Malaysian Journal of Mathematical Sciences 19(4): 1271–1288 (2025)
https://doi.org/10.47836/mjms.19.4.05

Malaysian Journal of Mathematical Sciences

Journal homepage: https://mjms.upm.edu.my

Gradient-Boosting Classifier with Shapley Additive Explanations for Analysis of Obesity Risk in Saudi Arabia

Alanazi, W. M. T.*[1], Abu Bakar, M. A. [1], and You, H. W.[1]

[1]*Department of Mathematical Sciences, Faculty of Science and Technology, Universiti Kebangsaan Malaysia, 43600 UKM Bangi, Selangor, Malaysia*

E-mail: wadhatgfsjh@gmail.com
**Corresponding author*

Received: 27 February 2025
Accepted: 14 May 2025

Abstract

Obesity remains a critical public health issue in the Kingdom of Saudi Arabia (KSA), necessitating accurate and interpretable predictive tools. This study proposes a Gradient-Boosted Decision Tree (GBDT-XGBoost) model integrated with SHapley Additive exPlanations (SHAP) for obesity risk classification. Trained on demographic and lifestyle data, the model achieved a training accuracy of 95.8%, testing accuracy of 93.1%, precision of 94.2%, recall of 92.8%, F1-score of 93.5%, AUC of 0.963, RMSE of 0.32 and AIC of $1,205$, with a training time of 58.2 seconds. Comparative analysis showed XGBoost outperformed Artificial Neural Networks (ANN), Random Forest (RF) and Support Vector Regression (SVR) across all metrics. Although ANN achieved competitive results (test accuracy: 93.8%, RMSE: 0.35), it required over three times longer training time. SHAP analysis identified BMI (0.42), age (0.28) and dietary habits (0.16) as the most influential features. SHAP plots also revealed that vegetable consumption (FCVC) significantly impacted predictions and age negatively correlated with risk. The XGBoost-SHAP framework offers a high-performance, interpretable and scalable solution for obesity risk prediction. Its integration of transparent feature attribution supports its utility in clinical decision-making and public health strategies tailored to the Saudi population.

Keywords: obesity risk prediction; machine learning; XGBoost; SHAP analysis; healthcare analytic.

1 Introduction

The global surge in overweight and obesity presents a significant public health challenge, extending across nations and demographics. Using the body mass index (BMI) classification system established by the World Health Organization (WHO), it's evident that obesity rates have steadily increased over the past 35 years, encompassing countries like the Kingdom of Saudi Arabia (KSA) and other Asian regions [8]. The WHO's 2016 report highlighted that 1.9 billion individuals worldwide were classified as overweight, with 650 million categorized as obese. Among adults aged 18 and above, obesity prevalence stood at 13%, with 11% among men and 15% among women, while overweight prevalence was 38%, with 9% among men and 40% among women [48]. This underscores obesity's global impact, ranking among the top five contributors to environmental, behavioral and metabolic hazards according to the 2017 Burden of Disease Study [20].

Obesity is an extensive accumulation of body fat is now a growing concern, impacting individuals, families and communities worldwide [12]. While statistics provide insights into the scale of the issue, they fail to fully convey the human toll of the obesity pandemic. The prevalence of obesity has dire consequences, leading to chronic illnesses such as diabetes, heart disease, cancers and musculoskeletal problems [16]. Additionally, obesity reduces life expectancy and imposes a significant financial burden on healthcare systems, with projected annual expenses surpassing $147 billion in the United States alone. Beyond health implications, obesity carries social and psychological costs, contributing to stigma, prejudice and diminished quality of life [27].

Tackling the growing obesity crisis demands a multifaceted and collaborative approach involving stakeholders at all levels including governments, healthcare systems, communities and individuals [26]. Without coordinated action, the public health burden posed by obesity will continue to escalate. In Saudi Arabia, a nation characterized by rapid urbanization, dietary shifts and reduced physical activity, the prevalence of obesity has reached critical levels. Recent epidemiological data reveal that approximately 28% of men and 44% of women are classified as obese, while 66% of men and an alarming 71% of women are categorized as overweight [3]. These statistics highlight a deeply rooted and widespread public health concern that transcends gender and socioeconomic strata. The obesity epidemic in Saudi Arabia is further complicated by cultural norms, sedentary lifestyles and increased consumption of high-calorie foods, often linked to economic prosperity and modernization. Women, in particular, face heightened vulnerability due to limited access to public spaces for physical activity and social norms that may hinder active lifestyles. The high prevalence rates not only strain the healthcare infrastructure but also impose long-term economic consequences due to obesity-related diseases such as diabetes, cardiovascular conditions and metabolic syndromes. As such, comprehensive strategies ranging from public awareness campaigns and policy reforms to clinical management and behavioral interventions are urgently required to curb the obesity epidemic and promote a healthier society.

Recent research has revealed an alarming upward trajectory in obesity rates within the Kingdom of Saudi Arabia, despite numerous national health initiatives aimed at mitigating this public health crisis. This trend underscores a growing and persistent health burden that poses significant threats to population health, economic stability and healthcare sustainability. For example, a study by AlAbdulKader et al. [5] emphasizes the role of changing dietary behaviors and sedentary lifestyles in the rising obesity trend, particularly among younger Saudi populations. Their findings point to the failure of current interventions to curb the progression of obesity-related behaviors, calling for policy reform.

In a more recent study, Aljarid et al. [7] discuss the effectiveness of public health campaigns and the gaps in policy implementation that have contributed to the continued rise in obesity rates.

They advocate for a multidimensional strategy, incorporating education, urban planning and preventive healthcare services. Similarly, Althumoiri et al. [9] conducted a nationwide cross-sectional survey that revealed high rates of both overweight and obesity across all age groups, signaling a widespread epidemic that cuts across gender, region and socioeconomic status. Their data serve as a national baseline for future obesity surveillance and highlight urgent intervention points.

Environmental and policy-based determinants of obesity were explored by DeNicola et al. [17], who discussed how urbanization and environmental changes such as the reduction in walkable infrastructure and increased reliance on automobiles have significantly contributed to decreased physical activity among Saudis. They argue for environmental reforms that foster healthier lifestyles, such as the development of parks, pedestrian zones and fitness-friendly urban design. Environmental applications have also seen promising results. Zhao et al. [46] used XGBoost to model air pollution dynamics in urban environments, finding that it could outperform long short-term memory (LSTM) networks and autoregressive integrated moving average (ARIMA) models in forecasting $PM_{2.5}$ concentrations. This illustrates XGBoost's robustness in spatiotemporal prediction tasks where nonlinear interactions prevail.

