

Personalised dynamic super learning: an application in predicting hemodiafiltration's convection volumes

Arthur Chatton,^{1,*} Michèle Bally,^{1,2} Renée Lévesque,³ Ivana Malenica,⁴ Robert W. Platt^{5,6,7} and Mireille E. Schnitzer^{1,7,8}

¹Faculté de Pharmacie, Université de Montréal, Montréal, QC, Canada, ²Département de Pharmacie et Centre de Recherche, Centre Hospitalier de l'Université de Montréal, Montréal, QC, Canada, ³Service de Néphrologie, Département de Médecine, Université de Montréal, Montréal, QC, Canada, ⁴Department of Statistics, Harvard University, Cambridge, MA, USA, ⁵Department of Pediatrics, McGill University, Montréal, QC, Canada, ⁶Center for Clinical Epidemiology, Lady Davis Institute, Jewish General Hospital, Montréal, QC, Canada, ⁷Department of Epidemiology, Biostatistics, Occupational Health, McGill University, Montréal, QC, Canada and ⁸Département de Médecine Sociale et Préventive, Université de Montréal, Montréal, QC, Canada

*Correspondence: Arthur Chatton. Faculté de pharmacie, Université de Montréal. Pavillon Jean-Coutu, 2940, chemin de Polytechnique, Montréal, QC, H3T-1J4, Canada. email:arthur.chatton@umontreal.ca

FOR PUBLISHER ONLY Received on Date Month Year; revised on Date Month Year; accepted on Date Month Year

Abstract

Obtaining continuously updated predictions is a major challenge for personalised medicine. Leveraging combinations of parametric regressions and machine learning approaches, the personalised online super learner (POSL) can achieve such dynamic and personalised predictions. We adapt POSL to predict a repeated continuous outcome dynamically and propose a new way to validate such personalised or dynamic prediction models. We illustrate its performance by predicting the convection volume of patients undergoing hemodiafiltration. POSL outperformed its candidate learners with respect to median absolute error, calibration-in-the-large, discrimination, and net benefit. We finally discuss the choices and challenges underlying the use of POSL.

Key words: Cross-validation, Dynamic prediction, Machine learning, Panel data, Personalised medicine, Stacking.

1. Introduction

Dynamic prediction models provide predicted outcome values that can be updated over time for an individual as new measurements become available (Proust-Lima and Blanche, 2016; Jenkins et al., 2021). Such models are crucial for providing updated predictions, giving users the most useful information at the current time. For instance, Fournier et al. (2019) developed a dynamic prediction model incorporating repeated creatinine measurement to inform kidney transplant recipients of their prognosis and disease progression. Sabathé et al. (2023) proposed a dynamic prediction model for guiding the early decision to switch from first- to second-line therapy in multiple sclerosis. Outside of clinical fields, dynamic models have been used, for example, to predict geological sharp slope deformation (Li et al., 2021) and growth/decline of tourism demand in Australia (Athanasopoulos et al., 2009).

Previous approaches to prediction were mainly based on parametric models, but there is a current trend towards an algorithmic culture (Breiman, 2001a; de Hond et al., 2022) using more flexible machine learning (ML) approaches. Indeed, many ML approaches do not rely on parametric assumptions about the underlying data-generating process and have the potential to predict more accurately. Nevertheless, there is still no free-lunch method, even with ML. Their improved flexibility is at the cost of an increased risk of over-fitting, needing more data to converge, and the final performance still depends on the particular data-generating process (Hand, 2006).

Ensemble methods aim to overcome these issues by combining a diverse and rich set of algorithms or “learners” into one final prediction (Breiman, 1996; Naimi and Balzer, 2018). For instance, the super learner (SL, van der Laan et al., 2007) takes a weighted sum of the individual predictions from each incorporated learner (hereafter candidate learners) to produce a final prediction that is theoretically as accurate as the best-performing candidate learner (Dudoit and van der Laan, 2005). Moreover, including parametric or semi-parametric learners in addition to non-parametric ones allows us to incorporate some knowledge about the data-generating process

(Keil and Edwards, 2018). Recent applications of SL include seasonal influenza hospitalisation (Gantenberg et al., 2023), mortality in intensive care units (Pirracchio et al., 2015), and prison violence (Bačák and Kennedy, 2019). SL was recently expanded to joint outcome modelling (Tanner et al., 2021; Devaux et al., 2022) and time-series (Benkeser et al., 2018; Malenica et al., 2023). Here, we will focus on the personalised online super learner (POSL) from Malenica et al. (2023), applying this methodology to predict a continuous outcome in panel/longitudinal data dynamically.

In this paper, we adapt POSL to predict a repeated continuous outcome dynamically, propose a new way to validate the performance of dynamic or personalized prediction models (subsection 4.2), and finally discuss the choices and challenges underpinning the use of the POSL. The paper is organised as follows. We introduce the motivating application in the next section. The third section describes the POSL methodology in our longitudinal context. We describe the procedures for model assessment in the fourth section and report the results from the application in the fifth section following the TRIPOD statement (Moons et al., 2015). We discuss the choices and challenges in the sixth section before concluding in the last section.

2. Motivating application: prediction of the individual convection volume for hemodiafiltration

End-stage kidney disease is an important cause of morbidity and mortality worldwide, with a prevalence estimated at almost 5 million individuals by 2030 (Liyanage et al., 2015). The only cure is kidney transplantation, but there is a severe shortage of transplant kidneys. In 2021, approximately one-third of transplant candidates received a kidney transplant in Canada, while the others remained on a renal replacement therapy (Canadian Institute for Health Information, 2022).

Amongst the different renal replacement therapies, we focus on hemodiafiltration (HDF), which combines diffusive clearance and convective removal of solutes. The primary outcome of HDF is the convection volume, defined as the sum of the replacement volume and the intradialytic net weight loss achieved over the entire HDF session. Recent randomised trials and an individual patient data meta-analysis suggested that one should aim for a convection volume of at least 23 or 24L per session to reach the highest levels of efficacy and thus both survival and quality of life (Blankestijn et al., 2023; Lévesque et al., 2015; Peters et al., 2016). Achieving high convection volumes time after time and across patient populations requires clinical skills and readily accessible information and data. Nephrologists must continually re-assess multiple parameters refreshed with each HDF session and consider time-varying clinical status changes, which is daunting in busy dialysis centres.

Therefore, we illustrate the POSL through the development of a personalised dynamic prediction framework for people undergoing chronic HDF in the Centre Hospitalier de l'Université de Montréal, Québec, Canada, to predict the individual convection volume with prediction made prior to starting each HDF session. We used an open cohort of patients undergoing chronic HDF at the outpatient dialysis clinic or the Centre externe De Gaspé. We considered HDF as chronic if there were at least 28 consecutive sessions. Patients with chronic HDF were followed from March 1st 2017 to December 1st 2021. Patients were permanently censored if they died, underwent renal transplantation, or were transferred to another dialysis facility. Table S1 (available online) lists the study variables with their misclassification and measurement error risks. The session-specific predictors are the following: year, season, time since the last session, hemoglobin, albumin, central venous catheter changes, intercurrent hospitalisation, dalteparin dose, access by central venous catheter, and excess weight. The baseline predictors are time since the first chronic HDF session, age, gender, hypertension, diabetes, peripheral vascular disease, congestive heart failure, arrhythmia, acute myocardial infarction, chronic pulmonary disease, liver disease, valvular disease, cancer, metastatic solid tumour, cerebrovascular disease, dementia, peptic ulcer disease, hemiplegia or paraplegia, and rheumatic disease. We also considered some individual recent sessions' history as predictors: the outcomes (convection volumes) averaged over the three last months, the alteplase dose averaged over the last week, and the number of inversion lines averaged over the last week.

3. Personalised super learner

3.1. Super learning

To introduce the SL as described in van der Laan et al. (2007), we first consider a time-fixed setting. Let $\mathcal{O} = (\mathbf{X}, Y)$ denote the matrix of observed variables, where \mathbf{X} is the vector of baseline predictors, and Y is the outcome. Let also \mathbf{K} denote the set of C candidate learners, indexed by $c = 1, \dots, C$. A learner is a function taking as inputs \mathbf{X} to provide a prediction \hat{Y} . We use uppercase letters for random variables, lowercase letters for possible values of variables or constants, bold letters for vectors, and calligraphic font for matrices.

The SL takes three inputs: the observed dataset with n draws of \mathcal{O} , the user-specified set of candidate learners \mathbf{K} , and a loss function (or a performance measure) for quantifying the prediction error. The set \mathbf{K} should include different algorithms (*e.g.*, tree-based, regression-based, parametric, non-parametric). It can also include multiple versions of the same algorithm with different tuning parameters or different subsets of data (Phillips et al., 2023).

The SL involves fitting each candidate learner to the data using cross-validation (CV) to ultimately determine either the best learner or a weighted combination of learners. The first option, the discrete SL (dSL), selects the best-performing learner amongst the C candidate learners. The latter option is an ensemble SL (eSL), producing the best combination (described below) of the predictions of the C candidate learners.

First, CV proceeds as follows. The n individuals are randomly partitioned into V folds. Each fold defines a specific validation set, while the remaining observations define the related training set. Each of the C candidate learners is trained on each of the V training sets, and C predictions (one per learner) are obtained for each individual in the corresponding V validation sets, resulting in the $n \times C$

matrix $\hat{\mathcal{Y}} = \begin{pmatrix} \hat{\mathbf{y}}_{11} & \cdots & \hat{\mathbf{y}}_{1C} \\ \vdots & \ddots & \vdots \\ \hat{\mathbf{y}}_{v1} & \cdots & \hat{\mathbf{y}}_{vC} \end{pmatrix}$, where $\hat{\mathbf{y}}_{vc}$ is the vector of predictions using learner c for all individuals in the validation fold v . We thus define the meta-level dataset $\mathcal{M} = (\hat{\mathcal{Y}}, Y)$, which includes the matrix of outcome predictions and the corresponding observed outcomes.

Second, the risk with respect to the selected loss function is empirically optimised over \mathcal{M} . The dSL is identical to the candidate learner with the best cross-validated performance (*i.e.*, yielding the lowest estimated risk (lowest average loss) in \mathcal{M}). The eSL uses a meta-learner to obtain a vector of length C of optimal weight values (*e.g.*, coefficients from a constrained linear regression) α , one coefficient per candidate learner, which is determined by finding the values of α that optimise the expected value of the selected loss function over all individuals.

Third, the C learners are fit on the whole observed dataset \mathcal{O} to obtain C n -vectors of predictions.

Last, we obtain the final prediction from the fitted values obtained in the previous step. The dSL prediction is that of the candidate learner selected in the second step. The eSL prediction is the linear combination of the predicted outcomes from each candidate learner using the coefficient values estimated with the meta-learner. See Phillips et al. (2023, Figure 2) for a visual illustration.

3.2. Personalised online super learning for longitudinal prediction

POSL (Malenica et al., 2023) leverages two sources of information: the individual i for which we want to predict the outcome (hence *personalised*) dynamically in real-time (or *online*) and an ‘‘historical’’ pool of H individuals used for the prediction, indexed by h ($h = 1, \dots, H$). Suppose nephrologists want to predict the HDF convection volume for a given patient. This specific patient will be denoted by i . The H other patients previously dialysed constitute the historical pool. The task can be viewed as an external prediction in a cross-sectional setting, where i is the individual for whom the prediction is made, with a score developed using the H historical individuals in addition to patient i ’s past information.

