

Learning Patient-Level Prediction Models Across Multiple Healthcare Databases: Evaluation of Ensembles for Increasing Model Transportability

Jenna Marie Reps (✉ jreps@its.jnj.com)

Janssen (United States)

Ross D. Williams

Erasmus MC

Martijn J. Schuemie

Janssen (United States)

Patrick B. Ryan

Janssen (United States)

Peter R. Rijnbeek

Erasmus MC

Research Article

Keywords: Ensemble learning, model transportability, prognostic model, observational data, patient-level prediction

Posted Date: February 7th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1217376/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Title: Learning patient-level prediction models across multiple healthcare databases: Evaluation of ensembles for increasing model transportability

Corresponding author: Jenna Marie Reps, jreps@its.jnj.com, +1 732 715 6300, Janssen Research and Development, Raritan, NJ, USA

Authors: Jenna Marie Reps^{1*}, Ross D. Williams^{2*}, Martijn J. Schuemie¹, Patrick B. Ryan¹, Peter R. Rijnbeek²

¹Janssen Research and Development, Raritan, NJ, USA;

²Department of Medical Informatics, Erasmus University Medical Center, Rotterdam, The Netherlands

* Co-first authors

Keywords: Ensemble learning; model transportability; prognostic model; observational data; patient-level prediction

ABSTRACT (322)

Background: Prognostic models that are accurate could help aid medical decision making. Large observational databases often contain temporal medical data for large and diverse populations of patients. It may be possible to learn prognostic models using the large observational data. Often the performance of a prognostic model undesirably worsens when transported to a different database (or into a clinical setting). In this study we investigate different ensemble approaches that combine prognostic models independently developed using different databases (a simple distributed learning approach) to determine whether ensembles that combine models developed across databases can improve model transportability (perform better in new data than single database models)?

Methods: For a given prediction question we independently trained five single database models each using a different observational healthcare database. We then developed and investigated numerous ensemble models (fusion, stacking and mixture of experts) that combined the different database models. Performance of each model was investigated via discrimination and calibration using a leave one dataset out technique, i.e., hold out one database to use for validation and use the remaining four datasets for model development. The internal validation of a model developed using the hold out database was calculated and presented as the ‘hypothetical optimum’ for comparison.

Results: In this study the fusion ensembles generally outperformed the single database models when transported to a previously unseen database and the performances were more consistent across unseen databases. Stacking ensembles performed poorly in terms of discrimination when the labels in the unseen database were limited. Calibration was consistently poor when both ensembles and single database models were applied to previously unseen databases.

Conclusion: A simple distributed learning approach that implements ensemble techniques to combine models independently developed across different databases for the same prediction question may improve the discriminative performance in new data (new database or clinical setting) but will need to be recalibrated using the new data. This could help medical decision making by improving prognostic model performance.

BACKGROUND

Big observational healthcare databases, such as insurance claims data or electronic healthcare records, often contain data on large and diverse populations. One area where these datasets may benefit healthcare is in the application of machine learning to develop prognostic models. Prognostic models aim to predict a patient's risk of experiencing some future event (e.g., cardiovascular illnesses) [1] based on their current and historic health. In general, a prognostic task can be decomposed into three parts, the target population/index, the outcome, and the time-at-risk [2]. The target population is the set of patients for whom you attempt to predict the risk of some future outcome and the index is the point in time you want to make the prediction. The outcome is the medical event you want to predict, and the time-at-risk is the time interval (relative to the index) you want to predict the outcome occurring within. Prognostic models are learned from observational healthcare databases by finding patients in the database who historically match the target population, determining features such as age, gender, and medical history at index for each patient and then observing whether they had the outcome during the time-at-risk. Supervised learning, such as binary classification, is then applied to learn the differences between the people who had the outcome during the time-at-risk vs the people who did not. Often the aim is to develop a model using the historical data but apply the model to current patients to calculate a probability of whether they will have the outcome during the future time-at-risk. Such models could improve healthcare by informing medical decision making, but only if these models perform sufficiently well when implemented in their intended setting. For example, a model intended to be used by a family medicine doctor to help them decide which patients should be given preventative medicine may be developed using a large insurance claims database but needs to transport well into the family medicine setting. The performance in a new database (transportability) of a model is initially assessed by externally validating a model across diverse datasets with different patient case mixes [2,3]. It is common for a model's performance to deteriorate when transported to a different database [2]. The deterioration in performance may be due to the model or the differences between the development and validation populations [4]. A model that transports well to other databases is much more valuable in clinical practice. The question is how to best develop models with high transportability?

Big observational healthcare datasets only contain a sample of the population. This is frequently a non-random sample, for example the data may over sample (or only contain) certain ethnicities, genders, ages or patients with low/medium/high wealth. If a database used to develop a prognostic model contains a non-random sample of the target population then this will most likely negatively affect its performance if applied on the full population. However, different datasets, with varying patient case mixes, may give diverse perspectives when developing prognostic model for the same prediction task. Learning models across different healthcare datasets (e.g., a US insurance claims database, a UK primary care database and a US electronic healthcare record database), known as distributed learning [5-8], may lead to more transportable models. There are numerous approaches to distributed learning: i) combine the datasets and develop a model using the combined data, ii) apply a distributed algorithm that iterates across

datasets or iii) combine models developed using different datasets. The first option is generally limited as sharing patient-level data between researchers is often not possible due to privacy restrictions and therefore it is not possible to train a single model using the combination of different datasets. The second option is limited due to the administration required if the algorithm needs to communicate with each dataset (held at different physical locations by different owners) multiple times. Although some distributed algorithms that only require access to each database once, termed ‘one-shot distributed algorithms’, exist for certain generalized linear models [8]. The one-shot approach is not suitable for most machine learning models. The third option is most feasible, as it is possible for researchers to easily share prognostic models, they develop using their own data and these models could be combined via ensemble techniques (ensemble modeling is the common machine learning approach used to combine binary classification models). This prompts the question; can we implement the third option and combine models developed using diverse datasets to improve model transportability in new data (e.g., in a clinical setting)?

