Mitigating Bias in Fitbit Data: A Comprehensive Analysis and Model Enhancement using Ensemble learning.

MSc Research Project
Artificial Intelligence for Business

Olayinka Jayeoba
Student ID: X22178759

School of Computing
National College of Ireland

Supervisor: Anh Duong Trinh

National College of Ireland

MSc Project Submission Sheet

School of Computing

Student Name:	Olayinka Mayowa Jayeoba		
Student ID:	X22178759		
Programme:	M.sc Artificial Intelligence for business	**Year:**	Jan 2023
Module:	Msc Research Praticum		
Supervisor:	Anh Duong Trinh		
Submission Due Date:	January 5, 2024.		
Project Title:	Mitigating Bias in Fitbit Data: A Comprehensive Analysis and Model Enhancement using Ensemble learning		
Word Count:	**14 Page** **4393**		

I hereby certify that the information contained in this (my submission) is information about research I conducted for this project. All information other than my contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project. ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature: ……………………………………………………………………………………

Date: 4th January, 2024 ……………………………………………………………………

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST

Attach a completed copy of this sheet to each project (including multiple copies)	□√
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	□√
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	□√

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Mitigating Bias in Fitbit Data: A Comprehensive Analysis and Model Enhancement using Ensemble learning.

Olayinka Mayowa Jayeoba
X22178759

Abstract

Over time, wearable technology has revolutionized the way people manage their health and fitness. However, there are still some worries about potential biases in the data these devices produce, especially when it comes to the result conclusion Fitbit device gives.
This paper undertakes a multifaceted exploration to detect, understand, and mitigate biases in Fitbit data, with a specific focus on calorie expenditure and sleep patterns. The study aims to not only identify and understand biases in Fitbit calorie and sleep data but also propose and implement effective solutions, culminating in the enhancement of predictive models using linear regression and random forest models to train an improved predictive model for calorie expenditure and sleep tracking. This would also ensure the improved accuracy of health and fitness data, contributing to a more reliable and trustworthy wearable technology ecosystem.

Keywords: *Fitbit, Wearable Technology, Bias Detection, Bias Mitigation, Linear Regression, Random Forest, Calorie Tracking, Sleep Pattern*

1 Introduction

The integration of wearable devices into daily life has ushered in a new era of health and fitness monitoring, providing users with real-time insights into various physiological parameters. Fitbit, a prominent player in the wearable technology market, has gained widespread popularity for its ability to track key health metrics, including calorie expenditure and sleep patterns. As individuals increasingly rely on these devices to inform their lifestyle choices, the need for accurate and unbiased data becomes paramount for ensuring the efficacy of health and fitness interventions. There have been multiple studies with toddlers [1], older adults [2], students [3], and athletes [4] that can be enriched with wearable devices and their increasing number of measured parameters.
However, despite the advancements in wearable technology, concerns persist regarding the accuracy and potential biases inherent in the data generated by these devices. The intricacies of human physiology, coupled with variations in user demographics, activities, and device usage patterns, contribute to the complexity of data acquisition and interpretation. This complexity is particularly evident in the realms of calorie tracking and sleep monitoring, where small discrepancies in data accuracy can have significant implications for users striving to optimize their health and well-being.
The problem at the heart of this study revolves around the potential biases present in Fitbit-generated data, specifically focusing on calorie expenditure and sleep patterns. While Fitbit devices have proven valuable in providing users with actionable health insights, the accuracy of these insights remains a subject of scrutiny. Variations in user demographics, activity types, and even device models may introduce biases that compromise the reliability of the data, raising concerns about the validity of health-related decisions based on such information.

The identification and mitigation of biases in Fitbit calorie and sleep data are essential for several reasons. Firstly, inaccurate data can mislead users, leading to suboptimal health choices and potentially undermining the effectiveness of health and fitness interventions. Secondly, biases may contribute to disparities in the performance of predictive models, impacting the precision and reliability of algorithms that users depend on for informed decision-making. Lastly, the ethical implications of relying on biased data in the context of health and fitness necessitate a thorough investigation and transparent reporting to maintain user trust and confidence in wearable technology.
This study seeks to address these challenges by conducting a comprehensive analysis of biases in Fitbit-generated data, particularly in the domains of calorie expenditure and sleep patterns. Through a detailed examination of factors contributing to biases and the implementation of targeted solutions, the research aims to contribute to the enhancement of data accuracy in wearable technology, ensuring that users can confidently leverage the insights provided by Fitbit devices for their health and fitness goals.

