2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 14-18, 2024. Abu Dhabi, UAE

Automatic Dietary Monitoring Using Inertial Sensor in Smartwatch

K. Pavlov[1], V. Tsepulin[2], N. Lutsyak[3], R. Khasianov[4], E. Simchuk[5], A. Perchik[6], E. Volkova[7]

Abstract— **This paper investigates the problem of eating activity detection using motion data from an off-the-shelf smartwatch. The development and integration of the algorithm for detecting eating activity will make it easier for users to monitor their eating habits. For development of a such algorithm about 27500 hours of data were collected from 91 participants. Moreover, a reliable and interpreted approach with adjustable tolerance for model quality estimation in real-world conditions is proposed in this work. The algorithm based on end-to-end neural network (NN) for eating events detection with special postprocessing was developed by our research group. It recognizes eating events with 1 minute delay from the beginning of food intake. For a such tolerance it achieves F1-score of 0.90 in average (at "free-living" scenario test) for users wearing smartwatches either on dominant or on non-dominant hand. To the best of authors' knowledge, the algorithm provides the best performance of any existing solution or described in the literature.**

I. Introduction

In this study, we developed a new technological feature for wearables to increase their interactivity by adding new functionality of automatic dietary monitoring. This new feature will promote a healthy lifestyle and motivate users to maintain it. That will lead to users' health improvement and extend life expectancy and quality.

Most wearables available on market carries onboard microelectromechanical inertial measurement units (IMU) to monitor user's motions. These devices compose of three-axis accelerometer and three-axis gyroscope. Collected data is mostly used for physical activity characterization [1] and in fitness exercise recognition applications. One of the most useful applications is a fall detection and prevention due it is able to save the lives of elderly people [2]. Physical activity measurements with consumer market wearables have performance comparable with professional devices [3] Raw IMU data holds much more useful information about the user. Activity of daily living (ADL) can be accessed through analysis of accelerometer and gyroscope data. Gait patterns recognition from motion signals can be used as a biomarker for several medical conditions [4] or for user authentication [5].

[1,2,3,4,5,6]All the authors were with Samsung Research where the present study was undertaken,
12, Dvintsev street, Moscow, Russian Federation
[1] k.pavlov@samsung.com
[2] v.tsepulin@gmail.com
[3] nikolaylucyak@gmail.com
[4] khasyanovrr@yandex.ru
[5] egor.sim4uk@gmail.com
[6] a.perchik@samsung.com
[7] v.elena@samsung.com

Our work is aimed to develop a method of user-independent eating behavior characterization by analysis of motion data from IMU sensor integrated into smartwatch. This method will be used in a mobile app such as Samsung Health that manages person's eating behavior to motivate user to lead a healthier life. These applications are aimed to solve global problems like the increase of obesity and prevention of metabolic disorders [6]. Automatization of the dietary data input is essential for user retention since conventional text-based input mode requires extra efforts.

The main contributions of this paper are:

- The largest dataset for eating activity recognition was collected, see **Table I**.
- Interpretable method for model evaluation with flexible tolerance setting, see **Section V-A**.
- Algorithm for non-dominant hand (related to micro-gestures accompanying eating process) provides very high performance (F1-score > 0.83 for eating events detection) at both "non-free-living" and "free-living" test scenarios, see **Section V-C**.

II. Related work

In this paper, we propose a method that has the potential to automate eating activity detection. Firstly, we observe recognition of eating events in free-living settings. Such a problem statement is the closest to our main goal of automatic dietary monitoring. After that we describe results in detection of actions, gestures or micromovements during a meal. Our goal is to outperform the existing investigations of eating moment detection with inertial sensors.

A. Eating activity detection with inertial sensors

Eating activity detection was studied in recent years using various body-worn devices. Previous works showed promising results in this area. For example, Thomaz *et al*. [8] and Dong *et al*. [7] collected datasets for predicting eating events. Nevertheless, results of these papers showed the high rate of false positives occurring during non-eating activities, that was also noticed in [12] with a detailed categorization of confounding gestures.

Dong *et al*. [7] presented the first study using body-worn sensors to automatically detect eating activities in free-living conditions. Their approach was focused on a segmenting and classifying periods as eating and non-eating activities. The authors achieved an accuracy of 81%, whereas an F1-score of 0.34.