Ensemble learning is the process of producing multiple models, potentially pruning the set of models, and then combining the remaining models [9]. Often the ensemble increases both model accuracy and performance stability compared to any single classifier [10]. Ensembles either combine homogeneous models (same learning algorithm) or heterogeneous models (different learning algorithms). Homogeneous ensembles use the same learning algorithm but modify the perspective by using different training data (e.g., different instances, different features or by adding noise), different metrics or using different model settings (e.g., hyper-parameter values). Heterogeneous models take a different perspective as each learning algorithm makes different assumptions about the data. Combining the models is often done by fusing the classifiers [11], stacking [12] or using a mixture of experts [13]. Examples of simple fusing classifiers include i) majority vote, the combination technique used by random forest ‘bagging’ [14], ii) calculating the mean prediction value across classifiers or iii) weighted mean of the classifier predictions based on performance measures. Weighing each classifier’s prediction based on performance is better than taking the mean of all classifier predictions when the classifier performances differ (e.g., one classifier is better than the others) [11]. A mixture of experts is similar to weighted mean but instead of using universal weights across the instances, the weights are assigned per instance [13]. These ensembles are considered independent ensemble frameworks, as the models are trained independently and then combined [10]. A more advanced independent ensemble framework is known as ‘stacking’. Stacking is a meta-combination method that uses the set of models’ predictions as features and trains a new model that learns to predict the outcome using these prediction features [12]. A limitation of stacking is that it requires additional labeled data to learn how to best combine the individual models. Alternatively, ‘dependent ensemble’ frameworks train classifiers sequentially and each classifier depends on the output of the prior classifier [10]. Boosting is a dependent fusion ensemble framework as models are sequentially trained, and weights are assigned to the objective function of each model during training based on prior models’ mistakes [15]. The above examples are just a selection of the commonly used combination methods and there are numerous other ways to combine the models [10].

objective

This paper aims to determine whether prognostic model ensembles that combine regularized logistic regression models independently developed across different healthcare databases

perform better in new data (more transportable) than each individual database prognostic model (single dataset model). A model with improved transportability is likely to also perform better when used clinically for decision making.

METHODS

The Observational Health Data Science and Informatics (OHDSI) PatientLevelPrediction framework is used throughout this paper [2] for developing prognostic models using observational healthcare data.

databases

Four US claims and an EHR databases are explored, see Table 1.

Table 1- Summary of the five databases used in this study

Name	Type	Description	Start	End	Size (million lives)
IBM Medicare Supplemental Beneficiaries (MDCR)	US Claims	Patients aged 65 or older with supplemental healthcare.	2000-01-01	2019-12-31	10.115
IBM Medicaid (MDCD)–	US Claims	Patients with government subsidized healthcare.	2006-01-01	2018-12-31	28.777
Optum® De-Identified Clinformatics® Data Mart Database (Optum Claims)	US Claims	Patients of all ages	2000-05-01	2019-12-31	84.310
IBM Commercial Claims and Encounters (CCAЕ)	US Claims	The patients in this database are aged 65 or younger. They are employees who receive health insurance through their employer and their dependents.	2000-01-01	2019-12-31	152.96
Optum® de-identified Electronic Health Record Dataset (Optum EHR)	US EHR	Patients of all ages	2006-01-01	2019-03-31	96.505

The five databases in this study contain retrospectively collected deidentified data. The use of IBM and Optum databases were reviewed by the New England Institutional Review Board (IRB) and were determined to be exempt from broad IRB approval.

All datasets used in this paper were mapped into the OHDSI Observational Medical Outcomes Partnership Common Data Model (OMOP-CDM) version 5 [16]. The OMOP-CDM was developed to enable researchers with diverse datasets to have a standard database structure. This enables analysis code and software to be shared among researchers which facilitates external validation of prediction models.

prediction problem

As an example, the problem: “Amongst patients with pharmaceutically-treated depression, which patients will develop <an outcome> during the 1-year time interval following the start of the depression episode?” is investigated.

The target population of pharmaceutically treated depressed patients is defined as: patients with a condition record of major depressive disorder and the index date was the first record date.

Inclusion criteria are:

- Antidepressant recorded within 30 days before to 30 days after the target population index date
- No history of psychosis
- No history of dementia
- No history of mania
- ≥ 365 days prior observation

Twenty-one models predicting 21 different outcomes occurring for the first time between 1 day after index until 1 year after index are developed. The 21 outcomes are: acute liver injury, acute myocardial infarction, alopecia, constipation, decreased libido, delirium, diarrhea, fracture, gastrointestinal hemorrhage, hyponatremia, hypotension, hypothyroidism, insomnia, nausea, seizure, stroke, sudden cardiac death, suicide and suicidal ideation, tinnitus, ventricular arrhythmia and vertigo.

The above definition of prediction problem is the same as used in reference [2].

In this study a random sample of 500,000 patients from the target population (>1 million patients in Optum claims, >2 million patients in Optum EHR and >2 million in CCAE) were used throughout the study. This improved the efficiency of model development and also resulted in some low outcome counts, enabling the investigation into whether the outcome count impacts the ensemble performance.