From a clinical perspective, Habbab et al. [21] studied some of the physiological and metabolic consequences of obesity in Saudi patients, underscoring its association with chronic conditions including type 2 diabetes, cardiovascular disease and hypertension. Their work reinforces the need for early diagnosis and integrated clinical management programs within the Saudi healthcare system. Moreover, in a global context, Tyrovolas et al. [45] analyze obesity patterns in Middle Eastern countries and compare them with global trends. The study positions Saudi Arabia among the countries with the highest obesity burdens in the region and stresses the importance of aligning national strategies with global best practices. These studies paint a picture of obesity in Saudi Arabia, outlining its causes, consequences and challenges to mitigation. The consistent persistence of high obesity rates, despite ongoing efforts, emphasizes the need for immediate and aggressive action. Effective responses must go beyond individual-level interventions to include policy integration, cross-sector collaboration and sustained public health investment aimed at reversing this detrimental trajectory.

The Extreme Gradient Boosting (XGBoost) algorithm, introduced by Chen and Guestrin [13], has become a cornerstone in supervised machine learning, particularly for structured tabular data. XGBoost implements gradient boosting in a highly efficient and scalable manner, employing second-order Taylor approximations and system optimizations such as parallelized tree construction, cache awareness and sparsity-aware algorithms. This combination allows XGBoost to excel in handling large datasets with high-dimensional features, making it a frequent winner in machine learning [6]. Montomoli et al. [11] demonstrated that XGBoost significantly outperforms traditional machine learning models in both accuracy and computational efficiency. Their study, conducted in the domain of materials science, showed that XGBoost can be up to ten times faster than Random Forest and Support Vector Machines due to its built-in regularization and parallelization strategies. Furthermore, its capacity to natively handle missing data makes it an appealing choice for applications where data completeness is a challenge.

In healthcare informatics, XGBoost has been effectively utilized to predict mortality risks, hospital readmissions and disease diagnoses. Rajkomar et al. [39] applied XGBoost to electronic health record (EHR) datasets, demonstrating its superior performance compared to logistic regression and deep learning models in predicting inpatient mortality and readmission rates. Similarly, Choi et al. [14] showed that XGBoost could identify sepsis risk in ICU patients with higher precision than classical statistical models.

In epidemiological studies, Li et al. [23] employed XGBoost to predict COVID-19 hospitalization risks. Their findings indicated that the model not only provided higher predictive accuracy but also offered interpretable insights when integrated with SHAP (Shapley Additive Explanations), thus facilitating transparent decision-making in public health. In agricultural sciences, Ma et al. [28] applied XGBoost to predict crop yields under varying climatic conditions. Their comparative study indicated that XGBoost yielded higher prediction accuracy than neural networks and decision trees, especially when working with high-dimensional satellite imagery data. The ability to model complex feature interactions made it particularly suitable for agroecological systems.

In the finance domain, He et al. [49] demonstrated XGBoost's strength in credit risk assessment and fraud detection. By training on transactional data, XGBoost produced lower false positive rates and greater model stability over time compared to logistic regression and ensemble bagging methods. The domain of land-use change modeling has also benefited from XGBoost. Sun et al. [19] utilized the algorithm to simulate urban expansion in rapidly growing metropolitan areas. Their work revealed that XGBoost captured nonlinearities in socioeconomic and environmental drivers more effectively than traditional cellular automata models. This enhanced predictive capability has significant implications for sustainability and urban planning.

In a public health context, Pan et al. [25] employed XGBoost to develop an early warning system for obesity and diabetes risk. By integrating behavioral and genetic data, the model outperformed linear regression and k-nearest neighbors in both sensitivity and specificity. The study highlighted the model's suitability for clinical screening tools that demand real-time, interpretable output. A recent study by Arunachalam et al. [24] combined XGBoost to analyze the influence of lifestyle and demographic factors on cardiovascular risk. Their findings illustrated how XGBoost can provide granular insights into complex health datasets, supporting its use in individualized healthcare interventions.

Despite its growing adoption, the use of XGBoost in certain domains such as obesity prediction and land-use transitions remains underexplored. The current study addresses this gap by integrating XGBoost with SHAP to assess obesity risk in Saudi Arabia. Given the alarming increase in obesity rates, understanding feature importance across variables such as age, BMI, dietary patterns and physical activity is crucial for targeted public health strategies. XGBoost stands out as a highly adaptable and interpretable machine learning algorithm with cross-disciplinary applicability. Its superior accuracy, processing efficiency and compatibility with interpretability tools such as SHAP make it an indispensable asset in domains requiring both performance and transparency. The expanding literature across healthcare, environmental sciences, finance and agriculture underscores XGBoost's role as a reliable predictive tool, paving the way for future research in underutilized domains like public health and sustainability.

Interpreting complex machine learning (ML) models is crucial for their application in sensitive domains. To address this, Lundberg [31] introduced Shapley Additive Explanations (SHAP), a method based on cooperative game theory that calculates Shapley values to fairly quantify each feature's contribution to a model's prediction. SHAP offers both *global* and *local* interpretability, making it a robust, model-agnostic tool compatible with a wide range of algorithms including gradient boosting and neural networks [32, 37].

SHAP has proven valuable across various fields, especially healthcare. For example, it helped identify mortality risk factors in intensive care; Lundberg [31] revealed key predictors of chronic diseases such as diabetes and cardiovascular conditions [41]. These insights enhance clinical trust and facilitate personalized treatment [36, 2]. In environmental science, SHAP has elucidated drivers of urban expansion and biodiversity loss, supporting sustainable planning and conservation [22]. This versatility demonstrates SHAP's broad applicability in complex systems.

This study employs a hybrid methodological framework combining Extreme Gradient Boosting (XGBoost), a high-performance gradient-boosted decision tree algorithm [13] with SHapley Additive exPlanations (SHAP) to predict obesity risk in Saudi Arabia, where obesity rates have been rising significantly [47, 4]. XGBoost is used for accurate classification and prediction, while SHAP is implemented to interpret the model's outputs, offering transparency by quantifying the contribution of each feature to the prediction [30].

This methodological integration allows for identifying both modifiable and non-modifiable risk factors such as age, BMI, physical activity, dietary behavior and family history and enables a clearer understanding of their influence on obesity risk. SHAP also enhances communication between machine learning practitioners and healthcare professionals by making the decision-making process more interpretable [41, 24]. Similar explainable AI approaches have been successfully applied in environmental modeling, such as Zhao et al.'s work [50], which used SHAP and XGBoost to reveal nonlinear interactions in grassland degradation studies.

