Tailored Bayes: a risk modelling framework under unequal classification costs

Solon Karapanagiotis[1,2], Umberto Benedetto[3], Sach Mukherjee[4], Paul Kirk[1], Paul Newcombe[1]

[1] MRC Biostatistics Unit, University of Cambridge, [2] The Alan Turing Institute, [3] Bristol Heart Institute, University of Bristol, [4] German Center for Neurodegenerative Diseases (DZNE), Germany

Background

- Risk prediction models are widely used in many areas.
- The most common type of risk prediction model is based on binary outcomes (y), with class labels 0 (negative) and 1 (positive).
- Such models are often constructed using methodology which assumes the costs of different misclassification errors are equal.
- In many applications, especially in healthcare, this assumption is not valid.
- For instance, in a diagnostic setting, the cost of misdiagnosing a person with a life-threatening disease as healthy may be larger than the cost of misdiagnosing a healthy person as a patient.
- Here, we present Tailored Bayes (TB), a novel Bayesian inference framework which "tailors" model fitting to optimise predictive performance with respect to unbalanced misclassification costs.

Preliminaries

- The benefits/harms of correct/incorrect classifications can be summarised into a single number, the target threshold, defined as

$$t = \frac{U_{TN} - U_{FP}}{U_{TN} - U_{FP} + U_{TP} - U_{FN}} = \frac{H}{H+B} = \frac{1}{1+\frac{B}{H}} \quad (1)$$

- U are utility functions assigning a value to each of the four possible classification-outcome combinations stating how beneficial/costly each classification is.
- Eq. (1), tells us we only need to specify t, which corresponds to specific benefits/cost (B/H) ratio, without having to explicitly specify the relevant utilities, U.

Methods

- Data: $D = \{(y_i, \mathbf{x}_i) : i = 1, \dots, n\}$
- Objective: Learn $\boldsymbol{\beta}$, the regression coefficients.
- The TB log-likelihood/loss is

$$\log(L(D \mid \boldsymbol{\beta})) = \sum_{i=1}^{n} w_i l_i(D \mid \boldsymbol{\beta})$$

  - $w_i = \exp\{ -\lambda(p_u(\mathbf{x}_i) - t)^2 \}$
  - $l_i(D \mid \boldsymbol{\beta})$ is the standard logistic log-likelihood function
  - $p_u(\mathbf{x}_i) = p(y_i = 1|\mathbf{x}_i)$; needs to be estimated (data splitting)
  - $\lambda \geq 0$ is a tuning parameter (cross-validation)
- Setting $w_i := 1$ for all i, we recover the Standard Bayes (SB) solution.

