A Distributed Ensemble Approach for Mining Healthcare Data under Privacy Constraints

Yan Li^a, Changxin Bai^a, Chandan K. Reddy^a

^a *Department of Computer Science, Wayne state University, Detroit, MI.
Corresponding author: red dy@cs.wayne.edu*

Abstract

In recent years, electronic health records (EHRs) have been widely adapted at many healthcare facilities in an attempt to improve the quality of patient care and increase the productivity and efficiency of healthcare delivery. These EHRs can accurately diagnose diseases if utilized appropriately. While the EHRs can potentially resolve many of the existing problems associated with disease diagnosis, one of the main obstacles in effectively using them is the patient privacy and sensitivity of the medical information available in the EHR. Due to these concerns, even if the EHRs are available for storage and retrieval purposes, sharing of the patient records between different healthcare facilities has become a major concern and has hampered some of the effective advantages of using EHRs. Due to this lack of data sharing, most of the facilities aim at building clinical decision support systems using limited amount of patient data from their own EHR systems to provide important diagnosis related decisions. It becomes quite infeasible for a newly established healthcare facility to build a robust decision making system due to the lack of sufficient patient records. However, to make effective decisions from clinical data, it is indispensable to have large amounts of data to train the decision models. In this regard, there are conflicting objectives of preserving patient privacy and having sufficient data for modeling and decision making. To handle such disparate goals, we develop two adaptive distributed privacy-preserving algorithms based on a distributed ensemble strategy. The basic idea of our approach is to build an elegant model for each participating facility to accurately learn the data distribution, and then can transfer the useful healthcare knowledge acquired on their data from these participators in the form of their own decision models without revealing and sharing the patient-level sensitive data, thus protecting patient privacy. We demonstrate that

our approach can successfully build accurate and robust prediction models, under privacy constraints, using the healthcare data collected from different geographical locations. We demonstrate the performance of our method using the Type-2 diabetes EHRs accumulated from multiple sources from all fifty states in the U.S. Our method was evaluated on diagnosing diabetes in the presence of insufficient number of patient records from certain regions without revealing the actual patient data from other regions. Using the proposed approach, we also discovered the important biomarkers, both universal and region-specific, and validated the selected biomarkers using the biomedical literature.

Keywords: Privacy-preserving data mining, Ensemble learning, Electronic health records, Boosting, Machine learning, Healthcare

1. Introduction

In recent years, electronic health records (EHRs) have been widely adapted at many healthcare facilities in an attempt to improve the quality of patient care and increase the productivity and efficiency of healthcare delivery. Typically, EHRs are generated and maintained within a particular healthcare facility, such as a hospital, clinic or physician's office. These EHRs can not only aid in various daily healthcare operations but can also help in accurately diagnosing diseases, if utilized appropriately [35].

While the EHRs can resolve many of the existing problems associated with disease diagnosis, one of the main obstacles in effectively using them is the patient privacy and sensitivity of the medical information available in the EHR. EHRs generally contain patient information about demographics, diagnostics, medications, living habits and other health-related information. It is needless to say that most of these information is extremely sensitive [34]. EHR systems must abide the Health Insurance Portability and Accountability Act (HIPAA) [6] to preserve the privacy of patient information. Thus, the public EHR products must be certified by certain institutes such as Certification Commission for Healthcare Information Technology (CCHIT), and Office of the National Coordination for Health Information Technology (ONC) [41]. Due to these legal concerns and medical ethics [22, 4], physicians are always keen about maintaining the highest possible standards while protecting patient privacy [39].

Due to these concerns, even if the EHRs are available for storage and

retrieval purposes, sharing of patient records between different healthcare facilities has become a major concern and has hampered some of the effective advantages of using EHRs. Due to this lack of data sharing, most of the facilities do not have any other option but to build their own clinical decision support systems with limited amount of patient data available in their own EHRs for important diagnosis related decisions. In addition, it becomes quite infeasible for a newly established healthcare facility, small hospital, or rural hospital to build a robust decision making system due to lack of sufficient patient records. However, to make effective decisions from clinical data, it is indispensable to have large amounts of data to train the decision making models. As a result, these small and/or rural hospitals are less motivated, and in 2011, only 20.8% of them were using EHR systems [15]. In this regard, it is clear that there is a conflicting objective of maintaining patient privacy and having sufficient data for modeling and decision making.

The problem statement of the work that is being developed in this paper is as follows. Given healthcare data collected from multiple facilities, how do we obtain a decision model that leverages the knowledge from all the facilities without revealing any patient specific information from any of the individual facilities. In other words, the goal is to develop a knowledge based data integration mechanism in a privacy-preserving context.

To deal with these challenges, we propose a privacy-preserving data mining framework based on horizontally distributed healthcare data. The goal of our work is to build an effective decision making system that can utilize the knowledge from multiple facilities (or geographical locations) without revealing any patient-level information [27]. In our approach, an elegant model for each participating facility to accurately learnt for approximating the local data distribution, and the useful healthcare knowledge acquired on such local data is then transferred in the form of their own decision models without revealing and sharing the sensitive data, thus, protecting the patient privacy. Transferring knowledge between multiple locations is a common practice with the goal of improving the prediction power [25]. Our approach can successfully build accurate and robust prediction models, under privacy constraints, with healthcare data collected from different geographical locations. Each participator shares its own local model with others and builds a specific integrated model based on its own specifications.

Merely trying to acquire the entire knowledge available from all the participators without carefully accounting for the distribution differences can potentially degrade the performance of the classifier [36]. While the overall

data distribution for a particular disease must be similar within different hospitals since we are dealing with the same disease, there will still be certain significant differences that arise due to the differences in the demographics of the patient population. Such distribution differences will play a vital role in building robust integrated decision making models that are specific to the population group. However, in the horizontal distribution privacy-preserving problem, there is no access to the original EHR data, one can not directly measure the data distribution differences among participators but can only approximately say how large the difference might be by analyzing the difference between the models trained by each participator.

Most of the state-of-the-art privacy-preserving data mining methods for horizontally distributed data are built based on a star network, where an untrusted third-party is needed [45, 14]. Each participator shares their statistical data distribution with a centralized agent, and this central agent is responsible for building the integrated decision model [27]. However, this framework suffers from two important issues: 1) participators need frequent information exchange with the central agent and hence the communication cost is high; 2) the decision model is built on the central agent whose aim is to provide a decision support for all participators, but it ignores the data distribution differences between the participators.

In addition, in the horizontal distributed privacy preserving data mining literature, there is no prior work on preventing “**negative impact during integration**”. In our work, as each participator will build a specific integrated model based on their own specifications, they can selectively decide which of the models from other participators can be used to build the integrated model. We build local prediction models to represent data, detect the data distribution difference, and transfer knowledge. In our work, the ADABOOST algorithm is chosen to be the local learner for each participator because it is a simple yet effective classifier which is extensively studied in the literature.

The main contributions of our work are summarized as follows:

- Propose an adaptive distributed privacy-preserving data mining technique based on an ensemble strategy which can successfully acquire knowledge from multiple healthcare facilities without gaining access to the sensitive patient data.
- Propose a new integration scheme which does not require a third-party agent, and each participator can build its own integrated model based

on its particular needs and can selectively prevent the “negative impact” during the integration process.

- Demonstrate the performance of the proposed distributed privacy-preserving ensemble method using type-2 diabetes patient records gathered from multiple sources from all the fifty states in the U.S. Evaluate our method on diagnosing diabetes in the presence of insufficient number of patient records from certain regions without revealing the actual patient data from other regions.
- Identify the important global and region-specific biomarkers for Type-2 diabetes and validate the selected biomarkers using the biomedical literature.

The rest of this paper is organized as follows: In section 2, we provide important relevant works on privacy-preserving data mining and also information about the Type-2 diabetes disease. In Section 3, we propose our adaptive distributed privacy-preserving ensemble method. Our experimental results along with the important biomarkers discovered using the proposed work are presented in Section 4. Finally, Section 5 concludes our discussion with some future research directions.