We denote the matrix of the P time-varying predictors measured until time $t = 1, \dots, T$ by $\mathcal{L}_i = \begin{pmatrix} L_{i11} & \cdots & L_{i1P} \\ \vdots & \ddots & \vdots \\ L_{it1} & \cdots & L_{itP} \end{pmatrix}$. Owing to the outcome being repeatedly measured over time, we denote it by $\mathbf{Y}_i = (Y_{i,1}, \dots, Y_{i,T})$. Let $\mathcal{O}_i = (\mathbf{X}_i, \mathcal{L}_{i,t \leq \tau}, \mathbf{Y}_{i,t < \tau})$ denote the observed data for the unique individual i with which one wants to predict Y_τ , with t being the time of the observation and τ being the time of prediction.

Let $\{\mathcal{O}_h\}_{h=1}^H$, where $\mathcal{O}_h = (\mathbf{X}_h, \mathcal{L}_h, \mathbf{Y}_h)$, denote the observed data for the historical pool of individuals. Each individual h has a series of measurements index by $t = 1, \dots, T_h$, a time-varying predictor matrix of the same P predictors $\mathcal{L}_h = \begin{pmatrix} L_{h11} & \cdots & L_{h1P} \\ \vdots & \ddots & \vdots \\ L_{hT_h1} & \cdots & L_{hT_hP} \end{pmatrix}$, and a vector of historical outcomes $\mathbf{Y}_h = (Y_{h,1}, \dots, Y_{h,T_h})$.

Leveraging the longitudinal structure of the data, POSL predicts $Y_{i,\tau}$ using *both* the covariates measured at previous times ($t < \tau$; $t = 1, \dots, T$) from the individual i and the entire trajectory of the individuals $h = 1, \dots, H$ (even if $t \geq \tau$ since τ does not refer to calendar time).

Two sets of candidate learners can be defined depending on the set of individuals used for training. Let \mathbf{K}^{ind} and \mathbf{K}^{hist} be the sets of C^{ind} and C^{hist} candidate learners trained on the individual i or individuals $h = 1, \dots, H$, respectively. The historical learners \mathbf{K}^{hist} use more information and provide more data for predicting the earlier outcomes. In contrast, the individual learners \mathbf{K}^{ind} can be more accurate at later times since they are based on personal trajectory. \mathbf{K}^{ind} and \mathbf{K}^{hist} can include the same learners or different ones.

Contrary to the SL described previously, POSL takes four inputs: the observed datasets \mathcal{O}_i , the observed historical dataset $\{\mathcal{O}_h\}_{h=1}^H$, the user-specified set $\mathbf{K} = (\mathbf{K}^{ind}, \mathbf{K}^{hist})$ of C candidate learners ($C = C^{ind} + C^{hist}$), and a loss function for quantifying the prediction error. Note that $\{\mathcal{O}_h\}_{h=1}^H$ is not strictly required if \mathbf{K}^{hist} is already trained.

First, the historical learners \mathbf{K}^{hist} are trained on the historical cohort $\{\mathcal{O}_h\}_{h=1}^H$ if it has not yet been done. Second, the individual data \mathcal{O}_i are partitioned into V folds according to specific CV schemes described in the next subsection. The V training sets are used to train the C^{ind} individual learners only. Then, one obtains C predictions (one per learner, whether individual or historical) for the individual i in the corresponding V validation sets. All of the predictions from the V validation sets are merged to define the meta-level dataset $\mathcal{M} = (\hat{\mathcal{Y}}_i, \mathbf{Y}_{i,t < \tau})$, where $\hat{\mathcal{Y}}_i$ is the matrix of learner-and-fold-specific predictions as in the single time-point SL. Note that $\hat{\mathcal{Y}}_i$ includes predictions from both individual and historical learners such that $\hat{\mathcal{Y}}_i = (\hat{\mathcal{Y}}_i^{ind}, \hat{\mathcal{Y}}_i^{hist})$, where $\hat{\mathcal{Y}}_i^{ind}$ and $\hat{\mathcal{Y}}_i^{hist}$ denote the matrix of fold-specific predictions from \mathbf{K}^{ind} and \mathbf{K}^{hist} , respectively. Third, \mathcal{M} is used for optimising the estimated risk based on the selected loss function. The dSL is the candidate learner that minimises the risk. As in the single time-point SL, the eSL uses a meta-learner to obtain a vector of C weights $\alpha = (\alpha^{ind}, \alpha^{hist})$ that minimise the risk over \mathcal{M} . For instance, one can use as meta-learner a convex non-negative least squares (NNLS) regression formalised as follows: $\hat{\mathbf{Y}}_{i,t < \tau} = \alpha \hat{\mathcal{Y}}_i$, such that $\sum \alpha = 1$, $\alpha \geq 0$, and the statistical unit is the person-time. Note that the non-negative coefficients are estimated first and then scaled with respect to the convex constraint. The related loss function is the cumulative squared error defined as $\sum_{v=1}^V (Y_{i,v} - \hat{y}_{i,v})^2$, where $Y_{i,v}$ and $\hat{y}_{i,v}$ are respectively the observed outcome and the prediction obtained (*i.e.* $\hat{y}_{i,v}$ is the v th element of $\hat{\mathbf{Y}}_{i,t < \tau}$) in the validation fold v . The meta-learner and the loss function can be weighted to give more importance to recent times over older ones. For instance, Malenica et al. (2023) suggested setting time-specific weights ω to 1 for the most recent time and the others to $(1 - \delta)^{\tau - t}$, where δ is a user-defined constant representing how quickly past times should be down-weighted. Fourth, the candidate learners \mathbf{K} are fit on the trajectory of the individual i up to time τ to obtain C predictions. Last, we obtain the final prediction from the fitted values obtained in the previous step. The dSL prediction is that of the candidate learner minimising the loss computed in step 2. The eSL prediction is the prediction from the meta-learner taking the fitted values (fourth step) as inputs. Therefore, the prediction evolves over time and will be based on the - possibly different - best learners at each time.

3.3. Rolling cross-validation

Owing to the temporal dependency in \mathcal{O}_i , the CV procedure cannot randomly assign person-time observations to either the training or validation set. One must avoid using future observations when training individual learners because these data do not yet exist to make the prediction. Therefore, instead of using a V -fold CV scheme, one must use a hold-out CV scheme (Tashman, 2000) where a subset of the timeline of data is reserved for validating the candidate learners (Figure 1).

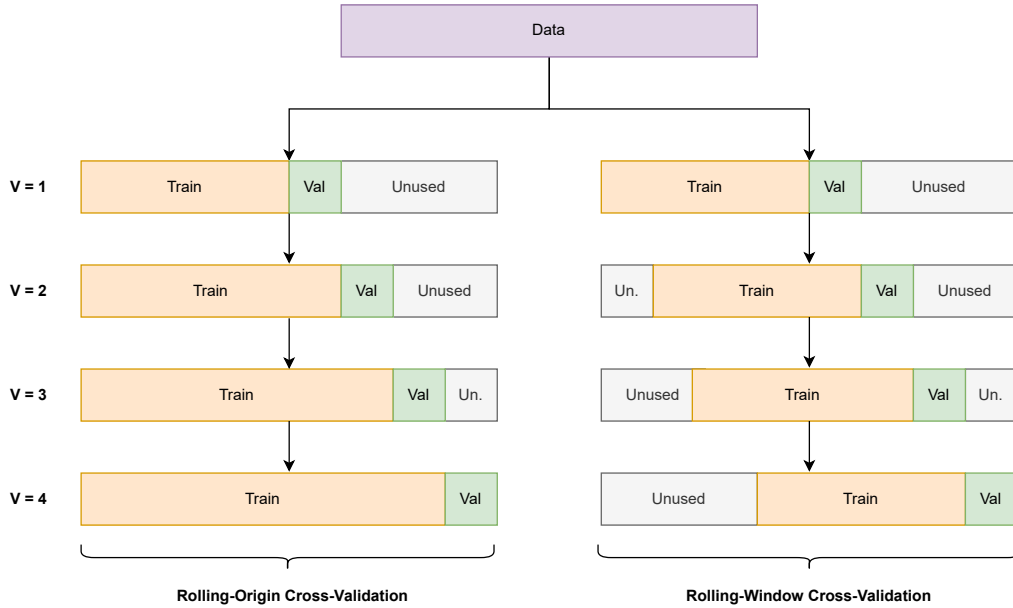


Fig. 1: Rolling cross-validation schemes across four folds v . Abbreviations: Un., Unused; and Val, Validation.

The global idea is to start with a small subset of data as the first training fold and validate in the following time (first validation fold). The validation fold is then iteratively added to the training fold to define the next training fold, and the subsequent time becomes the new validation fold. We refer to this scheme as Rolling-Origin CV (ROC) according to several authors (*e.g.*, Malenica et al., 2023; Hyndman and Athanasopoulos, 2018), though this term has been used to describe different CV schemes (*e.g.*, Tashman, 2000). Alternatively, one can prune the oldest time at each update for “cleaning out” patterns from one fold to the next, obtaining a so-called rolling-window CV (RWC). See Swanson and White (1997) for a discussion on the advantages and pitfalls of RWC.

4. Development and validation of the POSL for predicting HDF convection volumes

4.1. Model development

Since the strength of the SL is based on the diversity of the set of candidate learners (Phillips et al., 2023), we used the following set of algorithms as both individual and historical learners: an unpenalised linear model with main terms only, ridge linear model, lasso linear model, multivariate adaptive regression splines (MARS), and extreme gradient boosting trees (XGBoost). We augmented this set of candidate learners with these same learners incorporating a preliminary step of predictor screening with a random forest importance measure (Breiman, 2001b), taking only the top 5 predictors at time t . In addition, we considered the arithmetic mean of the outcome as a candidate learner to further reduce the risk of overfitting (Balzer and Westling, 2021). We refer the readers to James et al. (2021) for an introduction to these approaches. Software and related references are presented in Table S2. The individual learners are trained through ROCV and RWC (with a window size of 10 sessions), doubling the number of individual learners. Note that a fold is disregarded when at least one candidate learner did not converge (Malenica et al., 2023).

We selected a squared error loss function (or its cumulative sum for the dSL), leading to an NNLS meta-learner for the eSL (Tanner et al., 2021). Two meta-learners were considered (yielding two different eSLs): a convex and a non-convex NNLS. Convex NNLS is recommended for SL, but Malenica et al. (2023) reported better results with non-convex NNLS, likely because the convex combination was enforced after estimating the coefficients rather than as part of the estimation process itself. The meta-learners were weighted according to the time as described in the previous section. Specifically, the weight ω_t assigned to the squared loss measured at time t , was defined as

$$\omega_t = \begin{cases} 1 & \text{if } t \geq \tau - 5 \\ (1 - 0.1)^{\tau - t} & \text{otherwise.} \end{cases}$$

We truncated the predictions within $[0, 50]$ to respect clinically plausible values.