The work of Thomaz *et al*. [8] investigated smartwatch based system for the detection of food intake gestures and identification of eating moments. Initially, authors recognized

TABLE I

INFORMATION ABOUT DATASETS PRESENTED IN PAPERS

Source reference	Hands	Subjects #	Eating events #	Total, h	Eating, h	Annotation
Dong *et al.* [7]	Dominant	43	116	449	22.4	Activity
Thomaz *et al.* [8]	Dominant	28	120	470	25	Actions
Kyritsis *et al.* [9] ("non-free living" scenario)	Dominant	12	21	4	4	Micromovements
Kyritsis *et al.* [10] ("free living" scenario)	Dominant	6	16	77	77	Micromovements
Rouast *et al.* [11]	Both	202	202	N/A	N/A	Actions
This paper ("non-free living" scenario)	Both	91	**4165**	**27500**	**796**	Activity
This paper ("free living" scenario)	Both	6	149	780	21	Activity

gestures by Random Forest classifier [13] and then clustered them using DBSCAN [14] algorithm to detect eating activities.

F1-scores of 0.76 and 0.40 for recognition of eating events in 60-min and 15-min time segments respectively were achieved.

B. Eating activity detection with non-inertial sensors

Chun *et al*. [15] presented an approach for detection of eating activities with a necklace that captures head and jawbone movements. To assess the system in real-world settings, they conducted a 1-day in-the-wild study with 19 participants. The authors achieved an F1-score of 0.75 for 10-minute segments.

Zhang *et al*. [16] proposed a one-class-classification approach for eating event spotting using electromyogram monitoring eyeglasses. The authors trained oc-SVM [17] model on eating data, without specifying a model for non-eating data. They collected 122 hours of data in a free-living setting from 10 participants. The authors achieved an F1-score of 0.95.

Although non-inertial sensors provides high performance of eating detection, a form-factor requirements and privacy issues have severely limited commercialization of such approaches.

C. In-meal gesture detection

The work of Kyritsis *et al*. [9] presented an algorithm for automatic detection of in-meal food intake cycles from smartwatch IMU. Food intake cycle was defined as a sequence of 4 micromovements: pick food, upwards, downwards, mouth. The authors estimated the probability distribution of each wrist micromovement by applying a CNN [18] and using an LSTM [19] network to classify sequences of windows as food intake cycles. The authors achieved an F1 detection score of 0.913 from Leave One Subject Out experiments.

The work of Zhang *et al*. [20] proposed a framework for detecting and counting the number of feeding gestures during an eating episode using a six-axis inertial wrist-worn IMU. The authors achieved an average F1-score of 0.75 in-lab and 0.3 in test of feeding gesture count.

III. EXPERIMENT DESIGN

The dataset was collected from 112 subjects well distributed by age (37.5 ± 11.1 years), gender (51% of females, 49% of males) and body mass index (26.9 ± 5.0 kg/m^2). Prior to tests, all subjects gave their informed consent. The study protocol was approved by institutional review board (IRB) of Saratov State Medical University. The data was kept anonymized and it was used only for the intended research purpose.

The most of subjects (106 volunteers) were in controlled conditions (henceforth referred as "non-free living" group) and were required to follow the rules:

1) consume three different types of menus (developed by professional nutritionist),
2) do not consume any other food,
3) to be recorded on video during food intakes (so annotated the start and end times of eating could be checked).

In total "non-free living" datasets included 32000 hours of time-series data and 4200 eating events. Four eating scenarios were used for the "non-free living" group: standing position + eating with fingers; seating position + metal tableware or plastic tableware or eating with fingers.

A smaller group of subjects (6 volunteers) did not follow any rules and led their usual lifestyle (henceforth referred as "free-living" group). In total "free living" datasets included totally 910 hours of time-series data and 203 eating events.

Compared to other datasets presented in previous research papers on eating activities detection, our dataset is the largest in terms of number of meals and the total recording time (see **Table I** for details).

Raw IMU sensor data were collected with Galaxy Watch (Samsung, South Korea) smartwatch with custom software with $f_s = 20$ Hz sampling rate. All subjects wore smartwatch devices on both left & right hands and labeled all eating and drinking events.

