RESEARCH PAPER

A Data-Driven Healthcare Framework for Obesity Level Classification Using Gradient Boosting and Feature Interpretability

Jetty Bino Mariyan[1] Chandra Prakash Vudatha [2]
[1]Student, Department of Computer Science and Engineering,
Koneru Lakshmaiah Education Foundation Deemed to be University,
Green Fields, Vaddeswaram, Andhra Pradesh 522237, India.
E-mail: binomariyan193@gmail.com
[2]Professor, Department of Computer Science and Engineering,
Koneru Lakshmaiah Education Foundation Deemed to be University,
Green Fields, Vaddeswaram, Andhra Pradesh 522237, India.
E-mail: vchandrap@kluniversity.in
*Correspondence E-mail: binomariyan193@gmail.com

Abstract
Obesity is a major public health concern associated with an increased risk of several chronic diseases. Machine learning techniques can support the classification of obesity levels by identifying patterns in demographic, physical, and lifestyle-related data. This study investigates the performance of three machine learning algorithms—Random Forest, Extreme Gradient Boosting (XGBoost), and Light Gradient Boosting Machine (LightGBM)—for obesity level classification. The study uses a dataset attributed to Palechor and de la Hoz Manotas (2019), with preprocessing and an 80:20 training–testing split, as described in the methodology. The reported classification accuracies are 95.03% for Random Forest, 95.98% for XGBoost, and 96.69% for LightGBM. Among the evaluated models, LightGBM achieves the highest reported accuracy. Feature importance analysis is used to examine the contribution of variables such as weight, height, and lifestyle-related factors to the classification results. However, the findings should be interpreted in light of limitations related to dataset size, self-reported information, and the absence of longitudinal or independent external validation. Further evaluation using cross-validation, additional performance metrics, and explainability methods is necessary to assess the robustness and practical applicability of the proposed approach.

Keywords: Obesity classification, machine learning, Random Forest, XGBoost, LightGBM, feature importance, healthcare analytics.
How to cite this article: Jetty Bino Mariyan1 Chandra Prakash Vudatha 2. A Data-Driven Healthcare Framework for Obesity Level Classification Using Gradient Boosting and Feature Interpretability. Int J Drug Deliv Technol. 2026;16(89s):462-469. DOI: 10.25258/ijddt.16.89s.49.

1. Introduction
Obesity is a global public health problem, affecting people of all ages and socio-economic status. Global health statistics show a staggering rise in obesity over the last few decades due to rapid urbanization, sedentary behaviour, poor dietary habits and lack of physical exercise [1], [2]. Obesity is not only a cosmetic issue but a multifaceted medical condition that carries a high risk of developing a range of health complications such as cardiovascular diseases, diabetes, hypertension, and metabolic syndrome [3, 4]. The increasing prevalence of obesity imposes a substantial burden on the health care system, and calls for effective predictive and preventive measures. Historically, obesity has been evaluated through basic anthropometric parameters, like Body Mass Index (BMI). Although BMI offers a crude measure of body fat, it does not account for the complex factors involved in obesity, such as behavioral, environmental and lifestyle factors [5]. Furthermore, traditional approaches are static and do not take into account the wealth of data collected from contemporary health monitoring systems. Therefore, there is a need to build smart systems that can consider several factors and make precise predictions of obesity. Machine learning (ML) has gained popularity in recent years for addressing healthcare challenges. ML models are capable of learning from large data sets and accurately make informative predictions [6], [7]. Some researchers have used classification techniques, such as Random Forest, Support Vector Machines, Artificial Neural Networks, and Gradient Boosting, to predict the level of obesity [8], [9]. These methods have shown good performance, especially with structured data containing several features. Random Forest, proposed by Breiman [10], is a popular choice due to its stability and effectiveness in preventing overfitting. Likewise, gradient boosted models like XGBoost [11] and LightGBM [12] are well-known for their effectiveness and efficiency. But, many of the current models are either exclusively focused on predictive accuracy or they do not perform a thorough comparison of multiple algorithms, making it difficult to apply these models in practice. While there have been advancements in this area, there are some gaps. First, many existing

