



Published in final edited form as:

Sleep Health. 2025 April ; 11(2): 166–173. doi:10.1016/j.sleh.2024.10.003.

Performance of machine learning-based methodology using dynamical features to detect non-wear intervals in actigraphy data in a free-living setting

Jyotirmoy Nirupam Das^a, Linying Ji^b, Yuqi Shen^a, Soundar Kumara^a, Orfeu M. Buxton^a, Sy-Miin Chow^a

^aThe Pennsylvania State University, State College, PA

^bMontana State University, Bozeman, MT

Abstract

Goal and Aims: One challenge using wearable sensors is non-wear time. Without a non-wear (e.g., capacitive) sensor, actigraphy data quality can be biased by subjective determinations confounding sleep/wake classification. We developed and evaluated a machine learning algorithm supplemented by dynamic features to discern wear/non-wear episodes.

Focus Technology: Actigraphy data from wrist actigraph (Spectrum, Philips-Respironics)

Reference Technology: the built-in non-wear sensor as ‘ground truth’ to classify non-wear periods using other data, mimicking features of Actiwatch 2.

Sample: Data was collected over one week from employed adults (n=853).

Design: Extreme gradient boosting (XGBoost), a tree-based classifier algorithm, was used to classify wear/non-wear, supplemented by dynamic features calculated over various time windows.

Core analytics: The performance of the proposed algorithm was tested over 30-sec epochs. Additional analytics and exploratory analyses: Evaluation of the SHapley Additive exPlanations (SHAP) values to find the effectiveness of the dynamic features.

jfd5895@psu.edu .

Conflicts of interest

Dr. Orfeu M. Buxton: Outside of the current work, Dr. Orfeu M. Buxton discloses that he received subcontract grants to Penn State from Proactive Life (formerly Mobile Sleep Technologies), doing business as SleepSpace (National Science Foundation grant #1622766 and NIH/National Institute on Aging Small Business Innovation Research Program R43AG056250, R44 AG056250), received honoraria/travel support for lectures from Boston University, Boston College, Tufts School of Dental Medicine, New York University, University of Miami, University of South Florida, University of Utah, University of Arizona, Harvard Chan School of Public Health, Eric H. Angle Society of Orthodontists, consulting fees for SleepNumber, and receives an honorarium for his role as the Editor in Chief of the journal *Sleep Health*.

Dr. Chow: Dr. Chow reports grants from National Institutes of Health, during the conduct of the study.

No Disclosures to report: JND, LJ, YS

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Core Outcomes: The XGBoost classifier yielded substantial improvements in balanced accuracy, sensitivity, and specificity, including dynamic features and comparison to default actiwatch classification algorithms.

Important supplemental outcomes: The proposed classifier effectively distinguished between valid and invalid days, and the duration of contiguous periods of non-wear correctly identified.

Core conclusion: Our findings highlight the potential of XGBoost using dynamic features of varying activity levels across the time series to provide insights on wear/non-wear classification using a large dataset. The methodology provides an alternative to laborious manual benchmarking of the data for similar devices that do not have a non-wear sensor.

Keywords

non-wear detection; machine learning; actigraphy; sleep

Introduction

Data from wrist motion sensors for classifying sleep-wake patterns has grown exponentially due to ease of implementation and cost-effectiveness¹. Wearable-derived digital biomarkers generally provide activity (movement) and light levels or other data to capture passive data over long periods² in free-living participants non-obtrusively.

One of the crucial challenges of studies using wearable sensors is non-wear time, a form of missing data potentially introducing bias. When not worn, the devices may still register low activity levels during non-wear, which may be confounded with or indistinguishable from sleep or sedentary periods based on visual inspection. In devices without non-wear sensors, manually marking non-wear episodes is time-consuming and could introduce further bias. Common automated non-wear detection algorithms in the current literature use tri-axial accelerometers and other signals, such as skin temperature or respiration rate, to detect non-wear intervals accurately. Another approach is setting a time interval with low activity counts for non-wear detection. For example, using data from a cohort of children aged 11 years, ten different non-wear criteria were evaluated; it was concluded that time intervals of 45–60 minutes with zero activity counts were sufficient to detect non-wear intervals³, consistent with similar time interval thresholds for different cohorts and accelerometers in adults^{4,5}. Lok & Zeitzer, 2020⁵ used a temporal threshold to distinguish wear and non-wear periods after manually scoring the periods as wear and non-wear.