IV. DETECTION ALGORITHM

A. Preprocessing

Data for single IMU record can be presented as six vectors: accelerometer $\mathbf{a}_x$, $\mathbf{a}_y$, $\mathbf{a}_z$ and gyroscope $\mathbf{g}_x$, $\mathbf{g}_y$, $\mathbf{g}_z$. Where $\mathbf{a}_i = (a_i^1, \dots, a_i^n)$ and $\mathbf{g}_i = (g_i^1, \dots, g_i^n)$, $i \in \{x, y, z\}$, $n = T \cdot f_s$, where T is the total time of record in seconds and f_s is a sampling frequency of IMU.

Firstly, we apply smoothing to each axes using exponentially moving average approach with halflife $h = 20$ to clean signal from noise and too strong emissions: $\hat{a}_i^1 = a_i^1, \quad i \in \{x, y, z\}$
$\hat{a}_i^t = \alpha \cdot a_i^t + (1 - \alpha) \cdot \hat{a}_i^{t-1}, \quad i \in \{x, y, z\},\ t = \overline{1, n}$
$\alpha = 1 - \exp\left[\log(0.5)/h\right]$ After that, we filter signal with Chebyshev type I filter and downsample to $f_s' = 5$Hz. Such

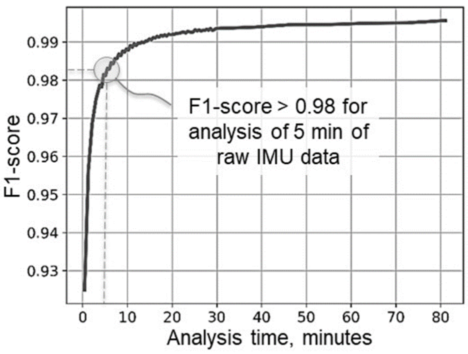

Fig. 1. Performance of supplementary hand detector algorithm depending on analysis time

frequency was selected because higher frequencies are not relevant to eating gestures (average eating gesture duration in our dataset is 3.5 seconds). At the same time, it will low down the power consumption during signal collection and reduce the computational cost. Finally, we normalize signal with $l2$-norm, so that at each timestamp t: $(a_x^t)^2 + (a_y^t)^2 + (a_z^t)^2 = 1$ and $(g_x^t)^2+(g_y^t)^2+(g_z^t)^2 = 1$. These preprocessing steps show their robustness for eating detection tasks [8], [21].

B. Automatic hand detection (left or right)

This is a common situation when a user wears smartwatch on a non-dominant hand, but uses mainly a dominant hand for food intake. The most of related studies are focused on dominant hand monitoring to detect eating gestures directly. But, as we found, modern IMU sensors are suitable to detect non-dominant hand's micro-gestures accompanying the eating gestures of dominant hand. It is reasonable to train two separate ML algorithms for eating gestures (dominant hand) and for accompanying micro-gestures (non-dominant hand).

According to [22] the rate of left-handed people in the world population is about 10%, so it is essential for our algorithm to know on which hand the user wears the smartwatch. Figure 1 shows a performance of our supplementary hand detector algorithm based on XGBoost [23] model with accelerometer and gyroscope features. It is shown that 5 minutes of smartwatch wearing allows to detect right / left hand with F1-score > 0.98.

C. Eating activity detection

We assign each segment belonging to eating event as "1", otherwise – "0". Since an eating activity lasts less than 1% of total dataset time – the data is imbalanced from the view point of classification task. A dynamically balanced batches technique is used to solve of imbalanced classes. Each batch in training process is composed in a special way to sample random segments from a dataset, and, at the same time, maintain constant ratio between amount of "1" samples (eating event) and "0" samples.