Figure 1: This is the Fitbit logo

2 Related Work

In recent years, there has been a concerted effort to enhance the clinical validation of data generated by smart and wearable devices, leading to the development of numerous datasets addressing diverse health issues and sociological considerations. Notably, the most extensive cohort examined to date is documented by Althoff et al[5]. in 2017, encompassing 68 million days of physical activity data from 717,000 mobile phone users across 111 countries. Their study aimed to establish correlations between inactivity and obesity within and across countries; however, the data's applicability to other research questions is limited, as only aggregated figures were disclosed. Similarly, Chan et al[6]. focused exclusively on asthma, creating a mobile application for daily questionnaires on asthma symptoms with data collected from 6,000 users. The MyHeartCounts Cardiovascular Health Study[7], concentrating on cardiovascular health, recruited 50,000 participants in the US within six months, but the mean app engagement per user was merely 4.1 days. Due to the exclusive focus on one health problem and data collection for that specific condition, such datasets often have limited utility for other researchers. In addressing this gap, this paper introduces the LifeSnaps dataset—a comprehensive, multi-modal, longitudinal, space and time-distributed collection. It includes anthropological data, such as physical activity, sleep, temperature, heart rate, blood oxygenation, stress, and mood, collected from 71 participants between mid-2021 and early-2022, designed for broad applicability in research inquiries.

Numerous investigations have explored biases within wearable technology, revealing the complexities involved in accurately capturing health metrics. In [8], an approach to predict sleep stages (light sleep, deep sleep, REM sleep, and wakefulness) is outlined, utilizing data collected from Fitbit and a standard questionnaire. This prediction involved assessing 21 static and dynamic features (e.g., age, sex, Pittsburgh sleep quality index, total sleep time) from 23 healthy adults, including 9 women. The study employed nine models, employing various combinations of resampling techniques and machine learning algorithms like Naïve Bayes, random forest, and support vector machine with a linear kernel, as well as the gradient boosting decision tree. The evaluation utilized metrics such as accuracy, Cohen's Kappa, Matthew's correlation coefficient, with the most favorable outcomes achieved by the model combining support vector machine with XGBoost.

In [9], an approach employing machine learning to predict sleep efficiency (categorized as poor or good sleep efficiency) based on physical activity data is outlined. The experiment utilized a dataset gathered with actigraphy sensors from 92 adolescents monitored over a week. The predictive model for sleep quality was created using physical activity data during awake time, while the evaluation used data from sleep time. Various neural network architectures, including multilayer perceptron, convolutional neural network, simple Elman-type recurrent neural network, long short-term memory, and a time-batched version of LSTM-RNN, were comparatively assessed. Results indicate that the convolutional neural network (CNN) yields the most favorable outcomes in terms of specificity and sensitivity.

In [10], a study is presented with the goal of developing prediction models utilizing data from Fitbit to medically grade sleep/wake stages. Different combinations of classifiers and resampling methods, such as decision tree and random forest in conjunction with four resampling methods, are scrutinized based on sensitivity, specificity, and area under the ROC curve (AUC) to identify the most effective ones. Experimental results reveal that decision tree algorithms perform best, being less influenced by the resampling method, and random up-sampling emerges as the most effective method for balancing training sets. The dataset for these experiments was collected using a Fitbit Charge 2 device and a medical device from 23 healthy adults, including 9 females, with total sleep hypnogram epochs ranging between 418 and 1208. Additionally, demographic information and subjective data on sleep quality from subjects were collected using the Pittsburgh Sleep Quality Index.

Our objective is to improve individuals' well-being by addressing issues related to calorie imbalances and weight management, leveraging gathered data to create a comprehensive health application. Utilizing a chart depicting an individual's food consumption, sensors play a pivotal role in calculating both calorie intake and expenditure during workouts. The person sets daily caloric goals, enabling effective monitoring of their intake. Recognizing the limitations of traditional methods in accurately assessing physical activity, our approach includes data on calorie expenditure behind the scenes, requiring a more nuanced data categorization [11]. In this context, our paper proposes an innovative machine-learning-based method to predict calorie content from images of food. This method, relying on supervised machine learning, involves identifying the food item, estimating its quantity, and predicting its calorie content, proving notably effective compared to conventional approaches that directly forecast calorie amounts from images [12].