Machine learning techniques have been used in non-wear detection studies, such as a logistic regression model to distinguish non-wear patterns during sleep and wake from 23 participants from the NHANES study⁶. A similar regression tree-based method classified wear/nonwear periods using movement and temperature sensor data⁷ from different time windows as an input feature. Temperature changes increase their model's performance. Some studies generally set a threshold over activity levels or accelerometry sensors over an interval such as 30, 60, or 90 minutes^{8,9}. Manual threshold constraints can be arbitrary and exclude the possibility of detecting shorter non-wear episodes.

Other studies have identified additional features from their data to improve the performance of their algorithms. A study using Random Forest for a sleep and non-wear classification study implemented in a clinical setting with 36 features extracted from tri-axial accelerometer data, including 12 statistical measures (e.g., mean, median, standard deviation, t increased model accuracy¹⁰. Additional opportunities to maximize features not used in these studies include the stability of the time series, shifts in the variable level, and variability across multiple time resolutions.

In this paper, we propose and evaluate, consistent with best practice for the evaluation of sleep algorithms¹¹, a machine learning-based methodology to detect non-wear time epochs (30-s) from actigraphy data. An eXtreme Gradient Boosting (XGBoost) model was applied to a dataset containing 853 participants, which led to a generalized machine learning-based approach for non-wear classification, including dynamic features, efficient tuning of hyperparameters, and visualization of results to facilitate interpretation of critical dynamic and time scale information. Due to the minimal data requirements, the method developed is versatile enough to be applied to a wide range of sleep wearable studies. Our methodology also addresses the limitations of thresholding used in earlier work, in which a hard threshold has to be chosen in terms of physical activity level and/or time interval for inactiveness. The optimal values for such fixed thresholds vary greatly from dataset to dataset, and do not capture the complexity of a dataset with varied non-wear intervals.

Methods

Sample

Data were collected from employed adults from two different industries, as previously described^{12–14}, including an extended care provider (nursing home direct patient care workers; n=545 individuals, 92% female, Mean_{age}: 38 years), and an information technology company (n=308 participants, 44% female, Mean_{age}: 46 years). We excluded other participants from the larger dataset with activity counts during off-wrist periods over-written by an early, defective version of the manufacturer's software and set to 'n/a,' so 'truth' was unavailable. Each participant was instructed to wear a wrist actigraph (Spectrum, Philips-Respironics, Murrysville PA) on their non-dominant wrist. The actiwatch provided wrist actigraphy measurements, including activity levels, white light, red light, blue light, and green light collected in 30-second epochs for a week. Non-wear periods were identified with the built-in wrist capacitive sensor, which, if the device was worn correctly, measured skin contact/proximity or its absence. The capacitive sensor could provide, in rare circumstances, both false on-wrist and false off-wrist values in stereotypical and visually identifiable ways, such as nearing battery failure, visually identified by two trained scorers. Scorers visually reviewed the entire recording for each file to better understand an individual's activity pattern and intensity levels across days and examined points of sudden change in activity levels and light levels in the recordings. After each scorer inspected the file independently, inter-rater reliability was assessed, and any discrepancies were resolved among the scoring team. The total number of off-wrist was 7.4% of epochs, and on-wrist was 92.5%, leading to an imbalanced dataset.

Machine learning approach

Data Preparation. In this paper, we explore the use of a tree-based classifier known as XGBoost to classify wear/non-wear epochs¹⁵. Wear is labeled as 1, and non-wear has been labeled as 0 from the Spectrum data. Mimicking the data of the closely related and commonly used Actiwatch 2 model (which has no non-wear sensor), the two key inputs taken from the actiwatch data are the white light and the activity levels. A correlation study between the different light signals confirmed that the white light correlated at least 0.9 with all other light signals; hence, only the white light level was taken as input. To leverage information concerning short- and long-term time dependencies, dynamical features were extracted using the tsfeatures package in R^{16,17}. Dynamical features created using the tsfeature package are spectral entropy, maximum level shift, maximum variance shift, maximum Kulback-Leiber divergence, hurst, spike, heterogeneity, the standard deviation of the first derivative of the time series, and activity lags. Spectral entropy, hurst, heterogeneity, and standard deviation of the first derivative of the time series capture if the time series is random/noisy or if it is stable and has a constant trend. Maximum level shifts and spikes keep track of any sudden changes in the time series. Maximum Kulback-Leiber divergence and maximum variance level shift detect any changes in data distribution in consecutive time windows. Dynamical features were created to provide statistical information to the XGBoost model and improve its ability to discern trends or patterns in the time series. The features have been made on the activity levels with a window size of 20 minutes and 1 hour. The 20-minute window captures shorter naps, and the 1-hour window size captures the average nap duration¹⁸. Furthermore, activity lags were created over 5 minutes, 10 minutes, 15 minutes, and 20 minutes. Next, we processed the time series to generate several temporal features such as minute, hour, quarter, day of the week, and time of the day. This information helps to interpret the timestamp in a useful representation, such as if the time of the day is night or the hour is early morning, then the person is more likely to be asleep. In addition, for each participant, the temporal features were converted into sinusoidal curves to capture prominent periodicities in the data. For each participant, employer information was added as a binary variable (1: extended care and 0: information technology). After processing each participant's data, the data were combined to form a larger dataset. This dataset was split into a 70:30 ratio to create the training and test sets. The random split was formed such that the training set has 597 unique participants and the test set has 256 unique participants.

