

Table SI. Model-specific tuning grids and fixed hyperparameter settings.

Algorithm	Implementation	Tuning/fixed settings	Resampling/seed
Logistic regression	caret::train, method="glm"	Binomial logistic regression; no model-specific tuning parameter	Outcome-stratified 5-fold CV; common folds seed=520; model seed=52
SVM	caret::train, method="svmRadial"	sigma={0.01, 0.1}; C={0.1, 1}; selected by CV ROC	Outcome-stratified 5-fold CV; common folds seed=520; model seed=52
GBM	caret::train, method="gbm"	n.trees=100; interaction.depth=2; shrinkage=0.1; n.minobsinnode=10	Outcome-stratified 5-fold CV; common folds seed=520; model seed=52
Neural network	caret::train, method="nnet"	size=5; decay=0.01	Outcome-stratified 5-fold CV; common folds seed=520; model seed=52
Random forest	caret::train, method="rf"	mtry={2, 3}; selected by CV ROC	Outcome-stratified 5-fold CV; common folds seed=520; model seed=52
XGBoost	xgboost::xgb.train	nrounds=50; max_depth=3; eta=0.1; gamma=0; colsample_bytree=0.8; min_child_weight=1; subsample=0.8	Manual evaluation using the same 5 training folds (fold-generation seed=520); no explicit model-specific seed recorded in the final native call
KNN	caret::train, method="kknn"	kmax=5; distance=1; kernel="optimal"	Outcome-stratified 5-fold CV; common folds seed=520; model seed=52
AdaBoost	caret::train, method="AdaBoost.M1"	mfinal=20; maxdepth=2; coeflearn="Zhu"	Outcome-stratified 5-fold CV; common folds seed=520; model seed=52
LightGBM	lightgbm::lgb.train	learning_rate=0.1; num_leaves=7; max_depth=3; feature_fraction=0.8; bagging_fraction=0.8; bagging_freq=1; lambda_l1=0.1; lambda_l2=0.1; nrounds≤50; early_stopping_rounds=5	Predefined settings; held-out set used as the early-stopping monitor; no explicit model-specific seed recorded
CatBoost	catboost::catboost.train	iterations=200; learning_rate=0.05; depth=3; l2_leaf_reg=10; loss=Logloss; eval_metric=AUC	random_seed=123; fixed iterations; no held-out early stopping

The same six consensus predictors were used for all candidate models. For caret models, the outcome-stratified folds were fixed across algorithms to support comparable internal resampling. The common five-fold partitions were generated with seed=520. CatBoost used random_seed=123. No explicit model-specific random seed was recorded for the final native XGBoost or LightGBM implementation, and no seed was retrospectively assigned. The held-out set was inspected during candidate-model evaluation and was also used as the early-stopping monitor for LightGBM; therefore, the manuscript describes it as a held-out test set rather than an external or untouched validation cohort. CV, cross-validation; GBM, gradient boosting machine; KNN, k-nearest neighbors; ROC, receiver operating characteristic; SVM, support vector machine.