3 Research Methodology

The initial stage in our methodology involves data pre-processing, wherein raw data undergoes transformation into a format suitable for input into machine learning algorithms. This encompassing step comprises several sub-steps, namely: (i) data acquisition, involving the procurement of a pertinent dataset; (ii) handling missing values, focusing on identification and management of missing data within the dataset; (iii) data splitting, which involves dividing the dataset into two sets – one designated for training and the other for testing, in our case, utilizing an 80/20 split ratio; and (iv) feature scaling, where independent variables are standardized using the min-max normalization method, placing all data within the range [0, 1]. For our experiments, we utilized the Life Snaps dataset, collected over four months from 71 individuals through Fitbit, Surveys, and SEMA3. Specifically, we extracted data related to calories burned, heart rate (measured in beats per minute), and sleep efficiency, all recorded with a Fitbit Sense smartwatch. To address missing values, we replaced them with the mean value of the respective feature.

4 Design Specification

Fig 1

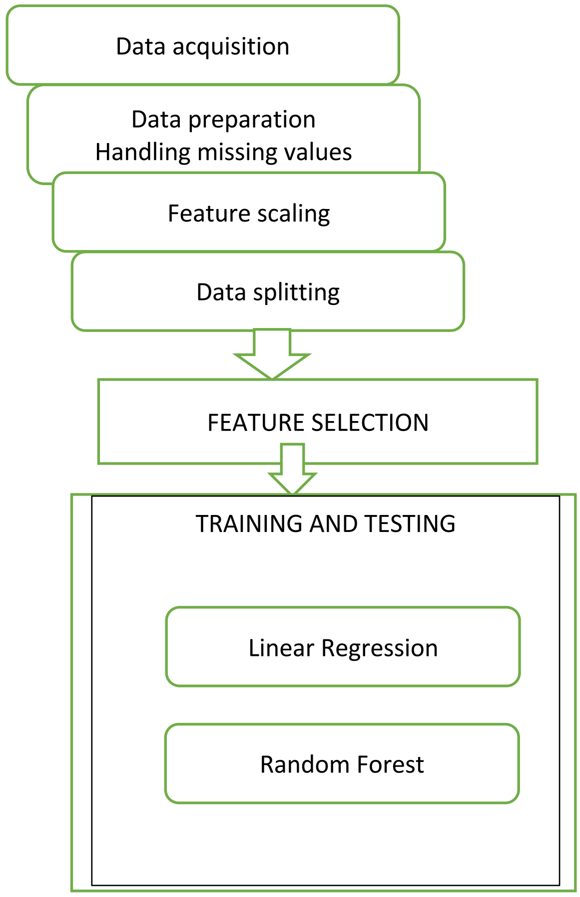

Fig 1. Our approach for predicting an individual's sleep quality and burned calories involves several key steps. The second step, feature selection, is aimed at reducing the number of input features utilized in training our machine learning algorithms by eliminating irrelevant ones. This step is

crucial for enhancing training efficiency, accuracy, and mitigating overfitting. The correlation matrix is employed for feature selection, providing insights into how features are correlated with each other or the target variable.

Moving on to the third step, which encompasses training and testing, the training phase involves utilizing training data to fit and fine-tune our models. The evaluation of these models employs metrics such as accuracy score and mean squared error loss. In the testing phase, the learned model is applied to testing data to make predictions. These steps collectively form a comprehensive method to predict both sleep quality and burned calories, encompassing feature selection and rigorous training and testing procedures.

5 Implementation

The output produced was conducted in Python programming language where libraries like panda, numpy, shap, sklearn, and statmodels were used for data transformation and manipulations and model developments.

6 Evaluation

This section provides a comprehensive overview of the dataset employed in our experiments and outlines the experimental results obtained for predicting sleep efficiency, burned calories, and the activity performed by an individual.