Hyperparameter tuning and model training. The parameters for the XGBoost model were determined by using hyperparameter optimization and cross-validation. The Hyperopt Python library¹⁹ was used to efficiently explore a range of hyperparameter values for the XGBoost model (see Fig. 1). The parameters were optimized on the average balanced accuracy across five folds of cross-validation training data. The initial ranges of hyperparameter values were expanded and re-evaluated when any hyperparameter value was observed near the search space's boundaries. The range of the search space of the parameters is given in Table 1. The final optimized hyperparameter values and parameter descriptions are summarized in Table 2. We used the histogram tree method, which buckets continuous features into bins to yield approximate tree-splitting solutions to enhance computational speed²⁰. To further speed up the algorithm gbtrees booster and GPU Quadro

K4000 were used. Here we tuned the `scale_pos_weight` as part of the hyperparameter values to be optimized via Hyperopt to allow for heavier weighing of the rare instances of non-wear under imbalanced data (i.e., with the inflated presence of wear data) in the loss function²¹. The two regularization parameters helped reduce overfitting. From Fig. 1. We can see that after fitting the XGBoost model to the training set, we used Youden's J statistic²² to select the probability threshold for classifying data as non-wear under imbalanced data optimal threshold was found and computed the final average balanced accuracy across the validation folds. The entire modeling process is shown in Fig.1. Once we got the classification results from the XGBoost model based on the optimum parameters, we manually changed the probability threshold to get the final performance metrics on the training dataset. We manually changed the threshold instead of relying on Youden's J statistic because Yoden's J statistic increased the specificity at the expense of Negative Predictive Value (NPV). The effects of using different probability thresholds for non-wear classification on the performance metrics are shown in Fig 2. We observe that at probability threshold 0.55, there is a balance between the specificity and NPV value. Increasing or decreasing the probability threshold would lead to unequal tradeoffs in the performance metrics; hence, 0.55 was set as the optimum probability threshold for our dataset.

To further improve our model, we employed a temporal smoothing function to address abrupt (and rare/unlikely 1-epoch) changes in the output at an epoch level. This smoothing function checks neighborhood labels and assigns the label that has the majority in that time window. For example, in a window size of 3 epochs, if there are two non-wear labels and one wear label, all epochs are labeled as non-wear. A range of window sizes from 3 to 20 epochs were used to find the optimal window size. The optimal window size was found to be 3 for the training dataset. This window size gave the maximum specificity; as we increased the window size, specificity decreased.

Classification performance evaluation. The main aim of this paper was to distinguish non-wear epochs in actigraphy data. The XGBoost model provided probabilities at an epoch level and we determined the optimal probability threshold level based on different classification metrics. The classification metrics studied in the paper are sensitivity, specificity, balanced accuracy, F1 score, positive predictive value (PPV), and negative predictive value (NPV). Area under the receiver operating characteristic curve (AUC) was computed for the initial set of results (Table 3). All other results are based on the optimal probability threshold.

To further evaluate our methodology, we compared our result with a published methodology addressing the same general problem⁵, which showed that 95% of the inactive episodes are of length shorter than 20 minutes in an Actiwatch 2 with a resolution of data of 60s. To create these non-wear /inactive episodes automatically, they used R studio to label any consecutive epochs with zero activity counts as an episode of non-wear. The episodes were then visually labeled as wear and non-wear. They then generated predictions on the episodes by calculating the length of each episode and classifying it as non-wear if the episode had a length of more than 20 minutes. The optimal time threshold was computed from the receiver operating curve. Following their methodology, we derived an optimal threshold of 1 minute, which gave the highest accuracy in detecting non-wear episodes in our dataset. They

used activity levels to classify wear/non-wear, and we followed similar approaches, adding dynamic features and white light to predict non-wear and wear epochs.