2. Related Work

In this section, we will provide the related literature in the field of privacy-preserving data mining and then discuss more about the Type-2 diabetes disease.

2.1. Privacy-Preserving Data Analysis

Recently, health data privacy has become an important area of research. As a result, privacy-preserving data mining has gained much more attention especially in the context of healthcare data analysis. The state-of-the-art methods for privacy-preserving data mining can be categorized under three types, namely, *randomization*, *k-anonymization*, and *distributed* privacy-preserving data mining [1]. In randomization approach, some techniques are used during the data collection process to induce perturbation in the selected attributes of the original data records for concealing certain sensitive information [3]. These techniques include swapping attribute values [17], principal component analysis (PCA) based techniques [21], adding

random components which follow a known probability distribution [2], etc. However, in some applications such as healthcare, randomization methods are not sufficient. It has been shown that by incorporating certain publicly available data, the hidden sensitive information in the processed records can be recovered and hence the privacy will be compromised. The k -anonymity procedure [37] was proposed as a solution for this indirect de-identification issue. By using generalization and suppression techniques, the k -anonymity method guarantees that any attribute value can be indistinguishably back-tracked to at least k records. In [8], k -Optimize algorithm has been developed to efficiently solve the problem of optimal k -anonymization in most cases.

The third category, which is the distributed privacy preservation, is closely related to secure distributed computation which is a research topic in cryptography [33]. In this scenario, non-fully trusting collaborators can jointly compute useful aggregate statistics over the entire dataset without sharing the original data. Thus the privacy of the individual dataset can be protected from various other participants. Based on the partitioning mechanism, distributed privacy-preserving data analysis can be grouped into two sub-categories: vertically partitioned distributed privacy preservation and horizontally partitioned distributed privacy preservation [13]. In the vertically partitioned distributed privacy preservation problem, each participant has a different set of attributes for the same group of data objects and the primary goal is to build a robust model that can leverage the entire set of attributes without revealing the individual details of the features to other participants. More detailed work about this approach can be found in [16]. In the horizontally partitioned distributed privacy preservation, each participant has different subset of records with the same set of attributes. The work proposed in this paper falls into this horizontally partitioned category where the patient data is being collected at multiple facilities and each patient record will consist of the same set of attributes. One of the first works in this category presented a strategy that extends the popular ID3 algorithm to two-party privacy preservation [26]. Then, there were a number of approaches proposed for horizontally partitioned distributed privacy preservation using various machine learning algorithms such as naive Bayes classifier [24], support vector machine [45], and ensemble based classifiers [19]. Amongst these methods, the ensemble-based approaches [7, 9, 12, 43, 42] have demonstrated good performance in terms of simultaneously preserving privacy and having good accuracy.

Gambs et al. [19] proposed MULTBOOST algorithm which is a boosting-

based privacy-preserving data mining method. However, the main problem with these above-mentioned models is that they need a third-party (central agent) that participates in the protocol to revise and aggregate the integrated final models. This also leads to an extra cost in processing. In addition, as each participator must have frequent contact with the third-party, a significant portion of the time is spent on the communication between the participators rather than the actual computation itself. Most importantly, those methods are not adaptable, i.e., a new integrated model has to be learned from scratch by repeating the entire learning process every time when a new participator is added. In this paper, we propose a distributed ensemble based horizontally partitioned privacy preservation algorithm which overcomes these limitations by making each participator compute their own local models independently from each other. Therefore, each participator can take advantage of the aggregate statistics (or a basic prediction model) without compromising their own privacy. Also the communication cost is substantially reduced because a third-party is not needed in our model, and the participators will have to only recompute the models for the newly integrated components once a new participator is added.

2.2. Type-2 Diabetes

Diabetes is the 7th leading cause of death in the United States [20] and the 8th leading cause of death in the world [44]. Around 29.1 million Americans (or 9.3% of the U.S. population) have diabetes [11]. In adults, type-2 diabetes (or insulin-dependent diabetes mellitus) accounts for approximately 90% to 95% of all diagnosed cases of diabetes. The risk for developing type-2 diabetes is mostly associated with older age, obesity, family history of diabetes, history of gestational diabetes, impaired glucose metabolism, physical inactivity, and race/ethnicity (African American, Alaska Native, American Indian, Asian American, Hispanic/Latino, or Pacific Islander American have a high prevalence of diabetes) [29, 11, 5].

There are several medical tests that can be performed to check if a person has diabetes or not. Currently, the American Diabetes Association (ADA) recommends four kinds of testing methods for diagnosing type-2 diabetes: 1) A1C test [30] which measures a person's average levels of blood glucose over a period of 3 months. 2) Fasting Plasma Glucose (FPG) test [31] which measures the blood glucose in a person who has fasted for more than 8 hours. 3) Oral Glucose Tolerance Test (OGTT) [28] measures the blood glucose after a person fasts for at least 8 hours and 2 hours after the person

drinks a liquid containing 75 grams of glucose dissolved in water. 4) Random Plasma Glucose (RPG) test [31] which is performed during a regular health checkup.

Though these medical diagnoses tests are relatively accurate (patients really have diabetes when these tests are positive), they typically have poor sensitivity (will miss a lot of cases). In 2012, 8.1 million (27.8% of 29.1 million) people with diabetes were undiagnosed [11]. In addition, in the data that we used in our study, only $1/8^{th}$ of the patients suffering from diabetes were diagnosed by those medical tests. In addition, such tests cause inconvenience to the patients. Hence, it is an important problem to diagnose these patients at early stages through their electronic health records (EHR) using data-driven methods.

To improve the diagnosis of type-2 diabetes, a machine learning community hosted a “Practice Fusion Diabetes Classification” challenge [23] in 2012 which aims at developing some advanced EHR based type-2 diabetes diagnostic support system. There were a total of 9,948 patients, and amongst them 1,904 patients suffer from diabetes. These patients come from all the 50 states and the District of Columbia in the U.S. and the Commonwealth of Puerto Rico. For each patient, the EHRs were comprehensively collected from the following primary sources of information [25]:

- **Demographic information** such as year of birth, gender, weight, and geographical location of each patient.
- **Diagnosis information** consists of the ICD9 Codes.
- **Allergy and immunization** consists of a list of allergies and vaccination records.
- **Laboratory information** consists of lab test observations for lab panels, and the lab test results received from lab facilities.
- **Medication and Prescription** consists of the medication history, where each medicine is identified by National Drug Code (NDC), and prescription records.
- **Patient smoking status** is an ordinal status variable maintained on a yearly basis.
- **Transcripts** consist of visited document records including allergy, medication, and diagnosis information.

We integrated all the information from 17 different sources and constructed a consolidated repository of diabetes EHR which contained 536 variables in total after combining all the data from multiple sources. It should be noted that, amongst all these features, only the gender and location are categorical in nature; the lab test results and personal information are real-valued attributes; certain status attributes such as smoking status, and emergency status have ordinal values. A significant majority of the remaining features are count variables.

3. Proposed work

In this section, we will explain the proposed framework for privacy-preserving data mining. Before describing the details of our proposed distributed ensemble based approach, we will first review the working of the standard ADABOOST algorithm [18]. The ADABOOST algorithm will be used as our local learner for each participator along with the decision stump which will be used as the weak learner. In our approach, we choose the ADABOOST algorithm because it is a simple, flexible and effective classifier, and it is constructed using a weighted combination of weak classifiers. Before going to the details of the proposed methods, we will first introduce the notations used in this paper in Table 1.

