Bayesian parametric mixture survival models in immuno-oncology applications

MSR53

Leveraging control arm observations to model heterogeneous response in the experimental arm

Daniel J. Sharpe[1], Tuli De[2], Georgia Yates[1], Sangita M. Baxi[2], Llenalia García-Fernández[3], Jackie Vanderpuye-Orgle[2]

[1]Parexel, Uxbridge, London, United Kingdom; [2]Parexel, Billerica, Massachusetts, United States; [3]Parexel, Madrid, Spain

Background

- Heterogeneous response to an intervention is a common feature of survival patterns in studies of immuno-oncology (IO) therapies
- Parametric mixture models (PMMs), which represent the cohort as a combination of two latent subpopulations with distinct survival curves, provide a natural framework to analyze heterogeneous response
- However, it is challenging to infer the proportion of responders, and hence the responder survival pattern, since this fraction is typically small and cohort sizes are limited
  - consequently, frequentist PMMs are often subject to high uncertainty, artefacts, and high sensitivity[1]
- We propose a Bayesian PMM (B-PMM) framework that is designed to make inferences in PMMs more tractable. Here, instead of using a “cure” assumption[1], which may be unrealistic for advanced cancers, we invoke the following core assumption:
  - non-responder patients in the IO arm exhibit a survival pattern similar to that of control arm patients

REFERENCES

[1] Chen T.-T. *BMC Med Res Methodol* 2016; **16**(1): 12

[2] Gelman A. *Entropy* 2017; **19**(10): 555.

[3] Herbst RS et al. *J Thorac Oncol.* 2021; **16**(10): 1718-1732.

Methods

- We consider mixtures of exponential, Weibull, gamma, log-normal, and log-logistic distributions, which are among those commonly considered in health technology assessments
- To impose the key assumption, informative priors are chosen for parameters of the non-responder survival functions. This is achieved by fitting standard parametric models to the control arm data. The maximum-likelihood estimates (MLEs) and standard errors define the corresponding B-PMM hyperparameters
- Weakly informative priors[2] are selected for the parameters of the responder survival functions. We use diffuse log-normal priors with hyperparameters chosen to yield appreciable probability density within the clinically plausible range: i.e., roughly bounded by an approximate general population mortality and overall population survival in the IO arm [Fig 1]
- A vague (uniform) prior is specified for the proportion of responders
- B-PMMs were fitted to digitized survival data for pembrolizumab [PEMBRO] (versus docetaxel [CHEMO]) in a subgroup of patients with advanced non-small cell lung cancer from the 5-year database lock of the KEYNOTE-010 study[3]

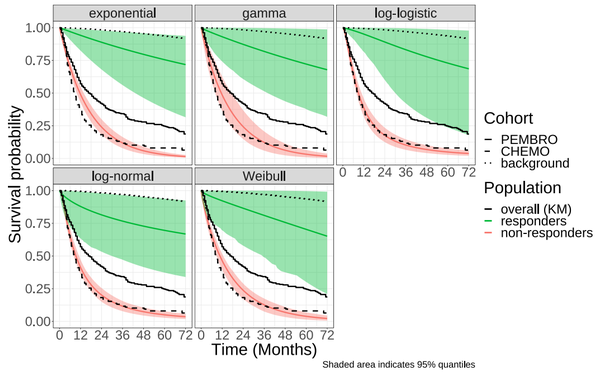

Figure 1: Kaplan-Meier [KM] survival functions for the KEYNOTE-010 trial and prior distributions for the B-PMM latency models

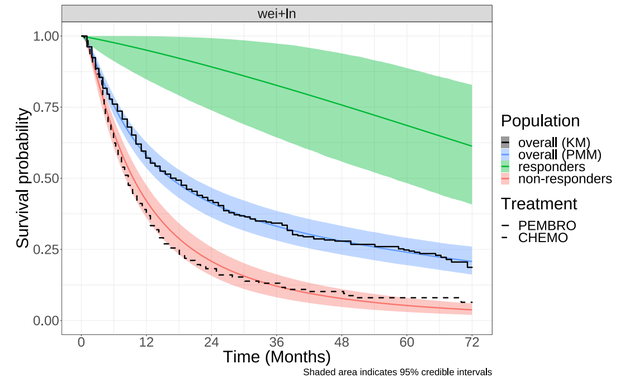

Figure 2: Predicted (B-PMM) and observed (Kaplan-Meier [KM]) survival probabilities for the KEYNOTE-010 trial

Results

- The range of posterior mean estimates for the proportion of responders across all 25 candidate models was 28.3-37.2%
  - this range is greatly reduced compared to the large inconsistency in MLEs from frequentist PMMs (range: 17.2-56.7%)
- The best-fitting B-PMM (based on Weibull and log-normal distributions for the responder and non-responder subpopulations, respectively) yields a plausible estimate for the proportion of responders (30.3% [95% credible interval: 19.0-44.0%]) [Figs 2 & 3]
  - in contrast, the best-fitting frequentist PMM yields very high uncertainty in the responder fraction estimate (37.6% [95% confidence interval: 14.5-68.2%])

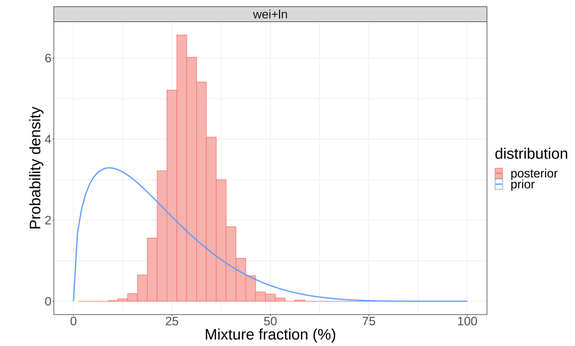

Figure 3: Prior and posterior probability distributions for the proportion of responders in the best-fitting B-PMM

Conclusions

- The proposed B-PMM approach overcomes the problem of limited observations in RCT data to reliably infer the proportion of responders from flexible survival models

Email: Daniel.Sharpe@parexel.com

www.parexel.com