4.2. Model validation

We split the whole sample into tuning and working samples. The former was used to define, through a 10-fold CV, the learners' hyperparameters used in the latter. We applied the model development procedure in the working sample for each individual (but one at a time). Therefore, \mathcal{O}_i corresponds to one individual and $\{\mathcal{O}_n\}$ to the remaining individuals. We summarise the sample-splitting process in Figure 2).

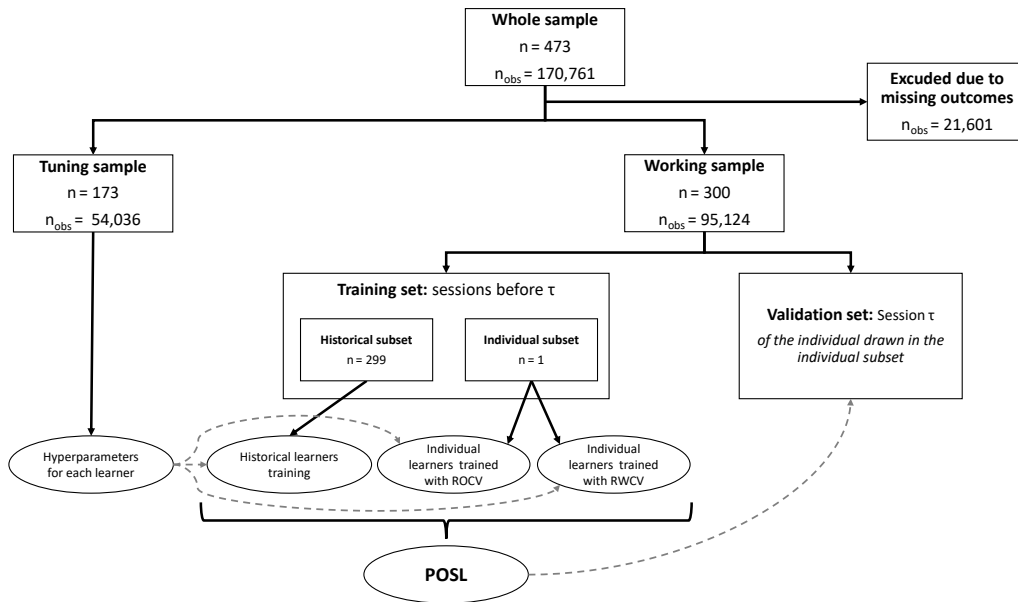


Fig. 2: Flow-chart of the tuning-training-validation process.

Abbreviations: n , Number of individuals; n_{obs} , Number of sessions; ROCV, Rolling-Origin Cross-Validation; RWCV, Rolling-Window Cross-Validation; POSL, Personalised Online Super Learner; and τ , session for training POSL.

To validate the prediction made by the POSL, we used a nested validation scheme called forward validation (Hjorth, 1982). Assuming we are interested in predicting the session's outcome $\tau + 1$, the sessions from 1 to τ are used as a training set for training the POSL with its own loop of ROCV or RWCV. Then the session $\tau + 1$ is the validation set in which we predicted the outcome (Figure 3). Predictions were made every session from the 12th ($\tau = 11$ gives ten sessions for training the POSL). Importantly, no session appears simultaneously in the training and validation sets, limiting over-fitting.

4.3. Model performance

For each individual i , we computed two accuracy measures (Hyndman and Athanasopoulos, 2018): the median absolute error (MdAE) and the mean squared error (MSE) defined, respectively, as $MdAE_i = \text{median}_{t=1}^{T_i} (|y_{ti} - \hat{y}_{ti}|)$ and $MSE_i = T_i^{-1} \sum_{t=1}^{T_i} (y_{ti} - \hat{y}_{ti})^2$. Note that the MSE is sometimes called Brier Score in predictive modelling. We assessed individual-specific calibration through the calibration intercept, the calibration slope and a flexible calibration curve (Van Calster et al., 2019). We used a generalised additive model with integrated smoothness estimation for obtaining the flexible curve (Wood et al., 2016). We assessed the discrimination by first dichotomising the predictions according to a clinically relevant threshold according to the literature (*i.e.*, $\geq 24L$ vs $< 24L$). Then, we computed the area under the ROC curve (AUROC) for each individual. Furthermore, we compared the performances of POSL to those of the candidate learners.

In addition, we evaluated the evolution over time of the MdAE, calibration slope, and calibration intercept averaged (using the median to circumvent extreme predictions) over all individuals for POSL, the individual linear model, and the historical XGBoost. Finally, we evaluated the clinical utility by computing the net benefit, defined by net benefit = true positives rate - false positives rate $\times p / (1 - p)$, where p is the binarised outcome prevalence (Peirce, 1884), at a range of binarisation thresholds.

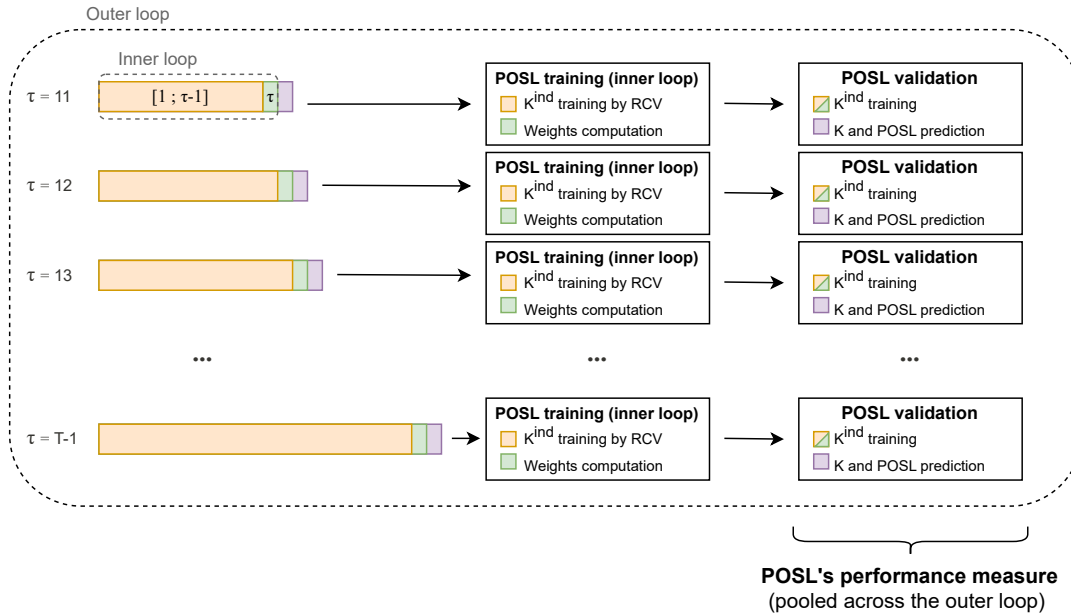


Fig. 3: POSL's forward cross-validation.

Abbreviations: K , Candidate learners set; K^{ind} , Candidate individual learners set; RCV, Rolling cross-validation (either origin or window); POSL, Personalised Online Super Learner; T , last session; and τ , session for training POSL. Adapted from Li et al. (2021)

4.4. Sample size

Because there is no current baseline performance data, the published literature is not informative, and no minimal sample size formula was available using a machine learning approach, we did not calculate a minimal sample size. For this study, we used convenience samples of all patients treated with chronic HDF.

4.5. Missing data

Missing predictors were imputed with the median or the mode of the individuals' previous sessions for continuous and discrete predictors, respectively; avoiding temporal data leakage (Roberts et al., 2017). Missingness indicators were added in the predictors' set as the pattern of predictor missingness might be informative for predicting the outcome (Sperrin et al., 2020). We excluded sessions with a missing outcome (*i.e.*, either cancelled or incomplete). Table S3 presents the characteristics of the included and excluded sessions.

4.6. Sensitivity analysis

We performed a sensitivity analysis with a binary outcome (at least 24L was successful; otherwise unsuccessful) where missing outcomes were imputed as unsuccessful. The procedure was almost identical, except that we evaluated the clinical utility through a decision curve analysis (Vickers and Elkin, 2006), and the loss function for the dSL was the negative log-likelihood to favour calibration and accuracy against discrimination. Furthermore, the negative log-likelihood loss function may even achieve better AUROC than an AUROC-maximising loss function, especially in smaller sample sizes (LeDell et al., 2016).

5. Results

5.1. Sample description

The study sample comprised 149,160 observations (*i.e.*, sessions) for 473 individuals with, on average, 315 sessions per individual (min: 1, max: 776). The convection volume ranged from 8L to 40L, with a mean of 27.4L. Seventy-two percent of the sessions presented a convection volume greater or equal to 24L. Characteristics of the tuning and working samples are described in Table 1. Note that the sample presented only a few missing data, with 0.8% of subjects missing excess weight and 0.6% missing albumin. Recall that three POSL implementations were considered: (a) a convex NNLS eSL, (b) a non-convex NNLS eSL, and (c) a dSL. The POSLs did not converge for two individuals in the working sample (0.67%) due to an insufficient number of recorded sessions (*i.e.*, less than twelve). Among the 2,295,825 predictions, 1583 (0.07%) were trimmed due to values outside of the bounds. No historical learners led to such predictions. Individual MARS (regardless of random forest screening) predicted extreme values the most often (0.50%), followed by the individual linear model (0.22%) and the individual lasso (0.15%). RF screening reduced their percentage of extreme predictions to 0.054% and 0.049, respectively. Among the POSLs, the dSL estimated less often extreme predictions than the two eSLs (0.04% vs 0.06%).

Table 1. Description of the tuning and working samples.