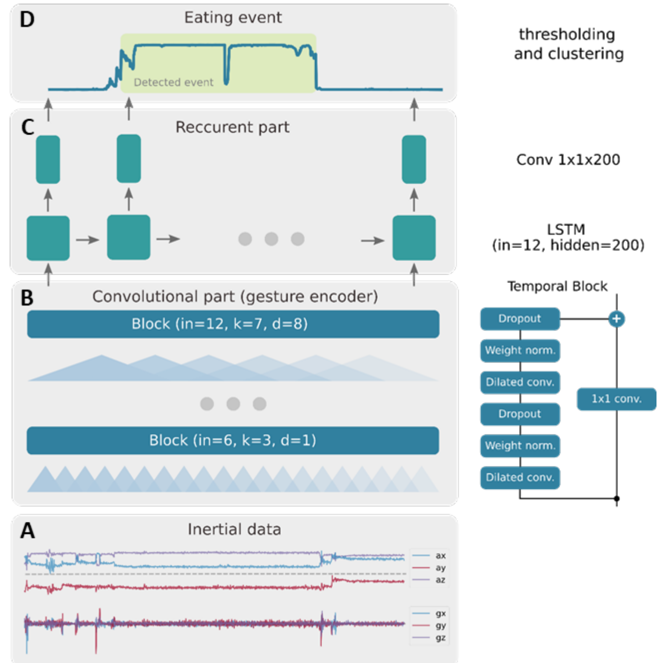

Fig. 2. ML pipeline of automatic eating events detection (from preprocessed IMU data on the bottom to detected event on the top of diagram)

The best performance is achieved with a end-to-end neural network (NN) containing convolutional and recurrent parts (see figure 2). Our ML pipeline is an improved combination of approaches described in [10] and [24]. Preprocessed input data from inertial sensors is represented by sequence of vectors containing 6 items each (block **A** in figure 2). This sequence is fed into the first temporal convolutional block (block **B** in figure 2). It contains dilated convolutional layers, weight-normalizations, and dropout layers. There is also a skip-connection in the block which improves neural network training convergence. There are 4 such blocks in the convolutional part of neural model. Output of previous layer is passed into input of the following layer. As a result of all computations that are performed in convolutional part of the model, we receive sequence of vectors describing human gestures recorded in data. The recurrent part of the model (block **C** in figure 2) is responsible for storing the time context and also learning the structure in sequences of gestures. The core block of this part is LSTM cell which is storing information related to recent user activity in it's state vector. The last layer of the network converts the cell state into the probability of the specific user activity. Then obtained sequence of probability values is converted into the time range of the activity. To perform this conversion, a threshold applied to probability values and then all point, where probability is greater than a threshold, are clustered (with DBSCAN) into activity interval. Such approach gives robust and accurate predictions of eating activity intervals. DBSCAN clustering algorithm was selected as this algorithm is best suited for combining close predictions into continuous one [8].

V. Evaluation

To evaluate the model we estimate the difference between the output probability of the model and our annotated ground truth. Since the algorithm consists of several steps we, firstly, optimize the architecture of end-to-end NN model, fit the model and adjust its hyperparameters on train data, secondly, optimize the threshold and DBSCAN parameters on validation data and, finally, estimate a score value on test samples of the dataset. Snack intakes (eating events with duration less than 5 minutes) were excluded from consideration The detailed information of validation structure can be found in **Table II**.

TABLE II

Data splitting information

	Train ("non-free-living")	Validation ("non-free-living")	Test1 ("non-free-living")	Test2 ("free-living")
Subjects	63	14	14	6
Eating events	1964	449	440	149
Hours	18900	4200	4100	780
Hours (eating)	547	125	124	21

Many researchers have suggested different ways to estimate the model quality. However, previously proposed metrics are not applicable to our tasks. Thus, we develop a new approach for reliable estimating the true quality of a model with the ability to select the desired prediction error. In this paper we introduce our metric, which is developed for comparison true and predicted food intake periods and takes into account both start and stop of eating process.

A. Proposed metric

In our metric we consider single ground truth food intake as one event. Besides, we measure the overlap between predicted and true segments.

To estimate the overlap we utilize an *Intersection over Union* (IoU) [25] metric. This metric is usually used for image detection, but we apply it for 1D task. Suppose we have two segments $s = (t_1, t_2)$ – true food intake period and $\hat{s} = (t_3, t_4)$ – predicted food intake period. Thus $\text{IoU}(s, \hat{s}) = |s \cap \hat{s}||s \cup \hat{s}|$, where $|s|$ is the length of the segment.