Results

Core analytics and main outcome variables

Epoch-by-Epoch Performance: In this paper, using data from 853 participants, 597 in the training dataset, and 256 in the test dataset. Each participant contributed ~ 20,000 epochs. Sensitivity refers to correctly predicting wear epochs whereas specificity gives the performance of the model on the non-wear epochs. The results from the optimum XGBoost model (Table 3) show that the balanced accuracy in the training is 0.881 and 0.876 in the test set. Due to the class imbalance in the data, we have a high sensitivity of 0.991 as opposed to a specificity of 0.771 in the training set. To balance this, we reclassify the probabilities generated by the XGBoost model using 0.55 as the new threshold and add the temporal smoothing function. The updated results are shown in Table 4, here, we can see that our specificity rises to 0.822 in the training data. The specificity is lower in the test at 0.791 but is still comparable to the training accuracy. Thus, we do not overfit or underfit our model. The sensitivity decreases from 0.991 to 0.985 in the training set but is still high enough to be determinantal in classifying wear time points. The NPV comes down when we use the probability threshold, but that is the best tradeoff we found in the model against increasing the specificity. The results of the comparison study can be seen in Table 4. We can see that our proposed model outperforms the temporal threshold method on every classification metric. This evaluation supports the claim that machine learning algorithms such as XGBoost can be an automation tool to classify wear and non-wear.

To validate the model further, we matched the invalid days in the dataset against the invalid days predicted by our algorithm. A day is defined as invalid if the person does not wear the actiwatch for more than 4 hours. We label invalid days as 1 and valid days as 0. From Table 5, we can calculate that we are predicting valid days with an accuracy of 0.997 and invalid days with 0.921 accuracy in the training dataset. We get comparable accuracy in the test set, as can be seen in Table 5. Table 6 shows a comprehensive list of classification metrics for predicting invalid and valid days. Most of the metrics are above 0.9; hence, our proposed methodology does not lead to a mismatch between valid and invalid days.

Additional analytics and exploratory analyses

To visualize the predictive roles of the dynamic features, we plotted the SHAP (SHapley Additive exPlanations)²³ values of the 20 most important features (Fig. 3) to identify how a change in each observed value of a feature would change the prediction of the probability of wearing the device in the model. A high standard deviation in the first derivative of the activity levels in a time window of one hour leads to a higher probability of wear. That is, higher growth in activity levels was associated with an increase in the probability of wearing the device. In contrast, lower maximum level shift for that hour (denoted as max_level_shift_hr), maximum variance shift for that hour (denoted as max_var_shift_hr), and number of sharp spikes for that hour (spike_hr) were observed to play multiple roles: they were associated in higher probability for wear in some instances, but also higher

probability for non-wear in other instances. In other words, while prolonged instances with little spikes, level, and variance shifts in an hourly window might be indicative of non-wear, these dynamical features might also capture episodes of sustained high physical activity. The abbreviation “short” in the feature labels in Fig. 3 corresponds to dynamical features created on the activity levels over a time window of 20 minutes. Hence, we have two-time windows of 20 minutes and one hour for each dynamic feature.

The standard deviation of the activity levels, spikes, and maximum variance shift over the shorter time window also contributed to the accuracy of the model in a similar pattern as the same features over a one-hour time window. Both time windows work in tandem to give short and long-term information to the model. The shorter time window reflected changes in the recent history of the time series, while the hourly time window captured the general trend of the feature. A high Hurst value suggested high persistence in an individual activity level, the most likely occurrence of which might be during individuals’ periods of non-wear. In contrast, a low Hurst value, suggesting low persistence, reflected sustained continuous changes in the process of interest, namely ongoing fluctuations in the individual’s physical activity during wear. For the Hurst exponent feature, the SHAP value plot gives similar interpretations for both time windows. Volatile activity levels during physical activities are indicated by high entropy levels over the hour (entropy_hr), whereas low entropy levels suggested no change in activity levels, which in turn, might be suggestive of non-wear. The maximum variance shift, maximum level shift, and spikes capture the sudden changes in the activity levels, in contrast, the Hurst and entropy indicate the flow of the activity levels. The SHAP value plot in Fig 3. shows that the light signal, when compared to the dynamic features, was not important in facilitating non-wear classification . However, we kept the light signal as a feature since it did not significantly affect computational time. The above analysis demonstrates the relevance and insights dynamic features provide to the XGBoost model.