	Overall N=149,160	Tuning N=95,124	Working N=54,036	SMD	Missing (%)
Time between HDF sessions, days (mean (sd))	2.45 (5.39)	2.46 (5.39)	2.45 (5.39)	0.003	0.0
Time in the RTWL, days (mean (sd))	583.07 (399.26)	581.25 (398.40)	586.27 (400.75)	0.013	0.0
Hemoglobin, g/L (mean (sd))	108.21 (13.35)	107.92 (12.93)	108.71 (14.05)	0.058	0.0
Albumin, g/L (mean (sd))	37.20 (3.86)	37.17 (3.88)	37.25 (3.82)	0.021	0.6
Dalteparin dose, IU (mean (sd))	5,067.45 (2,484.22)	4,946.77 (2,429.68)	5,279.89 (2,563.69)	0.133	0.0
Age at baseline, years (mean (sd))	67.17 (13.92)	67.29 (13.86)	66.96 (14.03)	0.024	0.0
Excess weight, kg (mean (sd))	1.74 (1.29)	1.72 (1.28)	1.79 (1.32)	0.060	0.8
Convection volume history [†] , L (mean (sd))	27.36 (2.91)	27.40 (2.84)	27.28 (3.03)	0.040	0.0
Inversion of HDF lines history [†] (mean (sd))	0.02 (0.09)	0.02 (0.09)	0.02 (0.09)	0.003	0.0
Alteplase doses history [†] (mean (sd))	0.05 (0.26)	0.05 (0.25)	0.06 (0.27)	0.038	0.0
Year (N, (%))				0.057	0.0
2017	25,827 (17.3)	16,334 (17.2)	9,493 (17.6)		
2018	39,870 (26.7)	24,657 (25.9)	15,213 (28.2)		
2019	41,597 (27.9)	26,941 (28.3)	14,656 (27.1)		
2020	36,196 (24.3)	23,497 (24.7)	12,699 (23.5)		
2021	5,670 (3.8)	3,695 (3.9)	1,975 (3.7)		
CVC change since last session (N, (%))	388 (0.3)	250 (0.3)	138 (0.3)	0.001	0.0
Hospitalisation (N, (%))	8,013 (5.4)	5,662 (6.0)	2,351 (4.4)	0.072	0.0
Access by CVC (N, (%))	70,851 (47.5)	44,351 (46.6)	26,500 (49.0)	0.048	0.0
Male (N, (%))	94,723 (63.5)	60,909 (64.0)	33,814 (62.6)	0.030	0.0
Hypertension (N, (%))	90,435 (60.6)	57,891 (60.9)	32,544 (60.2)	0.013	0.0
Diabetes (N, (%))	59,296 (39.8)	35,304 (37.1)	23,992 (44.4)	0.149	0.0
Peripheral vascular disease (N, (%))	32,449 (21.8)	20,597 (21.7)	11,852 (21.9)	0.007	0.0
Congestive heart failure (N, (%))	28,583 (19.2)	19,963 (21.0)	8,620 (16.0)	0.130	0.0
Cardiac arrhythmia (N, (%))	28,962 (19.4)	18,753 (19.7)	10,209 (18.9)	0.021	0.0
Acute myocardial infarction (N, (%))	22,214 (14.9)	15,766 (16.6)	6,448 (11.9)	0.133	0.0
Chronic pulmonary disease (N, (%))	20,563 (13.8)	14,126 (14.9)	6,437 (11.9)	0.086	0.0
Liver disease (N, (%))	17,501 (11.7)	11,783 (12.4)	5,718 (10.6)	0.057	0.0
Valvular disease (N, (%))	15,130 (10.1)	10,657 (11.2)	4,473 (8.3)	0.099	0.0
Cancer (N, (%))	12,369 (8.3)	6,788 (7.1)	5,581 (10.3)	0.113	0.0
Metastatic solid tumour (N, (%))	1,943 (1.3)	1,054 (1.1)	889 (1.6)	0.046	0.0
Cerebrovascular disease (N, (%))	11,904 (8.0)	7,361 (7.7)	4,543 (8.4)	0.025	0.0
Dementia (N, (%))	4,931 (3.3)	3,485 (3.7)	1,446 (2.7)	0.056	0.0
Peptic ulcer disease (N, (%))	4,383 (2.9)	2,041 (2.1)	2,342 (4.3)	0.124	0.0
Hemi- or paraplegia (N, (%))	4,470 (3.0)	1,830 (1.9)	2,640 (4.9)	0.164	0.0
Rheumatic disease (N, (%))	3,268 (2.2)	2,744 (2.9)	524 (1.0)	0.140	0.0
Week-end HDF session (N, (%))	24,959 (16.7)	16,009 (16.8)	8,950 (16.6)	0.007	0.0
Season (N, (%))				0.014	0.0
fall	37,047 (24.8)	23,481 (24.7)	13,566 (25.1)		
spring	38,388 (25.7)	24,501 (25.8)	13,887 (25.7)		
summer	36,513 (24.5)	23,468 (24.7)	13,045 (24.1)		
winter	37,212 (24.9)	23,674 (24.9)	13,538 (25.1)		

[†] Averaged over three months (convection volume) or one week (inversion lines and alteplase)

Abbreviations: CVC, Central venous catheter; HDF, Hemodiafiltration; IU, International unit; N, Number of sessions; RTWL, Renal transplant waiting list; SMD, Standardised mean difference.

5.2. Candidate learner selection

A nearly identical distribution of candidate learners was found between the convex and non-convex eSLs. Globally, the historical learners – especially XGBoost – were more important in the eSLs and the dSL than individual learners (Figure 4, panels A-C). Individual learners trained through ROCV were more present in the meta-learners than those trained through RWCV. The random forest screening step did not seem to improve performance since most of these learners were rarely present in the meta-learners, a notable exception being the historical XGBoost. The historical and individual lasso were seldom selected, consistent with the lack of improved performance observed in previous simulation studies (Van Calster et al., 2020; Riley et al., 2021). Furthermore, the eSLs included, on average, three individual and one historical learners with a constant trend over time. However, their relative importance fluctuated over time. Historical learners

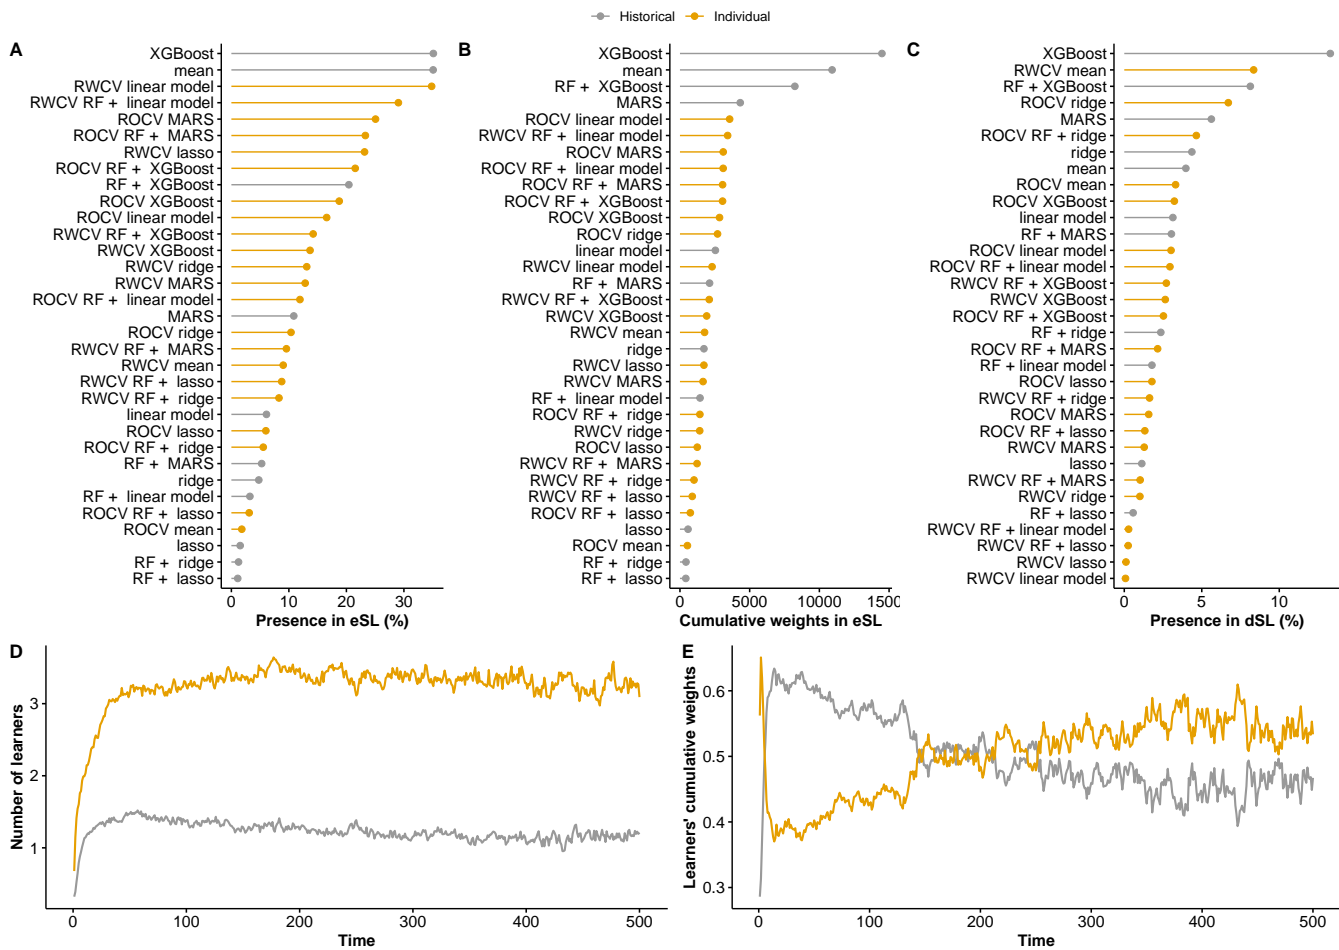


Fig. 4: Candidate learner distributions in the non-convex eSL and the dSL. (A) Percentage non-null weights in the eSL, (B) Sum of weights across time and individuals, (C) Percentage selection by dSL, (D) Number of learners composing the non-convex eSL over time, and (E) Sum of weights over time averaged between individuals. The X-axis of panels D and E is bound to 500 due to an insufficient number of observations after, and time 1 corresponds to the first predicted session.

Abbreviations: dSL, discrete super learner; eSL, ensemble super learner; MARS, multivariate adaptive splines regression; RF, random forest screening before training; ROCV, rolling-origin cross-validation; RWCV, rolling-window cross-validation; and XGBoost, extreme gradient boosting tree.

had more influence in the beginning until, approximately after the 200th session, the individual learners dominated (Figure 4, panels D-E).

5.3. Performance of POSL against its candidate learners

The accuracy of these three POSLs was similar (Figure 5), with a slight advantage for the non-convex eSL in line with previous results of Malenica et al. (2023). The calibration-in-the-large was near-perfect for all POSLs, but the convex eSL presented a higher variability. In contrast, the mean calibration was underoptimal for all POSLs, albeit convex eSL seemed slightly better. Historical learners almost always outperformed individual learners, especially in terms of mean calibration. Individual and historical learners were well-calibrated-in-the-large, but individual learners had a lower variance (except XGBoost). Individual MARS achieved poorer accuracy, while ridge regression yielded the best performance amongst individual learners. Additional numerical results are presented in Table S4. All historical learners were calibrated-in-the-small, while individual learners and POSLs underestimated the high convection volume and overestimated the low convection volume (Figure S1 available online). Individual XGBoost yielded poorer calibration-in-the-small, especially for the lowest outcomes. The eSLs achieved the highest discrimination with an AUROC of 0.825 (95%CI from 0.821 to 0.829) and 0.821 (95%CI from 0.817 to 0.825) for the non-convex and the convex eSLs, respectively (Figure S2). Individual learners achieved greater AUROCs than historical learners and dSL, except for MARS and XGBoost. Random forest screening did not seem to improve any performance measures significantly. Finally, non-convex eSL achieved the best net benefit around the threshold previously identified in the literature (Blankestijn et al., 2023; Lévesque et al., 2015), while the historical XGBoost achieved a lower net benefit than POSLs and the individual linear model regardless of the threshold (Figure S3).