ADABOOST is a well-known high-performance ensemble learning method which iteratively generates a strong classifier from a pool of weak hypotheses. In each iteration, a base classifier is learned, and based on the correctness of the prediction in each iteration, the weights of the training examples are updated. Thus, in the subsequent iteration, the base classifiers will focus on the misclassified samples in the previous iteration. The final ensemble classifier is a weighted combination of these base classifiers, where for each base classifier, the weight depends on the corresponding error rate. In other words, a classifier with lower error rate gets higher weight. Simple learners such as decision stumps (decision trees with only two leaf nodes) often perform well for ADABOOST [10]. Typically, a decision stump can be expressed using three parameters: (i) the name of the attribute to be tested (fn), (ii) the test threshold (θ), and (iii) the sign of the test ($\delta \in \{+1, -1\}$). Note that a decision stump can only handle one attribute (feature) at a time and hence the weights of weak classifiers can also be used to measure the ability to distinguish the corresponding features conditional on the learning task. The final boosted model will be a linear combination of these independent decision

Table 1: Notations used in this paper.

Notation	Description
$n^{(p)}$	sampling size of p^{th} participator
$D_{n^{(p)}}^p$	dataset of p^{th} participator
X	data matrix
Y	corresponding labels $\in \{+1, -1\}$
fn	name of the attribute
θ	test threshold
δ	sign of the test $\in \{+1, -1\}$
$h^{p(t)}(\cdot)$	the weak classifier trained at t^{th} iteration in p^{th} participator
$\alpha^{p(t)}$	the weight of $h^{p(t)}(\cdot)$
$H^p(\cdot)$	the classifier trained in the p^{th} participator
ϵ_p	training error rate of p^{th} participator
$\epsilon_p^{(q)}$	testing error rate of q^{th} model in the p^{th} participator's training data
$S^{(p)}$	the index set of selected models for p^{th} participator
λ_p	instance proportion of p^{th} participator
σ	a factor of the integrated weight

stumps, and this independence would not require each participator to have exactly the same set of features which is another primary advantage of our proposed model. Finally, this will allow us to select the common important features among all participators, and each participator can take advantage of the useful aggregate model statistics without sharing the original data.

3.1. Adaptive Distributed Privacy Preserving using Boosting

In this section, we will discuss the proposed boosting based adaptive model for privacy-preserving data mining with horizontally partitioned data. In this case, each participant has different set of records with both common features and local unique attributes. Let $D_{n^{(1)}}^1, D_{n^{(2)}}^2, \dots, D_{n^{(M)}}^M$ denote the datasets of M participators, the dataset of p^{th} participator contains a total of $n^{(p)}$ samples, and it can be represented as

$$D_{n^{(p)}}^p = \{(X_1^p, y_1^p), (X_2^p, y_2^p), \dots, (X_{n^{(p)}}^p, y_{n^{(p)}}^p)\},$$

where each example $X_i^p = (x_{i1}^p, x_{i2}^p, \dots, x_{iK^{(p)}}^p)$ is a covariate vector with $K^{(p)}$ components and each label $y_i^p \in \{+1, -1\}$ for the binary decision making that

is being considered here. In our model, we employ ADABOOST algorithm to build an ensemble classifier for each participator, and these ensemble classifiers can be used to represent the data distribution of each participator. By sharing these local models with each other, all the participators can build their individual integrated model which takes advantage of other participators' knowledge without direct access to the datasets. Hence, the privacy of the individual datasets can be preserved.

Firstly, the standard ADABOOST algorithm is applied to each participator for T boosting iterations, and set of T weak classifiers will be learned. In the decision stump case, each weak classifier can be represented by the variable set (fn, θ, δ) , where fn denotes the name of the selected attribute, θ is the decision threshold of the attribute fn , and δ is the sign of the decision. For any instance X_i , the hypothesis $h(X_i)$, which is made by decision stump, is either $+1$ or -1 . For the p^{th} participator, the ensemble classifier $H^p(\cdot)$ is a set of T weak classifiers:

$$\{h^{p(1)}(\cdot)\alpha^{p(1)}, h^{p(2)}(\cdot)\alpha^{p(2)}, \dots, h^{p(T)}(\cdot)\alpha^{p(T)}\}$$

where $h^{p(t)}(\cdot)$ is the weak classifier at t^{th} iteration and $\alpha^{p(t)}$ is the corresponding weight of that weak classifier. For a particular test instance X_i , the binary prediction made by the p^{th} participator's local model can be defined as:

$$\mathcal{H}^p(X) = \text{sign} \{H^p(X)\} = \text{sign} \left\{ \sum_{t=0}^T \alpha^{p(t)} h^{p(t)}(X_i) \right\} \quad (1)$$

where $\text{sign} \{\cdot\}$ is the signum function. Once the local models are built, all the participators share their own models and sample sizes with each other, and hence each participator will receive $M - 1$ models from other participators. Then, each participator will build an integrated model independently based on their specific requirements. For example, some facilities might want to maximize the sensitivity of their models instead of the area under curve (AUC) value for their clinical decision support system. Using our ensemble based approach, they can use their own cost matrix to select the decision threshold for the integrated model that is being built in their own facility. The P2P network based distributed framework of the proposed ensemble system is shown in Figure 1, and we can see that no third-party (central agent) is needed to build an integrated model in our proposed framework.

As each participator will build their local model independently, when a new participator is added, it will learn a local model by itself and share

this trained model with other participators. Then, each existing participator will decide whether it needs to update the integrated model based on the new local model’s performance on their own data. We can see that this P2P network based framework is more adaptable and can provide more autonomy for participators than the star network based distributed privacy-preserving framework. Each participator integrates the final model independently, and hence the integrated model will take care of the specific needs of their own local data. In summary, having an individualized integrated model separately at each participator will tailor the model to their own local data. Only the integration component will have to be selectively computed again when a new participator is added, rather than re-training the entire model from scratch.

Figure 1: Block diagram representing various participators, the local and the integrated models built at each participator.

Each participator receives $M - 1$ models and another important issue needs to be resolved for a specific participator. That is, how to integrate those local models and its own model in an appropriate manner so that the integrated model performs better than the local model. Also, there is a slight risk that by integrating various unwanted (or poor) models, the performance of the integrated model might deteriorate. To prevent such “negative impact” during integration, when the model integration is performed for the p^{th} participator, first one has to decide which of the other local models can be used to build the final integrated model.

The basic idea here is to not include the models from the other participators whose data distribution is very different from the data distribution of this p^{th} participator. Since looking into the original data distribution of the other participators is infeasible due to patient privacy, we resort to applying all the other local models on the training set of p^{th} participator, $D_{n(p)}^p$, and

compare the error rate of each local model with the training error rate of p^{th} participator's trained model. Note that the models received from the other participators might include some decision stumps for their own unique local features. Hence, before using those models, the p^{th} participator should select a suitable subset of common decision stumps based on f_n . For the p^{th} participator, the error rate of q^{th} participator's model is given by

$$\epsilon_p^{(q)} = \frac{1}{n^{(p)}} \left[\sum_{i=1}^{n^{(p)}} I \left(\text{sign} \left\{ \hat{H}^q(X_i^p) \right\} \neq y_i^p \right) \right] \quad (2)$$

where $\hat{H}^q(\cdot)$ is the selected subset of decision stumps from $H^q(\cdot)$, the model trained by q^{th} participator, $I(\cdot)$ is the indicator function. The training error rate of the p^{th} participator's trained model is given by

$$\epsilon_p = \frac{1}{n^{(p)}} \left[\sum_{i=1}^{n^{(p)}} I \left(\text{sign} \left\{ H^p(X_i^p) \right\} \neq y_i^p \right) \right] \quad (3)$$

For every local model, we compute the difference between $\epsilon_p^{(q)}$ and ϵ_p . If $(\epsilon_p^{(q)} - \epsilon_p)$ is less than a certain threshold τ , then we can assume that the data distribution in p^{th} participator and q^{th} participator are similar, and we can use q^{th} participator's trained model to build the integrated model for p^{th} participator. The integrated model only uses other local models (not the other participators' data) and selects the ones with similar data distribution to that of the local data; thus, it indirectly strengthens the data distribution. By doing this, it avoids the overfitting problem in the presence of only a few samples in the local data. For our experiments, we set $\tau = 0.2$ which is a relatively a large threshold. This is because ϵ_p (the training error rate of ADABOOST) is very small; thus, $(\epsilon_p + 0.2)$ is a reasonable threshold for $\epsilon_p^{(q)}$ (the testing error rate of q^{th} model in the p^{th} participator's training data) as the q^{th} model is not trained from the training data of p^{th} participator.