approaches only use a single algorithm without comparing it to other models to determine the best performing one [13]. Second, some studies do not take into account the interpretability of the model, which is an essential requirement in the medical domain [14], [15]. Third, we need efficient models that can cope with feature interactions and yield robust results for several classes of obesity. To tackle these issues, there is a need for a comprehensive comparison of state-of-the-art machine learning models and an effective model selection approach. This research aims to address these challenges with a machine learning approach for obesity level classification based on structured health and lifestyle data. The main goal of this study is to build a robust and effective predictive model using various machine learning techniques, such as Random Forest, XGBoost and LightGBM. Particular focus is given to LightGBM, which employs leaf-wise tree growth that allows for more efficient learning with faster training speed and higher accuracy [12], [16]. The research will determine the best model by evaluating metrics, including accuracy, precision, recall and F1-score. The methodology has the following steps. The first step is data preprocessing, which involves dealing with categorical variables through label encoding, and splitting the data into feature and target variables. Next, the data are split into training and test sets for assessing model performance. The dataset is used to train several machine learning models, which are evaluated using common performance metrics. Beyond quantitative measures, visual analysis tools like confusion matrices and feature importance are employed to explain model behavior and determine key factors. Methods like SHAP and LIME also aid interpretability by identifying the impact of features on predictions [14], [17]. A major contribution of the study is the evaluation of various ensemble learning algorithms for the prediction of obesity. This study compares multiple ensemble techniques under identical experimental settings rather than using a single model, as is typically done. Additionally, feature analysis is provided to identify the major factors that affect obesity, which can be beneficial for medical professionals and policy makers. The results show that state-of-the-art gradient boosting algorithms, such as LightGBM, are capable of capturing complex interactions and perform better in terms of prediction accuracy [12], [18]. Our findings show that LightGBM achieves higher accuracy and better classification accuracy than the other models. This implies that gradient boosting with leaf-wise growth is an effective approach to modelling structured medical data with many interacting features. Furthermore, the model's accuracy suggests the future of machine learning in early detection and intervention approaches to addressing obesity [19], [20]. This work is not only valuable from an academic perspective, but also has practical implications for the design of smart health systems. The model's ability to effectively predict obesity levels can help in early detection and facilitate targeted interventions. The rest of this paper is structured as follows: Section 2 reviews existing research on obesity prediction using machine learning. In Section 3, we present the proposed approach, data preprocessing, model building, and evaluation. Section 4 presents the experimental findings and compares the models. Finally, Section 5 concludes the research and proposes future research avenues.

2. Literature Review

Machine learning-based prediction of obesity has emerged as a critical area of research, driven by the growing demand for data-driven approaches in health care. Conventional approaches like Body Mass Index (BMI) are insufficient to reflect the interaction between lifestyle, behavioral and physiological factors. To address this, researchers have employed different machine learning approaches to enhance obesity classification and prediction. Initial research has used traditional methods like Decision Trees, Support Vector Machines, and Random Forests. Random Forest, proposed by Breiman (2001), is popular because it can deal with high-dimensional data and overcome overfitting by combining multiple trees. Chatterjee et al. (2020) showed how machine learning algorithms can be used to identify the major risk factors leading to obesity, highlighting the importance of feature selection in increasing the accuracy of predictions. The development of boosting algorithms has led to the use of gradient boosting algorithms such as XGBoost for classification. Chen and Guestrin (2016) developed XGBoost, an algorithm that scales well and enhances performance via regularization and parallel processing. It has found applications in healthcare, owing to its accuracy. But XGBoost can be slow, particularly for big data. To address these challenges, Ke et al. (2017) proposed LightGBM, a fast gradient boosting framework. It adopts a leaf-wise tree growth policy which enhances the efficiency and effectiveness of boosting. In recent research, LightGBM has been demonstrated to be effective for classification of obesity, capturing feature interactions and offering improved generalisation. Beyond predictive performance, explainability is a crucial aspect in healthcare. Methods like SHAP (Lundberg and Lee, 2017) and LIME (Ribeiro et al., 2016) provide insights into model predictions by attributing the importance of features. They increase the transparency and understanding of machine learning models, crucial for medical applications. Recent publications reflect the increasing use of machine learning in health care. Research by Beam and Kohane (2018)