A. Dataset Description:
As detailed in the preceding section, our experiments utilized the LifeSnap dataset, a repository of lifelogging information pertaining to 71 individuals, collected over a span of 4 months through Fitbit, Surveys, and SEMA3. Specifically, we focused on data acquired from a Fitbit Sense smartwatch, including sleep.json, which encompasses details about sleep efficiency, and calorie.json, offering an intricate breakdown of burned calories. Table 1 presents a segment of the sleep efficiency dataset (i.e., sleep.json file) featuring the most pertinent variables employed in our experiments for predicting sleep efficiency. To detect sleep efficiency, we selected specific features based on the correlation matrix, such as the duration of sleep, time spent in bed, minutes awake, minutes asleep, and minutes after wakeup. Figure 2 illustrates the correlation heatmap and histograms corresponding to key variables, including duration, minutes asleep, minutes awake, minutesAfterWakeup, time in bed, and efficiency variables within this dataset.

Table 1:

Values for predicting sleep efficiency.

ep_durati	resting_hr	rmssd	spo2	tress_scor	age	ender_nur	bmi
31260000	62.07307	89.603		78	<30	2	<19
32880000	62.12148	94.303		80	<30	2	<19
33600000	62.264	119.212		84	<30	2	<19
37620000	62.3689	111.709		82	<30	2	<19
33660000	61.96541	103.034		81	<30	2	<19
32280000	62.67175	89.941		82	<30	2	<19
33300000	63.35972	92.763		81	<30	2	<19
34740000	63.12127	109.509		84	<30	2	<19
35040000	62.41962	117.717		84	<30	2	<19

Values for predicting calorie efficiency.

calories	distance	bpm	_active_m	ely_active_	active_mi	ntary_min	age	ender_nur	bmi
2351.59	6517.5	71.70157	149	24	33	713	<30	2	<19
2332.08	7178.6	70.5793	132	25	31	704	<30	2	<19
2262.3	6090.9	71.84258	112	27	31	710	<30	2	<19
2325.1	6653.1	71.72548	133	21	37	622	<30	2	<19
2586.76	9557.9	74.40103	136	42	54	647	<30	2	<19
3806.02	18809	83.73977	305	128	98	371	<30	2	<19
1968.24	2799.7	68.47526	113	9	0	763	<30	2	<19
2300.02	6822.5	69.31436	149	23	34	655	<30	2	<19
2227.17	6215.8	69.3099	106	30	29	691	<30	2	<19

Fig 2: Correlation heatmap and Histogram variables for calories and sleep

(a)Correlation heatmap

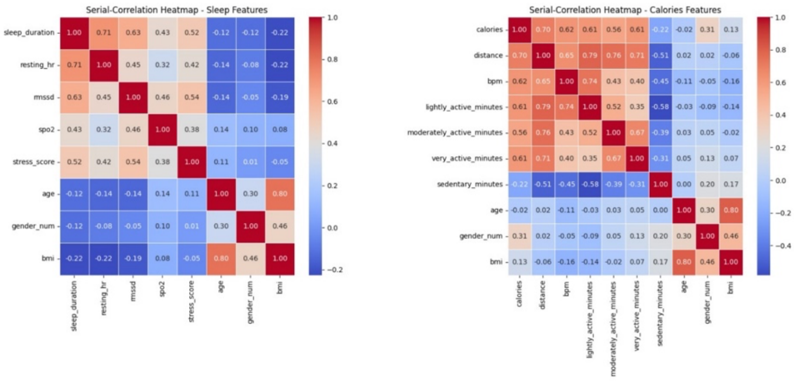

(b)Histogram Variables for Calories

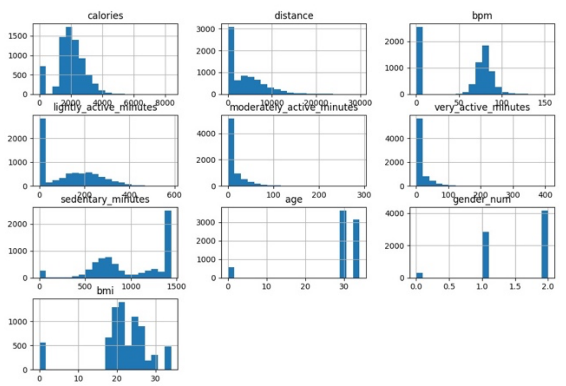

(c) Histogram Variables for Sleep

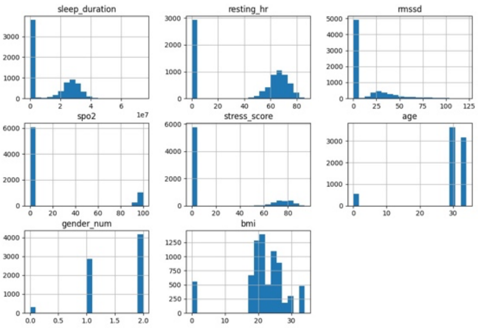

Table III – Model Accuracy.