Now that the model selection is complete, for p^{th} participator we get a index set of selected models ($S^{(p)}$) which will be used in the model integration. It should be noted that p also belongs to $S^{(p)}$. We will now explain the actual integration of the selected models. The model integration should not only take advantage of the aggregate statistics from other selected local models but also consider the specificity of the participator. To achieve this goal, we

establish the following three criteria for the proposed ensemble based model integration in the p^{th} participator:

1. More weight should be assigned to the participator's own local model compared to any other local model with same sample size. This is because of the fact that the data from other participators might follow a slightly different distribution and hence one cannot give more importance to the other local models. The assumption here is that by adding other local models, one can strengthen the local data distribution. This is important especially since the participator weights suffer from insufficient data.
2. The weight of the p^{th} participator's own local model is positively related to $n^{(p)}$. This is because the more the number of samples, the fewer the mistakes that are caused by insufficient sampling.
3. The weight of the p^{th} participator's own local model should have an upper bound. Otherwise, the weight of selected participator's model will become insignificant, and the final integrated model will be the same as p^{th} participator's own model and integration becomes unnecessary.

Considering these guidelines, the weight of each participator will be updated when building the integrated model for each participator. To illustrate this concept, let $\lambda_q \sigma_q^{(p)}$ denote the weight assigned to the q^{th} participator's model when we compute the integrated model for p^{th} participator, where $p \in \{1, 2, \dots, M\}$, $q \in S^{(p)}$, and $\lambda_q = \frac{n^{(q)}}{\sum_{s \in S^{(p)}} n^{(s)}}$ is the instance proportion of q^{th} participator. $\sigma_q^{(p)}$ is updated as follows:

$$\sigma_q^{(p)} = \begin{cases} \frac{\lambda_q}{\lambda_{max}^2} \cdot \lceil \frac{\lambda_{max}^2}{\lambda_{min}} \rceil & \text{if } q = p \\ 1 & \text{if } q \neq p \end{cases} \quad (4)$$

where $\lambda_{min} = \min(\lambda_q)$ is the minimum proportion among all the selected participators, $\lambda_{max} = \max(\lambda_q)$ is the maximum proportion, and $\lceil \cdot \rceil$ denotes the ceiling function. We will now prove that the proposed weight update mechanism $\lambda_q \sigma_q^{(p)}$ will satisfy all the three criteria mentioned earlier.

Theorem 1. *The weight update mechanism will assign more weight to p^{th} participator's own model compared to any other local models if they have the same sample size (criterion 1).*

Proof. If $q \neq p$ then $\sigma_q^{(p)} = 1$; that is to say the weight of other participator's local model is λ_q . Now let us consider the value of $\lambda_q \sigma_q^{(p)}$, the weight of p^{th}

participator's own model, i.e., when $q = p$. We note that the λ_q is bounded as follows. $\lambda_{min} \leq \lambda_q \leq \lambda_{max}$. We can see that the minimum possible value for $\sigma_p^{(p)}$ is obtained when $\lambda_p = \lambda_{min}$. Now, this value $\sigma_p^{(p)} = \frac{\lambda_{min}}{\lambda_{max}^2} \cdot \lceil \frac{\lambda_{max}^2}{\lambda_{min}} \rceil \geq 1 \Rightarrow \lambda_p \sigma_p^{(p)} \geq \lambda_p$. Thus, if $\lambda_p = \lambda_q$, we have $\lambda_p \sigma_p^{(p)} \geq \lambda_q \sigma_q^{(p)}$. This means that for any participator, in the model integration procedure, its own model will have more weight compared to any of the other selected participators once they have the same sample size. $\square$

Theorem 2. *The weight update mechanism will ensure that the weight of p^{th} participator's own model is positively related to $n^{(p)}$ (criterion 2).*

Proof. For p^{th} participator, since $S^{(p)}$ will be a constant set after model selection, λ_{min} and λ_{max} are also constant values. When $q = p$, $\lambda_p \sigma_p^{(p)} = \frac{\lambda_p^2}{\lambda_{max}^2} \cdot \lceil \frac{\lambda_{max}^2}{\lambda_{min}} \rceil$ is the weight of p^{th} participator's own model and contains only one variable $\lambda_p = \frac{n^{(p)}}{\sum_{s \in S^{(p)}} n^{(s)}}$, which is proportional to $n^{(p)}$ since the denominator is just a normalization factor. So, $\lambda_p \sigma_p^{(p)}$ is positively related to $n^{(p)}$. $\square$

Theorem 3. *The weight update mechanism will ensure that the weight of other participator's local model will not become insignificant thus making the model integration unnecessary (criterion 3).*

Proof. When $q = p$, $\lambda_p \sigma_p^{(p)}$, is the weight of p^{th} participator's own local model. We note that the λ_p is bounded as follows: $\lambda_{min} \leq \lambda_p \leq \lambda_{max}$. We observe that the maximum possible value for $\lambda_p \sigma_p^{(p)} = \lceil \frac{\lambda_{max}^2}{\lambda_{min}} \rceil$ is obtained when $\lambda_p = \lambda_{max}$. This limitation ensures that the weight of other participators' high-performance weak learner will not become insignificant. $\square$

Algorithm 1 outlines our proposed BOPPID (BOosting-based Privacy-Preserving Integration of Distributed data) algorithm. First, each participator employs standard ADABOOST algorithm to train their own local model and obtain the training error (lines 2-3). Then, they share their trained models with other participators (line 4). After the model sharing process, each participator will independently build its own integration model. In lines 7-13, each participator applies all the other local models to its own training data and selects a subset of "good" models from other participators for the subsequent integration process. The sample size proportion of the selected local models is calculated in line 14. Finally, the integration model is built using Eq. (4) (lines 16-23).

ALGORITHM 1: BOosting-based Privacy-Preserving Integration of Distributed data (BOPPID)

Input: The training sets of M participators $\{D_{n(1)}^1, \dots, D_{n(M)}^M\}$,
Number of boosting iterations (T), and

Threshold for error rate increase (τ).

Output: The set of M final classifiers $\{H^{1*}(\cdot), \dots, H^{M*}(\cdot)\}$.

```

1 foreach  $p$  in  $M$  participators parallel do
2    $H^p(\cdot) \leftarrow \text{ADABOOST}(D_{n(p)}^p, T);$ 
3    $\epsilon_p \leftarrow \frac{1}{n(p)} \left[ \sum_{i=1}^{n(p)} I(\text{sign}\{H^p(X_i^p)\} \neq y_i^p) \right];$ 
4   Share  $H^p(\cdot)$  with all other participators
5 end
6 foreach  $p$  in  $M$  participators parallel do
7   Initialize the Index set of selected participators  $S \leftarrow \{p\};$ 
8   for  $q \leftarrow 1$  to  $M$  do
9      $\epsilon_p^{(q)} \leftarrow \frac{1}{n(p)} \left[ \sum_{i=1}^{n(p)} I(\text{sign}\{\hat{H}^q(X_i^p)\} \neq y_i^p) \right];$ 
10    if  $(\epsilon_p^{(q)} - \epsilon_p) \leq \tau$  then
11      Add  $q$  in  $S^{(p)}$ 
12    end
13  end
14   $\Lambda = \left\{ \lambda_q = \frac{n(q)}{\sum_{s \in S} n(s)} \mid q \in S \right\};$ 
15  Find  $\lambda_{max}, \lambda_{min};$ 
16  foreach  $q$  in  $S$  do
17    if  $p = q$  then
18       $\sigma_q^{(p)} = \frac{\lambda_q}{\lambda_{max}^2} \cdot \lceil \frac{\lambda_{max}^2}{\lambda_{min}} \rceil;$ 
19    else
20       $\sigma_q^{(p)} = 1;$ 
21    end
22  end
23  return  $H^{p*}(\cdot) = \sum_{q \in S} \sum_{t=1}^T h^{q(t)}(\cdot) \cdot \alpha^{q(t)} \sigma_q^{(p)} \lambda_q$ 
24 end

