

Machine Learning Interpretability in Diabetes Risk Assessment

Turker Berk Donmez^{*a}, Mustafa Kutlu^a, Chris Freeman^b

^a*Mechatronics Engineering Department, Sakarya University of Applied Sciences
Sakarya, 54600, Turkey*

^b*Electronics and Computer Science, University of Southampton, Southampton SO17
1BJ, UK*

Abstract

Diabetes continues to be a complicated and prevalent metabolic illness, providing a serious burden to public health. While machine learning approaches like extreme gradient boosting (XGBoost) provide intriguing options for diabetes prediction, their 'black-box' nature typically limits clinical interpretability. To overcome this gap, our work applied SHapley Additive exPlanations (SHAP) to give insights into the XGBoost model's predictions. The dataset utilized in this research comprised of 253,680 patients and contained 21 parameters, such as General Health Status, High Blood Pressure Status, Age, and Body Mass Index. After feature selection using Recursive Feature Elimination (RFE), 15 important characteristics were discovered. In the test set, the XGBoost model obtained an accuracy of 86.6%, precision of 54.1%, recall of 17.0%, and an F1-score of 25.9% for the Original dataset. For the RFE dataset, the model displayed an accuracy of 86.6%, precision of 54.9%, recall of 16.5%, and an F1-score of 25.3%. SHAP analysis found that General Health Status, High Blood Pressure Status, Age, and Body Mass Index were the most important characteristics in both the Original and RFE datasets. This work provides as a platform for transparent and clinically applicable predictive modeling, assisting in early diabetes identification and preventive healthcare.

Keywords: Explainable AI, Diabetes, Recursive Feature Elimination

PACS: 0000, 1111

2000 MSC: 0000, 1111

1. Introduction

Diabetes, marked by chronic high blood sugar levels, is a pressing global health concern. Its prevalence continues to rise, leading to increased risks of cardiovascular diseases and other severe health issues. The World Health Organization (WHO) states that around 1.28 billion adults aged 30-79 years suffer from diabetes, predominantly in low- and middle-income nations (Diabetes - World Health Organization [1]). Yet, less than half receive proper diagnosis and treatment, with a mere one in five effectively managing their condition [2]. Such disparities highlight the pressing need for better diabetes prevention, diagnosis, and management.

In the last decade, the allure of machine learning in refining predictive models for diabetes has captivated researchers [3, 4, 5, 6, 7]. This has given birth to innovative methods for predicting various outcomes related to diabetes, from individualized risk evaluations to disease behaviour. These methods provide crucial insights into risk factors, pushing the frontier of tailored healthcare. However, many machine learning tools are criticized for their opaque "blackbox" methodologies [8]. For diabetes prediction, this opacity means healthcare professionals might struggle to grasp the biomarkers and elements that drive predictions. Considering the intricate relationships between diabetes risk factors, this lack of clarity can limit the models' clinical and public health value.

Several biomarkers critical to diabetes' onset and progression are recognized in medical literature [4]. Dyslipidemia, notable for its impact on cholesterol, is a recognized cardiovascular risk, and its ties with diabetes are well-documented [9]. Other factors, like electrolyte imbalances (high sodium, low potassium), iron metabolism, vitamin B12, and HbA1c levels, are also significantly linked with diabetes [10]. To enhance model transparency, the research community has pivoted towards explainable artificial intelligence (XAI) methodologies [10]. Notably, SHapley Additive exPlanations (SHAP) have risen in prominence which demystifies both global and local feature significance, offering a clearer understanding of diabetes risk factors, and aiding more informed medical and public health decisions [11].

This study delves deep into the existing literature on diabetes prediction, emphasizing transparent machine learning methods. SHAP is employed to elucidate an XGBoost model's predictions, utilizing a commonly available diabetes dataset from The Behavioral Risk Factor Surveillance System. Study involves 253,680 participants with their diabetes status and relevant clinical

cal data. This includes chronic illnesses, medications, and other markers. Firstly, one of the tree-based machine learning models called XGBoost is selected due to its better performance. This is followed by employing SHAP explanations to uncover novel links between diabetes risk and biomedical parameters. This work enriches the understanding of machine learning models in diabetes prediction and sets the stage for transparent, clinically pertinent predictive modeling in diabetes research and care.

The paper is structured as follows: The Methodology section delves into our research design, discussing data gathering, analytical methods, the XGBoost model, RFE(Recursive Feature Elimination) feature selection method and the use of SHAP for enhancing interpretability. The Results section shows our results, concentrating on the performance measures of the XGBoost model in both the Original and RFE datasets. It also digs into the SHAP analysis, emphasizing the most relevant factors like General Health Status, High Blood Pressure Status, Age, and Body Mass Index in both datasets. The section further covers the modest differences in accuracy, precision, recall, and F1-score between the Original and RFE sets, proving the efficiency of our feature selection approaches like RFE in diabetes prediction. A comprehensive discussion interprets our results, underlining the significance of transparent machine learning in diabetes research and healthcare choices in Discussion section. Finally, Conclusion section highlights the benefits of transparent machine learning in diabetes prediction, addresses constraints, and suggests future research directions to elevate diabetes risk evaluation, ultimately aiming for superior diabetes care and preventive medicine for better public health.

