Choice of Hyperparameters for Spike And Slab Prior, with Application to Dose Optimization in Oncology

Ramon I. Garcia[1], Jim Rogers[1]
[1]*Metrum Research Group, Tariffville, CT, USA*,

Objectives

Objectives: Dose optimization requires consideration of intrinsic factors to determine whether clinically identifiable subpopulations require dosing adjustments. In practice this involves the use of population pharmacokinetic and exposure-response models that include such factors represented as covariates. Regularized regression using spike and slab priors is appealing when the number of plausible covariate effects is too great to reliably estimate with maximum likelihood. Herein we propose heuristics to justify particular hyperparameter values when implementing spike and slab for logistic regression. We demonstrate our proposed reasoning in the context of valemetostat treatment for Adult T-cell Leukemia / Lymphoma (ATTLL) and Non-Hodgkin's lymphoma (NHL).

Methods

Spike and slab is a shrinkage method, in that it shrinks "weak" regression coefficients (covariate effect) values towards zero [1, 2, 3]. Each regression coefficient is modeled as coming from a mixture of two normal distributions with different variances: one with density concentrated around zero (the spike), and the other with density spread out over large plausible values (the slab). This method does not exclude any covariates, but instead stabilizes estimation by shrinking negligible (near zero) covariate effects towards the spike component of the prior, while applying minimal regularization to non-negligible effects. A prior distribution for the penalty parameter λ, which represents the probability of being in the spike as opposed to the slab prior, is also specified.

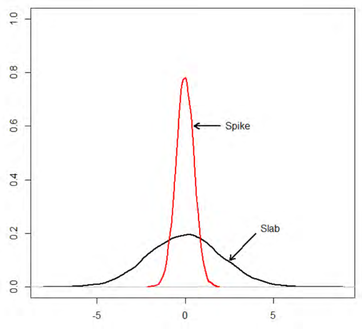

Figure 1. Spike and Slab Prior.

$$\text{logit}(p) = \alpha_0 + \beta_{\text{E}}\text{E} + \boldsymbol{X}^T\boldsymbol{\alpha}_x + \boldsymbol{X}^T\boldsymbol{\beta}_x\text{E} \qquad (1)$$

- p is the probability of having an event
- E is a standardized exposure metric
- X is a standardized covariate matrix

$$\begin{aligned} y &\sim \text{Bernoulli}(1, p) \\ \alpha_0 &\sim \text{N}(0, \sigma_{\text{slab}}) \\ \beta_{\text{E}} &\sim \text{N}(0, \sigma_{\text{slab}}) \\ \boldsymbol{\alpha}_x &\sim \lambda\,\text{N}(0, \sigma_{\text{spike}}) + (1-\lambda)\,\text{N}(0, \sigma_{\text{slab}}) \\ \boldsymbol{\beta}_x &\sim \lambda\,\text{N}(0, \sigma_{\text{spike}}) + (1-\lambda)\,\text{N}(0, \sigma_{\text{slab}}) \\ \lambda &\sim \text{Beta}(a, b) \end{aligned}$$

Exposure-response models for efficacy endpoints (overall response rate - central assessment, overall response rate - investigator assessment) and safety endpoints (anemia, neutrophil count decrease, platelet count decrease, any adverse event grade 3, dose interruption due to adverse event, dose reduction due to adverse event) were developed using unbound area under the concentration (AUC) curve at steady state as the exposure metric. Spike and slab priors were considered for the following covariates: age, weight, lactate dehydrogenase, Eastern Cooperative Oncology Group (ECOG) status, NHL disease stage, sex, race, country of origin, baseline hemoglobin, baseline platelets, and baseline neutrophils. Our strategy consisted of the following key elements:

Selection of Slab

$\sigma_{\text{slab}} = 1.5$ corresponds in a weakly informative prior for the probability of the response or event (inverse logit α_0) (**Figure 2c**) and odds ratio ($\exp(\beta_{\text{E}})$) (**Figure 2b**).