```

3.2. Complexity Analysis

The computational complexity of proposed BOPPID depends on the ADABOOST algorithm in line 2. For p^{th} participator, the overall cost of ADABOOST in all T iterations is $\Theta(K^{(p)}(T + \log n^{(p)}))$ if the weak learner is a decision stump, and the training error rate can be calculated in $\Theta(n^{(p)})$. In the integration procedure for each participator, the cost of selecting a subset of the beneficial local models is $\Theta(Mn^{(p)})$. In the worst case scenario, where all M participators are selected, the computational cost of updating the weights for all participators is $\Theta(M)$ and generating the final classifier $H^q(\cdot)$ takes $\Theta(MT)$ time. Thus, the total computational complexity of the p^{th} participator is $\Theta(K^{(p)}(T + \log n^{(p)}) + MT + Mn^{(p)})$. As each participator trains its local model and does the final integration independent from other participators, it is possible to perform these computations in parallel. In addition, note that the model integration can only start after all the participators complete the computation of their own local models, so the computational complexity of the proposed model depends on the participator with the largest number of samples. Hence, the total computational cost of BOPPID is $\Theta(K^{max}(T + \log n^{max}) + MT + Mn^{max})$, here n^{max} and K^{max} stands for the largest sample size and largest feature dimensions among all the M participators, respectively.

For the communication cost, let the network latency for transferring a model is f and the communication cost for sending and receiving a model are g and h respectively. For M participators, each of them will send their own model to $M - 1$ other participators, and receive $M - 1$ models from other participators. Hence, the total communication cost for BOPPID is $(M - 1)(f + g + h)$.

Now let us consider the scenario where a new participator is added, and the number of participators grows from M to $M + 1$. In our model, only the integration part has to be recomputed, and hence the computational cost for the existing M participators is $\Theta((M + 1)T + (M + 1)n^{max})$, and for the new participator, the computation cost is $\Theta(K^{max}(T + \log n^{new}) + MT + Mn^{max})$. Here n^{new} and K^{max} corresponds to the sample size and feature dimension of the new participator, respectively. Because all of the existing M participators will send their own models to the new participator and receive a new model from the new participator, the communication cost for the existing M participators is $(f + g + h)$ and for the new participator, the communication cost is $M(f + g + h)$. However, if a new participator is added in the existing methods such as MULTBOOST, a new integrated model

has to be retrained from the beginning. Therefore, the computation cost for all $M + 1$ participators is $\Theta(K^{(p)}(T + \log(\sum_{p=1}^{M+1} n^{(p)})) + MT)$ [19] and communication cost becomes $2T(f + g + h)$. Hence, we can clearly observe that our proposed model is more efficient when a new participator is added (which is a more practical scenario).

3.3. Adaptive Variant of Parallel Boosting

To strengthen our work, we also propose a new variant of an existing parallel boosting algorithm and make it applicable to the context of mining EHR data. In our previous work on parallel boosting [32], we developed a parallel variant of the boosting algorithm called ADABOOST.PL. Though originally designed for improving the computational efficiency, the overall layout of ADABOOST.PL can fit into the distributed privacy-preserving setting as it does not directly communicate the data from one participator to the others. The ADABOOST.PL algorithm follows the *MapReduce* workflows, where the original dataset is equally partitioned into M parts and mapped into M workers to learn M models which will then be reduced (integrated) into a final model. For each worker, its weak classifiers will be sorted with increasing order of $\alpha^{p^{(t)}}$ values (line 3 in Algorithm 2). The basic intuition of sorting the workers' weak classifiers with respect to their weights is to place the classifiers with similar correctness at the same sorted level. This is a critical component of the ADABOOST.PL since this will ensure that like-minded classifiers will be merged during each boosting iteration. After sorting, $h^{p^{*(t)}}(\cdot)$ is the weak learner of p^{th} participator at the t^{th} iteration, and $\alpha^{p^{*(t)}}$ denotes the corresponding weight. In this paper, we generalize the idea of ADABOOST.PL and propose a new variant, ADABOOST.PL.V2, which can handle the diversity of sample size of each participator in the privacy-preserving context. It should be noted that similar to ADABOOST.PL, the proposed ADABOOST.PL.V2 is also built on a star network, which is shown in Figure 2. Most of state-of-the-art horizontally distributed privacy-preserving methods are also built on this star network, so that we can see to it that there is only one integration model that is being built for all hospitals.

Thus, compared with BOPPID, ADABOOST.PL.V2 is less adaptive and cannot preserve the data characteristics of each participator. ADABOOST.PL.V2 discards the “map” part of the ADABOOST.PL because the datasets are distributed and stored by each participator, and α^t is the weighted average, (instead of the standard average used in ADABOOST.PL) of $\alpha^{p^{*(t)}}$ for all

possible p . The weights are based on the sample size proportion (line 9 in Algorithm 2).

Figure 2: The framework of star network based horizontal distributed privacy-preserving method.

The merged classifier, $h^{(t)}(\cdot)$ is a *ternary classifier*, a variant of weak classifier, which can return ‘+1’, ‘−1’, and ‘0’ as a way of abstaining from answering [38]. It is built by taking a simple majority vote among weak classifiers of the worker as follows:

$$h^{(t)}(X_i) = \begin{cases} \text{sign} \left(\sum_{p=1}^M h^{p* (t)} (X_i) \right) & \text{if } \sum_{p=1}^M h^{p* (t)} (X_i) \neq 0 \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

The ternary classifier will answer ‘0’ if equal number of positive and negative predictions are made by the workers’ weak classifiers. Otherwise, it will answer the majority prediction. In line 8, the weight of the ternary classifier are obtained by the weighted average of the corresponding classifier weights. Once all the ternary classifiers for T rounds are generated, the algorithm returns their weighted combination as the final classifier.

4. Experimental Results

In this section, we demonstrate the performance of the proposed approach using real-world EHRs of diabetes patients. We first present the clinical feature transformation performed on the EHR data. We will then compare the performance of our proposed model against the state-of-the-art distributed privacy preserving ensemble prediction models. In addition, we

ALGORITHM 2: ADABOOST.PL.V2

Input: Training sets of M participators

$\{D_{n(1)}^1, D_{n(2)}^2, \dots, D_{n(M)}^M\}$, and

Number of boosting iterations (T).

Output: The final classifiers $H(\cdot)$

```
1 foreach  $p$  in  $M$  participators parallel do
2    $H^p(\cdot) \leftarrow \text{ADABOOST}(D_{n(p)}^p, T);$ 
3    $H^{p*}(\cdot) \leftarrow$  the weak classifiers in  $H^p(\cdot)$  sorted w.r.t.  $\alpha^{p(t)}$ ;
4   Send  $H^{p*}(\cdot)$  to central agent;
5 end
6 Initialize  $\Lambda = \left\{ \lambda_p = \frac{n^{(p)}}{\sum_{m=1}^M n^{(m)}} \mid p = 1, 2, \dots, M \right\};$ 
7 for  $t \leftarrow 1$  to  $T$  do
8    $h^{(t)}(\cdot) \leftarrow \text{MERGE}(h^{1*(t)}(\cdot), \dots, h^{M*(t)}(\cdot));$ 
9    $\alpha^t \leftarrow \sum_{p=1}^M \alpha^{p*(t)} \lambda_p;$ 
10 end
11 return  $H(\cdot) = \sum_{t=1}^T \alpha^t h^{(t)}(\cdot)$ 
```

also perform a detailed study on the biomarkers selected by the proposed algorithm on this data. It shows that our proposed model can successfully select some important ‘universal’ and ‘region-specific’ biomarkers which can be used in the medical domain to improve the diagnosis of type-2 diabetes.