2. Methodology

In this work, two variations of the Behavioral Risk Factor Surveillance System (BRFSS) dataset, a yearly telephone survey conducted by the Centers for Disease Control and Prevention (CDC), were subjected to a thorough statistical analysis. The original dataset has two categorical outcome classes: '0,' which denotes the lack of diabetes or the presence of gestational diabetes alone, and '1,' which denotes either prediabetes or diabetes. There were originally 21 feature variables in this dataset, which included a variety of behavioral and health characteristics. Recursive Feature Elimination (RFE) was used as a feature selection strategy to narrow the feature space and improve the model's ability to forecast. A second dataset was created as

a result, however only 15 of the most useful features were kept. On both the original and the RFE-selected datasets, subsequent SHapley Additive exPlanations (SHAP) studies were performed to clarify the impact of each variable on the predictive model. This divided method provided a sophisticated assessment of feature relevance and offered insightful information on the characteristics most responsible for the likelihood of developing prediabetes or diabetes. These findings provide healthcare professionals with a strong analytical framework that enables more focused interventions based on the recognized important variables.

70% of the dataset is allocated for training, while the remaining 30% is used for testing, leading to a 70/30 data split. The 'Discussions' section also presents findings from a tenfold cross-validation, providing insights into the model's stability. Model interpretation is further enhanced using SHAP (SHapley Additive exPlanations).

2.1. Model Training

In this section, a detailed description of the classifier that has been implemented for the purpose of categorizing the diabetes dataset has been presented. A classifier, a fundamental machine learning algorithm, is utilized to partition the input data into predefined categories. In our specific context, the classifier employs the patient's features as input to discern the presence or absence of diabetes.

High levels of accuracy in predicting diabetes based on a wide range of attributes, including high blood pressure status, high cholesterol levels, cholesterol check status, Body Mass Index, stroke status, heart disease or attack status, physical activity status, heavy alcohol consumption, general health status, mental health status, difficulty walking or climbing stairs, gender, age category, education level, and income level are delineated in Table 1. Models like XGBoost, AdaBoost, Random Forest, and other tree-based standards provided excellent and comparable outcomes in the studies. XGBoost has been chosen as the main model despite alternative models having equivalent performance that has been demonstrated. XGBoost is regarded as the outstanding model due to its usability and practical efficiency, which are especially well-suited to the study objectives. It is the best option for addressing the complexity faced in our work in diabetes forecasting because of its exceptional computational speed mixed with strong predictive skills.

Model	Accuracy	Precision	Recall	F1-Score	Confusion Matrix	
XGBoost	0.8656	0.5413	0.1701	0.2588	64092	1513
Adaboost	0.8650	0.5402	0.1955	0.2871	8713	1786
Random Forest	0.8599	0.4788	0.4789	0.1748	63858	1747
Decision Tree	0.7979	0.2930	0.3289	0.3099	8446	2053
E. Boosting Machine	0.8654	0.5662	0.1664	0.2573	63608	1997
Naive Bayes	0.7718	0.3166	0.5647	0.4057	8663	1836
Logistic Regression	0.8652	0.5395	0.5395	0.1585	57274	8331
KNN	0.8468	0.3948	0.2066	0.2713	7045	3454
SVM	0.7798	0.2088	0.2184	0.2106	64266	1339
					8751	1748
					52809	12796
					4570	5929
					64184	1421
					8834	1665
					62279	3326
					8329	2170
					57117	8488
					8264	2235

Table 1: Model Metrics and Confusion Matrixes

2.2. *eXtreme Gradient Boosting (XGBoost)*

The XGBoost classifier builds upon the Gradient Boosting classifier, which also emphasizes speed and performance. One of the notable features of XGBoost is its regularized learning, which helps to smooth out the final learned weights and prevent overfitting. Overfitting occurs when the model performs well on the training data but poorly on the testing data due to learning both the information and noise from the training data [12].

The objective function $L(\theta)$ of the XGBoost can be written as:

$$L(\theta) = \sum_{i=1}^n l(y_i, \hat{y}_i) + \sum_{k=1}^K \Omega(f_k) \quad (1)$$

where n is the number of training instances, y_i is the actual value for the i^{th} training instance, \hat{y}_i is the predicted value for the i^{th} training instance, $l(y_i, \hat{y}_i)$ is the training loss, which measures the difference between the predicted and actual values for each training instance, K is the number of trees,

