

NeckSense: A Multi-Sensor Necklace for Detecting Eating Activities in Free-Living Conditions

SHIBO ZHANG, Northwestern University, United States

YUQI ZHAO, Northwestern University, United States

DZUNG TRI NGUYEN, Northwestern University, United States

RUNSHENG XU, Northwestern University, United States

SOUGATA SEN, Northwestern University, United States

JOSIAH HESTER, Northwestern University, United States

NABIL ALSHURAFA, Northwestern University, United States

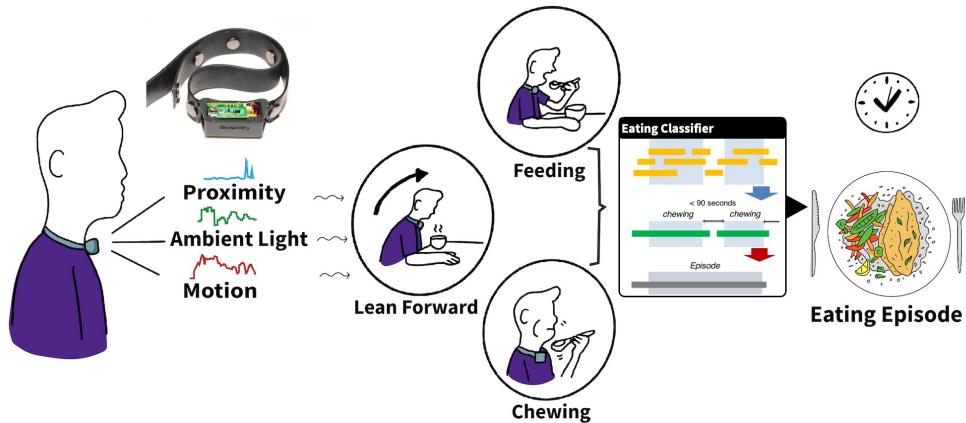


Fig. 1. Our non-contact, day-long battery life necklace, *NeckSense*, collects proximity, ambient light, and motion signals to detect chewing actions, feeding gestures, and lean forward motion, which allows detection of eating episodes occurring throughout the day. *NeckSense* enables long-term studies for monitoring eating behavior in free-living conditions.

We present the design, implementation, and evaluation of a multi-sensor, low-power necklace, *NeckSense*, for automatically and unobtrusively capturing fine-grained information about an individual's eating activity and eating episodes, across an

Authors' addresses: Shibo Zhang, Northwestern University, Evanston, IL, United States, shibo.zhang@northwestern.edu; Yuqi Zhao, Northwestern University, Evanston, IL, United States, yuqizhao2019@u.northwestern.edu; Dzung Tri Nguyen, Northwestern University, Evanston, IL, United States, dzung@u.northwestern.edu; Runsheng Xu, Northwestern University, Evanston, IL, United States, runshengxu2017@u.northwestern.edu; Sougata Sen, Northwestern University, Evanston, IL, United States, sougata.sen@northwestern.edu; Josiah Hester, Northwestern University, Evanston, IL, United States, josiah@northwestern.edu; Nabil Alshurafa, Northwestern University, 680 N Lake Shore Dr, Chicago, IL, 60611, United States, nabil@northwestern.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2474-9567/2020/6-ART72 \$15.00

<https://doi.org/10.1145/3397313>

entire waking day in a naturalistic setting. *NeckSense* fuses and classifies the proximity of the necklace from the chin, the ambient light, the Lean Forward Angle, and the energy signals to determine *chewing sequences*, a building block of the eating activity. It then clusters the identified chewing sequences to determine eating episodes. We tested *NeckSense* on 11 participants with and 9 participants without obesity, across two studies, where we collected more than 470 hours of data in a naturalistic setting. Our results demonstrate that *NeckSense* enables reliable eating detection for individuals with diverse body mass index (BMI) profiles, across an entire waking day, even in free-living environments. Overall, our system achieves an F1-score of 81.6% in detecting eating episodes in an exploratory study. Moreover, our system can achieve an F1-score of 77.1% for episodes even in an all-day-long free-living setting. With more than 15.8 hours of battery life, *NeckSense* will allow researchers and dietitians to better understand natural chewing and eating behaviors. In the future, researchers and dietitians can use *NeckSense* to provide appropriate real-time interventions when an eating episode is detected or when problematic eating is identified.