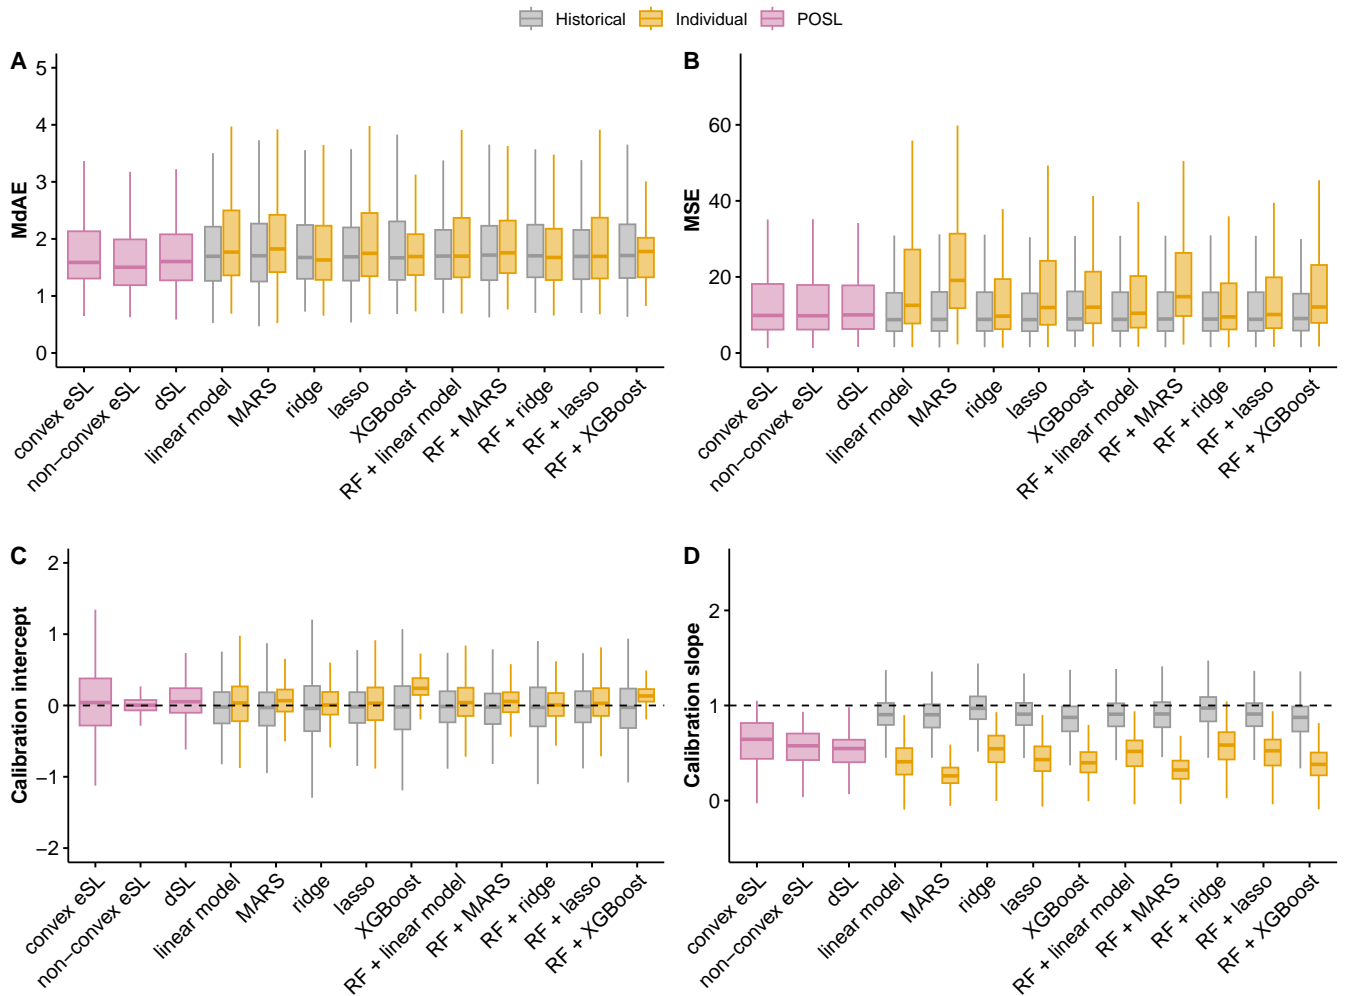


Fig. 5: Time-pooled performance measures of the different POSL implementations and their candidate learners. (A) Accuracy assessed by the MdAE, (B) Accuracy assessed by the MSE, (C) Calibration-in-the-large (or mean calibration) assessed by the calibration intercept, and (D) Weak calibration assessed by the calibration slope. The dashed line represents the median of the best-performing method for accuracy measures or the ideal value for calibration measures.

Abbreviations: dSL, discrete super learner; eSL, ensemble super learner; MARS, multivariate adaptive regression splines; MdAE, median absolute error; MSE, mean square error; POSL, Personalised Online Super Learner; RF, random forest screening before training; and XGBoost, extreme gradient boosting tree.

When looking at the evolution of the performances over time (Figure 6), the individual linear model achieved poorer accuracy in the early time periods (approximately until the 150th session). Interestingly, this time corresponds to the beginning of the equilibrium period between historical and individual learners in the eSL (Figure 4E). In contrast, POSLs consistently outperformed the candidate learners regarding MdAE, the non-convex eSL achieving the best MdAE over all time points. Historical XGBoost presented the lowest MSE until the hundredth session, from which it became equivalent to the two eSLs. Similarly, the historical XGBoost achieved the best calibration-in-the-large at the beginning before being slightly overtaken by the non-convex eSL. Surprisingly the convex eSL presented the reverse pattern. For the weak calibration, historical XGBoost outperformed the other approaches until the last times, followed by the convex eSL. The individual linear model caught up with the dSL and the non-convex eSL from approximately the 250th session, which corresponds to the time when the individual learners overtook the historical ones in the eSLs (Figure 4E). POSLs were highly miscalibrated-in-the-small at the first session, but the performance improved over time with an acceptable calibration-in-the-small achieved from the 100th session for the eSLs (Figure 6E).

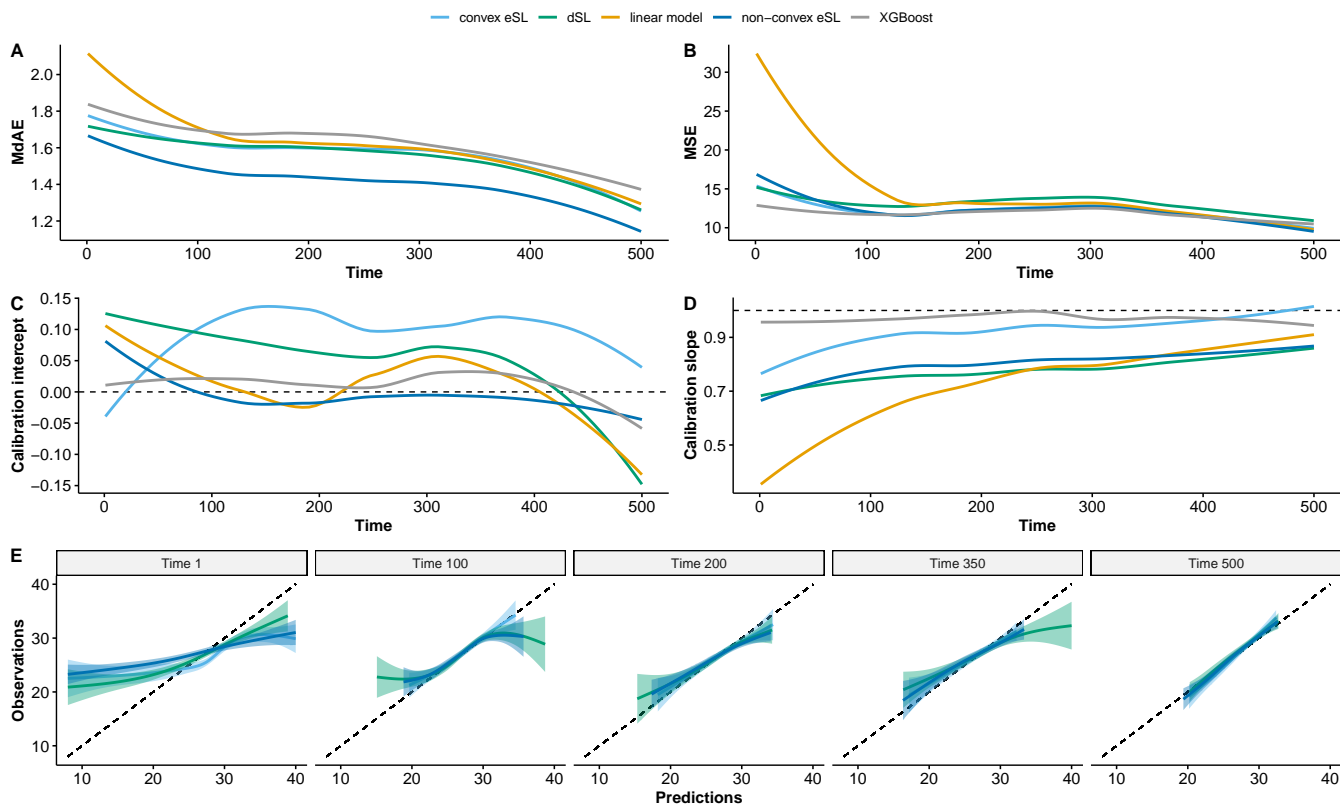


Fig. 6: Performance measures of the different POSL implementations, the individual linear model, and the historical XGBoost. (A) Accuracy assessed by the MdAE, (B) Accuracy assessed by the MSE, (C) Calibration-in-the-large (or mean calibration) assessed by the calibration intercept, and (D) Weak calibration assessed by the calibration slope. (E) Calibration-in-the-small (or moderate calibration) of the POSLs at five time points. The dashed line represents an ideal calibration. X-axes of panels 1-D are bound to 500 due to insufficient observations after. Curves of panels A-D were loess-smoothed. Time 1 corresponds to the first predicted session. Flexible curves of panel E were obtained through generalised additive models with integrated smoothness estimation.

Abbreviations: dSL, discrete super learner; eSL, ensemble super learner; MdAE, median absolute error; MSE, mean square error; POSL, Personalised Online Super Learner; and XGBoost, extreme gradient boosting tree.

5.4. Sensitivity analysis in a classification setting

First, note that the lasso was absent from the set of candidate learners since the best cross-validated penalisation in the tuning set was the absence of penalisation. As with a continuous outcome, the historical learners (especially XGBoost) were more present in the convex eSL than the individual learners (Figure S4). However, the relative importance of historical and individual learners was relatively constant over time, with a slight advantage for the historical learners until the 200th session and a clearer advantage for individual learners after the 250th session, approximately. Nevertheless, these advantages were less pronounced than those observed with a continuous outcome. In contrast, the non-convex eSL produced predictions greater than one in 13.1% of the cases because of overly extreme weights for the individual logistic models. All of these predictions were bound to one for computing the performance measures.