Variable Name	What It Indicates	Value Scales	Pt6	Pt156	Pt34	Pt265
Diabetes_binary	Diabetes Status	0=no diabetes, 1=prediabetes/diabetes	0.0	1.0	1.0	0.0
HighBP	High blood pressure status	0 = no high BP, 1 = high BP	1.0	1.0	1.0	1.0
HighChol	High cholesterol levels	0 = no high cholesterol, 1 = high cholesterol	0.0	1.0	1.0	0.0
CholCheck	Cholesterol check status	0 = no check in 5 years, 1 = yes	1.0	1.0	1.0	1.0
BMI	Body Mass Index	Continuous Scale	30.0	47.0	24.0	36.0
Smoker	Smoking status	0 = no, 1 = yes	1.0	1.0	1.0	1.0
Stroke	Stroke status	0 = no, 1 = yes	0.0	0.0	0.0	0.0
HeartDiseaseorAtc	Heart disease or attack stat	0 = no, 1 = yes	0.0	0.0	0.0	1.0
PhysActivity	Physical activity status	0 = no, 1 = yes	0.0	0.0	0.0	0.0
Fruits	Fruit consumption	0 = no, 1 = yes	0.0	1.0	0.0	1.0
Veggies	Vegetable consumption	0 = no, 1 = yes	0.0	0.0	0.0	1.0
HvyAlcoholCons	Heavy alcohol consumption	0 = no, 1 = yes	0.0	0.0	0.0	0.0
AnyHealthcare	Healthcare coverage	0 = no, 1 = yes	1.0	1.0	1.0	1.0
NoDocbcCost	Avoided doctor due to cost	0 = no, 1 = yes	0.0	0.0	0.0	0.0
GenHlth	General health status	1=excel,2=very good,3=good,4=fair,5=poor	3.0	3.0	2.0	3.0
MentHlth	Mental health status	1-30 days	0.0	0.0	0.0	0.0
PhysHlth	Physical health status	1-30 days	14.0	0.0	0.0	2.0
DiffWalk	Difficulty in walking	0 = no, 1 = yes	0.0	1.0	0.0	1.0
Sex	Gender	0 = female, 1 = male	0.0	0.0	0.0	0.0
Age	Age category	1 = 18-24, 9 = 60-64, 13 = 80 or older	9.0	11.0	12.0	11.0
Education	Education level	1 = Never attended, 6 = College graduate	6.0	6.0	3.0	5.0
Income	Income level	1 = less than \$10,000, 8 = \$75,000 or more	7.0	5.0	3.0	4.0

Table 2: Variable names, their indications, and value scales along with patient data.

f_k is the k^{th} tree, and $\Omega(f_k)$ is the regularization term for the k^{th} tree, which penalizes the complexity of the model to avoid overfitting.

The regularization term $\Omega(f_k)$ can be further defined as:

$$\Omega(f_k) = \gamma T + \frac{1}{2} \lambda \|w\|^2 \quad (2)$$

Where T is the number of leaf nodes in the k^{th} tree, w is the vector of scores on the leaf nodes in the k^{th} tree, and γ and λ are regularization parameters that control the complexity of the model. Large weights are penalized by the regularization term, which also encourages the model to have more streamlined and comprehensible structural elements. The goal of the model is to minimize this loss over the entire training set. Given that our dataset is also noisy, XGBoost is one of the appropriate classifiers for it [13]. The values of the hyperparameters are given in Table 3.

Hyperparameter	Value
eta (learning rate)	0.3
n_estimator (number of gradient-boosted trees)	100
Gamma (min split loss)	0
Max depth	6
Min child weight	1
Max delta step	0
subsample	1
Sampling method	Uniform

Table 3: Hyperparameter values for the model.

2.3. Recursive Feature Elimination (RFE)

Finding the ideal subset of features for a particular machine learning study is the goal of the feature selection method known as recursive feature elimination (RFE)[14]. When using RFE, a model is fitted to the data, and the features are then ranked according to how important or relevant they are to the prediction. The least significant characteristics are then gradually removed by RFE until the target number of features is obtained or a stopping requirement is satisfied. By lowering the dimensionality and complexity of the data, as well as by removing noise and multicollinearity among features, RFE can enhance the effectiveness and performance of machine learning models.[15]

Recursive Feature Elimination (RFE) is a feature selection algorithm commonly employed in machine learning for identifying a subset of most predictive features. Given a dataset $\mathcal{D} = \{(x_1, y_1), \dots, (x_n, y_n)\}$ with n samples, where each $x_i \in \mathbb{R}^d$ is a d -dimensional feature vector and y_i is the corresponding label, RFE aims to minimize a loss function $\mathcal{L}(f(x; \theta), y)$ by selecting the most relevant features. The optimization objective is $\theta^* = \arg \min_{\theta} \sum_{i=1}^n \mathcal{L}(f(x_i; \theta), y_i)$. The algorithm starts by fitting a predictive model $f : \mathbb{R}^d \rightarrow \mathbb{R}$ parameterized by θ to the entire feature set. It then ranks the features based on their importance, often calculated as $\text{Importance}(j) = \left| \frac{\partial \mathcal{L}}{\partial x_j} \right|$ for each feature j , and recursively eliminates the least important ones. This process continues iteratively until the desired number of features is retained. The computational complexity of RFE is generally $O(d \times p)$, where p is the complexity of the base estimator used for ranking features. RFE is particularly useful when model interpretability and simplicity are as crucial as predictive performance, although it can be computationally expensive for high-dimensional data.