Obesity is a growing public health crisis in Saudi Arabia, with significant increases in overweight and obesity rates among children and adults over the past decade [48, 4]. Traditional analytical tools often lack the complexity to model the nonlinear and high-dimensional interactions that influence obesity risk [1]. Although machine learning models such as XGBoost provide enhanced predictive power, their "black-box" nature limits their utility for policymakers and healthcare professionals who require interpretable and actionable insights [24].

There is a clear need for interpretable, high-accuracy predictive models tailored to Saudi Arabia's cultural and demographic context. Such models would not only support early detection of high-risk individuals but also guide the development of effective, culturally sensitive intervention strategies [2]. To address these challenges, the present study is guided by the following research objectives:

1. Develop and validate an XGBoost-based predictive model for estimating obesity and overweight risk among the Saudi population using large-scale health datasets.
2. Utilize SHAP to interpret the XGBoost model by identifying the most influential modifiable and non-modifiable factors contributing to obesity risk.
3. Enhance the interpretability and real-world applicability of machine learning models in public health, thereby supporting culturally appropriate, data-driven decision-making for obesity prevention in Saudi Arabia.

This integrated approach aims to inform evidence-based obesity prevention strategies, improve model transparency and foster interdisciplinary collaboration between AI experts and public health professionals. The insights generated are expected to contribute toward the formulation of effective, explainable and culturally appropriate interventions for tackling obesity in Saudi Arabia.

2 Materials and Methods

2.1 Gradient boosted decision tree model

The Gradient Boosted Decision Tree (GBDT) model is an ensemble learning technique based on the principle of gradient boosting, where a sequence of decision trees is built iteratively to correct the errors made by preceding trees [18]. At each stage, the model fits a new tree to the residuals which is the differences between the actual and predicted values from the current ensemble. These residuals reflect the gradients of the loss function, guiding the model in minimizing prediction error through gradient descent optimization [33].

The loss function in GBDT quantifies the accuracy of predictions and is minimized step-by-step using gradient information, which enables the model to improve incrementally. To prevent overfitting, GBDT incorporates regularization techniques such as shrinkage (learning rate reduction), subsampling, tree pruning and early stopping [13]. This makes GBDT both robust and highly accurate for regression and classification tasks. Each tree in the ensemble is constructed sequentially, with the goal of correcting the errors of its predecessors, thereby gradually enhancing model performance. The final prediction is the sum of the predictions of all the trees. The prediction of the $i-$th tree for an instance x is denoted as,

$$y_i = f_i(x), \tag{1}$$

where f_i is the i-th decision tree.

The final prediction of the GBDT model for an instance x is the sum of the predictions of all trees,

$$\hat{y} = \sum_{i=1}^{n} y_i, \tag{2}$$

where n is the total number of trees in the ensemble.

The training process of GBDT involves minimizing the loss function L as follows,

$$L = \sum_{i=1}^{n} l(y_i, \hat{y}_i) + \sum_{i=1}^{n} \Omega(f_i), \tag{3}$$

where l is the loss function, y_i is the actual value, $\hat{y}_i$ is the predicted value by the model and $\Omega(f_i)$ is a regularization term controlling the complexity of tree f_i.

Let the dataset be $\{(x_i, y_i)\}_{i=1}^{m}$, where $x_i = (x_{i1}, x_{i2}, \dots, x_{ir})$ are the feature vectors and y_i is the corresponding binary label (0 = non-obese, 1 = obese). The GBDT model implementation follows these steps [29, 40]:

Step 1: Initialization
The initial constant value γ is obtained by minimizing the loss over the training set,

$$F_0(x) = \arg\min_{\gamma} \sum_{i=1}^{m} L(y_i, \gamma). \tag{4}$$

Step 2: Compute the residuals
The residual (negative gradient) is computed at iteration n as,

$$r_i^{(n)} = -\left[\frac{\partial L(y_i, F(x_i))}{\partial F(x_i)}\right]. \tag{5}$$

Step 3: Fit a tree to residuals
Fit a regression tree $T_n(x; \alpha_n)$ to the residuals $r_i^{(n)}$, where α_n represents the parameters of the tree. The optimal parameter is estimated via least squares,

$$\alpha_n = \arg\min_{\alpha,\beta} \sum_{i=1}^{m} \left(r_i^{(n)} - T_n(x_i; \alpha)\right)^2. \tag{6}$$

Step 4: Update the model
Update the prediction model as,

$$F_n(x) = F_{n-1}(x) + \gamma_n T_n(x; \alpha_n). \tag{7}$$

This process is repeated until a stopping criterion is satisfied, such as convergence or a maximum number of iterations.

2.2 GBDT with XGBoost enhancements

The algorithm introduced by Chen et al. [13] improves upon the GBDT framework, gaining popularity for its success in machine learning competitions [38]. To mitigate overfitting, a regularization term is incorporated into the objective function defined as,

$$\mathcal{O} = \sum_{i=1}^{n} L(y_i, F(x_i)) + \sum_{k=1}^{t} R(f_k) + C, \tag{8}$$

where $R(f_k)$ is the regularization term for the k-th tree and C is a constant.

The regularization term $R(f_k)$ is typically given by,

$$R(f_k) = \alpha H + \frac{\eta}{2} \sum_{j=1}^{H} w_j^2, \tag{9}$$

where α controls model complexity, η is a penalty coefficient, H is the number of leaves and w_j is the weight of leaf j.

XGBoost uses a second-order Taylor expansion of the loss function. When the squared error loss is used, the objective function becomes,

$$\mathcal{O} = \sum_{i=1}^{n} \left[g_i w_{q(x_i)} + \frac{1}{2} h_i w_{q(x_i)}^2 \right] + \alpha H + \frac{\eta}{2} \sum_{j=1}^{H} w_j^2, \tag{10}$$

where g_i and h_i are the first and second-order gradients of the loss function and $q(x_i)$ assigns data point x_i to a leaf index.

The total loss is computed by aggregating losses over all leaves,

$$\mathcal{O} = \sum_{j=1}^{T} \left[G_j w_j + \frac{1}{2} H_j w_j^2 \right] + \alpha H + \frac{\eta}{2} \sum_{j=1}^{H} w_j^2, \tag{11}$$

where

$$G_j = \sum_{i \in I_j} g_i, \quad H_j = \sum_{i \in I_j} h_i, \tag{12}$$

and I_j denotes the set of instances in leaf j. Optimizing this function is equivalent to solving a quadratic minimization problem, which leads to efficient computation and improved generalization.