Regarding performance (Figures S5-S6, Table S5), historical learners achieved poorer MdAE and weak calibration than individual learners, but a better MSE. POSLs outperformed the historical learners regarding the MdAE and weak calibration but only outperformed the individual learners regarding the MSE. All methods achieved a correct and similar calibration-in-the-large. Between the POSLs, the non-convex eSL seemed slightly better, especially regarding the MdAE. In contrast, historical learners yielded the best calibration-in-the-small, while individual learners underestimated the convection volume more largely than the POSLs. Historical learners achieved a better AUROC than individual learners, solely outperformed by the eSLs (Figure S7). All POSLs presented a similar improved net benefit against the strategy of closely monitoring all patients when they are at least 25% chance of having a successful HDF session (Figure S8). The evolution of the performance over time of POSLs is relatively constant (Figure S9), although the eSLs were miscalibrated-in-the-small in the first time.

6. Choices and challenges

6.1. To combine or not to combine?

This paper describes the implementation of POSL (Malenica et al., 2023) to dynamically predict the individual HDF convection volume using individual trajectories and data available from a historical pool of other individuals. Furthermore, POSL can combine multiple parametric and non-parametric approaches to obtain the best combination of their predictions (or the best approach’s predictions) at each time. We have studied three POSLs: a dSL (the best candidate learner in terms of cumulative squared error), a convex NNLS eSL, and a non-convex NNLS eSL. We showed that their performances were close in our case, although the non-convex eSL was more discriminant and slightly more accurate than the two others. In addition, the non-convex eSL achieved correct calibration faster than the dSL in our context. These results confirm the previous findings of Malenica et al. (2023) about the weakness of imposing a convex scaling after estimating the NNLS coefficients.

Contrary to the asymptotic results (Dudoit and van der Laan, 2005; van der Vaart et al., 2006), some candidate historical learners outperformed the POSLs regarding MSE and weak calibration in this application. Nevertheless, Hixon and Evgeniou (2005) argued, “the advantage of combining [predictions] is not that the best possible combinations perform better than the best possible individual [prediction], but that it is less risky in practice to combine [predictions] than to select an individual [prediction] method.” Indeed, investigators don’t know before the analysis which method performs the best; using POSL ensures that they do not pick an inefficient method, albeit POSL may not perform the best in all cases. We consider this as the main strength of using POSL over its alternatives.

6.2. Which inputs to use for the POSL?

Like any ensemble method, POSL’s performance depends on its three inputs: (a) the data, (b) the library of candidate learners, and (c) the loss function.

For the first input, the dynamic nature of POSL requires careful thinking about the choice of predictors. In particular, lagged values of crucial predictors must be defined according to domain knowledge. We used summaries of the recent history of the outcomes and two session-specific predictors according to clinical knowledge. Nevertheless, the history of graft rejection was unavailable from our data source, which might constitute a limitation of this study. One must also keep in mind that the prediction model is developed for being *implemented and used* in practice. For instance, one might avoid including predictors that are difficult to collect in an automatic way (Vickers and Cronin, 2010). Similarly, predictors collected differently over time may induce model miscalibration (Luijken et al., 2020) even if RWCV might partially correct this. We used a window of 10 sessions to be able to predict sooner. However, increasing the window size may improve the RWCV-based candidate individual learners performance. We previously noted that no minimal sample size could be derived for such models. Therefore, one must assume that the collected sample is sufficient for training the models, which is a strong assumption since ML approaches are data-hungry (van der Ploeg et al., 2014). Furthermore, an *efficient* sample size must also be computed for defining a minimum number of CV folds (Phillips et al., 2023). For the POSL, it is unclear whether these sample sizes refer to the number of individuals or the number of observations for the individual i , and whether we should include the historical cohort in the calculation. In our application, we have observed that the individual learners became relevant from, on average, the 150th observation approximately and required around a hundred additional observations to be predominant. Another issue is the presence of missing data. We used a missingness indicator approach with median imputation because it is computationally non-intensive, and missing data patterns may be predictive by themselves (Sperrin et al., 2020). However, knowing that the missingness of a predictor is predictive may motivate the user to refrain from measuring that predictor, creating harmful feedback between the model and the user (Groenwold, 2020). Since POSL is a “black box”, this phenomenon may be less likely to occur. Although generally considered the gold standard, multiple imputation is difficult to use with POSL because it substantially increases the computational complexity. An under-exploited alternative could be relying on ML approaches that naturally handle missing data, such as RF or XGBoost (Nijman et al., 2022). However, including only tree-based approaches in the POSL restricts its flexibility and increases the risk of overfitting Balzer and Westling (2021).

Considering the choice of candidate learners, following the recommendation of Phillips et al. (2023), we incorporated a diverse library of candidate learners, including parametric models with and without penalisation, non-parametric approaches that do not enforce linear or monotonic outcome-covariate relationships, and the arithmetic mean. Moreover, we considered two different CV schemes for individual learners to limit the impact of potential distribution shifts. Finally, we added an RF screening procedure to reduce the number of predictors. However, Strobl et al. (2007) showed that RF screening might incorrectly favour the selection of continuous/multinomial predictors over binary ones in classification settings. In the context of POSL, it could be relevant to add time-series models, such as ARIMA (Hyndman and Athanasopoulos, 2018), to the pool of individual learners when the number of observations is sufficient. However, including too many candidate learners (as evoked above) might complicate the estimation of the meta-learner by falling into the curse of dimensionality. Indeed, the meta-level dataset could present many more predictors (*i.e.*, the candidate learner-specific predicted outcomes) than observations. Using a meta-screener before fitting the meta-learner could be a solution for reducing the dimensionality (*e.g.*, using a post-lasso NNLS (Belloni and Chernozhukov, 2013) as meta-learner). Alternatively, we can put different eSLs and their candidate learners in a dSL (Phillips et al., 2023), which will select a candidate learner over the eSL if the eSL does not converge. Nevertheless, this procedure requires an additional loop of CV to compute the eSL-loss needed by the dSL. Phillips et al. (2023) defined an efficient sample size of 500 as a rule of thumb for including as many learners as computationally possible in an SL. Further research is needed to assess this rule with POSL. SLs can also be considered as candidate learners, a historical SL being especially appealing for improving the first POSL predictions. However, this requires including a supplementary loop of CV in the POSL training process. The choice of the candidate learners implies defining their hyperparameters. Indeed, default values provided by statistical packages are unlikely the optimal ones (Probst et al., 2019). We applied a sample-splitting strategy by defining the hyperparameters and training the

models on mutually exclusive splits. We fixed our hyperparameters across time, which is unlikely to be optimal. Some individual learners' hyperparameters can be too aggressive regarding the amount of data available in the first CV folds, explaining the MARS' overfitting trend. Alternatively, one can compute time-specific hyperparameters by splitting the CV fold in tuning and training mutually exclusive sets Tashman (2000). Such a nested CV has been shown to improve performance by Varma and Simon (2006), while other authors advocated for the more pragmatic approach of estimating the hyperparameters only one time to reduce the computational burden at the cost of a small decrease in performance (Schumacher et al., 2007; Foucher and Danger, 2012). Another solution might be to include several versions of the same candidate learner but with different hyperparameter values. Thereby, POSL can select the best implementation at each time.

The final component is the loss function. We used the sum of the squared errors over time for the dSL and the related NNLS as a meta-learner for the eSLs. Note that non-parametric approaches can be used as alternative meta-learners (Phillips et al., 2023). Our loss function and meta-regressions were weighted to give less importance to the least recent sessions. We used weights similar to those proposed by Malenica et al. (2023), but they acknowledged that domain knowledge must drive this choice. Again, we could define several eSLs with different weighting schemes to put into an unweighted dSL. In addition, the predictions from the CV folds were not bounded to plausible values, which potentially biased the squared error at time t but also at the subsequent times since the loss function was the cumulative sum. This can explain why the dSL did not achieve the best MSE in our application. Using this loss function also implied disregarding the fold in which at least one learner did not converge to not bias the cumulative loss.

6.3. How to evaluate the POSL

Prediction score must be evaluated regarding overall performances, calibration, discrimination, and clinical usefulness (McLernon et al., 2023). Calibration refers to how well model predictions agree with the actual outcome frequencies in the population under study. Of note, calibration was seldom reported for ML-based prediction models (Andaur Navarro et al., 2023). Discrimination refers to how well the model predictions separate high- from low-risk patients. Overall performances mix calibration and discrimination. Clinical usefulness refers to how well the model predictions can improve clinical decision-making by targeting high-risk patients for additional intervention while not burdening the system and patients with excessive false positives. We used both the MSE and the MdAE as measures of overall performance. Nevertheless, these measures are time-scale dependent. If the time range is highly variable between individuals, one can prefer scale-independent measures such as the median absolute percentage error or the mean squared percentage error (Hyndman and Athanasopoulos, 2018). Calibration was assessed at three nested levels (Van Calster et al., 2019). First, calibration-in-the-large implies that the observed event rate is close to the average predicted risk. Second, weak calibration implies no overestimation or underestimation of risks. Third, calibration-in-the-small implies the predicted risks correspond to observed event rates (but not necessary for each predictor pattern). We used the calibration intercept, the calibration slope, and a flexible calibration curve, respectively. Discrimination and clinical usefulness are more difficult to evaluate with a continuous outcome. Indeed, clinical decision-making typically operates in a classification setting with a threshold defining high-risk and low-risk subgroups. We used a threshold of 24L according to the literature (Lévesque et al., 2015; Blankestijn et al., 2023) for the discrimination evaluated through the AUROC. Clinical usefulness can be evaluated through the net benefit solely or a decision analysis curve for a binary outcome (Vickers and Elkin, 2006). Although decision curve analysis allows for the investigation of the net benefit across a range of possible threshold probabilities, it should not be used to select the best threshold (Vickers et al., 2019). Note that if the envisaged intervention is dangerous or costly, we can incorporate this "harm" in the net benefit computation (Vickers and Elkin, 2006). We emphasise the importance of reporting all these performance measures to compare the methods fairly. Indeed, historical learners were better calibrated than individual learners and POSLs in our application, but they achieved lower discrimination and net benefit. These performance measures should be considered together, and their respective importance considered with respect to the intended use (Vickers and Cronin, 2010).

The traditional paradigm in prediction modelling involves developing a model in a dataset, validating its performances internally on the same dataset with a correction for overfitting (*e.g.*, through V-fold CV or bootstrap resampling), and validating externally in another population (Steyerberg and Harrell, 2016). The cost of personalising the prediction with POSL, because POSL is based on individual learners (*i.e.*, learners trained only on the previous data of the individual for which one wants a prediction), is that no general score is developed, making an external validation infeasible. Nevertheless, note that prediction models are seldom appropriately validated because this needs many external validations in various populations, including those of the predicted individual (Van Calster et al., 2023). We can only validate it (and use it) in the population from which the individual is drawn. As a partial remedy, we employed an internal-external validation as suggested by Steyerberg and Harrell (2016). Internal-external validation is close to a V-fold CV, but the splitting is non-random. If the data are clustered according to some characteristics, such as centres; each centre becomes a fold. Then, we develop the prediction model on all centres but one and validate it on the remaining centre. Finally, we average the performance over the validation folds. We used this approach by clustering by individual time series, using a forward CV for each cluster, and averaging over the clusters. Thereby, POSL can also be viewed as a dynamic refitting process or model revision (Jenkins et al., 2018). Finally, random-effects meta-analysis models could be used for verifying the homogeneity of the POSL's performance across individuals (Austin et al., 2016).

