.826	-	0.841	-	0.807	-	0.810	-
Overall[3]		43.16		42.53		42.81		42.04		42.11
%DLT[4]	5.5		5.9		5.8		6.2		6.1	

- New designs escalate faster to the highest dose

Public

All-Doses-Toxic Scenario

All Toxic Scenario										
	Original BLRM		Design 1		Design 2		Design 3		Design 4	
MTD	Frequency	N	Frequency	N	Frequency	N	Frequency	N	Frequency	N
AllToxic	0.855	-	0.851	-	0.855	-	0.838	-	0.838	-
10	0.064	4.06	0.055	4.00	0.064	4.06	0.056	3.83	0.057	3.88
25	0.060	1.47	0.065	1.58	0.060	1.47	0.064	1.94	0.063	1.83
50	0.019	0.47	0.022	0.57	0.019	0.46	0.030	0.86	0.029	0.86
100	0.001	0.08	0.005	0.20	0.002	0.08	0.009	0.25	0.010	0.25
200	0.000	0.00	0.001	0.03	0.000	0.01	0.002	0.07	0.002	0.07
400	0.000	0.00	0.000	0.01	0.000	0.00	0.000	0.02	0.000	0.02
800	0.000	0.00	0.000	0.00	0.000	0.00	0.000	0.00	0.000	0.00
NotFound	0.001	-	0.001	-	0.000		0.001	-	0.001	-
Overall		6.08		6.38		6.07		6.98		6.90
%DLT	47.2		47.4		47.2		47.6		47.6	

- New designs also terminate the study very fast

Public

H2H with BOIN and Keyboard Designs

- BOIN and keyboard designs as implemented in R packages *BOIN* and *keyboard* respectively
- Fundamental differences make it impossible to perfectly align simulation specifications
  - How doses are deemed overly toxic
  - How the MTD is determined/estimated
  - How early stopping is triggered
  - "MTD not found": possible outcome in BLRM-based designs, but not in BOIN and keyboard
- In simulation, the **average sample sizes** are aligned between BOIN/keyboard and new design 1
- Preliminary results are promising enough to merit a more comprehensive comparison

Public

H2H Results

TABLE 7 Simulation results: Comparison with BOIN and keyboard designs, TTI (0.16, 0.33)

	Original BLRM	Design 1		BOIN	Keyboard	
Scenario	MTD accuracy (*N* at MTD)	MTD accuracy (*N* at MTD)	Avg *N* total	MTD accuracy (*N* at MTD)	MTD accuracy (*N* at MTD)	Avg *N* total
Steep	0.318 (6.54)	0.470 (8.80)	26.30	0.444 (6.01)	0.445 (6.01)	27
S-shaped	0.561 (9.03)	0.749 (10.58)	31.67	0.834 (11.50)	0.834 (11.50)	33
Flat	0.524 (7.76)	0.685 (8.87)	32.09	0.721 (8.61)	0.721 (8.61)	33
All-safe[a]	0.956 (20.84)	0.960 (23.21)	42.53	0.990 (20.50)	0.987 (20.50)	42
All-toxic[b]	0.855 (-)	0.851 (-)	6.38	0.292 (-)	0.292 (-)	8.2

[a]For original BLRM and design 1, MTD accuracy is the frequency of the highest dose level being identified as MTD, or no MTD being claimed. For BOIN and keyboard designs, MTD accuracy is only the frequency of the highest dose level being identified as MTD. *N* at MTD is the number of patients treated at the highest dose level.

[b]For BOIN and keyboard designs, MTD accuracy is the frequency of the lowest dose level being eliminated due to toxicity. *N* at MTD is not applicable under this scenario.

Public

Parting Thoughts

- Various designs have respective strengths and limitations
- Model-based designs are more involved but flexible to handle sophisticated dose-finding settings
  - Covariates and drug combinations
  - Robust information-sharing across potentially heterogeneous strata
- Our idea applies to most BLRM extensions
- Can we improve the original BLRM design with more elegance, e.g., via Bayesian decision-theoretic?

Public

Commentaries and Rejoinder

- Commentaries by
  - Drs. Ying Yuan and Yixuan Zhao
  - Drs. Paul Frankel, Elizabeth Garrett-Mayer and Mark Kralio
  - Drs. Alexia Iasonos and Graham Wheeler
- Topics discussed include
  - Alternative approaches
  - Practical considerations
  - Model parsimony vs. flexibility

Public

Key References

- Neuenschwander, B., Branson, M., & Gsponer, T. (2008). Critical aspects of the Bayesian approach to phase I cancer trials. Statistics in medicine, 27(13), 2420-2439.
- Babb, J., Rogatko, A., & Zacks, S. (1998). Cancer phase I clinical trials: efficient dose escalation with overdose control. Statistics in medicine, 17(10), 1103-1120.
- Liu, S., & Yuan, Y. (2015). Bayesian optimal interval designs for phase I clinical trials. Journal of the Royal Statistical Society: Series C: Applied Statistics, 507-523.
- Ji, Y., Liu, P., Li, Y., & Nebiyou Bekele, B. (2010). A modified toxicity probability interval method for dose-finding trials. Clinical trials, 7(6), 653-663.
- Guo, W., Wang, S. J., Yang, S., Lynn, H., & Ji, Y. (2017). A Bayesian interval dose-finding design addressing Ockham's razor: mTPI-2. Contemporary clinical trials, 58, 23-33.
- Yan, F., Mandrekar, S. J., & Yuan, Y. (2017). Keyboard: a novel Bayesian toxicity probability interval design for phase I clinical trials. Clinical Cancer Research, 23(15), 3994-4003.
- Zhou, H., Yuan, Y., & Nie, L. (2018). Accuracy, safety, and reliability of novel phase I trial designs. Clinical Cancer Research, 24(18), 4357-4364.

MERCK