CCS Concepts: • **Human-centered computing** → **Mobile devices; Ubiquitous and mobile computing; Empirical studies in ubiquitous and mobile computing;** • **Applied computing** → **Law, social and behavioral sciences.**

Additional Key Words and Phrases: eating activity detection, automated dietary monitoring, human activity recognition, wearable, neck-worn sensor, sensor fusion, free-living studies

ACM Reference Format:

Shibo Zhang, Yuqi Zhao, Dzung Tri Nguyen, Runsheng Xu, Sougata Sen, Josiah Hester, and Nabil Alshurafa. 2020. NeckSense: A Multi-Sensor Necklace for Detecting Eating Activities in Free-Living Conditions. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 4, 2, Article 72 (June 2020), 26 pages. <https://doi.org/10.1145/3397313>

1 INTRODUCTION

Automatically and unobtrusively monitoring an individual's eating activity in free-living settings has been a long-standing objective of the research community [11]. The possibility of monitoring eating activity automatically will allow individuals, researchers, and ultimately clinicians to support various wellness goals in the form of interventions. For example, clinicians can design a system to trigger real-time interventions when people spend too much time eating [51], or a researcher can request for timely information about energy consumption in an individual's diet [12]. It has been well established that such interventions can help treat eating disorders in the long term and improve the quality of life [16]. However, it is difficult to provide eating-related interventions without automatically detecting the eating activity and fine-grained actions associated with it. Thus, it is extremely important for any eating activity monitoring system to detect the eating activity that occurs every day, with diverse food choices, at varying times and contexts during the waking day, and a myriad of environments.

The current gold-standard techniques for monitoring an individual's dietary habits are either asking the individual to maintain a food journal or for dietitians to perform 24-hour dietary recalls to record food items that were consumed by the individual in the previous 24 hours. As participants often record all their meals at the end of the day, such techniques are subject to user bias and forgetfulness [5, 25]. To complement 24-hour dietary recalls and enable more timely feedback, a large body of work around automatic dietary monitoring has emerged, with wearable sensors showing promise in automatically detecting eating behaviors and linking them to eating episodes [5, 27]. Researchers have explored the possibility of detecting eating activity using wrist-based inertial sensor data [18, 55, 61], on-ear- or on-throat-based audio sensor data [8, 13, 45], image-based information [41, 46, 56], or a combination of one or more of these techniques [33, 35, 50]. There is existing support among behavioral researchers, social scientists, and clinicians for mobile adaptive interventions such as ecological momentary interventions (EMIs) and just-in-time adaptive interventions (JITAIs) [29]. These interventions use contextual inputs, such as detection of an eating episode or number of mouthfuls consumed, to adapt the content and timing of interventions to participants. The behavioral researchers, or social scientists can use the fine-grained output of the automatic dietary monitoring systems to provide necessary interventions (e.g., providing intervention when they detect poor chewing habits [53]).

However, many of the existing automated eating detection systems are either obtrusive, or they have not been tested in a completely free-living setting with diverse BMI populations (the people most likely to benefit from such a technology) or have not been tested in longitudinal studies beyond a few days. Indeed, many of the existing dietary monitoring studies have been validated on a student population and in a university environment. This population might not be representative of a population that could benefit from real-time interventions intended to prevent problematic eating behaviors. It is currently unknown whether a model developed on people with a normal BMI will likely succeed when tested on people with obesity, a population most likely to benefit from such adaptive interventions. In our study, we recruited people with varying BMI levels and found that automatic dietary monitoring system models trained on participants without obesity have a worse performance when tested on individuals with obesity (details in Section 6). It is therefore necessary to build systems and models that can support real-time eating detection in various demographic groups.