To further analyze the results in terms of non-wear intervals of different durations, we created epoch intervals with different numbers of continuous non-wear epochs. An epoch interval length is defined as the length of the continuous set of non-wear epochs predicted by our methodology. Our nonwear epoch interval length varied from 2 to 19970 epochs, corresponding to nonwear intervals from 1 minute to 7 days for 30-second epochs. Next, we grouped the different epoch intervals based on their length with a bin size of 20 epochs (10 mins) as can be seen in Fig. 4, depicting the mean and median accuracy of the predicted non-wear epoch intervals against the actual number of non-wear epochs present in them. From the plot, we can observe that the mean accuracy gradually rises to 1 as the epoch interval length increases to 530 epochs. So for higher epoch interval length, the algorithm is better able to identify a set of continuous non-wear episodes. The median values approach 1 for epoch interval length of 110 epochs (55mins) to 530 epochs (nearly 4.5 hours). This suggests that the proposed methodology can discern between long naps or long periods of inactivity against non-wear epochs. Hence this methodology can be used in sleep studies where the main aim in preprocessing data is to identify valid days or long periods of non-wear in the dataset. To better illustrate the methodology’s performance on longer time windows, two data samples of length one week each have been shown in Fig. 5. The black bands represent non-wear episodes, and the white bands show wear episodes. For

each sample data, the ground truth from the capacitive sensor and the prediction from the methodology are depicted. The two bands from each sample are not identical, but the non-wear episodes on a longer horizon are effectively captured by the proposed methodology.

Discussion

In the current study, we addressed the problem of determining non-wear epochs in a sleep wearable (Actiwatch Spectrum) relevant to data from a different model of the same device (Actiwatch 2.0) that does not include a non-wear sensor feature, a commonly used device/data type in prior sleep studies. Addressing non-wear is particularly important because non-wear intervals near or during actual time in bed periods may be classified as sleep, leading to non-random bias in characterizing sleep and wake patterns. We used a large dataset using the Actiwatch Spectrum device, similar to the Actiwatch 2.0 but including a non-wear sensor that was used as the ‘ground truth’ for non-wear. We developed and evaluated an XGBoost model to classify wear and non-wear epochs from an actiwatch worn by a participant in a free-living setting. We extracted different dynamical features from the activity level data from the actiwatch and also expanded the time stamps into sinusoidal curves to better utilize the temporal component of the data. This expanded dataset with 37 features was then used as input into the XGBoost model. The optimum parameters for the model were selected by running a cross-validation and hyperparameter tuning algorithm on the XGBoost model. Our systematic evaluation revealed classification performance in classifying epoch level non-wear in the test dataset of AUC: 0.937, Sensitivity: 0.760, Specificity: 0.760, and Balanced Accuracy: 0.875, among other favorable evaluation metrics (Table 3). The results were further optimized by the application of manual probability thresholding and the temporal smoothing function, leading to a .887 balanced accuracy score in the test dataset. In comparison to the temporal thresholding methodology described in Lok & Zeitzer, 2020⁵ our methodology performed better in all performance metrics. Since we achieved high balanced accuracy on the test set, the proposed model and dynamic features can be implemented to detect nonwear on other datasets. As a practical outcome, if an ‘invalid day’ is deemed to require at least 24h of wear time or <4h of non-wear, we calculated that we predicted valid days with an accuracy of 0.997 and invalid days with an accuracy of 0.921. Hence, in studies where the determination of valid days plays a major role, our methodology can be used to identify the days efficiently. Importantly, we further evaluated the threshold of the duration of a non-wear period of consecutive non-wear epochs where classification approached 100% correct at a contiguous duration of ~ 1 hour. This evaluation supports the use of this non-wear detection algorithm for rigorous detection of sleep in the absence of a non-wear sensor on a wearable.