7. Conclusion

We demonstrated the potential of POSL to dynamically predict a repeated continuous outcome while optimising the prediction for its subject. We illustrated its performances on a real-life example in nephrology. However, this methodology can also be employed outside medical fields as long as sufficient observations are collected for the statistical unit. We then discussed several challenges involved in

its use and our choices. We hope this study will encourage researchers to use this approach more broadly and provide guidance on using ML-based approaches in anticipation of the publication of ML-specific guidelines (Collins and Moons, 2019; Collins et al., 2021). Furthermore, prediction being a necessary step for causality (Hernán et al., 2019; van der Laan and Rose, 2011), POSL presents a new tool for personalised treatment rules (Moodie and Stephens, 2018). Finally, prediction modellers should look beyond the statistical performances and technical challenges by considering the ethical issues of using ML (Wiens et al., 2019; Vollmer et al., 2020).

Competing interests

No competing interest is declared.

Author contributions statement

Conceptualisation: AC, MB, and MES; Data Curation: MB and RL; Formal Analysis & Investigation: AC; Funding Acquisition: AC, RWP, and MES; Methodology: AC, RWP, and MES; Project Administration & Resources: MES; Software & Visualisation: AC; Supervision: RWP and MES; Writing - Original Draft: AC; and Writing - Review & Editing: All authors.

Data Availability Statement

Data may be available upon request from MB (email: michele.bally.chum@ssss.gouv.qc.ca) conditional on approval from the ethics committee of the Centre Hospitalier de l'Université de Montréal. R codes are publicly available at <https://github.com/ArthurChatton/Personalised-dynamic-SL>.

Acknowledgments

AC was supported by an IVADO postdoctoral fellowship (#2022-7820036733). This project was approved by the Université de Montréal ethic committee (#2023-4344).

References

- C. L. Andaur Navarro, J. A. Damen, M. van Smeden, T. Takada, S. W. Nijman, P. Dhiman, J. Ma, G. S. Collins, R. Bajpai, R. D. Riley, K. G. Moons, and L. Hooft. Systematic review identifies the design and methodological conduct of studies on machine learning-based prediction models. *Journal of Clinical Epidemiology*, 154:8–22, 2023. doi: 10.1016/j.jclinepi.2022.11.015.
- G. Athanasopoulos, R. A. Ahmed, and R. J. Hyndman. Hierarchical forecasts for australian domestic tourism. *International Journal of Forecasting*, 25(1):146–166, 2009. doi: 10.1016/j.ijforecast.2008.07.004.
- P. C. Austin, D. van Klaveren, Y. Vergouwe, D. Nieboer, D. S. Lee, and E. W. Steyerberg. Geographic and temporal validity of prediction models: different approaches were useful to examine model performance. *Journal of Clinical Epidemiology*, 79:76–85, 2016. doi: 10.1016/j.jclinepi.2016.05.007.
- V. Bačák and E. H. Kennedy. Principled machine learning using the super learner: An application to predicting prison violence. *Sociological Methods & Research*, 48(3):698–721, 2019. doi: 10.1177/0049124117747301.
- L. B. Balzer and T. Westling. Demystifying statistical inference when using machine learning in causal research. *American Journal of Epidemiology*, Ahead of print, 2021. ISSN 0002-9262. doi: 10.1093/aje/kwab200.
- A. Belloni and V. Chernozhukov. Least squares after model selection in high-dimensional sparse models. *Bernoulli*, 19(2):521–547, 2013. doi: 10.3150/11-BEJ410.
- D. Benkeser, C. Ju, S. Lendle, and M. van der Laan. Online cross-validation-based ensemble learning. *Statistics in Medicine*, 37(2): 249–260, 2018. doi: 10.1002/sim.7320.
- P. J. Blankestijn, R. W. Vernooij, C. Hockham, G. F. Strippoli, B. Canaud, J. Hegbrant, C. Barth, A. Covic, K. Cromm, A. Cucui, A. Davenport, M. Rose, M. Török, M. Woodward, and M. L. Bots. Effect of hemodiafiltration or hemodialysis on mortality in kidney failure. *New England Journal of Medicine*, Ahead of print, 2023. doi: 10.1056/NEJMoa2304820.
- L. Breiman. Stacked regressions. *Machine Learning*, 24(1):49–64, 1996. doi: 10.1007/BF00117832.
- L. Breiman. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical Science*, 16(3):199–231, 2001a. doi: 10.1214/ss/1009213726.
- L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001b. doi: 10.1023/A:1010933404324.
- Canadian Institute for Health Information. Annual statistics on organ replacement in Canada, 2012 to 2021, 2022. URL <https://www.cihi.ca/en/annual-statistics-on-organ-replacement-in-canada-2012-to-2021>.
- G. S. Collins and K. G. M. Moons. Reporting of artificial intelligence prediction models. *The Lancet*, 393(10181):1577–1579, 2019. doi: 10.1016/S0140-6736(19)30037-6.
- G. S. Collins, P. Dhiman, C. L. Andaur Navarro, J. Ma, L. Hooft, J. B. Reitsma, P. Logullo, A. L. Beam, L. Peng, B. Van Calster, M. van Smeden, R. D. Riley, and K. G. Moons. Protocol for development of a reporting guideline (TRIPOD-AI) and risk of bias tool (PROBAST-AI) for diagnostic and prognostic prediction model studies based on artificial intelligence. *BMJ Open*, 11(7):e048008, 2021. doi: 10.1136/bmjopen-2020-048008.

- A. A. H. de Hond, A. M. Leeuwenberg, L. Hooft, I. M. J. Kant, S. W. J. Nijman, H. J. A. van Os, J. J. Aardoom, T. P. A. Debray, E. Schuit, M. van Smeden, J. B. Reitsma, E. W. Steyerberg, N. H. Chavannes, and K. G. M. Moons. Guidelines and quality criteria for artificial intelligence-based prediction models in healthcare: a scoping review. *npj Digital Medicine*, 5(11):1–13, 2022. doi: 10.1038/s41746-021-00549-7.
- A. Devaux, R. Genuer, K. Peres, and C. Proust-Lima. Individual dynamic prediction of clinical endpoint from large dimensional longitudinal biomarker history: a landmark approach. *BMC Medical Research Methodology*, 22(1):188, 2022. doi: 10.1186/s12874-022-01660-3.
- S. Dudoit and M. J. van der Laan. Asymptotics of cross-validated risk estimation in estimator selection and performance assessment. *Statistical Methodology*, 2(2):131–154, 2005. doi: 10.1016/j.stamet.2005.02.003.
- Y. Foucher and R. Danger. Time dependent roc curves for the estimation of true prognostic capacity of microarray data. *Statistical Applications in Genetics and Molecular Biology*, 11(6):Article 1, 2012. doi: 10.1515/1544-6115.1815.
- M.-C. Fournier, Y. Foucher, P. Blanche, C. Legendre, S. Girerd, M. Ladrière, E. Morelon, F. Buron, L. Rostaing, N. Kamar, G. Mourad, V. Garrigue, G. Couvrat-Desvergnès, M. Giral, E. Dantan, and D. Consortium. Dynamic predictions of long-term kidney graft failure: an information tool promoting patient-centred care. *Nephrology, Dialysis, Transplantation*, 34(11):1961–1969, 2019. doi: 10.1093/ndt/gfz027.
- J. R. Gantenberg, K. W. McConeghy, C. J. Howe, J. Steingrimsson, R. Van Aalst, A. Chit, and A. R. Zullo. Predicting seasonal influenza hospitalizations using an ensemble super learner: A simulation study. *American Journal of Epidemiology*, Ahead of print, 2023. doi: 10.1093/aje/kwad113.
- R. H. H. Groenwold. Informative missingness in electronic health record systems: the curse of knowing. *Diagnostic and Prognostic Research*, 4(1):8, 2020. doi: 10.1186/s41512-020-00077-0.
- D. J. Hand. Classifier technology and the illusion of progress. *Statistical Science*, 21(1):1–14, 2006. doi: 10.1214/088342306000000060.
- M. A. Hernán, J. Hsu, and B. Healy. A second chance to get causal inference right: A classification of data science tasks. *Chance*, 32(1):42–49, 2019. doi: 10.1080/09332480.2019.1579578.
- M. Hibon and T. Evgeniou. To combine or not to combine: selecting among forecasts and their combinations. *International Journal of Forecasting*, 21(1):15–24, 2005.
- U. Hjorth. Model selection and forward validation. *Scandinavian Journal of Statistics*, 9(2):95–105, 1982.
- R. Hyndman and G. Athanasopoulos. *Forecasting: Principles and Practice (2nd edition)*. OTexts, Melbourne, Australia, 2018. URL <https://otexts.com/fpp2/>.
- G. James, D. Witten, T. Hastie, and R. Tibshirani. *An Introduction to Statistical Learning (2nd edition)*. Springer Texts in Statistics. Springer New York, New York, NY, 2021.
- D. A. Jenkins, M. Sperrin, G. P. Martin, and N. Peek. Dynamic models to predict health outcomes: current status and methodological challenges. *Diagnostic and Prognostic Research*, 2(1):23, 2018. doi: 10.1186/s41512-018-0045-2.
- D. A. Jenkins, G. P. Martin, M. Sperrin, R. D. Riley, T. P. A. Debray, G. S. Collins, and N. Peek. Continual updating and monitoring of clinical prediction models: time for dynamic prediction systems? *Diagnostic and Prognostic Research*, 5(1):1, 2021. doi: 10.1186/s41512-020-00090-3.
- A. P. Keil and J. K. Edwards. You are smarter than you think: (super) machine learning in context. *European Journal of Epidemiology*, 33(5):437–440, 2018. doi: 10.1007/s10654-018-0405-9.
- E. LeDell, M. J. van der Laan, and M. Petersen. Auc-maximizing ensembles through metalearning. *The International Journal of Biostatistics*, 12(1):203–218, 2016. doi: 10.1515/ijb-2015-0035.
- H. Li, Q. Xu, Y. He, X. Fan, H. Yang, and S. Li. Temporal detection of sharp landslide deformation with ensemble-based lstm-rnns and hurst exponent. *Geomatics, Natural Hazards and Risk*, 12:3089–3113, 2021. doi: 10.1080/19475705.2021.1994474.
- T. Liyanage, T. Ninomiya, V. Jha, B. Neal, H. M. Patrice, I. Okpechi, M.-h. Zhao, J. Lv, A. X. Garg, J. Knight, A. Rodgers, M. Gallagher, S. Kotwal, A. Cass, and V. Perkovic. Worldwide access to treatment for end-stage kidney disease: a systematic review. *The Lancet*, 385(9981):1975–1982, 2015. doi: 10.1016/S0140-6736(14)61601-9.
- K. Luijken, L. Wynants, M. v. Smeden, B. V. Calster, E. W. Steyerberg, R. H. H. Groenwold, D. Timmerman, T. Bourne, and C. Ukaegbu. Changing predictor measurement procedures affected the performance of prediction models in clinical examples. *Journal of Clinical Epidemiology*, 119:7–18, 2020. doi: 10.1016/j.jclinepi.2019.11.001.
- R. Lévesque, D. Marcelli, H. Cardinal, M.-L. Caron, M. P. C. Grooteman, M. L. Bots, P. J. Blankestijn, M. J. Nubé, A. Grassmann, B. Canaud, and A. Gandjour. Cost-effectiveness analysis of high-efficiency hemodiafiltration versus low-flux hemodialysis based on the canadian arm of the contrast study. *Applied Health Economics and Health Policy*, 13(6):647–659, 2015. doi: 10.1007/s40258-015-0179-0.
- I. Malenica, R. V. Phillips, A. Chambaz, A. E. Hubbard, R. Pirracchio, and M. J. van der Laan. Personalized online ensemble machine learning with applications for dynamic data streams. *Statistics in Medicine*, 42(7):1013–1044, 2023. doi: 10.1002/sim.9655.
- D. J. McLernon, D. Giardiello, B. Van Calster, L. Wynants, N. van Geloven, M. van Smeden, T. Therneau, and E. W. Steyerberg. Assessing performance and clinical usefulness in prediction models with survival outcomes: Practical guidance for cox proportional hazards models. *Annals of Internal Medicine*, 176(1):105–114, 2023. doi: 10.7326/M22-0844.
- E. E. M. Moodie and D. A. Stephens. *Dynamic Treatment Regimes*, page 1–5. John Wiley & Sons, Ltd, 2018. doi: 10.1002/9781118445112.stat08040. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118445112.stat08040>.
- K. G. Moons, D. G. Altman, J. B. Reitsma, J. P. Ioannidis, P. Macaskill, E. W. Steyerberg, A. J. Vickers, D. F. Ransohoff, and G. S. Collins. Transparent reporting of a multivariable prediction model for individual prognosis or diagnosis (TRIPOD): Explanation and elaboration. *Annals of Internal Medicine*, 162(1):W1–W73, 2015. doi: 10.7326/M14-0698.