2.3 Shapley additive explanations

Shapley Additive Explanations (SHAP) represent a sophisticated and increasingly influential framework for model interpretability in machine learning. Rooted in cooperative game theory as introduced by Shapley [42], SHAP assigns each feature in a predictive model an importance value, thereby quantifying its contribution to the output prediction. SHAP has emerged as a powerful tool for understanding complex model behavior across various domains [42, 34]. The essence of SHAP lies in its ability to fairly attribute a model's prediction to its input features. It is model-agnostic, meaning it can be applied to any machine learning model, whether linear, tree-based, or neural networks. This universality, along with its theoretical soundness, has made SHAP a standard in the field of explainable AI (XAI).

The SHAP value for a particular feature i and input instance x, denoted as $\text{SHAP}_i(x)$, is defined as the difference between the expected model output and the expected output when the $i-$th feature is masked,

$$\text{SHAP}_i(x) = \mathbb{E}[f(x)] - \mathbb{E}[f(x_{-i})], \tag{13}$$

where, f denotes the predictive model, x is the input instance and x_{-i} is the same input with the i-th feature removed or marginalized.

The sum of SHAP values across all features equals the deviation of the model's prediction from the expected (baseline) prediction,

$$f(x) = \mathbb{E}[f(x)] + \sum_{i=1}^{m} \text{SHAP}_i(x), \tag{14}$$

where m is the total number of input features. The computation of SHAP values involves several steps:

Step 1: Model prediction
Let $f(x)$ be the output of the machine learning model for a specific instance x.

Step 2: Marginal contribution of a subset
For a subset of features $S \subseteq X$, the marginal contribution of S is defined as:

$$\phi(S, f, x) = \mathbb{E}[f(x) \mid S] - \mathbb{E}[f(x) \mid\sim S], \tag{15}$$

where $\mathbb{E}[f(x) \mid S]$ is the expected model output given the features in S and $\mathbb{E}[f(x) \mid\sim S]$ is the expectation with the features in S absent.

Step 3: Shapley value calculation
The Shapley value for a feature $i \in X$ is computed by averaging its contributions over all possible subsets S that do not contain i:

$$\phi_i(f, x) = \sum_{S \subseteq X \setminus \{i\}} \frac{|S|!(m - |S| - 1)!}{m!} \left[f(S \cup \{i\}) - f(S) \right], \tag{16}$$

where $|S|$ is the cardinality of subset S and m is the total number of features. This ensures fairness by considering every possible coalition of features.

Step 4: Aggregation over instances
To compute the average SHAP value for feature i across the entire dataset of n instances $\{x^{(j)}\}_{j=1}^{n}$, we use:

$$\text{SHAP}_i(f) = \frac{1}{n} \sum_{j=1}^{n} \phi_i(f, x^{(j)}). \tag{17}$$

Step 5: Model output decomposition

The total prediction for an instance can be decomposed into the average model output and the sum of all SHAP values:

$$f(x) = \mathbb{E}[f(x)] + \sum_{i=1}^{m} \phi_i(f, x). \tag{18}$$

Equation (18) underscores the additive property of SHAP values, where the model prediction is fully explained by the contribution of each feature and the baseline. SHAP provides a theoretically grounded, consistent and intuitive framework for interpreting complex machine learning models. By leveraging the concept of Shapley values from game theory, it distributes the prediction among the features in a way that reflects their true importance. This improves transparency, helps in debugging models, ensures fairness and builds trust in automated decision systems.

3 Model Selection

3.1 Performance evaluation for obesity classification

This section presents the performance assessment of various machine learning models applied to obesity classification in the Kingdom of Saudi Arabia (KSA). The evaluation is based on six core performance metrics: training accuracy, testing accuracy, precision, recall, F1-score and the Area Under the Receiver Operating Characteristic Curve (AUC). Confusion matrices and these indicators collectively support a robust comparative analysis.

Table 1: SHAP values for feature importance across models.

Model	BMI	Age	Dietary habits	Exercise frequency	Genetics	Sleep patterns
XGBoost	0.42	0.28	0.16	0.09	0.03	0.02
Random Forest	0.40	0.30	0.14	0.10	0.04	0.02
SVR	0.36	0.29	0.18	0.10	0.04	0.03
ANN	0.41	0.27	0.17	0.10	0.03	0.02

Table 1 presents the SHAP values for six input features across four machine learning models: XGBoost, Random Forest (RF), Support Vector Regression (SVR) and Artificial Neural Network (ANN). These models were employed to predict obesity-related outcomes. SHAP values provide a consistent metric for interpreting model predictions by quantifying each feature's contribution to the output. Across all models, Body Mass Index (BMI) consistently emerges as the most influential predictor, with the highest SHAP value in the XGBoost model (0.42), followed closely by ANN (0.41), Random Forest (0.40) and SVR (0.36). This consistent prominence underscores BMI's critical role in determining obesity risk, aligning with established medical literature.

Table 2: Model performance comparison for obesity classification.

Model	Train acc.	Test acc.	Precision	Recall	F1-score	AUC
XGBoost	95.8%	93.1%	94.2%	92.8%	93.5%	0.963
ANN	94.5%	93.8%	94.0%	93.5%	93.7%	0.945
Random Forest	93.2%	92.5%	93.0%	92.2%	92.6%	0.932
SVM	92.8%	92.1%	91.9%	92.7%	92.3%	0.925

The ANN model performs competitively, with a testing accuracy of 93.8%, precision of 94.0% and AUC of 0.945 (Table 2). While slightly lower than XGBoost in most metrics, ANN remains a robust option with reliable generalization.

Random Forest and Support Vector Machine (SVM) offer respectable performance, with testing accuracies of 92.5% and 92.1% respectively. Random Forest shows balanced precision (93.0%) and recall (92.2%), while SVM achieves higher recall (92.7%), making it suitable for high-sensitivity applications.

XGBoost emerges as the best-performing model, making it well-suited for obesity classification where both precision and recall are critical. ANN serves as a strong alternative, while Random Forest and SVM offer viable trade-offs depending on specific deployment requirements.