Ensemble Methods	Sleep/Calorie
Random Forest	0.9456
Linear Regression	0.8654

6.1 Experiment / Case Study 1

A. Experimental results evaluation

We have comparatively analyzed two unsupervised learning, namely, Random Forest and Linear Regression from scratch to see which is the most appropriate for predicting sleep efficiency, the burned calories, and the name of the activity performed by a person. In the training phase, we evaluated the learned models using two metrics, namely the accuracy score and the mean square error loss.

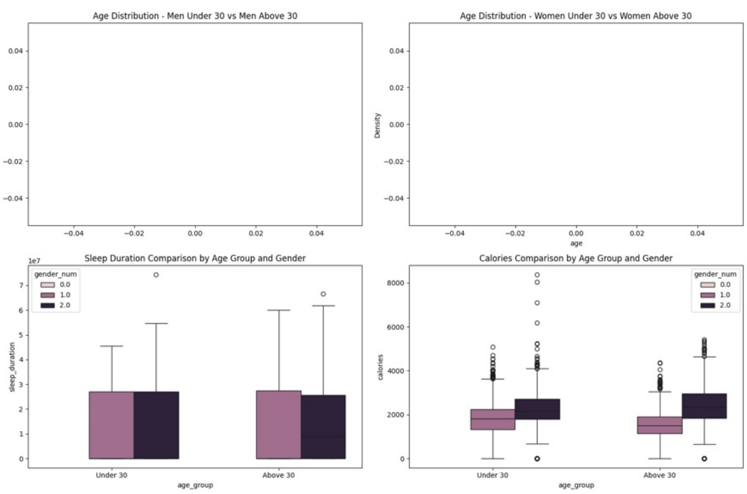

The left side of this image shows the prediction of sleep duration based on men under 30 and above 30. The predictions show that men under 30s get more sleep than above 30s.

[13] Nonetheless, sleep-related concerns are commonly voiced by older adults. In a comprehensive epidemiological study on sleep, Foley et al. [14] discovered that over 50% of elderly individuals reported complaints of insomnia, with chronic sleep disturbances primarily linked to indications of poor health [15]. Furthermore, older adults grappling with sleep issues often experience challenges related to balance, ambulation, and vision, even when accounting for medication use [16]. These difficulties contribute to an elevated risk of falls, a significant predictor for placement in nursing homes or assisted-living environments. A study involving over 3,000 older women revealed that sleeping less than 7 hours a night or having a sleep efficiency (defined as the amount of sleep relative to the time spent in bed) below 65% was associated with an increased risk of falls [17]. Moreover, sleep problems are correlated with higher mortality rates. Research has consistently shown that poor sleep efficiency, increased sleep latency, and reduced total sleep time are linked to a heightened risk of mortality, even after adjusting for relevant covariates [18].Even though during the past few decades, a lot has been discovered about sleep in older adults, many questions remain. What impact treatment has on daytime outcomes is the most urgent concern. There is less evidence indicating that getting more sleep lowers the risk, although epidemiological research consistently confirms that insomnia and poor sleep put people at risk for potentially dangerous outcomes[19].

On the right side of the experimental result evaluation, [1.0 is for women, 2.0 is for men]. It predicted that women have fewer calories than men. Since men need more energy than women due to men's duties and responsibilities including everyday work which consumes energy, this model shows that its prediction on both men's and women's calorie efficiency is okay[20].

6.2 Experiment / Case Study 2

The image below is another visualization of how the linear regression and random forest model predictions are.

The random forest calories feature states that from 0, the middle age residuals from age 12-20 plus have the highest calories, and the predicted image at the right bottom states that 29 plus to 32 are the highest predicted values. The linear prediction model predicted the sleep features showing that the residuals data states only 29plues to 40 years of age get adequate sleep.