Selection of Spike

$\sigma_{\text{spike}} = 0.1$ results in (0.82,1.21) as a 95% prior credible interval (CRI) for the odds ratio (**Figure 2a**) ($\exp(\alpha_{X,i})$) corresponding to a nominally negligible covariate effect obtained by solving for sigma as a function of the "negligible" odds ratio magnitude τ in Equation 2.

$$Pr(1/\tau < \exp(\alpha_X) < \tau) = Pr(|Z| < \log(\tau)/\sigma_{\text{spike}}) \implies \sigma_{\text{spike}} = \log(\tau)/1.96 \qquad (2)$$

Selection of λ

A Beta(a, b) prior has an expected value of a/(a+b). Beta$(1, 1)$, Beta$(1/2, 5)$, and Beta$(5, 1/2)$ priors for λ, correspond to approximately 50%, 9%, and 91% of the total covariate effects having negligible magnitudes, respectively (**Figure 2c**). The Beta(1/2, 5) prior results in less regularized estimates that more closely resemble maximum likelihood estimates from a full model. A Beta(1/2, 5) prior results in more regularized estimates that more closely resemble a base model, e.g., a model with just one covariate or exposure. Beta(1, 1) results in a weakly informative uniform prior. A Beta(1, 1) was used in the primary analysis, and Beta(1/2, 5) and Beta(1/2, 5) priors were used in sensitivity analyses.

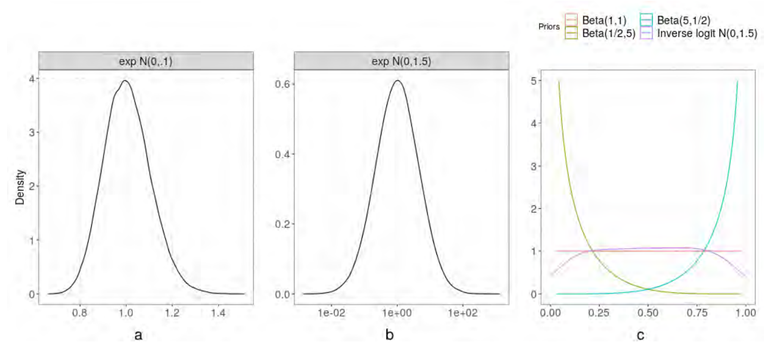

Figure 2. Prior Densities of Parameters Under Different Hyperparameter Values. Figure 2a is a spike prior for the odds ratio of a negligible covariate. Figure 2b is a slab prior for the odds ratio of a covariate with 95% prior CRI of (0.05, 19.0) which corresponds to a weakly informative prior. Figure 2c are different Beta priors for λ and a weakly informative prior of the overall probability (inverse logit α_0) of response or event.