Several machine learning models that offer some kind of feature importance metric can be used with recursive feature elimination (RFE). This comprises linear models with coefficients, tree-based techniques with feature significance scores, and support vector machines (SVMs) with their support vectors. It’s important to keep in mind, though, that RFE functions as a wrapper technique and requires the fitting of a new model for each subset of characteristics. This might require a lot of work and increase the risk of model overfitting. Table 2 provides more information on the 15 relevant parameters that were chosen for our investigation using RFE.

2.4. Model Interpretation with SHAP

To augment the transparency of our XGBoost model and provide clinically meaningful interpretations, we employed SHapley Additive exPlanations (SHAP) as an interpretability tool. SHAP is rooted in cooperative game theory, aiming to fairly allocate ”payouts” or contributions among features for a given prediction. Mathematically, the Shapley value ϕ_i of a feature i is computed as:

$$\phi_i(f) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|! \cdot (|N| - |S| - 1)!}{|N|!} [f(S \cup \{i\}) - f(S)]$$

where $f(S)$ is the prediction of the model for feature subset S , and N is the set of all features. These Shapley values represent the marginal contribution of each feature to a specific prediction, enabling us to decipher both global and local feature importance.

The application of SHAP to our XGBoost model adhered to a systematic methodology. Initially, specific data preprocessing steps tailored for SHAP analysis were executed. Utilizing the SHAP library, Shapley values for each feature across all data points were computed. These values were aggregated to form global feature importance metrics, which can be mathematically represented as:

$$\text{Global Importance}(j) = \frac{1}{n} \sum_{i=1}^n |\phi_{ij}|$$

where ϕ_{ij} is the Shapley value of feature j for data point i , and n is the total number of data points. This provided an overarching view of the most and least impactful features in diabetes prediction. Interestingly, our analysis revealed that certain biomarkers, traditionally not emphasized in diabetes literature, played a significant role in model predictions.

SHAP also excels in local interpretability, enabling us to dissect individual predictions. For example, in case studies of specific patients, the SHAP values highlighted the biomarkers that were instrumental in classifying them as high-risk or low-risk for diabetes. This local interpretability is invaluable for clinicians, as it provides personalized insights into each patient’s risk factors. Mathematically, local interpretability can be examined through the individual Shapley values ϕ_{ij} , providing a nuanced understanding of how each feature contributes to a specific prediction.

In XGBoost, SHAP values may be particularly useful for comprehending how various characteristics interact with one another. The authors’ modification of SHAP values to take interaction effects into consideration led to the idea of SHAP interaction values. These interaction values guarantee the explanations of each of XGBoost’s unique estimations for interaction effects. In order to highlight important interactions that XGBoost captures but that could otherwise go undetected, SHAP interaction values discriminate between main effects and interaction effects for specific model predictions.[16]

An efficient tool for understanding XGBoost models is the SHAP values. They constitute a significant theoretical breakthrough by overcoming consistency problems and permitting the separate assessment of main and

interaction effects for individual model predictions. This might highlight important interactions that XGBoost has recorded but that were otherwise missed. Data scientists are now better prepared to develop models that are more clear and intelligible thanks to the introduction of SHAP values in XGBoost.[17]

3. Results

3.1. Results of Classifications

First, two separate sets were created from the original dataset after a data preparation step. The original collection, referred to as the "Original Set," had 253,680 samples and kept all 21 features. The second set, referred to as the "RFE Set," was created by choosing 15 important features from the initial 253,680 samples using recursive feature elimination (RFE).

RFE was used to improve feature selection and reduce model complexity. This approach chose 15 characteristics, which are listed in Table `reftab:attributes`, including high blood pressure status, high cholesterol levels, cholesterol check status, body mass index, and others.

The two sets 'Original' and 'RFE' were each put through an XGBoost classifier separately to evaluate the precision and robustness of our feature selections, as shown also in Table `reftab:analysis-results`. Despite some variations in the measures, the classifier performed well on both sets of data.

The classifier's performance for the 'Original' set was accuracy of 86.6%, precision of 54.1%, recall of 17.0%, and F1-score of 25.9%.

Similar results were obtained for the 'RFE' set, where the classifier's accuracy was 86.6%, precision was 54.9%, recall was 16.5%, and F1-score was 25.3%.

These findings validate the efficacy of our feature selection approaches by showing that both feature sets are useful for classification, although with minor changes in accuracy and recall measures.