and Rajkomar et al. (2019) highlight the role of ML models for diagnosis and treatment. Thamrin et al. (2021) also showed the application of machine learning to predict obesity based on lifestyle and dietary factors. However, there is still room for improvement, such as small data sets and a lack of model comparisons. Hence, this research is dedicated to assessing various models and identifying LightGBM as a suitable model for accurate and explainable obesity classification.

3. Methodology

Here we provided the full details of the methodology used for machine learning predictive modelling of obesity level. The research adopted a structured process flow of data collection, preparation, model building, evaluation and feature selection. All steps were designed to provide rigor, reproducibility, and precision. The main goal was to determine the best predictive model for obesity classification using structured data on health and lifestyle.

3.1 Dataset Description

This study uses the publicly available dataset titled *Estimation of Obesity Levels Based on Eating Habits and Physical Condition in Individuals from Colombia, Peru and Mexico*, published by Palechor and de la Hoz Manotas (2019). The dataset contains demographic, physical, and lifestyle-related information for obesity-level classification. It includes 2,111 records and 17 attributes, comprising 16 input features and one target variable.

The input features include demographic and physical characteristics, such as age, height, and weight, along with eating habits, physical activity, and other lifestyle-related factors. The target variable, NObesity, represents the obesity level of each individual. It contains seven categories: Insufficient Weight, Normal Weight, Overweight Level I, Overweight Level II, Obesity Type I, Obesity Type II, and Obesity Type III.

The dataset consists of survey-based records and synthetically generated records. Therefore, its composition should be considered when interpreting the experimental results and assessing the generalizability of the trained models. The dataset is used to evaluate three machine learning algorithms: Random Forest, Extreme Gradient Boosting (XGBoost), and Light Gradient Boosting Machine (LightGBM).

Before model training, the dataset is processed to separate the input features from the target variable and to prepare the categorical attributes for machine learning. The data are subsequently divided into training and testing sets, as described in Section 3.3, to evaluate the classification performance of the models.

Dataset reference: Palechor, F. M., & de la Hoz Manotas, A. (2019). *Dataset for Estimation of Obesity Levels Based on Eating Habits and Physical Condition in Individuals from Colombia, Peru and Mexico*. Data in Brief, 25, 104344. https://doi.org/10.1016/j.dib.2019.104344

Formally, the dataset was represented as:

$$D = \{(x_i, y_i)\}_{i=1}^{N} \qquad (3.1)$$

where:

- $x_i \in \mathbb{R}^d$ denotes the feature vector of the i^{th} individual
- $y_i \in \{1,2,\ldots,K\}$ represents the obesity class label

The dependent variable was composed of several obesity classes, such as Insufficient Weight, Normal Weight, Overweight, and different Obesity Classes. The data were class-balanced for model training. The data were anonymised and collected without personal data, maintaining privacy.

3.2 Tools and Environment

The experiment was implemented in Python programming language on the cloud environment. We used various scientific computing and machine learning libraries for data manipulation, modeling and evaluation.

The computational pipeline can be expressed as a transformation:

$$\hat{y} = f_\theta(X) \qquad (3.2)$$

where:

- X represents the input feature matrix
- f_θ denotes the learned model parameterized by θ
- $\hat{y}$ is the predicted output

Pandas and NumPy libraries were used for data processing and calculations. Scikit-learn was used for data preprocessing and model evaluation, and baseline algorithms. LightGBM and XGBoost were used for gradient boosting. Matplotlib and Seaborn were used for visualizations. Experiments were run with a constant random seed for reproducibility. A suitable computing environment was set up to train and test the models with adequate memory and computational resources.