- A. I. Naimi and L. B. Balzer. Stacked generalization: an introduction to super learning. *European Journal of Epidemiology*, 33(5): 459–464, 2018. doi: 10.1007/s10654-018-0390-z.
- S. W. J. Nijman, A. M. Leeuwenberg, I. Beekers, I. Verkouter, J. J. L. Jacobs, M. L. Bots, F. W. Asselbergs, K. G. M. Moons, and T. P. A. Debray. Missing data is poorly handled and reported in prediction model studies using machine learning: a literature review. *Journal of Clinical Epidemiology*, 142:218–229, 2022. doi: 10.1016/j.jclinepi.2021.11.023.
- C. S. Peirce. The numerical measure of the success of predictions. *Science*, 4(93):453–454, 1884. doi: 10.1126/science.ns-4.93.453.b.
- S. A. E. Peters, M. L. Bots, B. Canaud, A. Davenport, M. P. C. Grooteman, F. Kircelli, F. Locatelli, F. Maduell, M. Morena, M. J. Nubé, E. Ok, F. Torres, M. Woodward, P. J. Blankestijn, and H. P. P. Investigators. Haemodiafiltration and mortality in end-stage kidney disease patients: a pooled individual participant data analysis from four randomized controlled trials. *Nephrology, Dialysis, Transplantation*, 31(6):978–984, 2016. doi: 10.1093/ndt/gfv349.
- R. V. Phillips, M. J. van der Laan, H. Lee, and S. Gruber. Practical considerations for specifying a super learner. *International Journal of Epidemiology*, Ahead of print, 2023. doi: 10.1093/ije/dyad023.
- R. Pirracchio, M. L. Petersen, M. Carone, M. Resche-Rigon, S. Chevret, and M. J. van der Laan. Mortality prediction in the ICU: can we do better? results from the super ICU learner algorithm (SICULA) project, a population-based study. *The Lancet. Respiratory medicine*, 3(1):42–52, 2015. doi: 10.1016/S2213-2600(14)70239-5.
- P. Probst, A.-L. Boulesteix, and B. Bischl. Tunability: Importance of hyperparameters of machine learning algorithms. *Journal of Machine Learning Research*, 20:1–32, 2019.
- C. Proust-Lima and P. Blanche. *Dynamic Predictions*, page 1–6. John Wiley & Sons, Ltd, 2016. doi: 10.1002/9781118445112.stat07876. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118445112.stat07876>.
- R. D. Riley, K. I. Snell, G. P. Martin, R. Whittle, L. Archer, M. Sperrin, and G. S. Collins. Penalization and shrinkage methods produced unreliable clinical prediction models especially when sample size was small. *Journal of Clinical Epidemiology*, 132:88–96, 2021. doi: 10.1016/j.jclinepi.2020.12.005.
- D. R. Roberts, V. Bahn, S. Ciuti, M. S. Boyce, J. Elith, G. Guillera-Arroita, S. Hauenstein, J. J. Lahoz-Monfort, B. Schröder, W. Thuiller, D. I. Warton, B. A. Wintle, F. Hartig, and C. F. Dormann. Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*, 40(8):913–929, 2017. doi: 10.1111/ecog.02881.
- C. Sabathé, R. Casey, S. Vukusic, E. Leray, G. Mathey, J. De Sèze, J. Ciron, S. Wiertlewski, A. Ruet, J. Pelletier, H. Zéphir, L. Michel, C. Lebrun-Frenay, X. Moisset, E. Thouvenot, J.-P. Camdessanché, S. Bakchine, B. Stankoff, A. Al Khedr, P. Cabre, E. Maillart, E. Berger, O. Heinzlef, K. Hankiewicz, T. Moreau, O. Gout, B. Bourre, A. Wahab, P. Labauge, A. Montcuquet, G. Defer, A. Maurousset, N. Maubeuge, D. Dimitri Boulos, H. Ben Nasr, C. Nifle, O. Casez, D.-A. Laplaud, and Y. Foucher. Improving the decision to switch from first- to second-line therapy in multiple sclerosis: A dynamic scoring system. *Multiple Sclerosis*, 29(2):236–247, 2023. doi: 10.1177/13524585221139156.
- M. Schumacher, H. Binder, and T. Gerds. Assessment of survival prediction models based on microarray data. *Bioinformatics*, 23(14): 1768–1774, 2007. doi: 10.1093/bioinformatics/btm232.
- M. Sperrin, G. P. Martin, R. Sisk, and N. Peek. Missing data should be handled differently for prediction than for description or causal explanation. *Journal of Clinical Epidemiology*, 125:183–187, 2020. doi: 10.1016/j.jclinepi.2020.03.028.
- E. W. Steyerberg and F. E. Harrell. Prediction models need appropriate internal, internal-external, and external validation. *Journal of Clinical Epidemiology*, 69:245–247, 2016. doi: 10.1016/j.jclinepi.2015.04.005.
- C. Strobl, A.-L. Boulesteix, A. Zeileis, and T. Hothorn. Bias in random forest variable importance measures: Illustrations, sources and a solution. *BMC Bioinformatics*, 8(1):25, 2007. doi: 10.1186/1471-2105-8-25.
- N. R. Swanson and H. White. Forecasting economic time series using flexible versus fixed specification and linear versus nonlinear econometric models. *International Journal of Forecasting*, 13(4):439–461, 1997. doi: 10.1016/S0169-2070(97)00030-7.
- K. T. Tanner, L. D. Sharples, R. M. Daniel, and R. H. Keogh. Dynamic survival prediction combining landmarking with a machine learning ensemble: Methodology and empirical comparison. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 184(1):3–30, 2021. doi: <https://doi.org/10.1111/rssa.12611>.
- L. J. Tashman. Out-of-sample tests of forecasting accuracy: an analysis and review. *International Journal of Forecasting*, 16(4):437–450, 2000. doi: 10.1016/S0169-2070(00)00065-0.
- B. Van Calster, D. J. McLernon, M. van Smeden, L. Wynants, and E. W. Steyerberg. Calibration: the Achilles heel of predictive analytics. *BMC Medicine*, 17(1):230, 2019. doi: 10.1186/s12916-019-1466-7.
- B. Van Calster, M. van Smeden, B. De Cock, and E. W. Steyerberg. Regression shrinkage methods for clinical prediction models do not guarantee improved performance: Simulation study. *Statistical Methods in Medical Research*, 29(11):3166–3178, 2020. doi: 10.1177/0962280220921415.
- B. Van Calster, E. W. Steyerberg, L. Wynants, and M. van Smeden. There is no such thing as a validated prediction model. *BMC Medicine*, 21(1):70, 2023. doi: 10.1186/s12916-023-02779-w.
- M. J. van der Laan and S. Rose. *Targeted learning: causal inference for observational and experimental data*. Springer series in statistics. Springer, New York, 2011.
- M. J. van der Laan, E. C. Polley, and A. E. Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1): Article 25, 2007. doi: 10.2202/1544-6115.1309.
- T. van der Ploeg, P. C. Austin, and E. W. Steyerberg. Modern modelling techniques are data hungry: a simulation study for predicting dichotomous endpoints. *BMC Medical Research Methodology*, 14(1):137, 2014. doi: 10.1186/1471-2288-14-137.
- A. W. van der Vaart, S. Dudoit, and M. J. van Der Laan. Oracle inequalities for multi-fold cross validation. *Statistics & Decisions*, 24(3):351–371, 2006. doi: 10.1524/stdn.2006.24.3.351.

- S. Varma and R. Simon. Bias in error estimation when using cross-validation for model selection. *BMC Bioinformatics*, 7(1):91, 2006. doi: 10.1186/1471-2105-7-91.
- A. J. Vickers and A. M. Cronin. Everything you always wanted to know about evaluating prediction models (but were too afraid to ask). *Urology*, 76(6):1298–1301, 2010. doi: 10.1016/j.urology.2010.06.019.
- A. J. Vickers and E. B. Elkin. Decision curve analysis: A novel method for evaluating prediction models. *Medical Decision Making*, 26(6):565–574, 2006. doi: 10.1177/0272989X06295361.
- A. J. Vickers, B. van Calster, and E. W. Steyerberg. A simple, step-by-step guide to interpreting decision curve analysis. *Diagnostic and Prognostic Research*, 3(1):18, 2019. doi: 10.1186/s41512-019-0064-7.
- S. Vollmer, B. A. Mateen, G. Bohner, F. J. Király, R. Ghani, P. Jonsson, S. Cumbers, A. Jonas, K. S. L. McAllister, P. Myles, D. Grainger, M. Birse, R. Branson, K. G. M. Moons, G. S. Collins, J. P. A. Ioannidis, C. Holmes, and H. Hemingway. Machine learning and artificial intelligence research for patient benefit: 20 critical questions on transparency, replicability, ethics, and effectiveness. *BMJ*, 368:l6927, 2020. doi: 10.1136/bmj.l6927.
- J. Wiens, S. Saria, M. Sendak, M. Ghassemi, V. X. Liu, F. Doshi-Velez, K. Jung, K. Heller, D. Kale, M. Saeed, P. N. Ossorio, S. Thadane-Israni, and A. Goldenberg. Do no harm: a roadmap for responsible machine learning for health care. *Nature Medicine*, 25(9):1337–1340, 2019. doi: 10.1038/s41591-019-0548-6.
- S. N. Wood, N. Pya, and B. Säfken. Smoothing parameter and model selection for general smooth models. *Journal of the American Statistical Association*, 111(516):1548–1563, 2016. doi: 10.1080/01621459.2016.1180986.