3.2. Explaining Model with SHAP

Machine learning algorithms have emerged as effective tools for disease prediction in the evolving healthcare landscape. However, the 'black-box' nature of these models often makes it challenging to interpret their predictions and understand the relative importance of different input factors. To address this issue, interpretability tools as SHAP have gained popularity. SHAP values provide a nuanced understanding of the factors influencing illness risk

No.	Attribution	Variable	RFE	Pt6	Pt156	Pt34	Pt265
1	High blood pressure	HighBP	Picked	1.0	1.0	1.0	1.0
2	High cholesterol levels	HighChol	Picked	0.0	1.0	1.0	0.0
3	Cholesterol check status	CholCheck	Picked	1.0	1.0	1.0	1.0
4	Body Mass Index	BMI	Picked	30.0	47.0	24.0	36.0
5	Smoking status	Smoker		1.0	1.0	1.0	1.0
6	Stroke status	Stroke	Picked	0.0	0.0	0.0	0.0
7	Heart disease or attack status	HeartDorA	Picked	0.0	0.0	0.0	1.0
8	Physical activity status	PhysActivity	Picked	0.0	0.0	0.0	0.0
9	Fruit consumption	Fruits		0.0	1.0	0.0	1.0
10	Vegetable consumption	Veggies		0.0	0.0	0.0	1.0
11	Heavy alcohol consmp.	HvyAlcoholC	Picked	0.0	0.0	0.0	0.0
12	Healthcare coverage	AnyHealthcr		1.0	1.0	1.0	1.0
13	Avoided doctor due to cost	NoDocbcCos		0.0	0.0	0.0	0.0
14	General health status	GenHlth	Picked	3.0	3.0	2.0	3.0
15	Mental health status	MentHlth	Picked	0.0	0.0	0.0	0.0
16	Physical health status	PhysHlth		14.0	0.0	0.0	2.0
17	Difficulty in walking or climbing stairs	DiffWalk	Picked	0.0	1.0	0.0	1.0
18	Gender	Sex	Picked	0.0	0.0	0.0	0.0
19	Age category	Age	Picked	9.0	11.0	12.0	11.0
20	Education level	Education	Picked	6.0	6.0	3.0	5.0
21	Income level	Income	Picked	7.0	5.0	3.0	4.0

Table 4: Attribution of the dataset and patient details with RFE Selection

by quantifying the impact of specific attributes on the model’s prediction. In this context, we explore the intricate network of health factors affecting the likelihood of developing hypertension, a common and complex cardiovascular disease. Our goal is to analyze the interaction of various health metrics, ranging from medication use and standard risk factors to less obvious influencers, and how they collectively influence the model’s predictions using patient-specific SHAP values. This investigation highlights how machine learning interpretability tools can enhance our understanding of disease prediction and potentially guide therapeutic decisions.

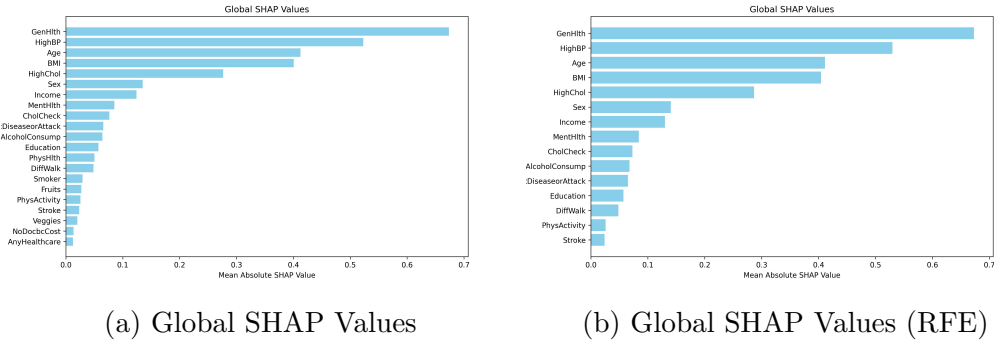



Figure 1: Comparison of Global SHAP Values

The study highlights the primary factors impacting hypertension across various dataset modifications through the analysis of global SHAP values. Several common trends and intriguing variances are observed. The average magnitude of the SHAP values for each feature over the entire dataset is evaluated, as depicted in Figure 3c.


To acquire a better understanding of how each feature effects the model’s predictions, SHAP (Shapley Additive Explanations) values were generated. Among all the characteristics, General Health Status, High Blood Pressure Status, Age, and Body Mass Index emerged as the most significant in both the Original and RFE datasets. In the Original dataset, the SHAP values were notably 0.674 for General Health Status, 0.523 for High Blood Pressure Status, 0.413 for Age, and 0.401 for Body Mass Index. Similarly, in the RFE dataset, the comparable SHAP values were 0.673 for General Health Status, 0.530 for High Blood Pressure Status, 0.411 for Age, and 0.404 for Body Mass Index as also shown in figure 1. These values provide as a quantifiable indication of the average influence each of these qualities has on the model’s output, therefore underlining their relevance in diabetes prediction.

These results validate the success of our feature selection techniques by showing the resilience of the chosen features in predicting diabetes, regardless of the feature set employed.