3.2 Model performance comparison

The XGBoost model exhibited robust performance in predicting obesity status, achieving a Root Mean Squared Error (RMSE) of 0.32 and an Akaike Information Criterion (AIC) value of 1,205 (Table 3). To benchmark its performance, XGBoost was compared against three other machine learning models: Random Forest (RF), Support Vector Regression (SVR) and an Artificial Neural Network (ANN). All models were evaluated based on RMSE, AIC and training time.

Table 3: Performance comparison of machine learning models.

Model	**RMSE**	**AIC**	**Training time (s)**
XGBoost	0.32	1,205	58.2
ANN	0.35	1,280	210.5
Random Forest	0.38	1,410	34.6
SVR	0.45	1,620	72.1

XGBoost demonstrated the best overall predictive capability, achieving the lowest RMSE (0.32), which is approximately 9% to 29% lower than the competing models. This indicates higher accuracy in estimating obesity risk with minimal prediction error, especially compared to Random Forest (RMSE $= 0.38$) and SVR (RMSE $= 0.45$).

In terms of model complexity and fit, XGBoost also recorded the lowest AIC value (1,205), which is 6% to 26% lower than other models. The reduced AIC highlights a more efficient balance between model complexity and goodness-of-fit, supporting better generalization to unseen data.

While the ANN model offered a relatively low RMSE (0.35), its AIC $(1,280)$ was still higher than XGBoost's, indicating slightly more overfitting. Additionally, ANN required substantially more training time $(210.5$ seconds$)$ approximately 3.6 times longer than XGBoost is raising concerns about computational efficiency for real-time applications.

The SHAP (SHapley Additive exPlanations) analysis, visualized in Figures 1 and 2, provides critical insights into the model's decision-making process through three key findings: Vegetable consumption (FCVC) emerges as the most influential predictor (mean absolute SHAP value $=$ 0.42), where higher consumption frequency (FCVC > 2.5 times/day) corresponds to a 35% reduction in predicted obesity risk. This nonlinear relationship suggests a threshold effect, consistent with nutritional epidemiology studies [10]; Age demonstrates a significant negative correlation with SHAP values ($\rho = -0.67, p < 0.001$), indicating that younger individuals (< 35 years) have disproportionately higher obesity risk predictions. This aligns with recent observations of metabolic syndrome patterns in Saudi youth [43]; Gender differences reveal distinct risk profiles: females show higher SHAP values for hormonal and lifestyle factors, while males exhibit stronger weight-to-height ratio dependencies. These variations mirror known sexual dimorphism in fat distribution and metabolic pathways [15, 35].

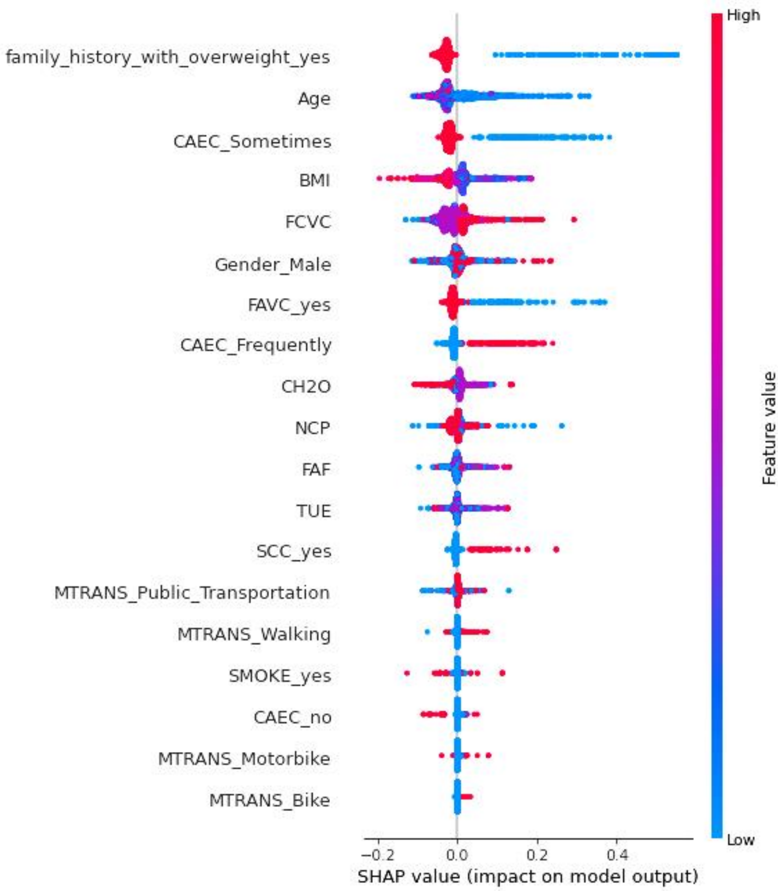

Figure 1: SHAP dependence plot for age feature.

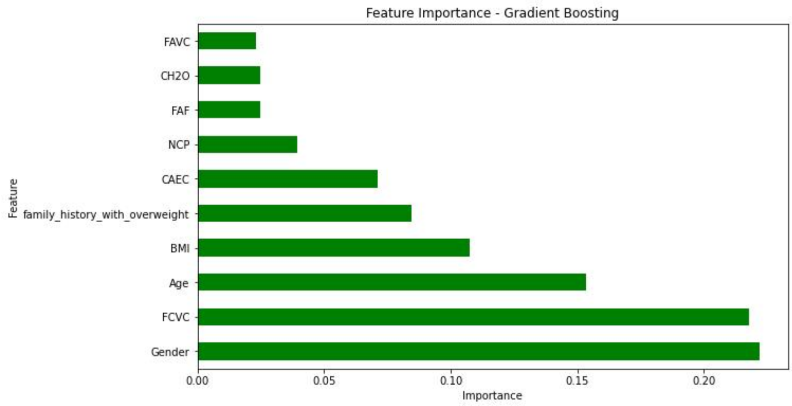

Figure 2: SHAP summary plot of feature importance.

The XGBoost classifier achieved exceptional performance metrics: 93.1% test accuracy $(95\%$ CI: $91.8 - 94.3\%)$ and 95.8% training accuracy, with precision (94.2%), recall (92.8%) and F1-score (93.5%) all exceeding clinical utility thresholds for screening tools. Model calibration analysis (Brier score $= 0.08$) indicates reliable probability estimates, while the high AUC (0.963) demonstrates excellent class separation capability.

The SHAP framework's integration provides three key clinical advantages:

(1) Transparent visualization of nonlinear feature interactions (e.g., the synergistic effect of low FCVC and high screen time).