4.1. Clinical Feature Pre-Processing

We now explain the clinical feature transformation procedure which transforms the raw EHR data to a format that is more suitable for data analysis. The original data contains patient records for a total of 9,948 patients, and each patient can be identified using a specific serial number. In Figure 3, we show the feature creation procedure from EHR data by considering a sample patient. In this example, we use several records from different categories for a particular patient (with PatientGuid “0113E5B3-66A5-44DA-91BE-425EF133651E”).

We extracted all the features for distinct anthropometric variables present in the data. In this example, we consider height, weight, Body Mass Index (BMI), systolic blood pressure (SystolicBP), and diastolic blood pressure (DiastolicBP). To tackle the problem of multiple anthropometric values for

the same patient, we represent each anthropometric feature using a set of summary statistics (maximum, minimum, and median). Note that, in this example, some anthropometric features are missing. It should be noted that certain missing values (like height in this example) are easy to fill. For example, we can confidently replace ‘NULL’ by 61 and calculate the corresponding BMI by using $\frac{\text{Weight}(\text{lb})}{(\text{Height}(\text{in}))^2} \times 703$. However, other missing values like weight cannot be filled-in, and hence we leave them empty. “HighLowBP” is the difference between the median of systolic blood pressure and the median of diastolic blood pressure. Lab variables are also generated using a similar approach. In addition, we also created a variable which counts the number of times the lab was conducted for the patient. ICD-9 code is a list code for International Statistical Classification of Diseases, and each code presents a disease description. With the acute indicator, two binary variables can be created to reflect whether the patient has a special disease or not, and if the disease is present, whether it is acute or not (“L2_ ” represents acute and “L1_ ” represents not acute). In this example, the codes 278 and 724.5 represent obesity and backache, respectively. Because this patient suffered from serious backache only once, we give the value of 1 to “L2_ BackPain” feature. For the medication information, we created two variables for each distinct medication (according to NDC) given to the patient; one feature represents the number of times the individual medication was given to the patient, and the other feature represents the dose. There are several states corresponding to smoking status which can be represented using ordinal numbers, and we build two new features, one for the current smoking state and another for the smoking history. When the patient has never smoked in the past, the status of “0 cigarettes per day” is denoted by 0.

In summary, based on the generated features which are listed in tables with pale blue color, one can observe that following our transformation procedure not only helps in immensely reducing the dimensionality and complexity of the raw EHR data, but also summarizes the complex EHRs into a succinct representation which is then used for building prediction models.

4.2. Experimental Setup

We will demonstrate the performance of the proposed algorithms on improving the diagnosis of the diabetes condition. Hence, for our experiments, we selected the top 14 states based on the total patient population in each state (which will be considered as a participator). Table 2 shows the demographic statistics of these 14 selected states.

Figure 3: Clinical feature transformation for a single patient.

We compare our proposed adaptive distributed privacy-preserving ensemble system with the MULTBOOST [19] and local ADABOOST (LOCAL_Ada) algorithms. LOCAL_Ada algorithm is a local model obtained by applying the standard ADABOOST on each participator independently. For all the algorithms, the number of boosting iterations is set to 100. In Table 3, we provide the performance results of AUC values from different algorithms using 10-fold cross validation.

In Table 4, we present the comparison of F-measure values which is calculated as follows:

$$F_measure = \frac{2 \times Precision \times Sensitivity}{Sensitivity + Precision} \quad (6)$$

where $Sensitivity = \frac{TP}{TP+FN}$, and $Precision = \frac{TP}{TP+FP}$; therefore, a high value of F-measure indicates that both precision and sensitivity are reasonably high. It should be noted that the positive individuals are diabetics and the negative individuals are non-diabetics. Thus, here TP corresponds to the number of diabetics correctly predicted by the classifier, FN corresponds to

Table 2: Demographic statistics of the top 14 states based on patient population.

State	Total No. of patients	No. of Diabetes Patients	No. of Non-diabetes Patients	Prevalence Rate (%)
Arizona (AZ)	295	56	239	18.98%
California (CA)	1917	258	1659	13.46%
Florida (FL)	804	158	646	19.65%
Illinois (IL)	412	75	337	18.20%
Michigan (MI)	344	61	283	17.73%
Missouri (MO)	702	104	598	14.81%
New Jersey (NJ)	575	119	456	20.70%
Nevada (NV)	495	135	360	27.27%
New York (NY)	557	167	390	29.98%
Ohio (OH)	515	111	404	21.55%
Pennsylvania (PA)	250	67	183	26.80%
South Dakota (SD)	325	40	285	12.31%
Texas (TX)	897	243	654	27.09%
Virginia (VA)	484	79	405	16.32%

the number of diabetics wrongly predicted as non-diabetics, and FP is the number of nondiabetics wrongly predicted as diabetics.

We observe that for all the 14 states, our proposed algorithm can diagnose type-2 diabetes more accurately compared with the other algorithms. This demonstrates that we have successfully transferred useful medical knowledge and clinical information within all the participators under the privacy constraints without revealing the data itself from each participator. This also means that those participators can take advantage of the aggregated global models without compromising the privacy of individual patients.

Figure 4 shows the 10-fold cross validation result of weighted average AUC values at different iterations of the algorithm. The weighted average AUC is defined as:

$$AUC_{avg} = \sum_{p=1}^M AUC^{(p)} \frac{n^{(p)}}{\sum_{m=1}^M n^{(m)}} \quad (7)$$

where $AUC^{(p)}$ is the AUC value of p^{th} participator, $n^{(p)}$ is the number of patients in the p^{th} participator. The figure shows that the proposed algorithm

Table 3: Performance comparison of different algorithms based on the AUC values (along with their standard deviation)

State	LOCAL_Ada	MULTBOOST	ADABOOST.PL.V2	BOPPID
AZ	0.799 $\pm$ 0.103	0.831 $\pm$ 0.085	0.845 $\pm$ 0.096	0.865 $\pm$ 0.088
CA	0.895 $\pm$ 0.046	0.844 $\pm$ 0.060	0.870 $\pm$ 0.064	0.902 $\pm$ 0.049
FL	0.882 $\pm$ 0.065	0.861 $\pm$ 0.059	0.878 $\pm$ 0.056	0.899 $\pm$ 0.053
IL	0.849 $\pm$ 0.071	0.876 $\pm$ 0.076	0.885 $\pm$ 0.053	0.896 $\pm$ 0.046
MI	0.881 $\pm$ 0.065	0.926 $\pm$ 0.038	0.923 $\pm$ 0.037	0.934 $\pm$ 0.038
MO	0.914 $\pm$ 0.034	0.886 $\pm$ 0.052	0.902 $\pm$ 0.048	0.921 $\pm$ 0.040
NJ	0.802 $\pm$ 0.040	0.835 $\pm$ 0.051	0.847 $\pm$ 0.042	0.857 $\pm$ 0.048
NV	0.863 $\pm$ 0.045	0.877 $\pm$ 0.033	0.891 $\pm$ 0.029	0.901 $\pm$ 0.019
NY	0.858 $\pm$ 0.060	0.857 $\pm$ 0.038	0.873 $\pm$ 0.061	0.891 $\pm$ 0.056
OH	0.885 $\pm$ 0.041	0.879 $\pm$ 0.066	0.892 $\pm$ 0.053	0.900 $\pm$ 0.050
PA	0.816 $\pm$ 0.082	0.809 $\pm$ 0.117	0.873 $\pm$ 0.071	0.885 $\pm$ 0.071
SD	0.892 $\pm$ 0.072	0.920 $\pm$ 0.067	0.929 $\pm$ 0.056	0.937 $\pm$ 0.047
TX	0.854 $\pm$ 0.040	0.838 $\pm$ 0.038	0.845 $\pm$ 0.027	0.871 $\pm$ 0.027
VA	0.789 $\pm$ 0.056	0.827 $\pm$ 0.045	0.838 $\pm$ 0.072	0.863 $\pm$ 0.059