A more sophisticated knowledge of the model’s prediction ability is revealed when the individual SHAP values are thoroughly analyzed and compared to the actual health situations of particular patients. Negative SHAP values for General Health Status were seen starting with Patient 6, who was correctly identified as not having diabetes (True Negative), as shown in the




(a) Local SHAP values for original dataset



(b) Local SHAP values for RFE dataset



(c) Local SHAP values for original dataset




(d) Local SHAP values for RFE dataset

Figure 2: Comparison of local SHAP values for Patient 6 and 156


Figure 2, -1.3621 in the "Original" set and -1.4022 in the "RFE" set. These data support the model's accurate prediction of a decreased risk of diabetes, together with a negative SHAP value for high blood pressure status (-0.5726 in "Original" and -0.5969 in "RFE").

Patient 156 had substantial positive SHAP values for BMI, 0.9521 in the "Original" set and 0.8867 in the "RFE" set, and was appropriately diagnosed as having diabetes (True Positive). As shown in the figure 2, positive SHAP values for High Blood Pressure Status (0.5819 in 'Original' and 'RFE') are confirmed to have a significant influence in the model's ability to diagnose diabetes accurately.


However, the model's forecasts did not show infallibility. Positive SHAP readings for High Blood Pressure Status were found in both sets for Patient 34, who was incorrectly diagnosed as not having diabetes (False Negative), as shown in the Figure 3, namely 0.5234 in the "Original" and 0.5418 in the




(a) Local SHAP values for original dataset



(b) Local SHAP values for RFE dataset



(c) Local SHAP values for original dataset



(d) Local SHAP values for RFE dataset

Figure 3: Comparison of local SHAP values for Patient 34 and 265

”RFE”. These were countered by a negative SHAP value for General Health Status (-0.5440 in ’Original’ and -0.4934 in ’RFE’), indicating that these competing signs may have caused the inaccurate prediction to be confused.

Similar positive SHAP readings as shown in the figure 3, for BMI (0.5051 in ’Original’ and 0.5029 in ’RFE’) and High Blood Pressure Status (0.5122 in ’Original’ and 0.4965 in ’RFE’) were reported in the case of Patient 265 who was misdiagnosed as having diabetes (False Positive). Despite these signs, the model overestimated the danger, which led to an inaccurate prognosis.

These distinct studies confirm the robustness of the feature selection while also highlighting potential areas where more model fine-tuning may be necessary for more precise and dependable predictions.

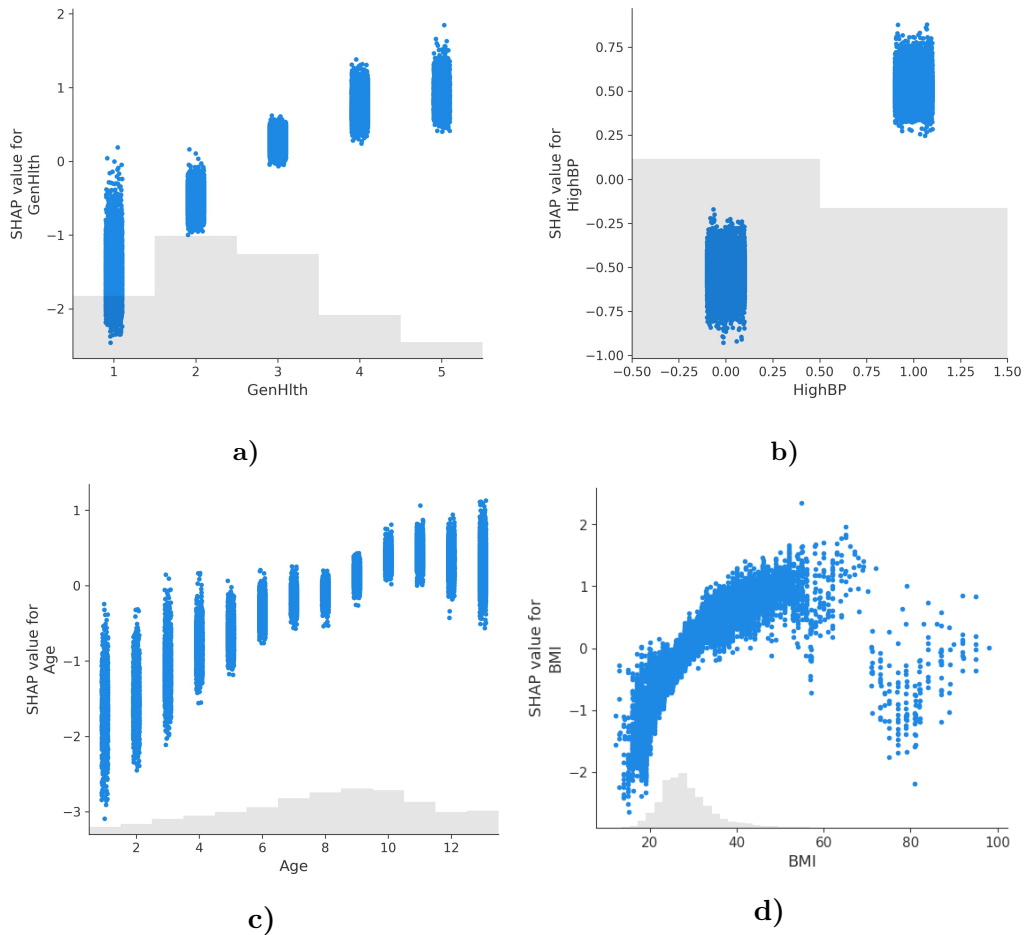


Figure 4: Dependence plots in original dataset for variables **a)** General health status, **b)** High blood pressure, **c)** Age, **d)** Body Mass Index(BMI)

3.3. SHAP Dependences for Variables

In the study, the clarification of how individual features influence the model's predictions was provided by SHAP dependence plots. Additionally, how these relationships are modulated by other interacting variables was indicated by the color of each point on the plots. A comprehensive understanding of the key variables and their interactions that drive the model's decisions was offered by these plots.