To bridge this gap, in this paper, we present the design and evaluation of *NeckSense*, a *multi-sensor necklace* for automatically detecting an eating episode. The goal of *NeckSense* is to ensure that it can *accurately, automatically, and unobtrusively* monitor its wearer's eating activities that occur *at any time* during the day, *in any setting*, while ensuring that the device has at least an entire waking day of *battery life*. However, most importantly, we want to ensure that *NeckSense* generalizes and detects eating accurately for a demographically diverse population, including people with and without obesity.

In this work, we assume that the combination of an individual's leaning forward action, performing the feeding gestures, and then periodicity of chewing behavior together constitutes an in-moment eating activity. We design and develop *NeckSense*, a necklace with an embedded proximity sensor, an ambient light sensor, and an Inertial Measurement Unit (IMU) sensor that can capture these aforementioned actions. Although researchers have previously explored using the proximity sensor to detect the eating activity by monitoring the jaw movement [15], in this work we demonstrate that we can detect the eating activity more accurately if we augment the proximity sensor data with the ambient light and IMU sensor data. To detect the eating activity, *NeckSense* fuses and classifies features extracted from these sensors. It then clusters the predicted eating activity to determine an eating episode. Figure 1 provides an overview of *NeckSense*. We evaluate the feasibility of *NeckSense* by conducting two user studies: a longer, intermittently-monitored free-living *Exploratory Study* (semi-free-living) and another completely *Free-Living Study*. The Exploratory Study allowed us to identify sensors that were useful in detecting the eating activity. It also allowed us to identify usability concerns with the necklace. The findings of the Exploratory Study allowed us to improve *NeckSense*. We evaluated the improved necklace on a diverse non-student population recruited for two full days in a completely free-living scenario, while participants carried out their everyday activities.

While designing, developing, and evaluating *NeckSense* to address practical challenges in free-living studies pertaining to (a) accurately monitoring the eating activity that occurs in diverse settings and with participants with in varied BMI range, (b) usability and comfort of wearing the device, and (c) collecting the sensor data energy efficiently, we make the following *key contributions*:

- (1) We describe the design and implementation of a multi-sensor necklace for detecting eating episodes. The necklace utilizes sensor data from its proximity, ambient light, and IMU sensors to determine eating activity. We utilize periodicity in chewing to apply a longest periodic subsequence algorithm on the proximity sensor signal, and then show the benefits of adding ambient light, IMU, and hour-of-day to improve proximity-based detection of eating episodes.
- (2) We evaluate the system in two studies: an exploratory semi-free-living and another completely free-living study to determine the possibility of detecting eating episodes in naturalistic settings. The studies involved participants with varied BMIs. The participants consumed 117 meals during this study period. Overall, we found that in a semi-free-living setting, the necklace could identify eating episodes at an ambitious

fine-grained per-second level with an average F1-score of 76.2% and at a coarse-grained per-episode level of 81.6%. This is an 8% improvement over using only data from the proximity sensor for eating episode detection. As expected, the fine-grained performance drops to 73.7% and the coarse-grained performance to 77.1% in a completely free-living setting. We also show improvement when comparing our algorithm to a prior system’s algorithm that uses solely proximity sensing for eating detection.

- (3) We evaluate the energy performance of the system during these studies and observed that on average the battery life of the device during the Exploratory Study was 13 hours, while the battery life improved to 15.8 hours in the Free-Living Study.
- (4) We will anonymize and make both our datasets, the source code, and the design files for the *NeckSense* hardware available for use by the community.¹ The dataset contains the sensor traces collected from 20 participants, tagged with ground truth labels generated from video and clinical standard labeling practices.

Overall, the necklace provides a practical solution for automatically monitoring eating activity in completely free-living settings. The system has been validated in both individuals with and without obesity, improving reliability of the eating detection system and laying the ground work for future mobile adaptive interventions to gain further insight on dietary habits. In the near future we anticipate that other researchers will use our dataset to validate their own methods for chewing sequence and eating episode detection.