Limitations and Future Perspectives

One of the limitations of our study can be observed in Fig 4, where we can see that the median accuracy is near zero for epoch lengths of less than 110 epochs (55 mins). Brief periods of time when the activity level is low, leading to misclassification of non-wear. Due to this, the proposed methodology has an NPV of 0.783. The median accuracy might be low due to short naps, which might be confused with non-wear epochs due to low activity levels. In addition to this short duration, epoch lengths might be inaccurately classified

due to a lack of knowledge of the recent history of the time series. Alternative machine learning variations, including long short-term memory networks and other recurrent neural network variations²⁴, and a wavelet scattering-based model²⁵ might be an approach for future work. Due to the unbalanced nature of the data, a high window length for the RNN model might be necessary to capture non-wear periods with the limitation of high computational resources, which might be infeasible in large-scale studies. In this manuscript, we extracted dynamic features from activity levels. One may further analyze the activity levels regarding the hierarchical organization of human physical activity. Depending on the time scale of a physiological process, it can be characterized in 4 ranges. Each range has specific characteristics that might add insights into the activity pattern²⁶. Similarly, if the data has sleep periods, one can find patterns in locomotor inactivity to study sleep dynamics²⁷. Such patterns can be used as additional data for the machine learning model but it might require high computational memory. Another limitations is due to the nature of the capacitive sensor on the wrist which may lead to brief false positives.

Core conclusion

An XGBoost model was built on a dataset to classify non-wear and wear epochs of a commonly used sleep wearable without a non-wear sensor using the data from a nearly identical version of the device with a non-wear sensor as a reference. A set of dynamic features was created to augment the dataset and substantially increase the accuracy of the model. Even though the non-wear classification algorithm proposed in this article was applied to data from one specific device, the key innovation resides in the extraction and use of general, device-agnostic dynamic features that summarize differences in the trends and patterns of movement data characterizing wear and non-wear episodes. Thus, the algorithm is also applicable to other accelerometry devices. When compared to a temporal thresholding method, our methodology performed better on five out of the six classification metrics used in this paper. Our methodology was able to detect invalid days in the dataset and contiguous non-wear periods over 55 minutes, meaning night sleep and longer naps can be detected using this algorithm in data from a sleep wearable without a non-wear sensor. Future work may involve using complex neural network models to improve accuracy.

Funding

- The source of this data, the Work Family and Health Study (www.WorkFamilyHealthNetwork.org), was funded by a cooperative agreement through the National Institutes of Health and the Centers for Disease Control and Prevention: Eunice Kennedy Shriver National Institute of Child Health and Human Development (U01HD051217, U01HD051218, U01HD051256, and U01HD051276); National Institute on Aging (U01AG027669); Office of Behavioral and Social Sciences Research and National Institute for Occupational Safety and Health (U01OH008788 and U01HD059773). Grants from the National Heart, Lung, and Blood Institute (R01HL107240), William T. Grant Foundation, Alfred P. Sloan Foundation, and the Administration for Children and Families have provided additional funding. Specific funding for this project is also provided by the Penn State Social Science Research Institute. The contents of this publication are solely the responsibility of the authors and do not necessarily represent the official views of these institutes and offices.

- National Institutes of Health grants U24AA027684, OT2 HL161847, and R01 DK134863; the Pennsylvania State University Quantitative Social Sciences Initiative and UL1TR002014–06 from the National Center for Advancing Translational Sciences