The model accuracy;
Random forest - 0.9456
Linear regression - 0.8654

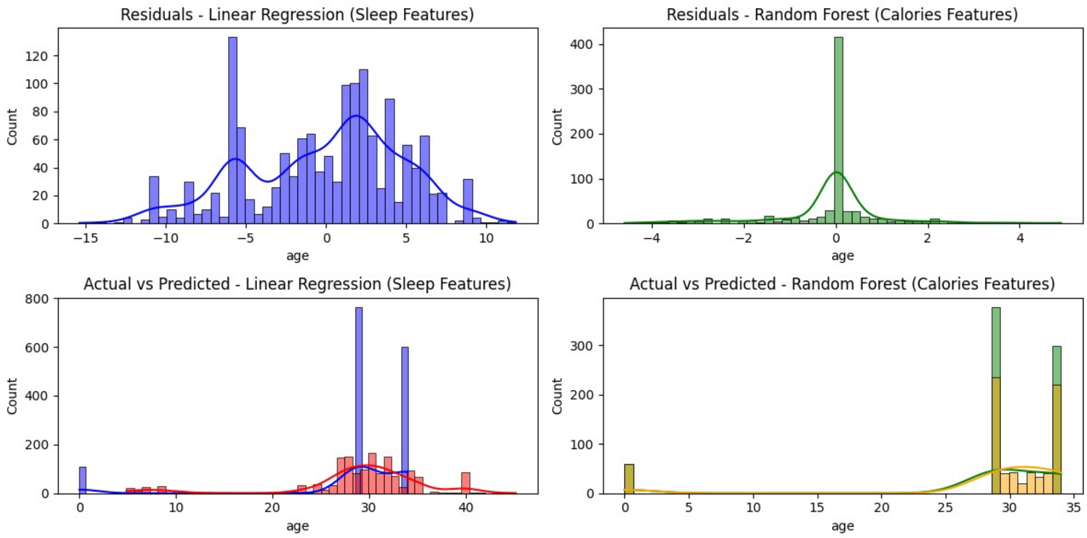

6.3 Experiment / Case Study 3

Separating the data into three categories training, validation, and test sets—is a conventional method of meeting the prescription for creating high-quality predictive models (for final model assessment). Over time, predictive models may change as data scientists or analysts get better at what they do—either by adding new data or refining what they already have (Raykar & Saha, 2015).

To avoid over-fitting, data is divided into distinct categories. To prevent overfitting and bias in model selection, Kumar (2021) divided the dataset into three sections. In addition, he recommended that the training set be the largest, with the development or testing set sharing the same percentages as the cross-validation set.

Predictions based on both findings in case studies 1 and 2, random forest model accuracy seems to be the best on both predictions. The number of instances or observations that were correctly classified by the Group Method of Data Handling (GMDH shell) classification approach, the number of instances or observations that were incorrectly classified, the x-measure, Accuracy, precision, Recall, F-measure, and RMSE are all divided into different split ratios, which include 50/50, 60/40, 70/30, 80/20, and 90/10, respectively. The results demonstrate that the split ratios of 70/30 and 90/10 had the best accuracy performance, both at 100%[21].

New features and targets will be selected, and prediction based on sex and age will be conducted on the 90/10 split methodology.

The below table shows the setup model score and how all features with sex and interaction act as targets and are nd labeled as features.

The table shows the age of females below and above 29 age.

Table I

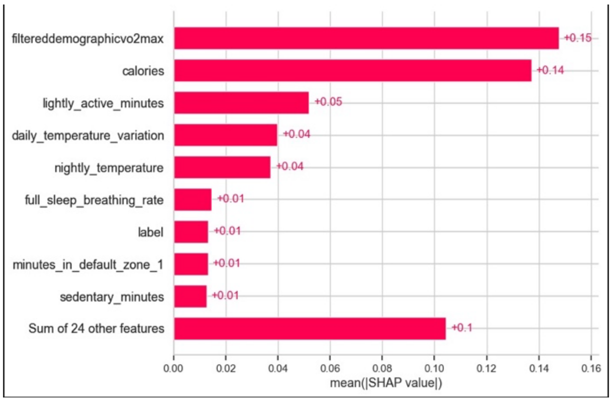

The table above is a prediction of how those under 29 years of age interact with the features.