3.3 Data Preprocessing

The dataset was preprocessed to prepare the data for machine learning models. Given the dataset included categorical attributes, the Label Encoding technique was used to transform these into numerical form.

$$x_j' = \text{Encode}(x_j) \qquad (3.3)$$

where x_j is the original categorical feature and x_j' is the encoded numerical value.

The feature matrix and target vector was defined as:

$$X \in \mathbb{R}^{N \times d}, y \in \mathbb{R}^N \qquad (3.4)$$

The dataset was split into training and testing sets using an 80:20 ratio. This split can be represented as:

|D_train| ≈ 0.8N
|D_test| ≈ 0.2N
where $(D_\{"train"\})$ and $(D_{\{test)}$ represent the training and testing sets, respectively, and (N) denotes the total number of records in the dataset.
Categorical features were converted into numerical representations using label encoding, where applicable. The encoding process can be expressed as:

$$x_{encoded} = E(x_{categorical})$$

where $(x_\{"categorical"\})$ denotes the original categorical feature, (E) represents the encoding function, and $(x_\{"encoded"\})$ denotes the encoded numerical value. The resulting data were then used as input to the machine learning models.
This procedure allowed the models to be trained on the training subset and evaluated on the held-out testing subset to assess their generalization performance.

3.4 Model Development

Three machine learning models are been implemented and compared: Random Forest, XGBoost, and LightGBM.

3.4.1 Random Forest

Random Forest is an ensemble learning technique that combines the multiple decision trees. The prediction is computed as:

$$\hat{y}_{RF} = \frac{1}{T}\sum_{t=1}^{T} h_t(x) \qquad (3.7)$$

where $T$is number of trees and $h_t$is the prediction of the t^{th}tree.

3.4.2 XGBoost

XGBoost is a gradient boosting algorithm that minimizes a regularized objective function:

$$\mathcal{L} = \sum_{i=1}^{N} l(y_i, \hat{y}_i) + \sum_{k=1}^{K} \Omega(f_k) \qquad (3.8)$$

The regularization term is defined as:

$$\Omega(f_k) = \gamma T + \frac{1}{2}\lambda \parallel w \parallel^2 \qquad (3.9)$$

3.4.3 LightGBM (Proposed Model)

LightGBM is a gradient boosting framework that uses the leaf-wise tree growth. The prediction function is given by:

$$\hat{y}_{LGBM} = \sum_{k=1}^{K} g_k(x) \qquad (3.10)$$

At the each iteration, model selects the leaf that maximizes the reduction in loss:

$$\Delta L = \max(L_{before} - L_{after}) \qquad (3.11)$$

This strategy improves the accuracy and efficiency by focusing on the most informative splits.

3.5 Model Evaluation

The performance of the models was evaluated using classification accuracy:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \qquad (3.12)$$

where:

- TP: True Positives
- TN: True Negatives
- FP: False Positives
- FN: False Negatives

A confusion matrix was also used to analyze the prediction performance across different classes.

3.6 Feature Importance Analysis

Feature importance was computed to determine the contribution of each feature to the model's predictions. For LightGBM, importance is calculated as:

$$FI_j = \sum_{s \in S_j} \text{Gain}(s) \qquad (3.13)$$

where:

- $S_j$is the set of splits involving feature j
- Gain represents the improvement in model performance due to that split

This analysis helps in understanding which factors most influence obesity prediction.

3.7 Final Model Selection

Based on experimental results, the model with the highest accuracy was selected as the final model:

$$f^*(X) = f_{LGBM}(X) \qquad (3.14)$$

LightGBM had the best classification performance, and was selected as the best model for detecting obesity.