Labelled data

We constructed labelled datasets for each database and outcome pair. For the i th patient in database k we used one-hot-encoding to create binary features indicating the presence of any medical condition or drug recorded prior to index (first record of major depressive disorder) and extracted the patient's gender and age at index (in 5-year bins). Let \mathbf{x}_i^k represent the feature vector for the i th patient in database k . Labels were determined per outcome, with y_{ij}^k corresponding to the presence ($y_{ij}^k = 1$) or absence ($y_{ij}^k = 0$) of outcome j in the year after index for patient i in database k . This resulted in 105 labelled datasets $\{(\mathbf{x}_i^k, y_{ij}^k)\}_i$.

statistical analysis

Binary Classifiers (Level 1 models)

For each database and outcome, a regularized logistic regression model with least absolute shrinkage and selection operator (LASSO) penalization was trained [17] using 80% of the data to develop the model and 20% of the data were held out to internally estimate the model performance (test set performance). Three-fold cross validation was applied in the 80% development data to learn the optimal regularization value. The final LASSO logistic regression coefficients were learned with the optimal hyper-parameter set using all of the 80% development data.

Let $f_{ij}(\mathbf{x}): R_m \rightarrow [0,1]$ correspond to the Level 1 logistic regression model that was developed using the i th database (database i) to predict the j th outcome (outcome j), where \mathbf{x} is the m -dimension feature vector for a patient. Given a patient's feature vector, the Level 1 model developed in database i predicts a value between 0 and 1 that corresponds to the probability that the patient will experience outcome j .

Performance Evaluation

Internal validation is when a model is developed and evaluated in the same database and external validation is when a model is developed and evaluated in different databases. For both internal and external validation, model discrimination and calibration were calculated. Model discrimination assesses how well a model ranks patients based on risk, this was calculated using the area under the receiver operating curve (AUROC). The AUROC is a ranking measure that corresponds to the probability that if a non-outcome patient was sampled and an outcome patient was sampled, the predicted risk assigned to the outcome patient is greater than the predicted risk assigned to the non-outcome patient. An AUROC of 0.5 corresponds to randomly predicting risk (no discriminative ability) and an AUROC of 1 corresponds to perfect prediction (a higher risk is predicted for all patients who will experience the outcome compared to those who will not). Calibration assesses how closely the predicted risk matches the true risk. For example, if a model is well calibrated, then if 10 patients are assigned a 10% risk, only 1 of them should experience the outcome. In this study, calibration was calculated using calibration-in-the-large [18] which compares the model's mean predicted risk in the population with the observed risk (a model is considered well calibrated if the mean predicted risk matches the observed risk in the population).

The internal validation of each Level 1 model (test set performance) provides a benchmark performance for the database and outcome pair. The internal validation of each Level 1 model,

trained in database k to predict outcome j , was determined by calculating the AUROC and calibration-in-the-large using the predicted risk $f_{ij}(\mathbf{x}_i^k)$ and the true label y_{ij}^k for each patient in the 20% held out set (test set).

Binary Ensemble Classifiers (Level 2 models)

The ensembles in this study combine the Level 1 models developed in the different databases that predict the same outcome. Generally, an ensemble that predicts outcome j is a function of the N Level 1 models that predict outcome j :

$$f_j(\mathbf{x}) = g(\{f_{ij}(\mathbf{x})\}_{i \in \{1,2,\dots,N\}})$$

Seven different ensemble approaches were investigated to combine the Level 1 models, that predict the same outcome (j) but are trained on N different databases ($\{f_{kj}\}_{k \in \{1,2,\dots,N\}}$), using different heuristics.

A weighted fusion ensemble to predict the outcome j combines the Level 1 models by assigning each Level 1 model a weight:

$$f_j(\mathbf{x}) = \sum_i w_{ij} f_{ij}(\mathbf{x})$$

where w_{ij} is the weight assigned to the Level 1 model trained using database i to predict outcome j . In this study different weighting heuristics are investigated:

1. Mean Ensemble (**mean**) – for a patient, their predicted risk is the mean of the predicted risks of the included Level 1 classifiers (equal weighting so $w_{ij} = 1/N$, where N is the number of models being combined)
2. AUROC Ensemble normalized weights (**auc1**) – for a patient, their predicted risk is a weighted mean of the predicted risks of the included Level 1 models, where each Level 1 model's weight is based on the model's internal area under the receiver operating characteristic curve (AUROC) that was calculated in the 20% held out data. The weights are scaled relative to an AUROC of 0.5 and normalized to ensure the total weight across models was 1 (AUROC performance weighting so $w_{ij} = \frac{|AUROC_{ij}-0.5|}{\sum_k |AUROC_{kj}-0.5|}$), where $AUROC_{ij}$ is the internal AUROC value for the Level 1 model developed in database i to predict outcome j .
3. AUROC Ensemble unnormalized weights (**auc2**) – similar to 2) a patient's risk is a weighted mean of the predicted risks of the included Level 1 models, where each Level 1 model's weight is based on the model's internal AUROC. The weights are scaled between 1 for perfect discrimination and -1 for models that predict the opposite labels perfectly ($w_{ij} = \frac{AUROC_{ij}-0.5}{0.5}$), where $AUROC_{ij}$ is the internal AUROC value for the Level 1 model developed in database i to predict outcome j .
4. Similarity Weighted Ensemble (**sim**)- for a patient, their predicted risk is a weighted mean of the predicted risks of the included Level 1 models, but weights are based on how similar the Level 1 model's development population mean value for each predictor are

compared to the population that the patient is in. The cosine similarity metric was used for the two vectors containing the mean values in the patient's dataset and the Level 1 model's development data (case mix similarity weighting $w_{ij} = \frac{\text{cosine}(\mathbf{d}, \mathbf{d}_i)}{\sum_k \text{cosine}(\mathbf{d}, \mathbf{d}_k)}$) where \mathbf{d} is an m-dimensional vector corresponding to the mean values of the features included in model f_{ij} in the database the ensemble is being applied to and \mathbf{d}_i is an m-dimensional vector corresponding to the mean values of the features included in model f_{ij} in database i .

