MSR82

Combating Sample Scarcity: A Novel Bayesian Approach to Pediatric Basket Trials

Emma Mackay,[1] Aaron Springford,[1] Bart Heeg,[2] Paul Arora,[3] Kristian Thorlund[4]
[1]Cytel, Toronto, Canada, [2]Cytel, Rotterdam, Netherlands, [3]Dalla Lana School of Public Health, University of Toronto, Toronto, Canada, [4]Cytel, Vancouver, Canada

Evaluating the efficacy of histology-independent therapies for rare indications is particularly challenging in pediatric settings due to the difficulty of enrolling enough patients to conduct a well-powered clinical trial. We propose a methodology that allows for information to be borrowed both across basket trial histologies and from adult basket trial populations to improve the precision of pediatric response rate estimates.

Background

Evaluating efficacy of novel treatments for rare diseases in pediatric populations presents a substantial challenge due to the difficulty of recruiting enough patients to conduct a well-powered clinical trial

These difficulties are particularly pronounced for pediatric basket trials evaluating the efficacy of histology-independent therapies (HIT) in oncology

Basket trials recruit patients with multiple disease subtypes which share a common mutation or biomarker targeted by the therapy

- A challenge is that different histologies may respond differently to treatment—should we pool each histology or analyze them separately?
- Basket trials typically have extremely limited sample sizes within each histology—especially for pediatric basket trials

There is guidance and precedent in both regulatory and health technology assessment (HTA) settings for using Bayesian methods to address some of these challenges arising in rare indications[1-4]

- Use of Bayesian borrowing methods to supplement limited pediatric sample sizes using adult trial data[5]
- Use of Bayesian hierarchical models (BHM) to partially pool information across histologies in basket trials based on the degree of observed cross-histology heterogeneity[1]

Objective

We propose a method which combines (1) a BHM approach for partial pooling of information across histologies and (2) Bayesian borrowing from adult basket trials into pediatric basket trials using power priors[6] to estimate histology-specific overall response rates (ORR).

Methods

Overview

We use a BHM which partially pools information across histologies, where the extent of pooling is dependent on the variability in response across histologies in the trial data.

We additionally augment the pediatric basket trial data by borrowing information from an adult basket trial investigating the same treatment via a power prior approach[6].

The power prior down-weights the adult data based on a fixed discount parameter which we vary from 0 (no borrowing) to 1 (complete pooling) and allows for a tipping point analysis[5,7] in which the sensitivity of any efficacy conclusions to the amount of borrowing is assessed (e.g. how much borrowing is needed to exceed an ORR threshold for concluding efficacy?)

We demonstrate the method using simulated data containing $n_P = 50$ patients in the pediatric data and $n_A = 100$ patients in the adult data, split across $K = 8$ prognostically important histologies.

Model Specification

We model patient responses $y_i \in \{0, 1\}$ for $i = 1, \dots, n_A + n_P$ patients using the following logistic regression specification:

$$y_i \sim \text{Bernoulli}(p_i)$$

$$\text{logit}(p_i) = \mu + Z_i'\gamma + X_i'\beta$$

where μ is an intercept term, Z_i is a one-hot-encoded vector specifying patient i's histology, $\gamma = (\gamma_1, \dots, \gamma_K)$ is a vector of histology-specific random effects with distribution $\gamma_k \sim N(0, \sigma^2)$ for all $k = 1, \dots, K$, X_i is an optional vector of regressors with corresponding fixed effect coefficient vector β.

To allow for down-weighting of the adult data, we use a static power prior for our parameter vector $\theta = (\mu, \gamma_1, \dots, \gamma_K, \beta_1, \dots, \beta_J, \sigma)$ which is based on the parameter likelihood on the external adult data raised to the power of a discount parameter $\alpha_0 \in [0, 1]$. For the external adult data $D_0 = \{(y_i, Z_i, X_i)\}_{i=1}^{n_A}$ and pediatric data $D = \{(y_i, Z_i, X_i)\}_{i=n_A+1}^{n_A+n_P}$, the power prior is as follows:

$$\pi(\theta|D_0, \alpha_0) \propto L(\theta|D_0)^{\alpha_0}\pi_0(\theta)$$

Where $L(\theta|D_0)$ is the likelihood for the external adult data and $\pi_0(\theta)$ is the prior for θ before observing either the adult or pediatric data.

By incorporating the likelihood for the pediatric data, $L(\theta|D)$, we can form our posterior for θ as follows:

$$\pi(\theta|D, D_0, \alpha_0) \propto L(\theta|D)\pi(\theta|D_0, \alpha_0)$$

$$\propto L(\theta|D)L(\theta|D_0)^{\alpha_0}\pi_0(\theta)$$

In our simulation/model implementation we ignore the $X_i'\beta$ term for simplicity. This assumes that, within each histology, adult and pediatric patients in our samples have similar prognosis at baseline. This is a strong assumption which can be relaxed but caution needs to be taken due to the risk of overparameterization.