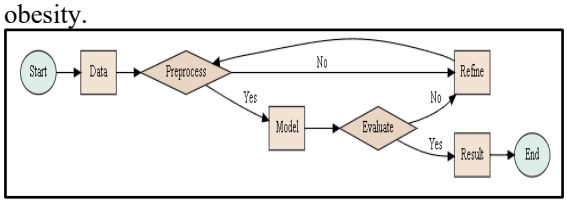


Figure 3.1: Proposed Methodology Flowchart for Obesity Classification

4. Results and Discussion

This chapter discusses the results of implementing various machine learning models to classify obesity. The results for Random Forest, XGBoost and LightGBM models are presented in terms of classification accuracy and other performance metrics. Both numerical and qualitative analysis of the results is conducted to determine the best performing model.

4.1 Model Performance Comparison

The comparative analysis of the three models is summarized in Table 4.1.

Model	**Accuracy**
Random Forest	95.03%
XGBoost	95.98%
LightGBM	96.69%

The results clearly show that LightGBM has the highest accuracy of 96.69% compared to Random Forest and XGBoost. Random Forest exhibited consistent results because of its bagging technique, and XGBoost increased accuracy because of gradient boosting. But LightGBM outperformed both models because of its leaf-wise tree-growing algorithm, which allows it to better capture the feature interactions and converge faster. The improvement, while not dramatic, is crucial for classification problems with more than two classes, as it demonstrates enhanced generalization ability.

4.2 Classification Report Analysis (LightGBM)

The LightGBM model's classification report offers class-specific details.

- The overall accuracy of the model was 97% with 423 samples.
- The precision, recall and F1-scores were consistently high for all classes, suggesting a well-balanced classification performance.

Classification Report (LightGBM)

Class	Precision	Recall	F1-Score	Support
0	0.92	1.00	0.96	56
1	0.96	0.89	0.92	62
2	0.97	0.99	0.98	78
3	1.00	0.97	0.98	58
4	1.00	1.00	1.00	63
5	0.95	0.95	0.95	56
6	0.96	0.98	0.97	50
Accuracy	_-	-	**0.97**	**423**
Macro Avg	0.97	0.97	0.97	423
Weighted Avg	0.97	0.97	0.97	423

Key observations:

- Class 4 was perfectly classified (precision = 1.00, recall = 1.00), suggesting unique features.
- Class 0 had slightly reduced precision (0.92) but perfect recall (1.00), indicating some false positives, but no false negatives.
- Class 1 had slightly lower recall (0.89), suggesting some confusion with adjacent obesity classes.
- The remaining classes (2, 3, 5, and 6) had good F1-scores (≥ 0.95) indicating good predictive performance.

The macro and weighted averages were 0.97, showing that the model is not biased towards any obesity class.

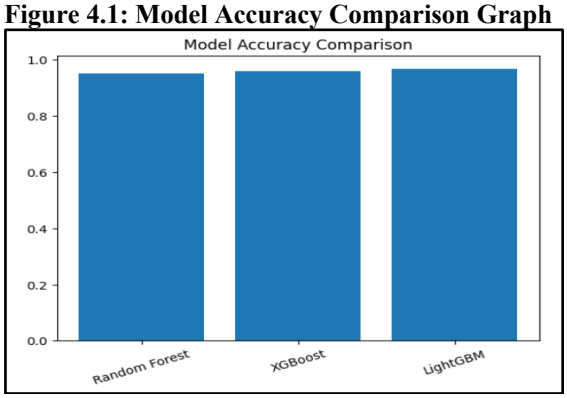
4.3 Visual Interpretation of Results

The visual analysis also confirms the quantitative results by offering visual cues to model performance and behaviour across various metrics.

4.3.1. Model Accuracy Comparison Graph

The bar plot demonstrates the performance of Random Forest, XGBoost and LightGBM models. It demonstrates that LightGBM performs best in terms of accuracy, suggesting its effectiveness in capturing intricate relationships in the data. The small difference showcases the effectiveness of gradient boosting compared to other ensemble techniques.