References

1. Danilenko KV, Stefani O, Voronin KA, et al. Wearable Light-and-Motion Dataloggers for Sleep/Wake Research: A Review. *Appl Sci*. 2022;12(22):11794. doi:10.3390/app122211794
2. Nakagata T, Murakami H, Kawakami R, et al. Step-count outcomes of 13 different activity trackers: Results from laboratory and free-living experiments. *Gait Posture*. 2022;98:24–33. doi:10.1016/j.gaitpost.2022.08.004 [PubMed: 36030707]
3. Aadland E, Andersen LB, Anderssen SA, Resaland GK. A comparison of 10 accelerometer non-wear time criteria and logbooks in children. *BMC Public Health*. 2018;18(1):323. doi:10.1186/s12889-018-5212-4 [PubMed: 29510709]
4. Berendsen BAJ, Hendriks MRC, Willems P, Meijer K, Schaper NC, Savelberg HHCM. A 20 min window is optimal in a non-wear algorithm for tri-axial thigh-worn accelerometry in overweight people. *Physiol Meas*. 2014;35(11):2205–2212. doi:10.1088/0967-3334/35/11/2205 [PubMed: 25341022]
5. Lok R, Zeitzer JM. A Temporal Threshold for Distinguishing Off-Wrist from Inactivity Periods: A Retrospective Actigraphy Analysis. *Clocks Sleep*. 2020;2(4):466–472. doi:10.3390/clockssleep2040034 [PubMed: 33198122]
6. Thapa-Chhetry B, Arguello DJ, John D, Intille S. Detecting Sleep and Nonwear in 24-h Wrist Accelerometer Data from the National Health and Nutrition Examination Survey. *Med Sci Sports Exerc*. 2022;54(11):1936–1946. doi:10.1249/MSS.0000000000002973 [PubMed: 36007161]
7. Vert A, Weber KS, Thai V, et al. Detecting accelerometer non-wear periods using change in acceleration combined with rate-of-change in temperature. *BMC Med Res Methodol*. 2022;22(1):147. doi:10.1186/s12874-022-01633-6 [PubMed: 35596151]
8. Choi L, Liu Z, Matthews CE, Buchowski MS. Validation of Accelerometer Wear and Nonwear Time Classification Algorithm. *Med Sci Sports Exerc*. 2011;43(2):357–364. doi:10.1249/MSS.0b013e3181ed61a3 [PubMed: 20581716]
9. Hecht A, Ma S, Porszasz J, Casaburi R, For The Copd Clinical Research Netw. Methodology for Using Long-Term Accelerometry Monitoring to Describe Daily Activity Patterns in COPD. *COPD J Chronic Obstr Pulm Dis*. 2009;6(2):121–129. doi:10.1080/15412550902755044
10. Sundararajan K, Georgievska S, Te Lindert BHW, et al. Sleep classification from wrist-worn accelerometer data using random forests. *Sci Rep*. 2021;11(1):24. doi:10.1038/s41598-020-79217-x [PubMed: 33420133]
11. De Zambotti M, Menghini L, Grandner MA, et al. Rigorous performance evaluation (previously, “validation”) for informed use of new technologies for sleep health measurement. *Sleep Health*. 2022;8(3):263–269. doi:10.1016/j.sleh.2022.02.006 [PubMed: 35513978]
12. Bray J, Kelly E, Hammer L, et al. An Integrative, Multilevel, and Transdisciplinary Research Approach to Challenges of Work, Family, and Health. RTI Press; 2013. doi:10.3768/rtipress.2013.mr.0024.1303
13. Olson R, Crain TL, Bodner TE, et al. A workplace intervention improves sleep: results from the randomized controlled Work, Family, and Health Study. *Sleep Health*. 2015;1(1):55–65. doi:10.1016/j.sleh.2014.11.003 [PubMed: 29073416]
14. Berkman LF, Kelly EL, Hammer LB, et al. Employee Cardiometabolic Risk Following a Cluster-Randomized Workplace Intervention From the Work, Family and Health Network, 2009–2013. *Am J Public Health*. 2023;113(12):1322–1331. doi:10.2105/AJPH.2023.307413 [PubMed: 37939328]
15. Chen T, Guestrin C. XGBoost: A Scalable Tree Boosting System. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM; 2016:785–794. doi:10.1145/2939672.2939785
16. Hyndman RJ, Wang E, Laptev N. Large-Scale Unusual Time Series Detection. In: *2015 IEEE International Conference on Data Mining Workshop (ICDMW)*. IEEE; 2015:1616–1619. doi:10.1109/ICDMW.2015.104
17. tsfeatures.pdf. Accessed April 4, 2024. <https://cran.r-project.org/web/packages/tsfeatures/tsfeatures.pdf>

18. Milner CE, Cote KA. Benefits of napping in healthy adults: impact of nap length, time of day, age, and experience with napping. *J Sleep Res.* 2009;18(2):272–281. doi:10.1111/j.1365-2869.2008.00718.x [PubMed: 19645971]
19. Bergstra J, Yamins D, Cox DD. Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures. In *International conference on machine learning*, pp. 115–123. PMLR, 2013.
20. Ke G, Meng Q, Finley T, et al. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. *Advances in neural information processing systems* 30 (2017).
21. XGBoost Parameters — xgboost 2.1.0-dev documentation. Accessed April 4, 2024. <https://xgboost.readthedocs.io/en/latest/parameter.html>
22. Youden WJ. Index for rating diagnostic tests. *Cancer.* 1950;3(1):32–35. doi:10.1002/1097-0142(1950)3:1<32::AID-CNCR2820030106>3.0.CO;2-3 [PubMed: 15405679]
23. Lundberg S, Lee SI. A Unified Approach to Interpreting Model Predictions. Published online November 24, 2017. Accessed February 21, 2024. <http://arxiv.org/abs/1705.07874>
24. Schmidt RM. Recurrent Neural Networks (RNNs): A gentle Introduction and Overview. Published online November 23, 2019. Accessed April 23, 2024. <http://arxiv.org/abs/1912.05911>
25. Sharma M, Lodhi H, Yadav R, Acharya UR. Sleep disorder identification using wavelet scattering on ECG signals. *Int J Imaging Syst Technol.* 2024;34(1):e22980. doi:10.1002/ima.22980
26. Búzás A, Makai A, Groma GI, et al. Hierarchical organization of human physical activity. *Sci Rep.* 2024;14(1):5981. doi:10.1038/s41598-024-56185-0 [PubMed: 38472275]
27. Winnebeck EC, Fischer D, Leise T, Roenneberg T. Dynamics and Ultradian Structure of Human Sleep in Real Life. *Curr Biol.* 2018;28(1):49–59.e5. doi:10.1016/j.cub.2017.11.063 [PubMed: 29290561]