Define $\mathbf{s} = (s_1, s_2, \dots, s_n)$ as a vector of ground truth food intakes bounds and $\hat{\mathbf{s}} = (\hat{s}_1, \hat{s}_2, \dots, \hat{s}_m)$ as a vector of predicted food intakes bounds. Computing the pairwise IoU between the predicted and true period, give us a matrix M: $M_{ij} = \text{IoU}(s_i, \hat{s}_j)\, \forall\, i \in \overline{1, n},\, \forall\, j \in \overline{1, m}$.

The following statistics, True Positives (TP), False Positives (FP) and False Negatives (FN) are defined with respect to the parameter τ, which we call *IoU threshold*:

- Ground truth food intake s is considered to be $\text{TP}(\tau)$ if there is *any* predicted food intake $\hat{s} : \text{IoU}(s, \hat{s}) > \tau$.
- Ground truth food intake s is considered to be $\text{FN}(\tau)$ if there is *none* predicted food intake $\hat{s} : \text{IoU}(s, \hat{s}) > \tau$.
- Predicted food intake $\hat{s}$ is considered to be $\text{FP}(\tau)$ if there is *none* ground truth food intake s : $\text{IoU}(s, \hat{s}) > \tau$.

The total number of TP, FP and FN are calculated as:
$\text{TP}(\tau) = \#\{i \mid \exists\, j \in \overline{1, m} : M_{ij} > \tau\}$
$\text{FP}(\tau) = \#\{j \mid M_{ij} < \tau\, \forall\, i \in \overline{1, n}\}$
$\text{FN}(\tau) = \#\{i \mid j \in \overline{1, m} : M_{ij} > \tau\}$

Precision, Recall and F1-score are derived from TP, FP an FN statistics, therefore should also be computed with respect to τ since these metrics are usually used for detection task and represents the quality of the model. $\text{Precision}(\tau) = \frac{\text{TP}(\tau)}{\text{TP}(\tau)+\text{FP}(\tau)}$, $\text{Recall}(\tau) = \frac{\text{TP}(\tau)}{\text{TP}(\tau)+\text{FN}(\tau)}$, $\text{F1}(\tau) = 2 \cdot \frac{\text{Precision}(\tau)\cdot\text{Recall}(\tau)}{\text{Precision}(\tau)+\text{Recall}(\tau)}$

B. Alternative metric

To compare with existed results, we use the metric proposed by Thomaz *et al.* [8], as this paper performs the task similar to ours. There is the only fundamental difference between our and Thomaz' metrics. While we consider predicted or ground truth food intake as one event, ref. [8] consider one time segment with specified duration as one event. This may lead to splitting of food intakes into several time segments and makes this metric less interpreted in real applications.

C. Results

We selected the best parameters of end-to-end NN on train split and adjusted parameters of thresholds, filter and clustering for postprocessing on validation part of dataset.

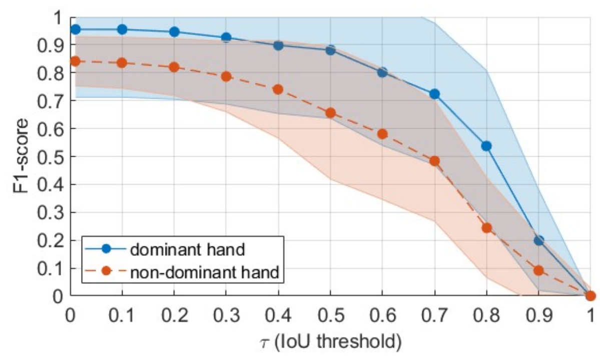

Fig. 3. F1-score with ±std shading for proposed metric evaluation. The values for different IoU were averaged over test ("non-free-living") users with respect to hand on which the smartwatch was worn

We computed our proposed metric for different values of τ (**Figure 3**) for Test1 split ("non-free-living"). Large τ is required for accurate prediction of both – start and stop of eating process. But our goal is to detect the fact of eating event as soon as possible after it starts. Our goal requires $\tau = 0.1$ (hereinafter F1-score at $\tau = 0.1$ is used for performance estimation). In this case, the end-to-end NN model provides F1-score of 0.955 ± 0.24 for 14 subjects ("non-free-living" scenario dataset) wearing smartwatch on the dominant hand and 0.836 ± 0.09 for the same subjects wearing smartwatch on the non-dominant hand during two weeks.