Two distinct sets were utilized for this analysis: the 'Original' set, which encompasses all 21 variables, and the 'RFE' set, a refined subset containing

15 selected features. In both sets, the identification of the most impactful features was made: General Health Status (GenHlth), High Blood Pressure Status (HighBP), Age, and Body Mass Index (BMI), listed in descending order of their importance. A more comprehensive understanding of the implications of these critical variables was enabled by the utilization of both sets, thereby enriching the interpretability of the model.

It must be noted that rigorous statistical analysis has not been applied to the SHAP values, which are intended to provide an initial understanding of the relationships between the variables and diabetes risk. For example, a greater likelihood of diabetes was suggested by higher SHAP values for General Health Status and High Blood Pressure Status in both sets, as illustrated in the dependence plots.

Regarding Age, an increase in the risk of diabetes as age increases was indicated by the SHAP values. This is consistent with broader medical understanding, in which older age is linked to a higher likelihood of developing chronic conditions, including diabetes.

Similarly, a link between elevated Body Mass Index (BMI) values and a higher risk of diabetes was established, corroborating existing medical literature that associates obesity with diabetes.

In summary, the validation of the efficacy of the feature selection process was achieved through the SHAP dependence plots generated from both the 'Original' and 'RFE' sets. Nuanced insights into how these selected variables interact to influence the model's predictive capabilities were also offered.

Discussion

In the current work, the efficacy of machine learning methods in predicting diabetes risk was examined, with a special emphasis on the XGBoost classifier. Two unique datasets were employed: the 'Original' set, encompassing all 21 variables, and the 'RFE' set, a refined subset of 15 chosen characteristics. The SHAP dependent charts were developed to give subtle insights into the correlations between various variables and the model's prediction capabilities. It was discovered that General Health Status (GenHlth), High Blood Pressure Status (HighBP), Age, and Body Mass Index (BMI) were the most impacting factors in both groups. These results accord with current research, strengthening the assumption that these characteristics are key markers of diabetes risk.

The investigation also demonstrated that the XGBoost classifier performed with a high degree of accuracy on both the 'Original' and 'RFE' sets. However, it should be noted that the model's performance indicators, such as accuracy and recall, differed marginally across the two sets. This implies that although feature selection strategies like RFE may be efficient in decreasing model complexity, they may not necessarily lead to a substantial gain in predictive accuracy. Therefore, the option between employing all accessible data or a subset of chosen features should be made carefully, considering the unique aims and restrictions of the study.

One of the disadvantages of this research is the absence of comprehensive statistical analysis done to the SHAP values. While the SHAP dependency plots give helpful early insights into the correlations between factors and diabetes risk, they are not a replacement for a complete statistical study. Future study might benefit from a more in-depth investigation, potentially adding additional statistical approaches to confirm the results further.

In conclusion, this work adds to the expanding corpus of research on the application of machine learning algorithms in healthcare. It reveals that XGBoost, when trained on well chosen characteristics, may be a strong tool for predicting diabetes risk. The insights gathered from the SHAP dependent plots give an extra layer of interpretability to the model, making it more transparent and perhaps more trustworthy in a clinical environment. Further study is required to confirm these results across diverse demographics and healthcare systems.

4. Conclusion and Future Work

4.1. Conclusion

In our comprehensive research journey into the domain of machine learning interpretability, an emphasis on diabetes risk assessment was pronounced, leveraging the usability of SHAPley analysis. Drawing from the Behavioral Risk Factor Surveillance System (BRFSS) dataset, the study precisely analyzed two distinct variations (original and feature selected) of the dataset, which is a renowned annual telephone survey collected by the Centers for Disease Control and Prevention (CDC). The dataset presented a separation in its outcome classes, defining the presence of diabetes, offering a granular understanding of this pervasive disease. Firstly, encompassing 21 varied feature variables, the dataset provided a kaleidoscopic view of a many of

behavioural and health characteristics. Identifying the potential for refinement, we harnessed RFE as a strategic tool for feature selection. This not only rationalised our feature space but significantly augmented our model’s prediction ability. In conclusion our efforts underscore the enormous benefits of adopting a transparent machine learning paradigm, especially in the diabetes prediction.