Table 4: Performance comparison of different algorithms based on the F-measure values (along with their standard deviation)

State	LOCAL_Ada	MULTBOOST	ADABOOST.PL.V2	BOPPID
AZ	0.541 $\pm$ 0.141	0.592 $\pm$ 0.084	0.612 $\pm$ 0.148	0.631 $\pm$ 0.151
CA	0.592 $\pm$ 0.092	0.498 $\pm$ 0.049	0.562 $\pm$ 0.079	0.608 $\pm$ 0.090
FL	0.664 $\pm$ 0.098	0.615 $\pm$ 0.085	0.643 $\pm$ 0.083	0.679 $\pm$ 0.092
IL	0.572 $\pm$ 0.102	0.643 $\pm$ 0.136	0.630 $\pm$ 0.110	0.659 $\pm$ 0.088
MI	0.619 $\pm$ 0.110	0.708 $\pm$ 0.087	0.707 $\pm$ 0.087	0.720 $\pm$ 0.091
MO	0.654 $\pm$ 0.069	0.608 $\pm$ 0.097	0.646 $\pm$ 0.094	0.667 $\pm$ 0.103
NJ	0.534 $\pm$ 0.059	0.583 $\pm$ 0.058	0.585 $\pm$ 0.056	0.589 $\pm$ 0.063
NV	0.668 $\pm$ 0.049	0.691 $\pm$ 0.058	0.705 $\pm$ 0.058	0.715 $\pm$ 0.043
NY	0.699 $\pm$ 0.086	0.682 $\pm$ 0.054	0.709 $\pm$ 0.099	0.731 $\pm$ 0.093
OH	0.657 $\pm$ 0.074	0.639 $\pm$ 0.104	0.688 $\pm$ 0.080	0.701 $\pm$ 0.084
PA	0.637 $\pm$ 0.090	0.623 $\pm$ 0.128	0.693 $\pm$ 0.092	0.721 $\pm$ 0.112
SD	0.666 $\pm$ 0.190	0.696 $\pm$ 0.159	0.691 $\pm$ 0.140	0.732 $\pm$ 0.107
TX	0.661 $\pm$ 0.065	0.628 $\pm$ 0.050	0.639 $\pm$ 0.035	0.675 $\pm$ 0.042
VA	0.493 $\pm$ 0.088	0.512 $\pm$ 0.075	0.560 $\pm$ 0.096	0.566 $\pm$ 0.079

Figure 4: Comparison of weighted average AUC values of Local AdaBoost, MULTBOOST, and BOPPID algorithms on the diabetes dataset for 100 iterations.

always works better than MULTBOOST and Local AdaBoost, and only a few selected important features (or iterations) can provide a good prediction model.

4.3. Biomarker Discovery

Biomarkers are important indicators used to diagnose a particular disease in a clinical setting. From Figure 4, we can see that the weighted average AUC of our proposed model achieves 0.88 after 10 iterations and approximately 0.89 after 20 iterations. This indicates that the top selected features in each state can well represent the corresponding data distribution. It should be mentioned that since the weak classifier used in our algorithm is a decision stump, each iteration will correspond to the usage of only one feature with a simple split. In Figure 5, we show the selected biomarkers for each state after 10 iterations (in the left) and 20 iterations (in the right). In the 10 iterations scenario, we find 39 important biomarkers, and in the 20 iterations case, this number becomes 60. For comparison, we present these two settings under a same coordinate system, where the vertical axis is the index of selected features and the horizontal axis corresponds to the list of states that contain this attribute in their top feature list. Figure 5 clearly shows the selected important features and the corresponding state.

From Figure 5, we can conclude that there are a total of 6 universal biomarkers: “YearOfBirth” which represents the age of patient, “High-LowBP” measures the difference between systolic blood pressure and di-

astolic blood pressure, “L2_HypertensionEssential” is a biomarker that reflects whether the patient has been diagnosed with hypertension or not, “L2_MixedHyperlipidemia” is an indicator that tells whether the patient has suffered with mixed hyperlipidemia or not, “TotDiagPerVisit” indicates the number of complications the patient had, and “medication_diag_NA” indicates whether the patient was previously diagnosed with type-2 diabetes through medical tests. These 6 selected features meet the well known common knowledge in the clinical domain. The clinical study suggests that hypertensive persons are more predisposed to the development of diabetes than normotensive persons. Moreover, up to 75% of cardiovascular disease (CVD) in diabetes is attributed to hypertension [40] and older people are more likely to suffer from type-2 diabetes [29]. Medical diagnoses tests are relatively accurate, i.e., most of the patients who were diagnosed with diabetes by the medical tests actually have diabetes. However, the sensitivity of these tests is low, i.e., a lot of diabetic patients remain undetected through these tests.

Our results can also reflect that overweight or obese is closely related to type-2 diabetes [29], but for different states it is represented in different ways. For most of the states the body mass index (BMI) (related to features 6-9) is more important than the direct body weight, while for few other states direct body weight plays an important role as well. It can be seen from our results, that our approach can provide a more accurate set of “region-specific” biomarkers in addition to the global markers. Such analysis on region-specific biomarkers in comparison to the overall biomarkers (which are typically known in the medical literature) can provide new medical insights into specific regional patient data which can thus lead to new discoveries in the clinical domain.

5. Conclusions and future work

Healthcare in the United States is currently undergoing a revolutionary transformation and EHRs are playing a vital role in improving the quality of patient care. However, due to the patient privacy and the sensitivity of the patient information that is being stored in the EHRs, there are considerable barriers to EHR proliferation. In this paper, we developed a novel distributed privacy-preserving integration method based on an ensemble based strategy for building robust prediction models from EHR data. Using our approach, each participator will need to reveal only limited model information

Figure 5: Selected biomarkers for each state.

built using the local data without revealing any patient-specific information. The local models built from each participator will be effectively combined to build an integrated model that is specific to each participator. This integrated model is built by taking advantage of aggregate knowledge from all the participators instead of the local participator’s data alone. Compared to other existing works, our proposed framework can prevent the “negative impact” during integration from multiple sources, which typically happens when the data distribution differences amongst the participators are very large, by building a specific integration model for each participator with only a few selected local models instead of taking all the local models. We demonstrated the performance of our proposed method using the Type-2 diabetes EHR data and have shown that constructing a global model which utilizes other local models performs better compared to the original local models built on each participator’s data separately. In addition, we also successfully identified some important universal/global and region-specific biomarkers.

In the future, we will try to find more accurate ways to measure the difference of data distributions among multiple participators under privacy constraints. We also plan to extend our work with the idea of k -anonymity, so that we can measure the data distributions more accurately without compromising on patient privacy.

Acknowledgements

This work was supported in part by the National Science Foundation grants IIS-1242304, IIS-1231742, IIS-1527827 and the National Institutes of Health grant R21CA175974.

Reference

- [1] C. C. Aggarwal and S. Y. Philip. *A general survey of privacy-preserving data mining models and algorithms*. Springer, 2008.
- [2] D. Agrawal and C. C. Aggarwal. On the design and quantification of privacy preserving data mining algorithms. In *Proceedings of the twentieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 247–255. ACM, 2001.
- [3] R. Agrawal and R. Srikant. Privacy-preserving data mining. *ACM Sigmod Record*, 29(2):439–450, 2000.