2 RELATED WORK

Large-scale, whole-population interventions [34], such as advertising campaigns targeted towards curtailing eating-related disorders, have had little or no success in addressing the obesity epidemic [38, 39, 43, 58]. Instead, researchers are proposing just-in-time interventions [19, 28, 36, 37] to test personalized interventions that are tailored to person-specific needs [47, 54]. Detecting eating automatically is the first step towards testing these personalized just-in-time interventions. However, several factors make automated eating detection challenging to implement. These factors range from identifying the right device, signals, or form-factor to validating eating activity using visual confirmation in real-world settings.

2.1 Eating Detection Techniques

Researchers have proposed several techniques to automatically detect the eating activity. We detail some of the proposed techniques, grouped based on signal type.

Audio and Video: Automated sensing platforms using image- and audio-based techniques with sensors placed around the throat or ear have shown promise in detecting eating [7, 13, 35, 40, 45, 56, 57]. However, the utility of these sensors is limited by short battery life (reducing autonomy) and security or privacy concerns. Eating detection systems designed without camera or audio components (as in this work) reduce privacy concerns and enable longer battery life [3].

Physiological: Various techniques have been used to indirectly detect chewing, which involve electromyography (EMG) [60] sensors on an eye-glass frame, in-the-ear microphones combined with photoplethysmography (PPG) sensors [42], and mechanoelectrical sensors placed directly around the neck [7], among others [30, 31]. While these sensors have shown promise in controlled environments, they have not been tested in free-living populations or over significant periods of time, and many need direct skin contact, which can be uncomfortable and affect adherence, thereby limiting their potential utility in longitudinal studies.

Inertial or Proximity: Several researchers have proposed techniques to automatically detect eating using inertial sensors embedded in wrist-worn devices and phones [18, 35, 50, 55]. However, wrist-worn sensors are limited by several confounding gestures [61]. More recently, researchers have explored the possibility of detecting

¹The dataset and code can be found at <http://doi.org/10.5281/zenodo.3774395>.

Table 1. Comparing the literature on in-wild eating detection to this work.

Year	Study	Sensors	On-body position	No. of participants	Avg hours per day	Validation video	Non-student	Obese
2014	Fontana et al. [22]	S1, S4, S6	Ear, wrist, chest	12	24.0	✗	✓	✓
2015	Thomaz et al. [55]	S1	Wrist	7+1	5.7/13.6	✗	✗	✗
2015	Bedri et al. [10]	S2, S5	Ear, head	6	6.0	✗	✓	✗
2016	Farooq et al. [21]	S4	Temple	8	3.0	✗	✓	✗
2017	Bedri et al. [9]	S1-S3, S5, S7	Neck, ear	10	4.5	✓	✓	✗
2017	Zhang et al. [60]	S8	Ear	10	6.1	✗	✗	✗
2017	Mirtchouk et al. [35]	S1-S3, S7	Ear, wrist, head	11	11.7	✗	✓	✗
2018	Sen et al. [49]	S1, S2, S10	Wrist	9	5.8	✗	✓	✗
2018	Chun et al. [15]	S5	Neck	17	4.6	✗	✗	✗
2018	Bi et al. [13]	S7	Ear	14	2.3	✓	✗	✗
2020	This work	S1-S3, S5, S9	Neck	10+10	4.9/9.5	✓	✓	✓

S1 - accelerometer, S2 - gyroscope, S3 - magnetometer, S4 - piezo, S5 - proximity, S6 - radio frequency, S7 - microphone, S8 - electromyography, S9 - light, S10 - camera

chewing bouts and eating episodes using only an on-neck proximity sensor, combined with a threshold-based algorithm [15]. We show that our multi-sensor method outperforms such a threshold-based method in recalling chewing sequences, even in a naturalistic setting.

Our method re-imagines **inertial or proximity** sensing modality by fusing them with data from an ambient light sensor, chewing-related periodic subsequence features, and time of day. We present an eating detection framework that uses this fusion and is tested in free-living settings.