5. Age Weighted Ensemble (**age**)– for a patient, their predicted risk is a weighted mean based on how similar the model development data population mean age was compared to the patient's population mean age (case age similarity weighting $w_{ij} = \frac{d(\text{age}, \text{age}_i)}{\sum_k d(\text{age}, \text{age}_k)}$), where age is the mean age in years of the patients in the dataset the model is being applied to, age_i is the mean age of the patients in database i and $d(\text{age}, \text{age}_i) = 1/(1+|\text{age} - \text{age}_i|)$.

The mixture of expert ensembles $f_j(\mathbf{x})$ use the equation: $f_j(\mathbf{x}) = \sum_i g_{ij}(\mathbf{x})f_{ij}(\mathbf{x})$ where g_{ij} is the gating function value for Level 1 model developed in database i to predicted outcome j .

6. Age Mixture of Experts Ensemble (**ageME**) – for a patient, their predicted risk is calculated using the Level 1 model developed using a population with a mean age that most closely matches the patient's age, the gating function is:

$$g_{ij}(\mathbf{x}) = \begin{cases} \mathbf{1}, & i \equiv \min_k (\text{age}_k - \text{age}) \\ \mathbf{0}, & \text{otherwise} \end{cases}$$

Where age_k is the mean age in years of the patients in database k and age is the age in years of the patient whose risk is being calculated.

Stacking ensembles involved training a Level 2 model that uses the Level 1 model predictions as features.

7. Stacking ensemble –a LASSO logistic regression model was trained as the Level 2 model that used the predicted risk from each Level 1 model as predictors (effectively this learned the Level 1 model weightings). The stacking ensemble requires labelled data in the validation dataset whereas the other ensembles do not require this. As it is often not possible to get large amounts of labelled data in the validation dataset or application dataset, it was investigated how well the stacking ensemble would do if i) only 1,000 patients (**s|1000**), ii) only 10,000 patients (**s|10000**) and iii) all available patients (**s|All**) were used to learn the weightings.

Model Transportability

$$\mathbf{pred}_j^k = \left(\sum_{i \neq k} \frac{f_{ij}(\mathbf{x}_1^k)}{n}, \sum_{i \neq k} \frac{f_{ij}(\mathbf{x}_2^k)}{n}, \dots, \sum_{i \neq k} \frac{f_{ij}(\mathbf{x}_m^k)}{n} \right)$$

The ground truth in database k is:

$$\mathbf{truth}_j^k = (y_{1j}^k, y_{2j}^k, \dots, y_{mj}^k)$$

The external AUROC and calibration metrics for the mean fusion ensemble applied to database k for outcome j is then calculated by comparing the predictions and ground truth labels.

$$\mathit{externalAUROC}_j^k = \mathit{AUROC}(\mathbf{pred}_j^k, \mathbf{truth}_j^k)$$

In general, the predictions for all Level 1 models and Level 2 ensemble models when transported to database k are:

$$\mathit{level1Pred}_j^k = (f_{ij}(\mathbf{x}_1^k), f_{ij}(\mathbf{x}_2^k), \dots, f_{ij}(\mathbf{x}_m^k)), \text{ where } i \neq k$$

$$\mathit{ensemblePred}_j^k = (g(\{f_{ij}(\mathbf{x}_1^k)\}_{i \neq k}), g(\{f_{ij}(\mathbf{x}_2^k)\}_{i \neq k}), \dots, g(\{f_{ij}(\mathbf{x}_m^k)\}_{i \neq k}))$$

To put the performance of the Level 1 and Level 2 models (that do not use database k) into context, the ‘hypothetical optimal’ performance that is achievable in database k was estimated. The ‘hypothetical optimum’ is defined as the internal validation performance (using a 20% test set $\{(\hat{\mathbf{x}}_i^k, \hat{y}_{ij}^k)\}_i$) of the Level 1 model developed in database k :

$$\begin{aligned} \mathit{internalPred}_j^k &= (f_{kj}(\hat{\mathbf{x}}_1^k), f_{kj}(\hat{\mathbf{x}}_2^k), \dots, f_{kj}(\hat{\mathbf{x}}_t^k)) \\ \mathit{truthTest}_j^k &= (\hat{y}_{1j}^k, \hat{y}_{2j}^k, \dots, \hat{y}_{tj}^k) \end{aligned}$$

The internal AUROC in database k for outcome j is then:

$$\mathit{internalAUROC}_j^k = \mathit{AUROC}(\mathit{internalPred}_j^k, \mathit{truthTest}_j^k)$$

Given sufficient data, the internal performance of a model can be considered the upper bound of achievable performance (conditional on the same features being available to internal and external model development). If a model transported to new data has an external performance close to the internal performance of a model developed using the data, then this can be considered to have transported well. Consequently, to determine how well a model transports the difference in performance between the internal validation AUROC of the Level 1 model trained using the left-out database, database k , and the external validation AUROC of models when applied to the left-out database k was calculated:

$$\mathit{AUROC_difference}_j^k = \mathit{externalAUROC}_j^k - \mathit{internalAUROC}_j^k$$

where $\mathit{externalAUROC}_j^k$ is the performance of the model in database k (trained without dataset k) in predicting outcome j and $\mathit{internalAUROC}_j^k$ is the Level 1 model predicting outcome j trained in database k 's performance on the 20% test set. To show how well each model transports in general, box plots were created to show the distribution

of $AUROC_difference_j^k$ across the different outcomes and databases. Distributions centered around 0 indicate excellent transportability and distributions with a small range indicate consistency.