- [4] M. Alther and C. K. Reddy. Clinical decision support systems. In C. K. Reddy and C. C. Aggarwal, editors, *Healthcare Data Analytics*. Chapman and Hall/CRC Press, 2015.
- [5] American Diabetes Association and others. Diagnosis and classification of diabetes mellitus. *Diabetes care*, 31(Supplement 1):S55–S60, 2008.
- [6] An Act. Health insurance portability and accountability act of 1996. *Public Law*, 104:191, 1996.
- [7] D. Baumgartner and G. Serpen. Performance of global–local hybrid ensemble versus boosting and bagging ensembles. *International Journal of Machine Learning and Cybernetics*, 4(4):301–317, 2013.
- [8] R. J. Bayardo and R. Agrawal. Data privacy through optimal k-anonymization. In *Data Engineering, 2005. ICDE 2005. Proceedings. 21st International Conference on*, pages 217–228. IEEE, 2005.
- [9] M. Bhardwaj and V. Bhatnagar. Towards an optimally pruned classifier ensemble. *International Journal of Machine Learning and Cybernetics*, pages 1–20, 2014.
- [10] R. Caruana and A. Niculescu-Mizil. Data mining in metric space: an empirical analysis of supervised learning performance criteria. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 69–78. ACM, 2004.
- [11] Centers for Disease Control and Prevention and others. National diabetes statistics report: estimates of diabetes and its burden in the united states, 2014. *Atlanta, GA: US Department of Health and Human Services*, 2014.
- [12] I. T. Christou, G. Gekas, and A. Kyrikou. A classifier ensemble approach to the tv-viewer profile adaptation problem. *International Journal of Machine Learning and Cybernetics*, 3(4):313–326, 2012.
- [13] C. Clifton, M. Kantarcioglu, J. Vaidya, X. Lin, and M. Y. Zhu. Tools for privacy preserving distributed data mining. *ACM SIGKDD Explorations Newsletter*, 4(2):28–34, 2002.

- [14] C. Clifton, M. Kantarciolu, A. Doan, G. Schadow, J. Vaidya, A. Elmagarmid, and D. Suciu. Privacy-preserving data integration and sharing. In *Proceedings of the 9th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery*, pages 19–26. ACM, 2004.
- [15] C. M. DesRoches, C. Worzala, M. S. Joshi, P. D. Kralovec, and A. K. Jha. Small, nonteaching, and rural hospitals continue to be slow in adopting electronic health record systems. *Health Affairs*, 31(5):1092–1099, 2012.
- [16] C. Dwork and K. Nissim. Privacy-preserving datamining on vertically partitioned databases. In *Advances in Cryptology–CRYPTO 2004*, pages 528–544. Springer, 2004.
- [17] S. E. Fienberg and J. McIntyre. Data swapping: Variations on a theme by dalenius and reiss. In *Privacy in statistical databases*, pages 14–29. Springer, 2004.
- [18] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- [19] S. Gambs, B. Kégl, and E. Aïmeur. Privacy-preserving boosting. *Data Mining and Knowledge Discovery*, 14(1):131–170, 2007.
- [20] D. L. Hoyert, J. Xu, et al. Deaths: preliminary data for 2011. *National vital statistics reports*, 61(6):1–51, 2012.
- [21] Z. Huang, W. Du, and B. Chen. Deriving private information from randomized data. In *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, pages 37–48. ACM, 2005.
- [22] A. K. Jha, C. M. DesRoches, E. G. Campbell, K. Donelan, S. R. Rao, T. G. Ferris, A. Shields, S. Rosenbaum, and D. Blumenthal. Use of electronic health records in us hospitals. *New England Journal of Medicine*, 360(16):1628–1638, 2009.
- [23] kaggle.com. *Practice Fusion Diabetes Classification: Identify patients diagnosed with Type 2 Diabetes*. 2012. <https://www.kaggle.com/c/pf2012-diabetes>.

- [24] M. Kantarcioglu, J. Vaidya, and C. Clifton. Privacy preserving naive bayes classifier for horizontally partitioned data. In *IEEE ICDM workshop on privacy preserving data mining*, pages 3–9, 2003.
- [25] Y. Li, B. Vinzamuri, and C. K. Reddy. Constrained elastic net based knowledge transfer for healthcare information exchange. *Data Mining and Knowledge Discovery*, 29(4):1094–1112, 2015.
- [26] Y. Lindell and B. Pinkas. Privacy preserving data mining. In *Advances in Cryptology-CRYPTO 2000*, pages 36–54. Springer, 2000.
- [27] G. Mathew and Z. Obradovic. Distributed privacy-preserving decision support system for highly imbalanced clinical data. *ACM Transactions on Management Information Systems (TMIS)*, 4(3):12, 2013.
- [28] M. Matsuda and R. A. DeFronzo. Insulin sensitivity indices obtained from oral glucose tolerance testing: comparison with the euglycemic insulin clamp. *Diabetes care*, 22(9):1462–1470, 1999.
- [29] National Diabetes Information Clearinghouse. Am I at risk for type 2 diabetes? taking steps to lower your risk of getting diabetes, June 2012. NIH Publication No. 12-4805.
- [30] National Diabetes Information Clearinghouse. The A1C test and diabetes, March 2014. NIH Publication No. 147816.
- [31] National Diabetes Information Clearinghouse. Diagnosis of diabetes and prediabetes, June 2014. NIH Publication No. 144642.
- [32] I. Palit and C. K. Reddy. Scalable and parallel boosting with mapreduce. *Knowledge and Data Engineering, IEEE Transactions on*, 24(10):1904–1916, 2012.
- [33] B. Pinkas. Cryptographic techniques for privacy-preserving data mining. *ACM SIGKDD Explorations Newsletter*, 4(2):12–19, 2002.
- [34] R. Rahman and C. K. Reddy. Electronic health records: A survey. In C. K. Reddy and C. C. Aggarwal, editors, *Healthcare Data Analytics*. Chapman and Hall/CRC Press, 2015.
- [35] C. K. Reddy and C. C. Aggarwal. *Healthcare Data analytics*. CRC Press, 2015.

- [36] C. K. Reddy and Y. Li. A review of clinical prediction models. In C. K. Reddy and C. C. Aggarwal, editors, *Healthcare Data Analytics*. Chapman and Hall/CRC Press, 2015.
- [37] P. Samarati. Protecting respondents identities in microdata release. *Knowledge and Data Engineering, IEEE Transactions on*, 13(6):1010–1027, 2001.
- [38] R. E. Schapire and Y. Singer. Improved boosting algorithms using confidence-rated predictions. *Machine learning*, 37(3):297–336, 1999.
- [39] S. R. Simon, R. Kaushal, P. D. Cleary, C. A. Jenter, L. A. Volk, E. J. Orav, E. Burdick, E. G. Poon, and D. W. Bates. Physicians and electronic health records: a statewide survey. *Archives of Internal Medicine*, 167(5):507–512, 2007.
- [40] J. R. Sowers, M. Epstein, and E. D. Frohlich. Diabetes, hypertension, and cardiovascular disease an update. *Hypertension*, 37(4):1053–1059, 2001.
- [41] US Department of Health, Human Services and others. Health information technology: Initial set of standards, implementation specifications, and certification criteria for electronic health record technology. *Federal Register*, 75(8):13, 2010.
- [42] X. Wang, A. Chen, and H. Feng. Upper integral network with extreme learning mechanism. *Neurocomputing*, 74(16):2520–2525, 2011.
- [43] X.-Z. Wang, R. Wang, H.-M. Feng, and H.-C. Wang. A new approach to classifier fusion based on upper integral. *Cybernetics, IEEE Transactions on*, 44(5):620–635, 2014.
- [44] World Health Organization. The top 10 causes of death. <http://www.who.int/mediacentre/factsheets/fs310/en/>, 2012.
- [45] H. Yu, X. Jiang, and J. Vaidya. Privacy-preserving svm using nonlinear kernels on horizontally partitioned data. In *Proceedings of the 2006 ACM symposium on Applied computing*, pages 603–610. ACM, 2006.