Results

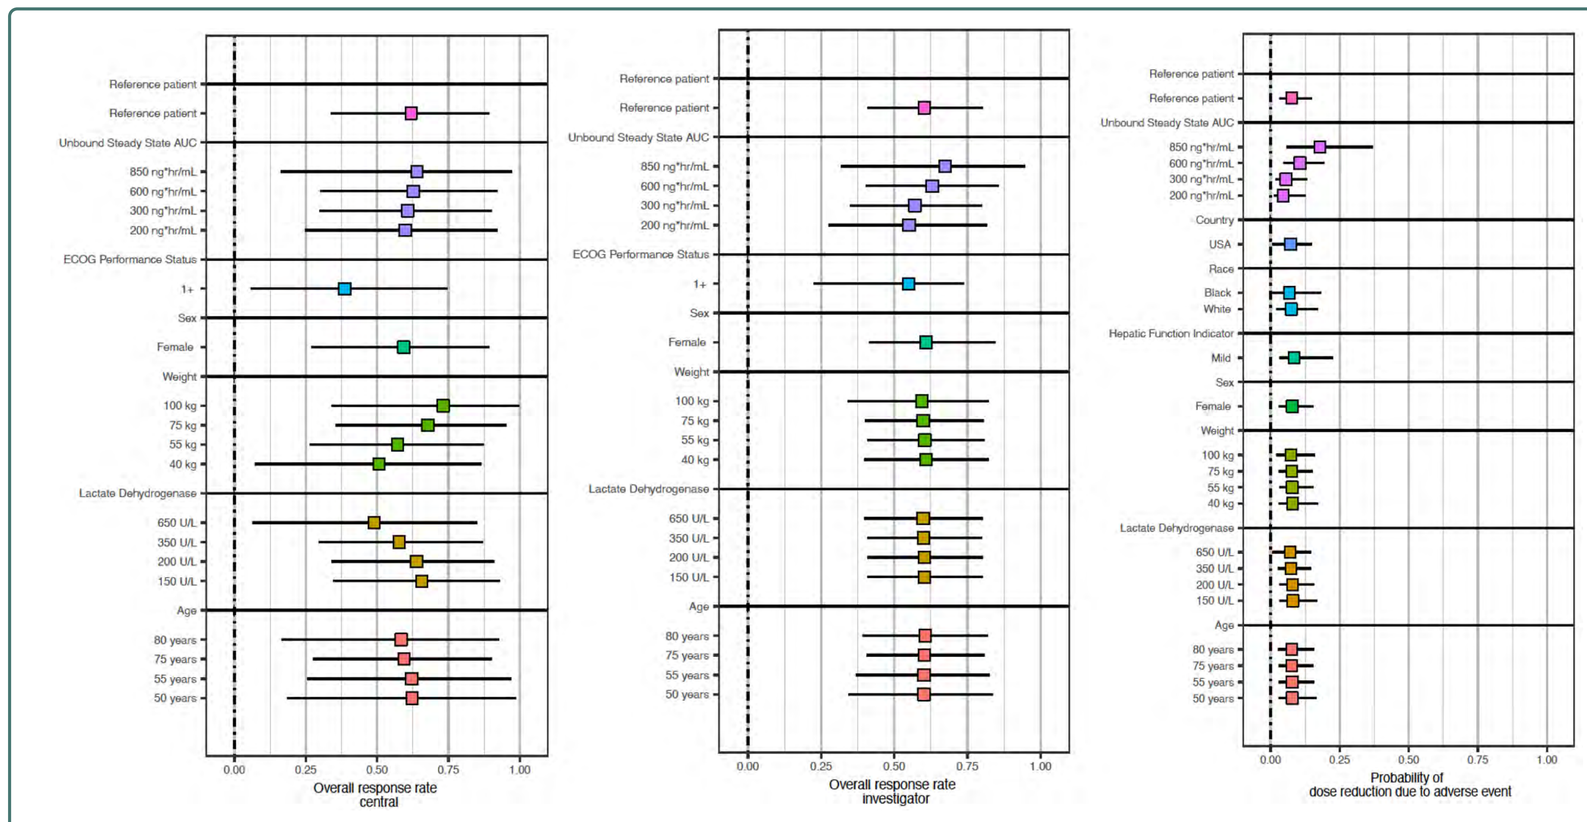

Figure 3. Conditional Predictions for Overall Response Rate - Central Assessment ($n = 25$), Overall Response Rate - Investigator Assessment ($n = 39$), and Dose Reduction Due to Adverse Event ($n = 102$). Point and interval are predictions at different covariate levels or percentiles of covariate for the reference patient.