4.2. Future works

The horizon of future endeavours, our study gazed into several areas of research. The profound depth of the BRFSS dataset hints at the potential for uncovering more intricate patterns and correlations, thereby exploring diabetes risk prediction. While our current methodologies have proven the usability, the dynamic nature of healthcare mandates continuous evolution. This could clear in the form of integrating more sophisticated machine learning algorithms, or perhaps investigating the distinctions of SHAP to loosen more complex model behaviors. Collaborative research performed with medical professionals could surface the real-world results for the application of our models, offering concrete benefits to patients. It is aspired by amalgamating feedback from the clinics, refining methodologies, and developing innovation, to redefine the landscape of machine learning in diabetes risk assessment.

References

- [1] W. H. Organization, et al., Classification of diabetes mellitus (2019).
- [2] D. A. Mihai, D. S. Stefan, D. Stegaru, G. E. Bernea, I. A. Vacaroiu, T. Papacocea, M. O. D. Lupuoru, A. E. Nica, O. Stiru, D. Dragos, et al., Continuous glucose monitoring devices: A brief presentation, *Experimental and Therapeutic Medicine* 23 (2) (2022) 1–6.
- [3] Y. Lin, Y. Li, X. Huang, L. Liu, H. Wei, X. Zou, Analysis of diabetes clinical data based on recurrent neural networks, *Computational Intelligence and Neuroscience* 2022 (2022).
- [4] F. Gómez-Peralta, C. Abreu, Clinical research on type 2 diabetes: A promising and multifaceted landscape (2022).
- [5] Y. Qin, J. Wu, W. Xiao, K. Wang, A. Huang, B. Liu, J. Yu, C. Li, F. Yu, Z. Ren, Machine learning models for data-driven prediction of diabetes

- by lifestyle type, *International Journal of Environmental Research and Public Health* 19 (22) (2022) 15027.
- [6] S. Shankaracharya, Diabetes risk prediction using machine learning: Prospect and challenges, *Bioinformatics, Proteomics and Immaging Analysis* 3 (2017) 194–195.
- [7] E. Afsaneh, A. Sharifdini, H. Ghazzaghi, M. Z. Ghobadi, Recent applications of machine learning and deep learning models in the prediction, diagnosis, and management of diabetes: a comprehensive review, *Diabetology & Metabolic Syndrome* 14 (1) (2022) 1–39.
- [8] J. Carreras, S. Hiraiwa, Y. Y. Kikuti, M. Miyaoka, S. Tomita, H. Ikoma, A. Ito, Y. Kondo, G. Roncador, J. F. Garcia, et al., Artificial neural networks predicted the overall survival and molecular subtypes of diffuse large b-cell lymphoma using a pancancer immune-oncology panel, *Cancers* 13 (24) (2021) 6384.
- [9] M. Gutch, S. Rungta, S. Kumar, A. Agarwal, A. Bhattacharya, S. M. Razi, Thyroid functions and serum lipid profile in metabolic syndrome, *Biomedical journal* 40 (3) (2017) 147–153.
- [10] M. A. Hasan, M. F. K. Chowdhury, S. Yasmin, S. Paul, T. Ahmed, S. Monira, G. S. Hasan, M. R. Amin, H. Sarker, Association between glycemic control and serum lipid profile in type 2 diabetic patients: Experience in a medical college hospital, *Bangabandhu Sheikh Mujib Medical University Journal* 14 (4) (2021) 138–143.
- [11] H. Gong, M. Wang, H. Zhang, M. F. Elahe, M. Jin, An explainable ai approach for the rapid diagnosis of covid-19 using ensemble learning algorithms, *Frontiers in Public Health* 10 (2022) 874455.
- [12] N. L. Fitriyani, M. Syafrudin, G. Alfian, J. Rhee, Hdpm: an effective heart disease prediction model for a clinical decision support system, *IEEE Access* 8 (2020) 133034–133050.
- [13] N. Shao, H. Hu, Exploring the path of enhancing ideological and political education in universities in the era of big data, *Journal of Environmental and Public Health* 2022 (2022).

- [14] Y. Zhang, X. Zhang, J. Razbek, D. Li, W. Xia, L. Bao, H. Mao, M. Daken, M. Cao, Opening the black box: interpretable machine learning for predictor finding of metabolic syndrome, *BMC Endocrine Disorders* 22 (1) (2022) 1–15.
- [15] X.-w. Chen, J. C. Jeong, Enhanced recursive feature elimination, in: *Sixth international conference on machine learning and applications (ICMLA 2007)*, IEEE, 2007, pp. 429–435.
- [16] Z. He, Y. Yang, R. Fang, S. Zhou, W. Zhao, Y. Bai, J. Li, B. Wang, Integration of shapley additive explanations with random forest model for quantitative precipitation estimation of mesoscale convective systems, *Frontiers in Environmental Science* 10 (2023) 1057081.
- [17] X. R. Gao, M. Chiariglione, K. Qin, K. Nuytemans, D. W. Scharre, Y.-J. Li, E. R. Martin, Explainable machine learning aggregates polygenic risk scores and electronic health records for alzheimers disease prediction, *Scientific Reports* 13 (1) (2023) 450.