eMethods: Statistical methods supplement

Adjusted AUC calculation

The AUC, which is defined as the area under the ROC curve, may be alternatively formulated as an estimate of the probability that a model correctly predicts a greater risk of mortality for a patient who actually died than for a patient who actually survived. Let A denote a randomly selected patient, and let B denote a randomly selected patient with the opposite outcome (i.e., if patient A died, patient B survived and vice versa). Then, let p_A and p_B be the predicted probabilities of in-hospital mortality for patient A and patient B, respectively. We say that the set of predictions $(A,\ B,\ p_A\ , p_B\ )$ are "correct" if the patient who died had a larger predicted probability of mortality, "tied" if the predicted probabilities are equal, and "incorrect" if the patient who survived had a larger predicted probability of mortality. Then, the AUC is the average of the function $\psi(A,\ B,\ p_A\ , p_B\ )$ defined below across all possible randomly selected pairs of predictions.

$$\psi(A,\ B,\ p_A\ , p_B\ ) = \begin{cases} 1 \ \ if \text{correct} \\ \frac{1}{2} \quad\ if \text{ tied} \\ 0 \ if \text{ incorrect} \end{cases}.$$

In the setting where the data includes patients from multiple hospitals and the mortality rate within each hospital is included in the model, the predicted probabilities for patients in a hospital with a higher mortality rate will already be larger than the predicted probabilities for patients in a hospital with a lower mortality rate. Therefore, there is built in discrimination between patients from different hospitals that has nothing to do with the clinical characteristics in the model. We only want to estimate the discrimination of the model based on the clinical characteristics included, without this artificial boost from including the hospital mortality rate, so we calculate an adjusted AUC instead, described below.

The key idea is that we modify the procedure described above so that we only compare pairs of patients (A, B) that are *in the same hospital*; as a result, any systematic differences between hospitals are irrelevant to the calculation of the adjusted AUC. More specifically, assume we have *m* hospitals and let N_i denote the number of patients who are in hospital *i* and *N* denote the total number of patients across all *m* hospitals. As before, let A denote a randomly selected patient. However, now let B denote a randomly selected patient *from the same hospital* as patient A where patient B has the opposite outcome. We keep all other notation the same as above.

We can now calculate $AUC^{(w)}$ as:

$$\begin{aligned} AUC^{(w)} &= \mathrm{E}[\psi(A,\ B,\ p_A\ , p_B\ )] \\ &= \textstyle\sum_{i=1}^{m} E[\psi(A,\ B,\ p_A\ , p_B\ )|\text{A in hospital i}]\mathrm{P}(\text{A in hospital } i) \\ &= \textstyle\sum_{i=1}^{m} E[\psi(A,\ B,\ p_A\ , p_B\ )|\text{A and B in hospital i}]\mathrm{P}(\text{A in hospital } i) \\ &= \textstyle\sum_{i=1}^{m} E[\psi(A,\ B,\ p_A\ , p_B\ )|\text{A and B in hospital } i]\,\frac{N_i}{N} \\ &= \textstyle\sum_{i=1}^{m} \frac{N_i}{N} AUC_i \end{aligned}$$

Where AUC_i denotes the within-hospital AUC for hospital i, i.e.,

$$\text{AUC}_i = E[\psi(A,\ B,\ p_A\ , p_B\ )|\text{A and B in hospital } i]$$

So, $\text{AUC}^{(w)}$ can be written as:

$$\text{AUC}^{(w)} = \sum_{i=1}^{m} w_i\, AUC_i$$

$$w_i = \frac{N_i}{N}.$$

Therefore, $\text{AUC}^{(w)}$ is a weighted average of the individual hospital AUC's with weights proportional to the hospital sample size.

Modeling for web application

We use the following technique to allow for approximate refitting of the model without sharing the proprietary dataset when we share the model in the web application. First, we fit the logistic regression model described on the mortality outcome. Next, we transformed the predictions using a logit transform. If you fit a linear regression model (OLS) on these transformed predictions, the coefficients of the OLS model are exactly the same as those in the original logistic regression model. The coefficients of an OLS model (β) can be calculated using matrix algebra with the design matrix (*X*) and the outcome vector (*y*):

$$\beta = (X'X)^{-1}(X'y)$$

We save the matrices *G = X'X and B = X'y* where *y* is the nx1 vector of logit transformed predictions from the logistic regression model and *X* is the design matrix.

Then, say that one wishes to estimate the coefficients for a model without creatinine. Let creatinine be the fourth column in *X* and let G^* be *G* with the 4th row and column removed and B^* be B with the 4th row removed. Then, we can estimate the coefficients β^* for a model without creatinine as:

$$\beta^* = (G^*)^{-1}(B^*)$$

Thus, we can estimate the coefficients of the OLS model with any subset of the variables that we included in our final model. This method allows us to still share a model that can be refitted and updated while maintaining data privacy since we save and share the *G* and *B* matrices rather than the raw data.

Odds ratio and the hospital mortality rate

Note that the estimated odds ratio of mortality in a hospital with mortality rate r, for a patient with other covariates X_1 as compared to a patient with other covariates X_2, where γ is the model coefficient for $r$and β are the coefficients for all other covariates is $e^{\gamma r+\ \beta\prime X_1}/e^{\gamma r+\beta\prime X_2} = e^{\beta\prime X_1-\beta\prime X_2}$. The hospital mortality rate is cancelled out in the odds ratio. Therefore, the estimated odds ratio between two patients in the same hospital can be calculated without knowing the hospital mortality rate.