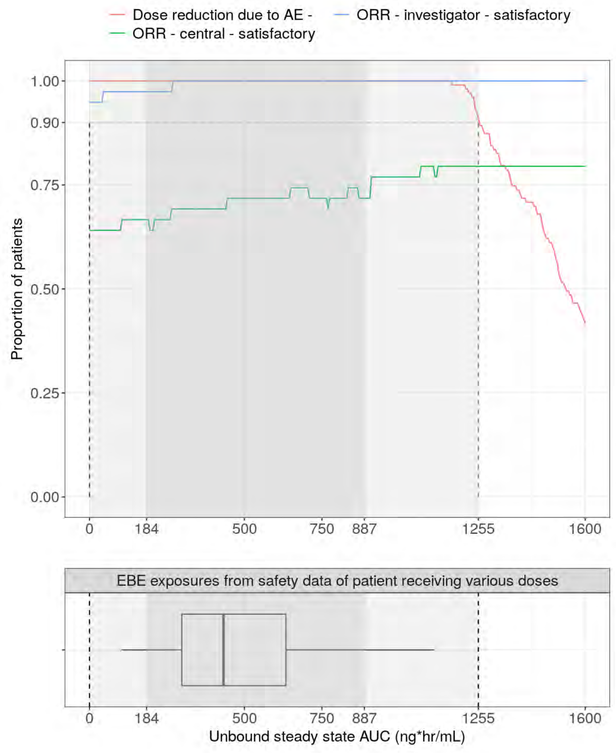

Figure 4. Estimated Region of Practical Equivalence (ROPE). A patient-specific probability of overall response of 30% or higher is proposed to define satisfactory efficacy, and a patient-specific probability of dose reduction due to adverse event of 50% or less is proposed to define acceptable safety. The vertical dashed lines represent the lowest exposure that is predicted to provide satisfactory efficacy for at least 50% of patients and the highest exposure that is predicted to provide acceptable safety for 90% of all patients. The light and dark gray areas indicate the ROPE and modified ROPE, respectively.

Spike and Slab Application

- Slab priors were set at N(0, 1.5), since this prior reflects a weakly informative prior of the probability of having an event and odds ratio of exposure.
- Spike priors were specified to correspond to a working definition of a negligibly small odds ratio (95% CRI = (0.82,1.21)).
- Spike component priors were not included for any covariates considered to have a very high prior plausibility of influencing the outcome variable (e.g., it was considered highly plausible that baseline neutrophils, baseline hemoglobin, and baseline platelets would affect the probability of a neutropenia adverse event, probability of anemia adverse event, and probability of a thrombocytopenia adverse event, respectively).
- Penalty parameter hyperparameters were varied by setting the prior expectation to the expected relative frequency of negligible covariates.
- For the Overall Response Rate - Central model, the posterior for λ departed only moderately from the Uniform prior (**Figure 3**), because the data were limited in distinguishing negligible effects from non-negligible effects.

Exposure - Response Analysis

- Slightly positive associations with exposure were observed for the efficacy endpoints but they were steeper for the safety endpoints.
- A steeper exposure-response relationship was estimated for Overall Response Rate - Investigator (Odds Ratio =1.22 per 250 ng*hr/mL, 95% CRI (0.63,2.87)) as compared to Overall Response Rate - Central (Odds Ratio = 1.08 per 250 ng*hr/mL 95% CRI (0.40,3.05)) (**Figure 3**).
- Rates of dose reduction due to adverse event were generally low (less than 10%) over most of the studied exposure range but began to increase steeply at higher exposures (**Figure 3**).

Dose - Optimization

- A ROPE of 0 to 1255 ng*hr/mL and modified ROPE with direct empirical support of 184 to 887 ng*hr/mL were estimated as target exposure ranges (**Figure 4**).
- Simulations at 50, 100, 150, and 200 mg once daily dosing suggest that 200 mg once daily dosing is most likely to achieve exposures within the modified ROPE for subpopulations of interest.

Conclusion

- Implementation of spike and slab priors allows for inclusion of all covariates and covariate effects to be estimated simultaneously through regularization of negligible effects.
- Hyperparameters can be specified in a general manner across multiple efficacy and safety endpoints.
- The efficacy and safety models can be used to simulate exposure regions of practical equivalence which achieve satisfactory efficacy and acceptable safety based on criteria specifing the probability of response or event and the minimum proportion of patient population achieving those probabilities.
- Analyses confirmed a positive exposure-response relationship for all endpoints and established a target exposure range (184 to 887 ng*hr/mL) that provides satisfactory efficacy and acceptable safety at the recommended dose (200 mg once daily) of valemetostat.

References

[1] George, E.I. and McCulloch, R.E. Variable Selection Via Gibbs Sampling. *J. Am. Stat. Assoc.* **88** (1993):881–889.

[2] Piironen, J. and Vehtari, A. Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electron. J. Stat.* **11** (2017):5018–5051.

[3] Fukae, M., Baron, K., Rogers, J., Garcia, R., Tachibana, M., Mondick, J. and Shimizu, T. Poster: Landmark and longitudinal exposure-response analyses for multiple efficacy and safety endpoints to justify the clinical dose of valemetostat for adult T-cell leukemia/lymphoma. In *American Conference on Pharmacokinetics* (2022).