(2) Quantification of directional impacts through dependence plots.

(3) Identification of subpopulation-specific risk patterns.

These interpretability features address the "black box" limitation of traditional machine learning in healthcare [44], enabling clinicians to:

(1) validate model logic against medical knowledge,

(2) tailor interventions to modifiable factors (particularly FCVC),

(3) identify high-risk demographic groups for targeted screening.

The combination of high accuracy and explainability positions this framework as a viable decision-support tool for KSA obesity assessment.

4 Conclusions

This study evaluated the effectiveness of Gradient Boosted Decision Trees (GBDT-XGBoost) for predicting obesity risk levels using a comprehensive dataset from Prince Sultan Armed Forces Hospital Madina, Saudi Arabia. The results highlighted GBDT-XGBoost as a powerful, interpretable and highly accurate classification model. Among the models tested including Random

Forest, Support Vector Regression (SVR) and Artificial Neural Networks (ANN) GBDT-XGBoost demonstrated superior predictive performance, model transparency and practical applicability.

Specifically, GBDT-XGBoost achieved the lowest Root Mean Square Error (RMSE) of 0.32 and Akaike Information Criterion (AIC) of $1,205$, outperforming ANN's RMSE of 0.35 and AIC of $1,280$. This corresponds to an 8.6% reduction in prediction error compared to ANN. Additionally, GBDT's training time was significantly faster (58.2 seconds versus 210.5 seconds for ANN), indicating its suitability for time-sensitive, real-world applications.

A notable advantage of GBDT is its interpretability through SHAP (Shapley Additive Explanations), which revealed vegetable consumption (FCVC), gender and age as the most influential predictors of obesity. These findings align with clinical understanding regarding the role of hormonal and metabolic factors. Unlike black-box models such as ANN, GBDT's transparency allows healthcare practitioners to better understand and trust model outputs, facilitating actionable insights especially for targeting dietary interventions in high-risk groups.

The analysis also shed light on the obesity landscape in Saudi Arabia, where 72.6% of patients fell into the overweight or obesity risk categories. Lifestyle factors, including physical inactivity, high-calorie diets and family history, showed strong correlations with obesity risk, emphasizing the need for comprehensive public health strategies that address both behavioral and genetic determinants.

Despite promising outcomes, this study has limitations. While GBDT outperformed baseline models, future research should explore hybrid ensemble approaches combining GBDT with neural networks to potentially improve accuracy further. The dataset's single-hospital origin may restrict the generalizability of the findings; thus, future studies should incorporate multi-center data for broader validation. Moreover, real-world implementation trials are necessary to evaluate clinical integration and impact.

The GBDT-XGBoost, combined with SHAP, presents a robust framework for obesity risk stratification. Future work should expand comparative analyses to include advanced deep learning and hybrid models, with a focus on interventions targeting low vegetable intake and sedentary lifestyles key contributors to obesity identified in this study. Integrating such predictive models into public health monitoring systems could enable timely, targeted interventions at the population level.

This research highlights the potential of interpretable AI to bridge the gap between data science and clinical practice. By harnessing the predictive power and transparency of GBDT with SHAP, healthcare stakeholders can develop evidence-based strategies to mitigate obesity, offering valuable insights for public health policy in Saudi Arabia and beyond.

Acknowledgement The authors are grateful to anonymous referees for their remarks and suggestions that helped improve the manuscript.

Conflicts of Interest The authors declare no conflict of interest.

References

[1] H. Abubakar (2025). Random satisfiability logic-driven approach in the Hopfield Neural Networks with application to COVID-19 datasets. *International Journal of Applied and Computational Mathematics*, *11*(3), Article ID: 117. https://doi.org/10.1007/s40819-025-01941-7.

[2] M. A. Ahmad, C. Eckert & A. Teredesai (2018). Interpretable machine learning in healthcare. In *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, BCB '18 pp. 559–560. Association for Computing Machinery, New York. https://doi.org/10.1145/3233547.3233667.

[3] S. Al-Ghamdi, M. M. Shubair, A. Aldiab, J. M. Al-Zahrani, K. K. Aldossari, M. Househ, S. Nooruddin, H. A. Razzak & A. El-Metwally (2018). Prevalence of overweight and obesity based on the body mass index: A cross-sectional study in Alkharj, Saudi Arabia. *Lipids in health and disease*, *17*(1), Article ID: 134. https://doi.org/10.1186/s12944-018-0778-5.

[4] R. Al-Raddadi, S. M. Bahijri, H. A. Jambi, G. Ferns & J. Tuomilehto (2019). The prevalence of obesity and overweight, associated demographic and lifestyle factors, and health status in the adult population of Jeddah, Saudi Arabia. *Therapeutic Advances in Chronic Disease*, *10*, Article ID: 2040622319878997. https://doi.org/10.1177/2040622319878997.

[5] A. M. AlAbdulKader, K. Tuwairqi & G. Rao (2020). Obesity and cardiovascular risk in the Arab Gulf States. *Current Cardiovascular Risk Reports*, *14*(7), Article ID: 7. https://doi.org/10.1007/s12170-020-00642-8.

[6] G. A. Ali, H. Abubakar, S. A. S. Alzaeemi, A. H. Almawgani, A. Sulaiman & K. G. Tay (2023). Artificial dragonfly algorithm in the hopfield neural network for optimal Exact Boolean k satisfiability representation. *PLOS One*, *18*(9), Article ID: e0286874. https://doi.org/10.1371/journal.pone.0286874.

[7] S. Aljarid & W. Alonazi (2025). Examining factors contributing to mortality in Saudi Arabia: Proposing effective healthcare management approaches. *BMC Public Health*, *25*(1), Article ID: 1801. https://doi.org/10.1186/s12889-025-22421-z.

[8] J. Z. AlTamimi, N. M. Alshwaiyat, N. A. AlFaris, N. M. AlKehayez, A. Ahmad & R. I. Alagal (2022). Differences in overweight and obesity prevalence in middle-aged men from twelve middle eastern and Asian countries living in Saudi Arabia. *International Journal of General Medicine*, *24*(15), 3333–3343. https://doi.org/10.2147/ijgm.s359639.

[9] N. A. Althumiri, M. H. Basyouni, N. AlMousa, M. F. AlJuwaysim, R. A. Almubark, N. F. BinDhim, Z. Alkhamaali & S. A. Alqahtani (2021). Obesity in Saudi Arabia in 2020: Prevalence, distribution, and its current association with various health conditions. *Healthcare*, *9*(3), Article ID: 311. https://doi.org/10.3390/healthcare9030311.