We use the following priors in forming $\pi_0(\theta)$:

$$\mu \sim \text{N}(0, 10^2)$$

$\gamma_k \sim N(0, \sigma^2)$ for all $k = 1, \dots, K$ (the random effects)

$$\sigma \sim \text{Half-Cauchy}(0, 1)$$

Simulation and Estimation Procedure

To demonstrate the modelling approach, we simulate adult and pediatric basket trial datasets under this model specification with parameter values $\mu = 0$, $\sigma = 0.5$, and $\boldsymbol{\beta} = \mathbf{0}$, with each patient having an equal probability of assignment to each of the $K = 8$ histologies

We estimate posterior medians and 95% credible intervals (CrI) for the ORR in each histology under a sliding scale of α_0 borrowing weights between 0 (no borrowing) and 1 (complete pooling) to assess the impact of borrowing from the adult population on pediatric ORR estimates.

Posterior estimates are computed using Markov chain Monte Carlo (MCMC) methods using the Stan No-U-Turn Sampler (NUTS) Hamiltonian Monte Carlo (HMC) implementation. Stan was invoked via the R programming language using the *rstan* package.

For each α_0 scenario, 40,000 MCMC iterations were run with a burn-in of 5,000 for each of 4 chains. MCMC convergence was assessed via $\hat{R}$ statistics for the ORRs.

Results

Table 1 reports the number and percent of pediatric and adult responders in the simulated data, as well as the "true" response rates that were used to simulate the data. We can see that there are very few patients within each histology—especially for the pediatric basket trial—and that observed ORRs can vary wildly from the "true" ORR due to sampling variability. Note in particular histology 2—where the sample pediatric and adult ORRs are 60% and 20%, respectively, compared to the true ORR of 47.2%—and histology 7 where 100% (5 / 5) of the pediatric patients were responders despite the true ORR of 66%.

Table 1. Pediatric and Adult Sample ORRs for Simulated Data vs. True ORR

Histology	Pediatric Responders (ORR in %)	Adult Responders (ORR in %)	True ORR
Histology 1	2 / 4 (50.0%)	4 / 12 (33.3%)	37.9%
Histology 2	6 / 10 (60.0%)	2 / 10 (20.0%)	47.2%
Histology 3	4 / 7 (57.1%)	9 / 18 (50.0%)	52.7%
Histology 4	2 / 4 (50.0%)	3 / 5 (60.0%)	58.9%
Histology 5	4 / 6 (66.7%)	10 / 16 (62.5%)	61.4%
Histology 6	4 / 5 (80.0%)	9 / 11 (81.8%)	64.7%
Histology 7	5 / 5 (100.0%)	9 / 16 (56.2%)	66.0%
Histology 8	8 / 9 (88.9%)	9 / 12 (75.0%)	76.0%

Abbreviations: ORR, overall response rate

Figure 1 plots point estimates (posterior medians) and 95% credible intervals (CrI) for the pediatric ORR for each histology under a series for increasing borrowing weight scenarios (increasing values of the α_0 discount parameter) from the adult trial population. We see that the width of the 95% CrIs tends to shrink with additional borrowing and that the ORR point estimates tend to (but not always) move closer to the true values.

The method yielded reductions in the width the 95% credible intervals (CrI) ranging from 7.8% to 28.2% for the 8 histologies when varying the power prior discount parameter from 0 to 1.

Figure 1. Pediatric ORR Estimates (Posterior Median and 95% CrIs) for Various Borrowing Weights (Red Line Indicates True ORRs)

