Supplementary Material 1. Flow chart of the XGBoost algorithm

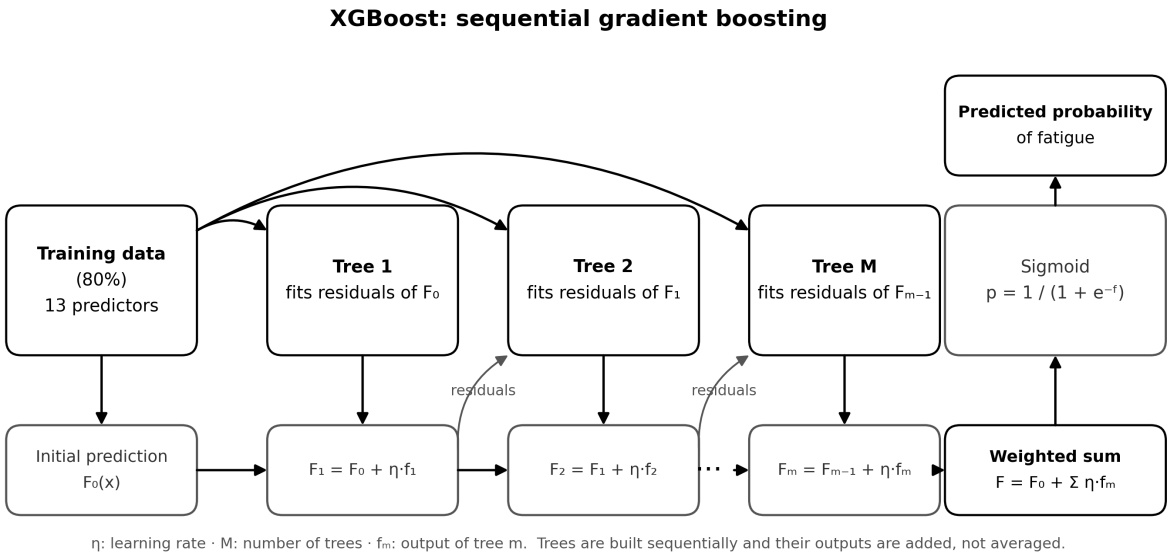

XGBoost builds decision trees sequentially. Starting from an initial prediction (the log-odds of fatigue prevalence in the training set), each new tree is fitted to the residuals (negative gradient of the loss) of the ensemble built so far, and its output, scaled by the learning rate η, is added to the running prediction. The final score is the weighted sum of all M trees, transformed by the logistic (sigmoid) function into a predicted probability of fatigue. This sequential, additive mechanism distinguishes gradient boosting from bagging methods such as Random Forest, in which independently grown trees are averaged.