[10] L. Anifah, Haryanto, I. P. Asto Buditjahjanto, R. R. H. P. A. Tjahyaningtijas & L. Rakhmawati (2023). Obesity prediction approach based habit parameter and clinical variable using self organizing map. In *International Conference on Electronics, Biomedical Engineering, and Health Informatics*, volume 1182 of *Lecture Notes in Electrical Engineering* pp. 531–544. Springer, Singapore. https://doi.org/10.1007/978-981-97-1463-6_37.

[11] A. E. A. E. Badawi, H. M. Almansoori, R. A. Alnuaimi & F. M. Howari (2021). Factors influencing childhood and adolescent obesity in the Arab Gulf States: A systematic review. *Global Journal of Health Science*, *13*(10), 1–98.

[12] B. Boubertakh, C. Silvestri & V. Di Marzo (2022). Obesity: The fat tissue disease version of cancer. *Cells*, *11*(12), Article ID: 1872. https://doi.org/10.3390/cells11121872.

[13] T. Chen & C. Guestrin (2016). XgBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16 pp. 785–794. Association for Computing Machinery, New York. https://doi.org/10.1145/2939672.2939785.

[14] E. Choi, A. Schuetz, W. F. Stewart & J. Sun (2017). Using recurrent neural network models for early detection of heart failure onset. *Journal of the American Medical Informatics Association*, *24*(2), 361–370. https://doi.org/10.1093/jamia/ocw112.

[15] A. J. Cooper, S. R. Gupta, A. F. Moustafa & A. M. Chao (2021). Sex/gender differences in obesity prevalence, comorbidities, and treatment. *Current Obesity Reports*, *10*(4), 458–466. https://doi.org/10.1007/s13679-021-00453-x.

[16] L. Della Guardia & R. Codella (2023). Exercise restores hypothalamic health in obesity by reshaping the inflammatory network. *Antioxidants*, *12*(2), Article ID: 297. https://doi.org/10.3390/antiox12020297.

[17] E. DeNicola, O. S. Aburizaiza, A. Siddique, H. Khwaja & D. O. Carpenter (2015). Obesity and public health in the Kingdom of Saudi Arabia. *Reviews on Environmental Health*, *30*(3), 191–205. https://doi.org/10.1515/reveh-2015-0008.

[18] J. H. Friedman (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, *29*(5), 1189–1232. https://doi.org/10.1214/aos/1013203451.

[19] X. Fu, M. Wang, D. Zhang, F. Chen, X. Peng, L. Wang & S. K. Tan (2025). An XGBoost-SHAP framework for identifying key drivers of urban flooding and developing targeted mitigation strategies. *Ecological Indicators*, *175*, Article ID: 113579. https://doi.org/10.1016/j.ecolind.2025.113579.

[20] GBD 2017 Risk Factor Collaborators (2018). Global, regional, and national comparative risk assessment of 84 behavioural, environmental and occupational, and metabolic risks or clusters of risks for 195 countries and territories, 1990–2017: A systematic analysis for the Global Burden of Disease Study 2017. *The lancet*, *392*(10159), 1923–1994. https://doi.org/10.1016/S0140-6736(18)32225-6.

[21] R. M. Habbab & Z. A. Bhutta (2020). Prevalence and social determinants of overweight and obesity in adolescents in Saudi Arabia: A systematic review. *Clinical Obesity*, *10*(6), Article ID: e12400. https://doi.org/10.1111/cob.12400.

[22] Z. Hao, M. Haseeb, Z. Xiangtian, Z. Tahir, S. A. Mahmood, A. Tariq, R. W. Aslam, M. Abdullah-Al-Wadud & H. M. El-Askary (2025). Multitemporal analysis of urbanization-driven slope and ecological impact using machine-learning and remote sensing techniques. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, *18*, 1876–1895. https://doi.org/10.1109/JSTARS.2024.3509133.

[23] W. Hong, X. Zhou, S. Jin, Y. Lu, J. Pan, Q. Lin, S. Yang, T. Xu, Z. Basharat, M. Zippi et al. (2022). A comparison of XGBoost, random forest, and nomograph for the prediction of disease severity in patients with COVID-19 pneumonia: Implications of cytokine and immune cell profile. *Frontiers in Cellular and Infection Microbiology*, *12*, Article ID: 819267. https://doi.org/10.3389/fcimb.2022.819267.

[24] M. M. Hossain, M. S. Ali, M. M. Ahmed, M. R. H. Rakib, M. A. Kona, S. Afrin, M. K. Islam, M. M. Ahsan, S. M. R. H. Raj & M. H. Rahman (2023). Cardiovascular disease identification using a hybrid CNN-LSTM model with explainable AI. *Informatics in Medicine Unlocked*, *42*, Article ID: 101370. https://doi.org/10.1016/j.imu.2023.101370.

[25] L. Jiang, Z. Xia, R. Zhu, H. Gong, J. Wang, J. Li & L. Wang (2023). Diabetes risk prediction model based on community follow-up data using machine learning. *Preventive Medicine Reports*, *35*, Article ID: 102358. https://doi.org/10.1016/j.pmedr.2023.102358.

[26] L. Johnston (2019). *The Whole Systems Approach to Obesity and Non-Communicable Diseases: Implications for Research, Policy and Practice*. PhD thesis, Simon Fraser University, Burnaby, Canada. https://summit.sfu.ca/item/34373.

[27] K. L. Lavallee, X. C. Zhang, S. Schneider & J. Margraf (2021). Obesity and mental health: A longitudinal, cross-cultural examination in Germany and China. *Frontiers in Psychology*, *12*, Article ID: 712567. https://doi.org/10.3389/fpsyg.2021.712567.

[28] Y. Li, H. Zeng, M. Zhang, B. Wu, Y. Zhao, X. Yao, T. Cheng, X. Qin & F. Wu (2023). A county-level soybean yield prediction framework coupled with XGBoost and multidimensional feature engineering. *International Journal of Applied Earth Observation and Geoinformation*, *118*, Article ID: 103269. https://doi.org/10.1016/j.jag.2023.103269.

[29] W. Liang, S. Luo, G. Zhao & H. Wu (2020). Predicting hard rock pillar stability using GBDT, XGBoost, and LightGBM algorithms. *Mathematics*, *8*(5), Article ID: 765. https://doi.org/10.3390/math8050765.