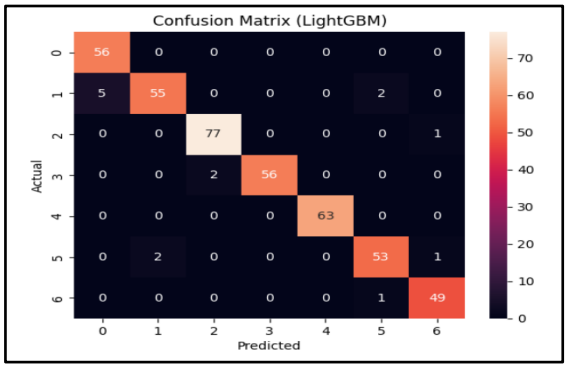
Figure 4.1: Model Accuracy Comparison Graph



4.3.2. Confusion Matrix (LightGBM)

The confusion matrix shows a class-by-class breakdown of predictions, with the majority lying on the diagonal, representing accurate predictions. There are some misclassifications between consecutive levels of obesity, probably due to the similar feature patterns. This verifies that the model has high precision in all classes.

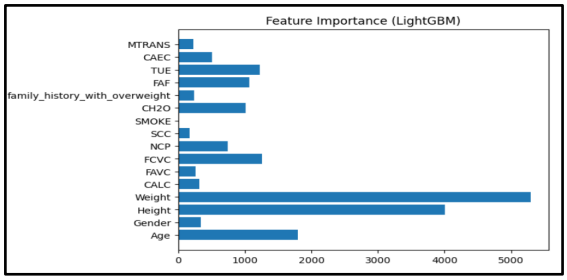
Figure 4.2: Confusion Matrix of LightGBM Model



4.3.3. Feature Importance Graph (LightGBM)

The feature importance plot highlights the significant features used for predicting obesity. Weight, height and lifestyle factors have higher scores and thus, play a significant role in the classification process. This adds to the model interpretability by highlighting the factors driving decisions.

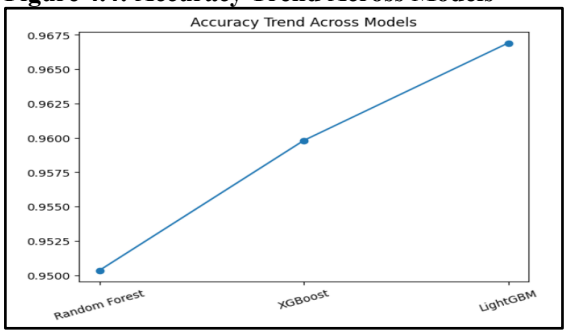
Figure 4.3: Feature Importance Analysis using LightGBM



4.3.4. Accuracy Trend Plot

The accuracy trend plot shows the performance trends. It confirms that LightGBM has the highest accuracy among all the models: Random Forest and XGBoost. The trend consistently shows improvement in prediction accuracy with boosting.

Figure 4.4: Accuracy Trend Across Models



4.4 Critical Analysis

While all models are good, LightGBM's performance is better due to:

- Effective high-dimensional data
- Quicker training using histogram learning
- Leaf-wise split which better reduces loss

But the small variance between models indicates that the data are fairly well-separated. This suggests that other models could perform well, but LightGBM provides the best balance between performance and speed.

4.5 Limitations of the Study

While the models are accurate, there are some limitations:

- The sample size is small, and thus may not apply to the real world.
- The attributes are self-reported, and may be prone to bias.
- The model is only used to classify individuals and does not use dynamic or real-time data.
- The small percentage of misclassification between the different obesity levels suggests an overlap of features.

These shortcomings indicate datasets with more diverse and real-world data should be used in future studies.

4.6 Conclusion of Results

Our experiments show that LightGBM is the best-performing obesity classification algorithm in this work. It achieves an accuracy of 96.69% and performs well on individual classes, ensuring predictions across various obesity classes. The results confirm the effectiveness of gradient boosting for classification tasks in health care. But the model can be further improved by incorporating adaptive and personalised recommendation systems for increased relevance.