RESULTS

The data sizes are presented in Table 2 and the database characteristics are displayed in Table 3. The smallest target population was the one extracted from the MDCR database, and this population were older and had higher rates of cancer and cardiovascular issues prior to index. The MDCD target population was the youngest and also had the highest rate of obesity recorded in the prior year. In general, the characteristics varied greatly across the datasets, indicating different patient case-mixes. The outcome count was generally greater than 100 except for delirium in MDCD and Optum Claims and Seizure in MDCR and Optum EHR.

Table 2 - The outcome counts and percentage of target population who develop the outcome during the time-at-risk

Outcome	CCAE (N ~ 499,678) (%)	MDCR (N ~160,956) (%)	MDCD (N ~ 469,302) (%)	Optum EHR (N ~ 499,881) (%)	Optum Claims (N ~ 499,753) (%)
Acute liver injury	14875 (3.35)	7226 (5.4)	21654 (5.47)	18535 (4.18)	18619 (4.31)
Acute myocardial infarction	1494 (0.3)	935 (0.59)	3800 (0.83)	816 (0.16)	1298 (0.26)
Alopecia	10672 (2.32)	7569 (5.64)	20597 (5.2)	16597 (3.69)	16571 (3.75)
Constipation	4170 (0.85)	6399 (4.39)	9210 (2.05)	10192 (2.13)	10282 (2.16)
Decreased libido	491 (0.1)	1080 (0.69)	905 (0.19)	287 (0.06)	708 (0.14)
Delirium	174 (0.03)	510 (0.32)	86 (0.02)	267 (0.05)	91 (0.02)
Diarrhea	1661 (0.34)	130 (0.08)	785 (0.17)	1210 (0.24)	1603 (0.32)
Fracture	509 (0.1)	963 (0.61)	894 (0.19)	381 (0.08)	758 (0.15)
Gastrointestinal hemorrhage	985 (0.2)	1298 (0.81)	1666 (0.36)	356 (0.07)	1021 (0.2)
Hyponatremia	19754 (4.65)	7824 (5.95)	33518 (9.82)	24043 (5.65)	23304 (5.67)
Hypotension	380 (0.08)	1153 (0.74)	636 (0.14)	230 (0.05)	683 (0.14)
Hypothyroidism	297 (0.06)	642 (0.4)	1056 (0.23)	162 (0.03)	333 (0.07)
Insomnia	3046 (0.62)	2086 (1.38)	2468 (0.53)	3049 (0.62)	4114 (0.85)
Ischemic stroke all inpatient	3120 (0.64)	1824 (1.19)	2655 (0.57)	2775 (0.56)	4139 (0.85)
Nausea	2722 (0.56)	4071 (2.77)	4033 (0.89)	4368 (0.9)	5846 (1.22)
Open angle glaucoma	6117 (1.33)	3853 (2.83)	5374 (1.22)	8786 (2.03)	9943 (2.33)
Seizure	184 (0.04)	67 (0.04)	307 (0.07)	94 (0.02)	199 (0.04)
Suicide and suicidal ideation	10221 (2.13)	993 (0.62)	21518 (5.09)	9957 (2.1)	8063 (1.67)
Tinnitus	2628 (0.53)	4276 (2.87)	5082 (1.12)	6920 (1.44)	7643 (1.62)

Ventricular arrhythmia and sudden cardiac death	20806 (4.91)	6846 (5.12)	27233 (6.92)	23655 (5.6)	23772 (5.89)
Vertigo	2577 (0.53)	748 (0.47)	2269 (0.49)	2341 (0.48)	2782 (0.57)

Table 3 - characteristics of the target population (patients with depression initiating treatment) per database

	CCAЕ	MDCD	MDCR	Optum Claims	Optum EHR
Mean Age	41	35	75	50	49
Male (%)	30.8	25.9	32.2	31.7	29.2
Mean number outpatient visits in prior year	16.3	31.2	26.8	16.6	32.4
Frequency of patients experiencing condition in prior year:					
Pain	0.60	0.74	0.74	0.66	0.57
Anxiety	0.41	0.50	0.28	0.42	0.43
Acute inflammatory disease	0.32	0.36	0.24	0.31	0.18
Neoplastic disease	0.22	0.14	0.46	0.27	0.17
Essential hypertension	0.25	0.31	0.69	0.40	0.37
Obesity	0.11	0.19	0.11	0.13	0.17
Heart disease	0.09	0.14	0.46	0.20	0.18
Diabetes mellitus	0.09	0.14	0.27	0.16	0.16
Urinary tract infectious disease	0.09	0.14	0.16	0.12	0.07
Anemia	0.07	0.12	0.20	0.12	0.11