Table II

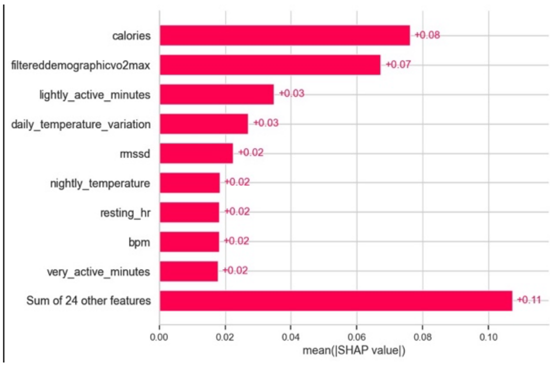

Visualization of prediction based on females above 29 is decent and the model score for both predictions is 1.0 the prediction also looks accurate compared to the prediction result in case studies 1 and 2.

Females are getting enough sleep predictions and females above 29 sleep patterns also look alright, pointing out the resting _hour, we can see that there is a boost in the visualization for both sleep and calorie efficiency compared to the one in Table I.

6.4 Discussion

Principal Findings

The findings concluded in the first experimental findings that the random forest model is the best model for Fitbit model to be built on, and the data splitting will consist of splitting the data set in two sets, one that will be used for training, and another one that will be used for testing. In our former app, roach the data set is split into a ratio of 80/20, but now since the prediction is based on new data, a new approach of splitting the data into 90/10 will be

infused. Th% of the data would be used for training the model, and 10% of the data would be used before testing.
As can be seen he Table III for predicting sleep efficiency and calorie expenditure. Random Forest provides the best result in our case.
According to the experimental result evaluation, Random Forest came out as the best model with a prediction of 0.9456 which is very close to 1.

7 Conclusion

Not as a replacement for conventional clinical polysomnography, Fitbit models are offered to the public to enable consumers to self-derive knowledge about their sleep quality. They seem helpful in researching the 24-hour sleep-wake cycle as well as in determining the length, pattern, and quality of sleep longitudinally—that is, over a long period of consecutive nights—in a typical home setting. In this sense, people can use the data collected by wristband trackers to enhance both their sleep and their calorie prediction. Primary care physicians and sleep experts may occasionally be able to obtain a cursory understanding of patients' sleep patterns. Even though not enough research has been done on modern sleep staging techniques, based on the few published research to date, it appears that Fitbit models perform better than actigraphy-related literature when it comes to distinguishing between wake and sleep intervals.

My next proposal is to design, build, deploy, and monitor an ML application that predicts Taxi demand in New York, using the deep learning framework, TensorFlow to predict based on Long, short, time memory, and time series. This brings in more revenue for the company and makes life easy for the drivers and passengers. This project makes me develop more skills and knowledge in the Machine learning path.

References

1. Nakagawa, M. et al. Daytime nap controls toddlers' nighttime sleep. Scientific Reports 6, 1–6 (2016).
2. Kekade, S. et al. The usefulness and actual use of wearable devices among the elderly population. Computer methods and programs in biomedicine 153, 137–159 (2018).
3. Wang, R. et al.Student lifee: Assessing mental health, academic performance and behavioral trends of college students using smartphones. 3–14, https://doi.org/10.1145/2632048.2632054 (Association for Computing Machinery, Inc, 2014).
4. Düking, P., Hotho, A., Holmberg, H.-C., Fuss, F. K. & Sperlich, B. Comparison of non-invasive individual monitoring of the training and health of athletes with commercially available wearable technologies. Frontiers inPhysiologyy 7, 71 (2016)
5. Althoff, T. et al. Large-scale physical activity data reveal worldwide activity inequality. Nature 547, 336–339, https://doi.org/10.1038/nature23018 (2017).
6. Chan, Y. F. Y. et al. Data descriptor: The asthma mobile health study, smartphone data collected usinga researchkit. Scientific Data 5, https://doi.org/10.1038/sdata.2018.96 (2018).