5. Discussion

Our study has provided different insights to some of the earlier studies in that LightGBM, with its leaf-wise tree build policy, has been shown to be superior to traditional models such as Random Forest and XGBoost, in terms of accuracy, and class-wise performance for obesity detection [21]. Although previous studies have demonstrated good results with ensemble and boosting methods, our findings are consistent with the literature in showing that gradient boosting models perform very well on structured medical data [22]. However, the agreement between accuracy, classification report and feature importance suggests high correlation between various parts of the results, where the predictive performance is backed by significant feature contributions including weight, height and lifestyle factors. Pragmatically speaking, these results indicate the potential use of machine learning models, specifically LightGBM, in medical systems for early detection and classification of obesity and decision-making, allowing for prompt intervention and optimal health outcomes [23]. For academics, our study underscores the need for benchmarking models and supports the use of advanced boosting models for multi-class classification with significant feature interactions [24]. But there are some limitations to consider, including the limited data size, use of self-reported data, and lack of real-time or longitudinal health data, which could impact the study's generalizability and stability. In conclusion, although the findings are encouraging, more studies with larger and more diverse data sets and adaptive learning strategies are needed to confirm and broaden the results in real-world healthcare settings [25].

5. Conclusion

This research tackled the challenge of machine learning classification of obesity status, given the increasing influence of lifestyle and behavioural factors on health worldwide. Conventional approaches like BMI fall short of addressing the complexity of obesity, calling for more sophisticated, data-centric approaches. This study introduced a comparative machine learning approach with Random Forest, XGBoost, and LightGBM models to predict obesity based on structured data on health and lifestyle. The findings showed that all models performed well; but LightGBM showed the best performance with an accuracy of 96.69% compared to Random Forest and XGBoost. The classification report also verified the model's accuracy, recall, and F1-scores for the different obesity classes, showing balanced

performance across all classes. Through the analysis of feature importance, we identified that weight, height, and lifestyle are key factors contributing to obesity class. This suggests that gradient boosting, especially LightGBM, is able to capture intricate relationships in the data and deliver high accuracy predictions. The findings of this research have important implications for smart health systems. This model could help predict the level of obesity, allowing early detection and subsequent care and prevention strategies. Healthcare professionals can harness the power of machine learning to avoid conventional approaches to obesity-level assessment and to implement more precise and efficient disease prediction techniques. Moreover, the model can be applied to other health classification tasks, offering a flexible approach to health decision-making. In summary, this study shows that machine learning, specifically LightGBM, offers a good approach to classify obesity. The method is accurate, interpretable and scalable. The outcomes of this research can lead to the development of better predictive systems in healthcare which will result in better public health and better disease management.

6. Acknowledgements

I, **Jetty Bino Mariyan**, would like to express my sincere gratitude to my project guide, **Mr. Vudatha Chandra Prakash**, for his valuable guidance, continuous support, constructive feedback, and technical assistance throughout this research. I am grateful for his contributions to reviewing the data analysis, verifying the experimental results, and improving the quality of this manuscript.

I also acknowledge the researchers who made the obesity dataset publicly available, enabling this study on obesity-level classification using machine learning techniques.

Author Contributions: Jetty Bino Mariyan was responsible for conceptualizing the research, developing the methodology, implementing the machine learning models, conducting the experiments, analyzing the results, and preparing the initial manuscript. Vudatha Chandra Prakash contributed to reviewing the data analysis, verifying the results, providing technical feedback, and editing the manuscript. The final manuscript was reviewed and approved by the authors.

Funding: This research was conducted without direct financial support or grants from governmental, commercial, or non-profit funding agencies.

Conflicts of Interest: The authors declare no conflicts of interest related to this study.