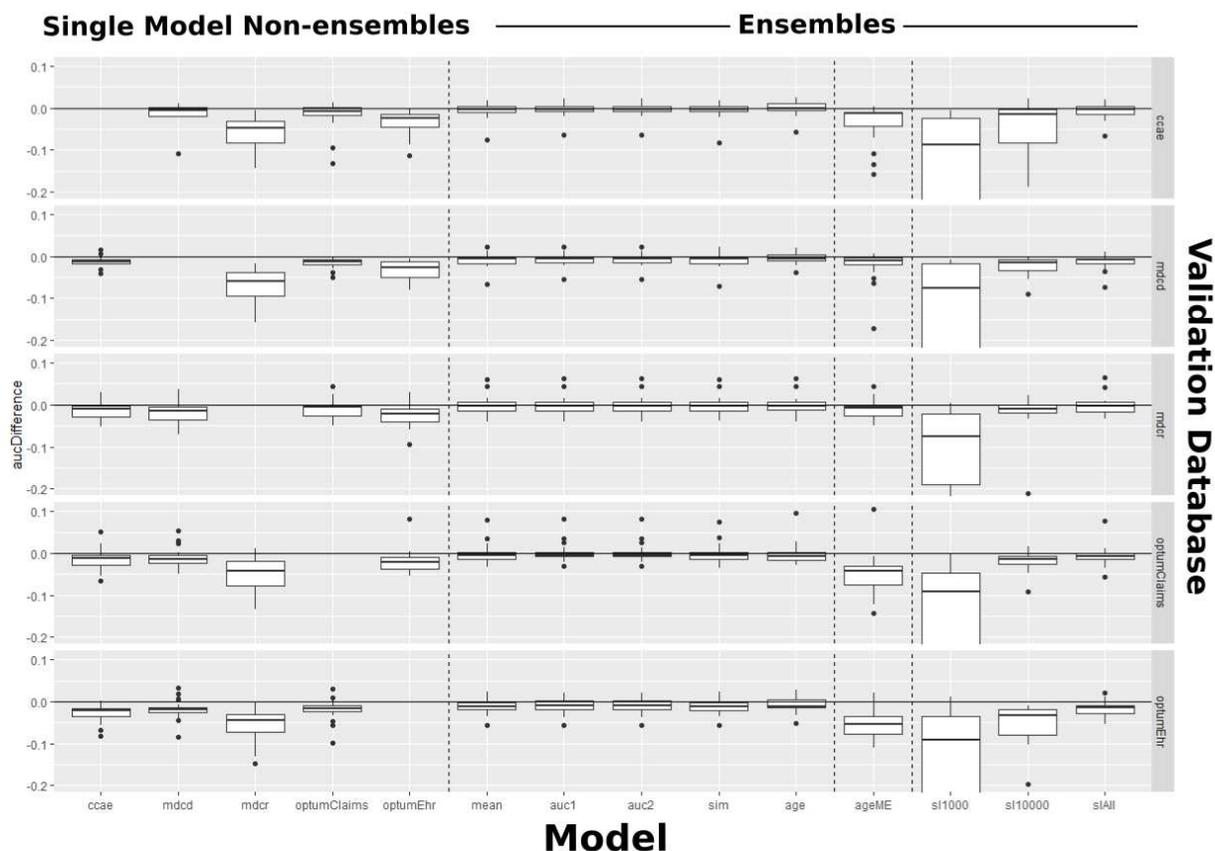


Figure 2 Box plots showing the difference between the external validation AUROC minus the internal validation AUROC per non-ensemble (Level 1 model) and ensemble method (Level 2 model) across the five databases. The rows represent the external database (the database that was excluded from the model/ensemble development) that was used to fairly evaluate the models/ensembles. The x-axis represents the model/ensemble technique. Box plots centered around 0 with a small range indicate highly transportable and consistent external discriminative performance. The dashed vertical lines separate the non-ensembles, the fusion ensembles, the mixture of expert ensembles and the stacking ensembles.

Figure 2 presents box plots of the $AUROC_{differences}$ per Level 1 model (non-ensemble) and Level 2 model (ensemble) when transported to each held out database across the 21 outcomes. The non-ensemble box plots show a lower median value and greater range of values compared to the fusion ensembles. The fusion ensembles achieved discriminative performances similar to the ‘hypothetical optimum’ when transported to new databases ($AUROC_{difference}$ values close to 0). The age-based mixture of expert and stacking ensembles that used 1,000 or 10,000 labels generally performed worse than the non-ensembles in terms of discrimination when transported. The stacking ensemble using all the labelled data available achieved external AUROC similar to the ‘hypothetical optimum’ but was not better than the fusion ensembles. The full external validation discrimination performance across the 21 outcomes and 5 databases for the non-ensembles and ensembles are presented in Appendix A.



Figure 3 - The calibration-in-the-large (mean predicted risk) for each non-ensemble (Level 1 model) and ensemble (Level 2 model) when externally validated. The rows represent the prediction problem (different outcomes) and the columns represent the external validation databases. The x-axis represents the different models/ensemble techniques. The solid horizontal line is the observed population risk in the external database. The dashed vertical lines divide the non-ensembles, the weighted ensembles, the mixture of expert ensemble and the stacker ensembles. A model is well calibrated when externally transported if the top of the bar is near to the solid horizontal line.



Figure 4. The calibration-in-the-large (mean predicted risk) for each non-ensemble (Level 1 model) and ensemble (Level 2 model) when externally validated. The solid horizontal line is the observed population risk in the external database. The dashed vertical lines divide the non-ensembles, the weighted ensembles, the mixture of expert ensemble and the stacker ensembles.

Each calibration in the large (the mean predicted risk) is presented in Figure 3 and Figure 4. The calibration in the large plots show the mean predicted risk per Level 2 model (ensemble) or Level 1 model (non-ensemble) and the dashed horizontal line is the observed population risk. A model is well calibrated if the mean predicted risk matches the observed population risk. Figures 3-4 show that the mean predicted risks did not often match the observed population risk, except for the stacking ensemble.