The performance of the developed algorithm was also estimated for 6 subjects at "free-living" scenario. **Figure 4** shows the particular metrics for each of subjects and for each hand. It should be noticed that Subject 1 in contrast to others

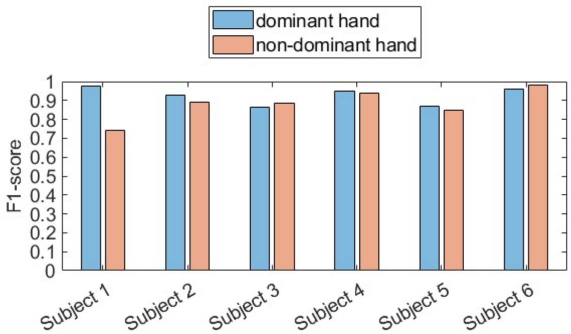

Fig. 4. F1-score metrics for each subject of "free-living" dataset

TABLE III
ALGORITHM PERFORMANCE AT TEST DATA SPLITS

(for $\tau = 0.1$)	At Test1 ("non-free-living")	At Test2 ("free-living")
For dominant hand		
F1-score	0.955	0.925
TP	419	135
FN	20	14
FP	19	8
For non-dominant hand		
F1-score	0.836	0.881
TP	370	126
FN	70	23
FP	75	11

is left-hander, so the ML model for dominant hand (related to eating gestures) was implemented for left hand to achieve F1-score = 0.974 (otherwise 0.937) and ML model for non-dominant hand (related to accompanying micro-gestures) was implemented for right hand to achieve F1-score = 0.742 (otherwise 0.554). Performance metrics with TP, FN and FP cases statistics are presented in the **Table III**.

VI. DISCUSSIONS AND CONCLUSIONS

This paper considers the problem of recognizing eating activities by passively tracking wrist motions. In our approach we use algorithm based on end-to-end NN model containing convolutional and recurrent parts to detect start of eating events. The algorithm is trained and optimized large and diverse dataset including 2413 eating events (different types of menu and eating scenarios) from 77 subjects (well distributed by parameters) with whole-day background data.

We estimate the performance of the detection algorithm by proposed evaluation metric on our test datasets of 20 participants (14 at "non-free-living" scenario, 6 at "free-living" scenario). Average F1-score is higher than 0.90 for both tests scenarios. The essence of the proposed algorithm is the top performance among existing solution in the area of smartwatch-based fully-automatic detection of eating events. Especially high performance is achieved for micro-gesture-based algorithm for non-dominant hand providing F1-score > 0.83.

These results show that the proposed approach is robust and can be used as a passive detection system of eating events in free-living settings. The proposed algorithm has some constraints:

- Short meals and snacks (duration < 5 min) are ignored,
- Eating only at seating position at table,
- No eating at transport.

For future work we aim to overcome these constraints and improve our detection system robustness for any user scenarios. The proposed approach of physical activity detection could be implemented to other applications and user scenarios (monitoring of drinking, smoking, tooth-brushing, hands and face washing, etc.). Monitoring of user's activities supported by health recommendation system will motivate users to achieve and maintain a healthy lifestyle.

REFERENCES

[1] C. C. Yang and Y. L. Hsu. A review of accelerometry-based wearable motion detectors for physical activity monitoring. *Sensors (Basel)*, 10(8):7772–7788, 2010.

[2] Shubham Ranakoti, Shagneet Arora, Shweta Chaudhary, Shakun Beetan, Amtoj Singh Sandhu, Padmavati Khandnor, and Poonam Saini. Human fall detection system over imu sensors using triaxial accelerometer. In Nishchal K. Verma and A. K. Ghosh, editors, *Computational Intelligence: Theories, Applications and Future Directions - Volume I*, pages 495–507, Singapore, 2019. Springer Singapore.