7. Hershman, S. G. et al. Physical activity, sleep, and cardiovascular health data for 50,000 individuals from the myheart counts study. Scientific Data 6, https://doi.org/10.1038/s41597-019-0016-7 (2019).
8. Zilu Liang and Mario Alberto Chapa-Martell, A Multi-Level Classification Approach for Sleep Stage Prediction With Processed Data Derived From Consumer Wearable Activity Trackers, FrontDigit. Health, 28 May 2021,https://doi.org/10.3389/fdgth.2021.665946
]9. Sathyanarayana A, Joty S, Fernandez-Luque L, et al. Sleep Quality Prediction From Wearable Data Using Deep Learning [published correction appears in http://mhealth.jmir.org/2016/4/e130/] [published correction appears in JMIR Mhealth Uhealth. 2016 Nov 25;4(4):e130]. JMIR Mhealth Uhealth. 2016;4(4):e125. Published 2016 Nov 4. doi:10.2196/mhealth.6562
10. Zilu Liang, Mario Alberto Chapa Martell, Predicting Medical-Grade Sleep-Wake Classification from Fitbit Data Using Tree-Based Machine Learning, Report number: IPSJ SIG Technical Report, 2019-UBI-62:14, 1-8, Project: Ubiquitous and personal computing for sleep health
11. Ingmar, "Forecasting Workout Burnt Calories Using Machine Learning," in Qatar Computing Research Institute, January 2016
12.GENERAL CONCERNS FOR DATA SCIENTISTS AND DATA ANALYSTS', LifeSnaps, a 4-month multi-modal dataset capturing unobtrusive snapshots of our lives in the wild [Preprint]. doi:https://doi.org/10.1038/s41597-022-01764-x.
13. Foley D, Ancoli-Israel S, Britz P, Walsh J: Sleep disturbances and chronic disease in older adults: results of the 2003 National Sleep Foundation Sleep in America Survey. J Psychosom Res 2004;56:497–502.
14. Foley DJ, Monjan AA, Brown SL, Simonsick EM, Wallace RB, Blazer DG: Sleep complaints among elderly persons: an epidemio- logic study of three communities. Sleep 1995; 18:425–432.
15. Foley DJ, Monjan AA, Brown SL, Simonsick EM, Wallace RB, Blazer DG: Sleep complaints among elderly persons: an epidemio- logic study of three communities. Sleep 1995; 18:425–432.
16.Foley DJ, Monjan A, Simonsick EM, Wallace RB, Blazer DG: Incidence and remission of insomnia among elderly adults: an epidemi- ologic study of 6,800 persons over three years. Sleep 1999;22:S366–S372.
17.Vitiello MV, Moe KE, Prinz PN: Sleep com- plaints cosegregate with illness in older adults: clinical research informed by and in- forming epidemiological studies of sleep. J Psychosom Res 2002;53:555–559.
18. Ohayon MM, Carskadon MA, Guilleminault C, Vitiello MV: Meta-analysis of quantitative sleep parameters from childhood to old age in healthy individuals: developing norma- tive sleep values across the human lifespan. Sleep 2004;27:1255–1273.
19.Shochat T, Martin J, Marler M, Ancoli-Israel S: Illumination levels in nursing home patients: effects on sleep and activity rhythms. J Sleep Res 2000;9:373–380.
20. Sleep Disorders in the Older Adult – A Mini-Review (2009). https://karger.com/ger/article-pdf/56/2/181/2835468/000236900.pdf (Accessed: January 4, 2024).
21. IDEAL DATASET SPLITTING RATIOS IN MACHINE LEARNING ALGORITHMS: GENERAL CONCERNS FOR DATA SCIENTISTS AND DATA ANALYSTS (no date). https://www.researchgate.net/profile/Ismail-Muraina/publication/358284895_IDEAL_DATASET_SPLITTING_RATIOS_IN_MACHINE_LEARNING_ALGORITHMS_GENERAL_CONCERNS_FOR_DATA_SCIENTISTS_AND_DATA_ANALYSTS/links/61fb97e711a1090a79cc1a8b/IDEAL-DATASET-SPLITTING-RATIOS-IN-MACHINE-LEARNING-ALGORITHMS-GENERAL-

CONCERNS-FOR-DATA-SCIENTISTS-AND-DATA-ANALYSTS.pdf (Accessed: January 4, 2024).