DISCUSSION

The results show that weighted fusion ensembles that combine multiple prognostic models developed in different databases appear to have more stable discriminative performances when transported to new databases compared to the Level 1 (single database) models. However, calibration appears to be an issue for all models that are transported to new databases (except stacking ensembles with sufficient labels).

This study showed that certain ensembles combining models developed independently across different databases transport better than the Level 1 single database models. The weighted fusion ensembles and stacking ensemble (that used all data) consistently achieved discrimination close to the ‘hypothetical optimum’ in the new data whereas the Level 1 single models generally performed slightly worse than the ‘hypothetical optimum’. The Level 1 single database models

were also less consistent across outcomes and certain database models did better than others (e.g., Optum claims models transported better than MDCR models). This variability may be due to each database containing diverse patient case-mixes, as seen in Table 3. The ensembles can combine the perspectives of the Level 1 models trained with different populations making them more robust to new populations. The calibrations of the transported models were generally poor, except the stacking model (using all data) as this used labelled data so was effectively recalibrated. If all the Level 1 single database models are mis-calibrated, then it makes sense that any ensemble combining them would also be mis-calibrated. This highlights the importance of model recalibrating before implementing them in new patient populations. It may be possible to recalibrate without labelled data by changing the intercept based on how common the outcome is in the target population the model is being applied to. If labels are available for some patients, then standard recalibration techniques can be implemented.

The results show the type of ensemble heuristic impacted transportability. The ensembles that performed the best in terms of discrimination when transported were the weighted fusion ensembles. The stacking ensemble did almost as well as the weighted fusion ensemble when there were sufficient labels, but it required labels in the new data it is being transported to whereas the weighted fusion ensembles did not. Requiring labels is a big disadvantage and therefore the weighted fusion ensembles are more useful. Interestingly, the simple mean fusion ensemble (uniform weighting) was comparable to the AUROC, age and database similarity weighted ensembles. Due to its simplicity, the mean fusion ensemble shows promise at being able to lead to more transportable prognostic models. If it worth noting, the age weighted ensembles may have benefitted in this study by the databases being similar (mostly US claims databases). For example, Optum claims is a mixture of patients that are similar to the patients in CCAE and MDCR, Therefore the age weighting may not perform well when the databases are more diverse. The weighted fusion ensembles and mixture of expert ensemble may have been impacted by the outcome rate differing between the databases. If the outcome is more common in a database, then a logistic regression model's intercept is likely to be greater and the model's mean predicted risk is likely to be higher than a model trained in data with fewer outcomes. This effectively may add more weighting to Level 1 models trained in databases that have a higher outcome percentage in the data.

The key advantage of this study is that we were able to compare the transportability of Level 1 models (developed in a single database) and ensembles combining Level 1 models developed in different databases across many prediction problems and across five datasets. In total we trained 21 (outcomes) x 5 (databases) single database models and created 21 (outcomes) x 5 (databases) x 7 (ensemble methods) ensemble models. The limitation of this study is the generalizability of findings as we only investigated one target population and we only used US data. In future work it would be useful to repeat this experiment across different target populations and externally validate the models (ensemble/non-ensemble) developed in this study across non-US databases. The OHDSI network and collaboration could be used to scale up this study across more diverse databases in future work [19]. In addition, there are numerous ways to combine the Level 1 models into an ensemble and we only investigated 7 simple approaches. However, these results provide a benchmark for comparing other ensembles techniques.

In this study 500,000 patients were sampled from each database as this provided a range of outcome sizes for the 21 outcomes investigated and enabled us to investigate the impact of outcome count in the study. Predicting rare outcomes is often an area of interest in healthcare and this may be where learning across multiple databases is more advantageous.

In future work it would be interesting to investigate whether rescaling the Level 1 models' predictions within the ensemble, to make the mean predicted risk for each Level 1 model within the ensemble equal, could improve the weighted fusion or mixture of expert ensembles. In addition, in this study we did not investigate pruning the Level 1 models within the ensembles, but this is an area of future research that may further improve transportability of an ensemble. In this study none of the Level 1 single database models achieved an AUROC ~ 0.5 , but it may make sense to prune such models if the situation arises.

CONCLUSION

In this study we performed a large-scale empirical evaluation to investigate the transportability of a simple and feasible distributed learning approach that uses ensemble learning to combine models developed independently in different databases. The results show that a mean fusion ensemble appears to transport to new data with higher discrimination compared to models developed in any single database. Consequently, developing a mean fusion ensemble of prognostic models developed using different databases (but for the same task) may lead to more clinically robust and useful prognostic models. However, recalibration is likely to be required.

Abbreviations

AUROC – Area under the receiver operating characteristic curve
CCAE - IBM Commercial Claims and Encounters (database)
CDM - common data model
EHR – electronic healthcare records
LASSO - least absolute shrinkage and selection operator
MDCD - IBM Medicaid (database)
MDCR - IBM Medicare Supplemental Beneficiaries (database)
OHDSI – Observational Health Data Science and Informatics
OMOP – Observational medical outcomes partnership
US – United States
UK – United Kingdom

Declarations

Ethics approval and consent to participate

All patient data included in this study were deidentified. The New England Institutional Review Board determined that studies conducted in Optum/IBM data are exempt from study-specific IRB review, as these studies do not qualify as human subjects research. No experiments were

conducted on humans in this study. The research methods were conducted in accordance with appropriate guidelines.