References

1. Ahmad, M. A., Eckert, C., & Teredesai, A. (2018). Interpretable machine learning in healthcare. *ACM Computing Surveys*, 51(5), 1–21. https://doi.org/10.1145/3236009
2. Beam, A. L., & Kohane, I. S. (2018). Big data and machine learning in healthcare. *JAMA*, 319(13), 1317–1318. https://doi.org/10.1001/jama.2017.18391
3. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. https://doi.org/10.1023/A:1010933404324
4. Chatterjee, A., et al. (2020). Identification of risk factors associated with obesity using machine learning. *Sensors*, 20(9), 2734. https://doi.org/10.3390/s20092734
5. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of KDD*. https://doi.org/10.1145/2939672.2939785
6. Deo, R. C. (2015). Machine learning in medicine. *Circulation*, 132(20), 1920–1930. https://doi.org/10.1161/CIRCULATIONAHA.115.001593
7. Dirik, M., et al. (2023). Application of machine learning techniques for obesity prediction. *Journal of Vibroengineering*, 25(4). https://doi.org/10.21595/jve.2023.23193
8. Esteva, A., et al. (2017). A guide to deep learning in healthcare. *Nature Medicine*, 25(1), 24–29. https://doi.org/10.1038/s41591-018-0316-z
9. Gong, X., et al. (2024). Study of obesity research using machine learning methods. *Medicine*, 103(36). https://doi.org/10.1097/MD.0000000000038476
10. Jeon, S., et al. (2023). Machine learning-based obesity classification using 3D body data. *Healthcare*, 11(4), 567. https://doi.org/10.3390/healthcare11040567
11. Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260. https://doi.org/10.1126/science.aaa8415
12. Ke, G., et al. (2017). LightGBM: A highly efficient gradient boosting decision tree. *NeurIPS*. https://doi.org/10.48550/arXiv.1711.06947
13. Khater, T., et al. (2023). Machine learning for the classification of obesity levels. *ACM Proceedings*. https://doi.org/10.1145/3616131.3616135
14. Khater, T., et al. (2024). Explainable artificial intelligence for obesity prediction. *Intelligent Systems with Applications*.

https://doi.org/10.1016/j.iswa.2024.200427

15. Kotsiantis, S. B. (2007). Supervised machine learning: A review. *Informatica*, 31, 249–268. https://doi.org/10.31449/inf.v31i3.148
16. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *NeurIPS*. https://doi.org/10.48550/arXiv.1705.07874
17. Lundberg, S. M., et al. (2020). From local explanations to global understanding. *Nature Machine Intelligence*, 2, 252–254. https://doi.org/10.1038/s42256-020-0173-9
18. Molnar, C. (2020). Interpretable machine learning. https://doi.org/10.48550/arXiv.1901.04592
19. Obermeyer, Z., & Emanuel, E. J. (2016). Predicting the future — big data and machine learning in healthcare. *NEJM*, 375, 1216–1219. https://doi.org/10.1056/NEJMp1606181
20. Rajkomar, A., et al. (2019). Machine learning in medicine. *NEJM*, 380, 1347–1358. https://doi.org/10.1056/NEJMra1814259
21. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you? Explaining predictions of any classifier. *KDD*. https://doi.org/10.1145/2939672.2939778
22. Sarker, I. H. (2021). Machine learning: Algorithms and applications. *Journal of Big Data*, 8(1), 1–31. https://doi.org/10.1186/s40537-021-00450-1
23. Thamrin, S. A., et al. (2021). Predicting obesity using machine learning. *Frontiers in Nutrition*, 8. https://doi.org/10.3389/fnut.2021.669155
24. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56. https://doi.org/10.1038/s41591-018-0300-7
25. Bzdok, D., et al. (2018). Machine learning for precision psychiatry. *Biological Psychiatry*, 84(3), 178–190. https://doi.org/10.1016/j.biopsych.2017.11.007
26. Palechor, F. M., & de la Hoz Manotas, A. (2019). Dataset for estimation of obesity levels based on eating habits and physical condition in individuals from Colombia, Peru and Mexico. *Data in Brief*, 25, 104344. https://doi.org/10.1016/j.dib.2019.104344.