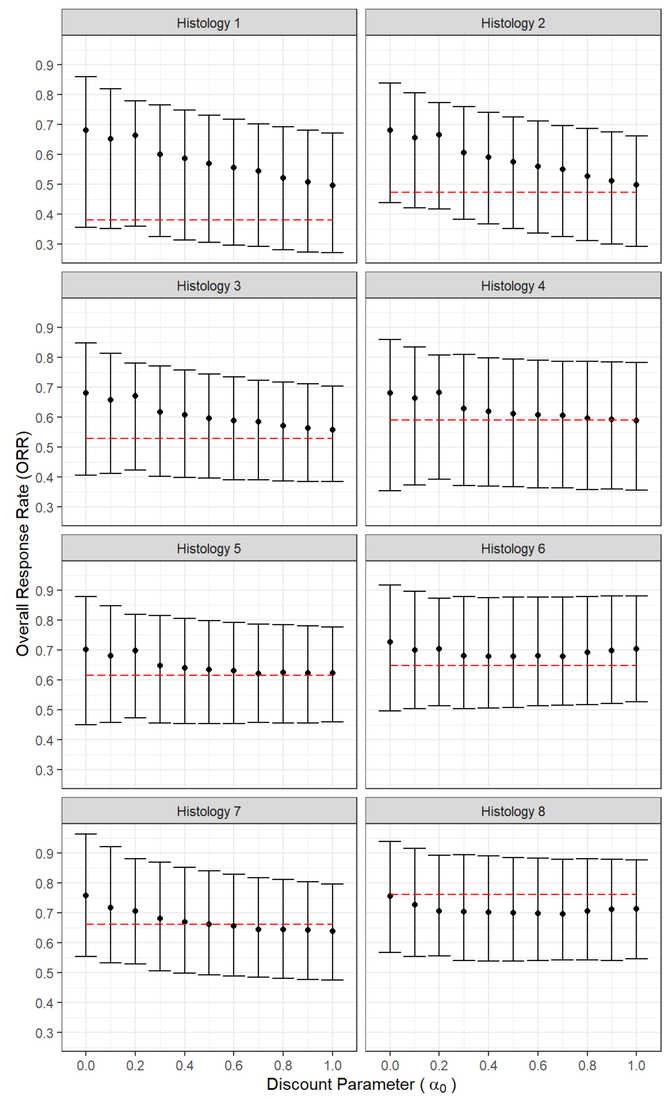

Note: MCMC convergence was not achieved for the $\alpha_0 = 0.2$ case so point estimates and 95% CrIs for this scenario should be regarded as unreliable.

Abbreviations: ORR, overall response rate; CrI, credible interval

Limitations

Bayesian borrowing methods have the potential to introduce bias and inflate type-I error where information borrowing is unwarranted—i.e. if adult trial patients differ systematically from pediatric patients in terms of their prognosis at baseline or treatment benefit then borrowing will come at the expense of some bias

BHM methods for modelling cross-histology heterogeneity in basket trials assume that histologies satisfy an exchangeability requirement.

While modifications to the model can be made to adjust for differences in prognosis at baseline, different cross-histology heterogeneity assumptions, or other clinical considerations, these modifications may rely on strong structural assumptions and may lead to issues of model over-parameterization where sample sizes are extremely limited.

Practitioners should be cognizant of the potential variance-bias trade-offs involved in incorporation of information across histologies or patient populations when using these methods; care should be taken in the selection of priors and sensitivity analyses should also be considered.

Conclusions

We propose a new method combining a BHM and Bayesian borrowing via a power prior to facilitate information sharing both across histologies and between adult and pediatric trial populations to improve the precision of ORR estimates for pediatric basket trials.

In conjunction with tipping point analysis and other sensitivity analyses, this approach can help surmount the challenges of evaluating the efficacy of HITs in pediatric populations where sample sizes are extremely limited.

References

1. Murphy P, Glynn D, Dias S, Hodgson R, Claxton L, Beresford L, Cooper K, Tappenden P, Ennis K, Grosso A, Wright K. Modelling approaches for histology-independent cancer drugs to inform NICE appraisals: a systematic review and decision-framework. 2. US Food and Drug Administration. Guidance for the Use of Bayesian Statistics in Medical Device Clinical Trials. Guidance for Industry and FDA Staff. 2010. https://www.fda.gov/media/71512/download. 3. US Food and Drug Administration. Adaptive Designs for Clinical Trials of Drugs and Biologics Guidance for Industry. 2019. https://www.fda.gov/media/78495/download. 4. US Food and Drug Administration. Interacting with the FDA on Complex Innovative Trial Designs for Drugs and Biological Products: Guidance for Industry. 2020. https://www.fda.gov/media/78495/download. 5. Food and Drug Administration. BLA 125370/s-064 and BLA 761043/s-007 multi-disciplinary review and evaluation benlysta®(belimumab) for intravenous infusion in children 5 to 17 years of age with SLE. 2021. https://www.fda.gov/media/127912/download. 6. Ibrahim JG, Chen MH. Power prior distributions for regression models. Statistical Science. 2000 Feb 1:46-60. 7. Best N, Price RG, Pouliquen IJ, Keene ON. Assessing efficacy in important subgroups in confirmatory trials: An example using Bayesian dynamic borrowing. Pharmaceutical statistics. 2021 May;20(3):551-62.

Disclosures

EM, AS, BH and KT are employees of Cytel. PA is an employee of the Dalla Lana School of Public Health at the University of Toronto, Canada.

Presented at ISPOR EU | Copenhagen, Denmark | November 12–15, 2023

Contact information: emma.mackay@cytel.com