Consent for publication

Not applicable

Availability of data and materials

The data that support the findings of this study are available from Optum and IBM but restrictions apply to the availability of these data, which were used under license for the current study, and so are not publicly available. The IBM CCAE, MDCD and MDCR data that support the findings of this study are available from IBM MarketScan Research Databases (contact at: <https://www.ibm.com/products/marketscan-research-databases/databases>). The Optum data that support the findings of this study are available from Optum (contact at: <https://www.optum.com/business/solutions/life-sciences/real-world-data.html>).

Competing interests

JMR, MJS, PBR are employees of Janssen Research and Development and shareholders of Johnson and Johnson. PRR and RDW work for a research group who received unconditional research grants from Boehringer-Ingelheim, GSK, Janssen Research & Development, Novartis, Pfizer, Yamanouchi, Servier. None of these grants result in a conflict of interest to the content of this paper.

Funding

PRR and RDW received funding from the European Health Data and Evidence Network (EHDEN) project of the Innovative Medicines Initiative 2 Joint Undertaking (JU) under grant agreement No 806968. The JU receives support from the European Union's Horizon 2020 research and innovation programme and EFPIA.

Authors' contributions

JMR, RDW and MJS contributed to the conception and design of the study. JMR and PBR contributed to the acquisition of data. JMR, RDW, MJS, PBR and PRR contributed to the analysis and interpretation of data. All authors contributed to drafting the article and revising it critically for important intellectual content. All authors contributed to the final approval of the version to be submitted.

Acknowledgements

None

REFERENCES

1. Farzadfar, F. Cardiovascular disease risk prediction models: challenges and perspectives. *Lancet Glob Health* 2019;7(10):e1288-e1289.

2. Reps JM, Schuemie MJ, Suchard MA, et al. Design and implementation of a standardized framework to generate and evaluate patient-level prediction models using observational healthcare data. *J Am Med Inform Assoc* 2018;25(8):969-75.
3. Debray TP, Vergouwe Y, Koffijberg H, et al. A new framework to enhance the interpretation of external validation studies of clinical prediction models. *J Clin Epidemiol* 2015;68(3):279-89.
4. Vergouwe, Y., Moons, K.G. and Steyerberg, E.W., External validity of risk models: use of benchmark values to disentangle a case-mix effect from incorrect coefficients. *Am J Epidemiol* 2010;172(8):971-980.
5. Jochems, A., Deist, T.M., Van Soest, J., Eble, M., Bulens, P., Coucke, P., Dries, W., Lambin, P. and Dekker, A., Distributed learning: developing a predictive model based on data from multiple hospitals without data leaving the hospital—a real life proof of concept. *Radiotherapy and Oncology* 2016;121(3):459-467.
6. Bogowicz, M., Jochems, A., Deist, T.M., Tanadini-Lang, S., Huang, S.H., Chan, B., Waldron, J.N., Bratman, S., O'Sullivan, B., Riesterer, O. and Studer, G., Privacy-preserving distributed learning of radiomics to predict overall survival and HPV status in head and neck cancer. *Scientific reports* 2020;10(1):1-10.
7. Deist, T.M., Dankers, F.J., Ojha, P., Marshall, M.S., Janssen, T., Faivre-Finn, C., Masciocchi, C., Valentini, V., Wang, J., Chen, J. and Zhang, Z., Distributed learning on 20 000+ lung cancer patients—The Personal Health Train. *Radiotherapy and Oncology* 2020;144:189-200.
8. Luo, C., Islam, M.N., Sheils, N.E., Reps, J.M., Buresh, J., Duan, R., Tong, J.M., Edmondson, M., Schuemie, M.J. and Chen, Y., 2020. Lossless Distributed Linear Mixed Model with Application to Integration of Heterogeneous Healthcare Data. *medRxiv*.
9. Tsoumakas, G., Partalas, I. and Vlahavas, I. A taxonomy and short review of ensemble selection. In *Workshop on Supervised and Unsupervised Ensemble Methods and Their Applications* 2008;1-6.
10. Rokach, L. Ensemble-based classifiers. *Artif Intell Rev* 2010;33(1–2):1-39. [doi:10.1007/s10462-009-9124-7](https://doi.org/10.1007/s10462-009-9124-7).
11. Fumera, G. and Roli, F., Performance analysis and comparison of linear combiners for classifier fusion. In *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)* Springer, Berlin, Heidelberg. 2002:424-432.
12. Wolpert, D.H., Stacked generalization. *Neural networks* 1992;5(2):241-259.
13. Masoudnia, S. and Ebrahimpour, R. Mixture of experts: a literature survey. *Artif Intell Rev* 2014;42(2):275-293.
14. Breiman L. Random forests. *Mach Learn* 2001;45:5–32
15. Freund Y, Schapire RE. Experiments with a new boosting algorithm. *Machine learning: proceedings of the thirteenth international conference* 1996:325–332
16. Voss EA, Makadia R, Matcho A, et al. . Feasibility and utility of applications of the common data model to multiple, disparate observational health databases. *J Am Med Inform Assoc* 2015;223:553–564.
17. Suchard MA, Simpson SE, Zorych I, et al. . Massive parallelization of serial inference algorithms for complex generalized linear models. *ACM Transact Model Comput Simulation* 2013;231:10–32.
18. Van Calster B, Nieboer D, Vergouwe Y, et al. A calibration hierarchy for risk models was defined: from utopia to empirical data. *J Clin Epidemiol* 2016;74:167-76.
19. Hripcsak, G., Duke, J.D., Shah, N.H., et al. Observational Health Data Sciences and Informatics (OHDSI): opportunities for observational researchers. *Stud Health Technol Inform* 2015;216:574.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [AppendixA.docx](#)