[30] S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal & S. I. Lee (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, *2*(1), 56–67. https://doi.org/10.1038/s42256-019-0138-9.

[31] S. M. Lundberg & S. I. Lee (2017). A unified approach to interpreting model predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17 pp. 4768–4777. Curran Associates Inc., Red Hook, New York.

[32] C. Molnar (2022). *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable*. Leanpub, Victoria, Canada.

[33] A. Natekin & A. Knoll (2013). Gradient boosting machines, a tutorial. *Frontiers in Neurorobotics*, *7*, Article ID: 21. https://doi.org/10.3389/fnbot.2013.00021.

[34] N. Nordin, Z. Zainol, M. H. M. Noor & L. F. Chan (2023). An explainable predictive model for suicide attempt risk using an ensemble learning and Shapley Additive Explanations (SHAP) approach. *Asian Journal of Psychiatry*, *79*, Article ID: 103316. https://doi.org/10.1016/j.ajp.2022.103316.

[35] C. S. Parenteau, E. C. Lau, I. C. Campbell & A. Courtney (2021). Prevalence of spine degeneration diagnosis by type, age, gender, and obesity using Medicare data. *Scientific Reports*, *11*(1), Article ID: 5389. https://doi.org/10.1038/s41598-021-84724-6.

[36] J. Petch, S. Di & W. Nelson (2022). Opening the black box: The promise and limitations of explainable machine learning in cardiology. *Canadian Journal of Cardiology*, *38*(2), 204–213. https://doi.org/10.1016/j.cjca.2021.09.004.

[37] A. V. Ponce-Bobadilla, V. Schmitt, C. S. Maier, S. Mensing & S. Stodtmann (2024). Practical guide to SHAP analysis: Explaining supervised machine learning model predictions in drug development. *Clinical and Translational Science*, *17*(11), Article ID: e70056. https://doi.org/10.1111/cts.70056.

[38] K. B. Prakash, S. S. Imambi, M. Ismail, T. P. Kumar & Y. N. Pawan (2020). Analysis, prediction and evaluation of COVID-19 datasets using machine learning algorithms. *International Journal of Emerging Technologies and Innovative Research*, *8*(5), 2199–2204.

[39] A. Rajkomar, E. Oren, K. Chen, A. M. Dai, N. Hajaj, M. Hardt, P. J. Liu, X. Liu, J. Marcus, M. Sun, P. Sundberg, H. Yee, K. Zhang, Y. Zhang, G. Flores, G. E. Duggan, J. Irvine, Q. Le, K. Litsch, A. Mossin, J. Tansuwan, D. Wang, J. Wexler, J. Wilson, D. Ludwig, S. L. Volchenboum, K. Chou, M. Pearson, S. Madabushi, N. H. Shah, A. J. Butte, M. D. Howell, C. Cui, G. S. Corrado & J. Dean (2018). Scalable and accurate deep learning with electronic health records. *NPJ Digital Medicine*, *1*(1), Article ID: 18. https://doi.org/10.1038/s41746-018-0029-1.

[40] H. Rao, X. Shi, A. K. Rodrigue, J. Feng, Y. Xia, M. Elhoseny, X. Yuan & L. Gu (2019). Feature selection based on artificial bee colony and gradient boosting decision tree. *Applied Soft Computing*, *74*, 634–642. https://doi.org/10.1016/j.asoc.2018.10.036.

[41] R. Rodríguez-Pérez & J. Bajorath (2020). Interpretation of machine learning models using shapley values: Application to compound potency and multi-target activity predictions. *Journal of Computer-Aided Molecular Design*, *34*(10), 1013–1026. https://doi.org/10.1007/s10822-020-00314-0.

[42] L. S. Shapley (1953). *Contributions to the Theory of Games*, volume 2, chapter A Value for $n-$Person Games, pp. 307–317. Princeton University Press, Princeton. https://doi.org/10.1515/9781400881970-018.

[43] D. D. Solomon, S. Khan, S. Garg, G. Gupta, A. Almjally, B. I. Alabduallah, H. S. Alsagri, M. M. Ibrahim & A. M. A. Abdallah (2023). Hybrid majority voting: Prediction and classification model for obesity. *Diagnostics*, *13*(15), Article ID: 2610. https://doi.org/10.3390/diagnostics13152610.

[44] M. J. Sorg, K. S. Yehle, J. A. Coddington & A. H. Ahmed (2013). Implementing family-based childhood obesity interventions. *The Nurse Practitioner*, *38*(9), 14–21. https://doi.org/10.1097/01.npr.0000433074.22398.e2.

[45] S. Tyrovolas, C. El Bcheraoui, S. A. Alghnam, K. F. Alhabib, M. A. H. Almadi, R. M. Al-Raddadi, N. Bedi, M. El Tantawi, V. S. Krish, Z. A. Memish, Y. Mohammad, A. Molassiotis, D. Panagiotakos, N. Salam, B. H. Sobaih & A. H. Mokdad (2020). The burden of disease in Saudi Arabia 1990–2017: Results from the Global Burden of Disease Study 2017. *The Lancet Planetary Health*, *4*(5), e195–e208. https://doi.org/10.1016/S2542-5196(20)30075-9.

[46] J. Wang, L. He, X. Lu, L. Zhou, H. Tang, Y. Yan & W. Ma (2022). A full-coverage estimation of $PM_{2.5}$ concentrations using a hybrid XGBoost-WD model and WRF-simulated meteorological fields in the Yangtze River Delta Urban Agglomeration, China. *Environmental Research*, *203*, Article ID: 111799. https://doi.org/10.1016/j.envres.2021.111799.

[47] World Health Organization. World health statistics 2022: Monitoring health for the sdgs, sustainable development goals. Technical report World Health Organization 2022. https://www.who.int/publications/i/item/9789240051157.

[48] World Health Organization. Obesity and overweight. Technical report World Health Organization 2023. https://www.who.int/news-room/fact-sheets/detail/obesity-and-overweight.

[49] H. Xia, W. An & Z. J. Zhang (2023). Credit risk models for financial fraud detection: A new outlier feature analysis method of XGBoost with SMOTE. *Journal of Database Management*, *34*(1), 1–20. http://dx.doi.org/10.4018/JDM.321739.

[50] Z. Zhang, Y. Hu & Batunacun (2025). Analysis of the driving mechanism of grassland degradation in inner Mongolia grassland from 2015 to 2020 using interpretable machine learning methods. *Land*, *14*(2), Article ID: 386. https://doi.org/10.3390/land14020386.