[3] Anis Davoudi, Amal Asiri Wanigatunga, Matin Kheirkhahan, Duane Benjamin Corbett, Tonatiuh Mendoza, Manoj Battula, Sanjay Ranka, Roger Benton Fillingim, Todd Matthew Manini, and Parisa Rashidi. Accuracy of samsung gear s smartwatch for activity recognition: Validation study. *JMIR Mhealth Uhealth*, 7(2):e11270, Feb 2019.

[4] D. Jarchi, J. Pope, T. K. M. Lee, L. Tamjidi, A. Mirzaei, and S. Sanei. A review on accelerometry-based gait analysis and emerging clinical applications. *IEEE Reviews in Biomedical Engineering*, 11:177–194, 2018.

[5] Thang Hoang, Deokjai Choi, and Thuc Nguyen. Gait authentication on mobile phone using biometric cryptosystem and fuzzy commitment scheme. *International Journal of Information Security*, 14(6):549–560, Nov 2015.

[6] N. Rigby. Eating and obesity–the new world disorder. *Nutrients*, 5(10):4206–4210, Oct 2013.

[7] Yujie Dong, Jenna Scisco, Mike Wilson, Eric Muth, and Adam Hoover. Detecting periods of eating during free-living by tracking wrist motion. *IEEE journal of biomedical and health informatics*, 18, 09 2013.

[8] Edison Thomaz, Irfan Essa, and Gregory D. Abowd. A practical approach for recognizing eating moments with wrist-mounted inertial sensing. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '15, pages 1029–1040, New York, NY, USA, 2015. ACM.

[9] K. Kyritsis, C. Diou, and A. Delopoulos. Modeling wrist micromovements to measure in-meal eating behavior from inertial sensor data. *IEEE Journal of Biomedical and Health Informatics*, pages 1–1, 2019.

[10] Konstantinos Kyritsis, Christos Diou, and Anastasios Delopoulos. A data driven end-to-end approach for in-the-wild monitoring of eating behavior using smartwatches. *IEEE Journal of Biomedical and Health Informatics*, 25(1):22–34, 2021.

[11] Philipp V. Rouast, Hamid Heydarian, Marc T. P. Adam, and Megan E. Rollo. Oreba: A dataset for objectively recognizing eating behavior and associated intake. *IEEE Access*, 8:181955–181963, 2020.

[12] Shibo Zhang, Rawan Alharbi, Matthew Nicholson, and Nabil Alshurafa. When generalized eating detection machine learning models fail in the field. In *UbiComp/ISWC Adjunct*, 2017.

[13] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.

[14] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Kdd*, volume 96, pages 226–231, 1996.

[15] Keum San Chun, Sarnab Bhattacharya, and Edison Thomaz. Detecting eating episodes by tracking jawbone movements with a non-contact wearable sensor. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2(1):4:1–4:21, March 2018.

[16] R. Zhang and O. Amft. Free-living eating event spotting using emg-monitoring eyeglasses. In *2018 IEEE EMBS International Conference on Biomedical Health Informatics (BHI)*, pages 128–132, March 2018.

[17] Larry M Manevitz and Malik Yousef. One-class svms for document classification. *Journal of machine Learning research*, 2(Dec):139–154, 2001.
[18] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
[19] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
[20] Shibo Zhang, William Stogin, and Nabil Alshurafa. I sense overeating: Motif-based machine learning framework to detect overeating using wrist-worn sensing. *Information Fusion*, 41:37 – 47, 2018.
[21] Edison Thomaz, Abdelkareem Bedri, Temiloluwa Prioleau, Irfan Essa, and Gregory D Abowd. Exploring symmetric and asymmetric bimanual eating detection with inertial sensors on the wrist. In *Proceedings of the 1st Workshop on Digital Biomarkers*, pages 21–26. ACM, 2017.
[22] Marietta Papadatou-Pastou, Eleni Ntolka, Judith Schmitz, Maryanne Martin, Marcus R Munafo, Sebastian Ocklenburg, and Silvia Paracchini. Human handedness: A meta-analysis, Apr 2019.
[23] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. *arXiv preprint arXiv:1603.02754*, 2016.
[24] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling, 2018.
[25] Wikipedia. Jaccard index — Wikipedia, the free encyclopedia. `http://en.wikipedia.org/w/index.php?title=Jaccard%20index&oldid=890384326`, 2019. [Online; accessed 12-April-2019].