Fig. 1. Hyperparameter tuning and k-fold cross-validation to optimize XGBoost model parameters.




Fig. 2. Performance measures associated with different probability thresholds for non-wear classification. The optimal probability threshold (.55) was selected since, at this probability, the best tradeoff between NPV and specificity was achieved.




Fig. 3.

The SHAP value plot is arranged according to the feature importance. Also, the SHAP values show the combined impact of each feature on the probability of wear and non-wear. Higher absolute SHAP values, defined as feature importance, conveying greater changes – either increases (in red) or decreases (in blue) in wear classification probability, are arranged in decreasing order from top to bottom according to their absolute SHAP value.




Fig 4.

Each epoch interval represents a non-wear episode of certain time period. Here [10,30) represents non-wear episodes ranging from 5 mins to 15 mins. The Mean and Median provide the model classification accuracy on these epoch interval bins. On the 2nd y-axis shows the number of epochs in each epoch intervals.




Fig. 5.

Two data samples over one week showcase the comparison between methodology-based prediction and ground truth.

Table 1.

Parameter Search Space

Parameter	Range
Maximum Bin	10–500
Colsample_bytree	0.01–1
Learning rate	0.00001–1
Maximum Depth	1–40
Reg_alpha	0–10
N_estimators	0–400
gamma	0–15
eta	0.01–1
Reg_lambda	0–10
Rate Drop	0–1
Scale pos weight	0.05–1

Table 2.

Optimal parameter values

Parameters	Value	Parameter Purpose
tree_method	Hist	Can handle larger datasets and has been optimized for speed
booster	gbtree	The gbtree method is more suitable to fit nonlinear boundaries
normalize_type	Forest	Stabilizes learning
device	Cuda	The model can be trained on GPU
sampling_method	uniform	No bias in selecting the samples for each tree
colsample_bytree	0.2674	Low value reduces the risk of overfitting
gamma	12.7667	High value reduces overfitting by reducing tree complexity
learning_rate	0.0083	Lower learning rate helps in getting optimal loss
max_bin	489	Buckets continuous features to make a detailed yet computationally efficient representation of features
max_depth	19	A deeper tree can better learn the hidden patterns
n_estimators	91	Denotes the number of trees and a higher number can avoid underfitting
rate_drop	0.4`	Makes the model more robust
reg_alpha	5.44	Makes the weights sparse
reg_lambda	1.51	Smoothens the impact of weights
scale_pos_weight	0.157	Gives more weight for the negative class in the loss function

Table 3.

Classification metrics on the training and test set.

	AUC	Sensitivity	Specificity	Balanced Accuracy	F1 score	PPV	NPV
Training dataset	0.981	0.991	0.771	0.881	0.986	0.981	0.879
Test dataset	0.937	0.990	0.760	0.875	0.985	0.981	0.855

Table 4.
Comparison Study on the Improved XGBoost methodology and the threshold Method

	XGBoost with Probability Threshold		Threshold Method	
	Train set	Test set	Train set	Test set
Sensitivity	0.985	0.983	0.957	0.947
Specificity	0.822	0.791	0.104	0.117
Balanced Accuracy	0.904	0.887	0.526	0.532
F1 score	0.985	0.983	0.965	0.967
PPV	0.985	0.983	0.984	0.989
NPV	0.824	0.783	0.032	0.025

Table 5.

Prediction of Invalid days on the train and test set

	Train set		Test Set	
	Predicted Invalid	Predicted Valid	Predicted Invalid	Predicted Valid
True Invalid	389	33	156	18
True Valid	8	3457	1	1505

Table 6.

Classification metrics for predicting invalid days

	Sensitivity	Specificity	Balanced Accuracy	F1 score	PPV	NPV
Training dataset	0.921	0.997	0.959	0.949	0.979	0.990
Test dataset	0.896	0.999	0.944	0.942	0.993	0.988