Results: Simulated Data

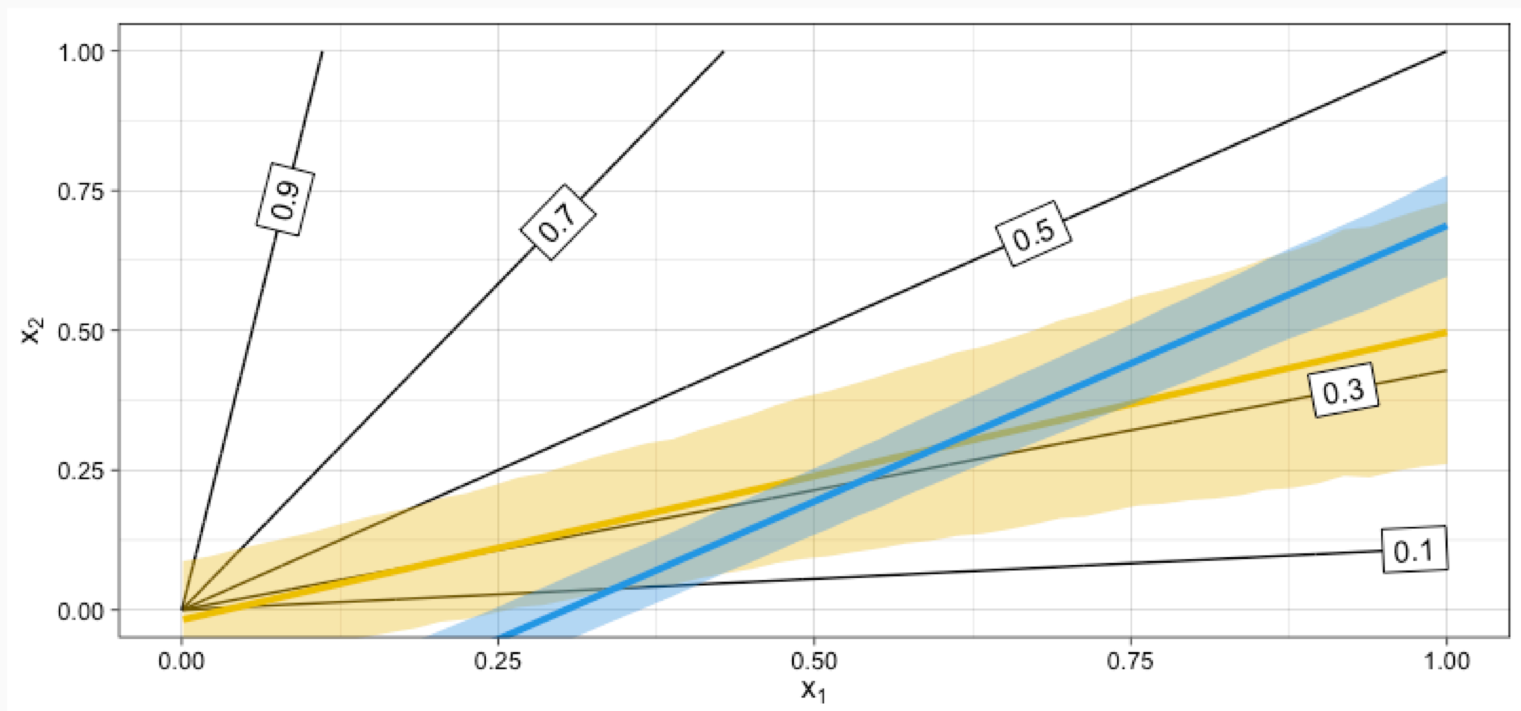

Fig. 1: Optimal decision boundaries (black lines) for $t = 0.1, 0.3, 0.5, 0.7, 0.9$. Posterior mean boundaries for SB (blue) and TB (yellow) when targeting the $t = 0.3$ boundary - corresponding to B/H=2.3:1 ratio (see eq. (1)). Shaded regions are 90% highest predictive density (HPD) regions. Data simulated from $p(y = 1|x_1, x_2) = \frac{x_2}{x_1+x_2}$, where $x_1, x_2 \sim \mathcal{U}(0, 1)$ and $n = 5000$.

Results: Cardiac Surgery Data

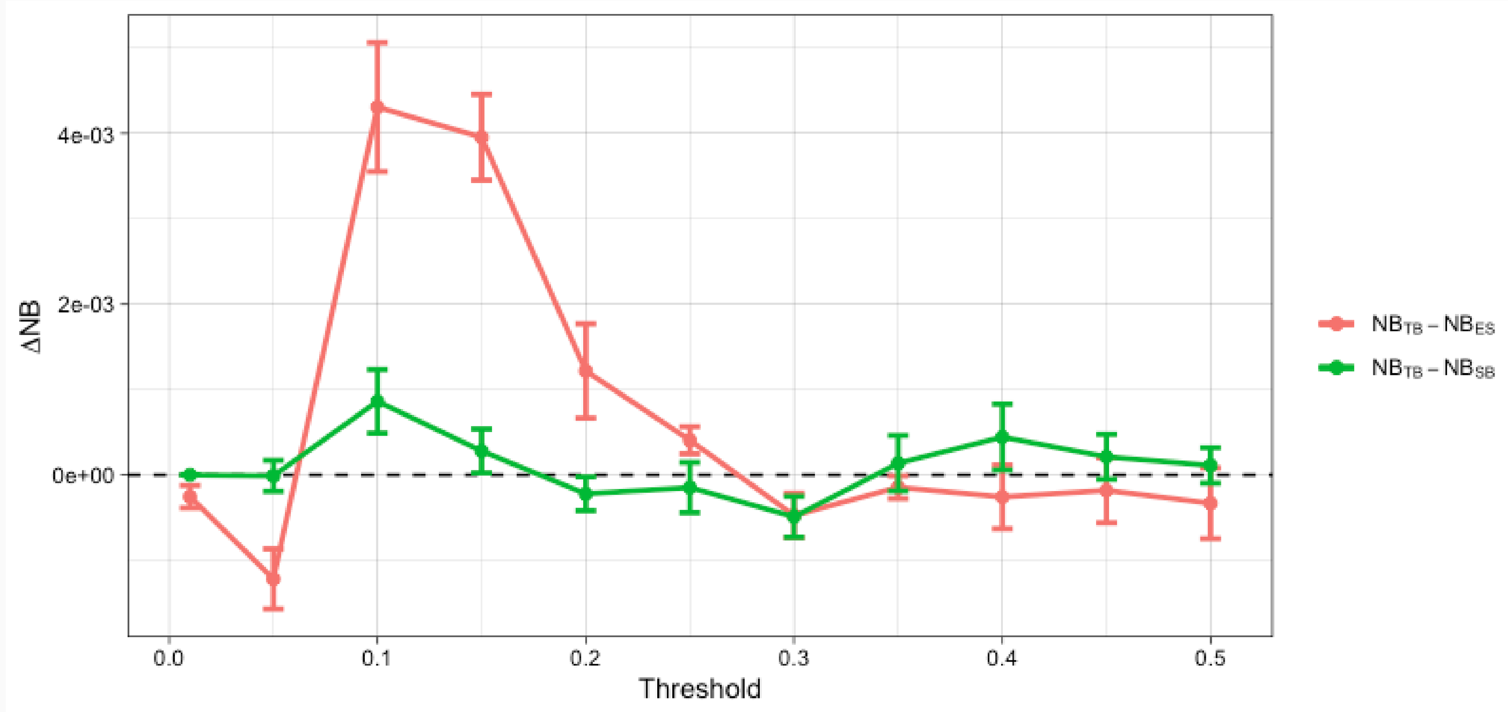

Fig. 2: Difference in net benefit (ΔNB) between TB and EuroSCORE (ES) (red), and between TB and SB (green). A positive difference means TB performs better. In clinical practice, treatement decisions are made under $t = 0.1$ or 0.2. $\text{NB} = \frac{TP}{n} - \frac{FP}{n}\frac{t}{1-t}$, where TP and FP are the true and false positives. This is a suitable model evaluation metric since it accounts for the different misclassification costs (through t). The units on the y axis may be interpreted as the difference in benefit associated with one patient who would die without treatment and who receives therapy.

Contributions

- The framework is flexible, easy-to-use, and widely applicable:
  - Robust to different choices of w_i, eg $w_i = \exp\{-h(p_u(\mathbf{x}_i), t)\}$ with h an arbitrary function.
  - Bayesian: hierarchical modelling and incorporation of external information.
- Generic:
  1. Implemented in any learning framework (not necessarily Bayesian).
  2. Not restricted to logistic likelihood/loss. The scheme can be used to adapt any likelihood/loss.
- R package: TailoredBayes

slides