2.2 Eating Detection in Naturalistic Setting

Researchers have explored the use of various sensors for detecting the eating activity outside laboratory settings [9, 10, 20–22, 42, 60], but the length of continuously recorded experimentation rarely exceeds more than few hours in a day, limiting its potential for longitudinal eating-related studies. Table 1 lists some existing research that utilizes various sensing modalities to detect the eating activity. Overall, we identified some key factors that have made these free-living studies hard to execute:

Validation Requirement: A straightforward validation technique that researchers commonly adopt is self-reporting; that is, the researchers requested participants to note down the start and end time of every meal that the participant consumes [17, 48, 49] or create a food journal [42, 60]. However, manually noting details about every meal can be burdensome and error-prone [25]. To reduce the burden on participants, several researchers have proposed the use of a front-, upward-, or downward-facing camera [9, 13, 26, 55, 56, 59] for validating their system’s performance. Researchers have tested various on-body camera positions and found that the shoulder camera is the best for privacy and easiest to wear [2]. This placement was also found to be best at capturing eating-related activities [4]. We thus use a shoulder-mounted camera for validating our device’s performance.

Diverse Population Requirement: While several sensing systems have shown promise in the wild, they have predominantly been validated within the student population [13, 15]. Before such systems can be truly generalizable, additional data are needed from a diverse population sample (especially including people belonging to various body mass ranges) that are not only student based or focused on people without obesity. People with varied body shapes may experience varied comfort levels and accuracy of any eating detection system, potentially confounding the system, but enabling deeper insight and translation of research to practice. To the best of our

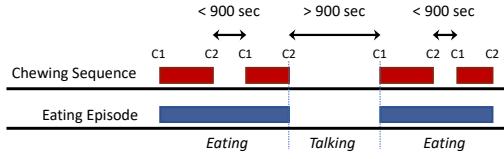


Fig. 2. Schematic of eating episodes that is composed of multiple chewing sequences. C1 and C2 correspond to start and end, respectively, of each chewing sequences. We use a data-driven approach to determine minimum interval between chewing sequences to identify episode boundaries.

knowledge, we are among the few to explicitly validate an automated eating detection system with a population of people with obesity.

3 SYSTEM DESIGN & IMPLEMENTATION

While keeping the challenges of free-living data collection in mind, we present our multi-sensor neck-worn device that is tolerant of varying sensor positions as captured in real-world settings, comfortable to wear, semi-autonomous (only requiring users to turn on/off and charge the device at the end of the day), and validated using a custom-designed ground truth camera with full-day battery life (allowing uninterrupted ground truth monitoring). We next present the design of the multi-sensor eating detection system and also describe the camera used for ground truth validation. These devices were used in two user studies: a semi-free living Exploratory Study and a completely Free-Living Study.

3.1 Defining Eating

We consider the act of mastication (i.e., the act of periodically bringing the jaw together to grind food) as a *chewing sequence*. In this paper, we define an *eating episode* as an aggregate of *chewing sequences* that occur within a short duration of time, and these chewing sequences are separated from other chewing sequences by a large time gap. An eating episode can represent either a snack or a meal. Figure 2 pictorially represents a chewing sequence and an eating episode.

Chewing Sequence: We define a chewing sequence as a combination of chews that occur together with breaks no greater than 3 seconds between subsequent chews. In this work we determine chews (detailed in Section 4) by applying a prominent-peak detection algorithm on the proximity sensor data, followed by running a longest period subsequencing algorithm, and finally extracting and classifying features from proximity, IMU, and ambient light sensors.

Eating Episode: We define an eating episode as a group of chewing sequences with inter-chewing sequence breaks no larger than δ seconds. Two adjacent chewing sequences with a gap longer than δ seconds are identified as two separate eating episodes. We use a data-driven approach to determine the value of δ . Figure 3 presents the cumulative distribution function (CDF) of the interval between subsequent chewing sequences. From the figure we can observe that a value for δ between 540 and 1100 seconds provides a clear boundary between

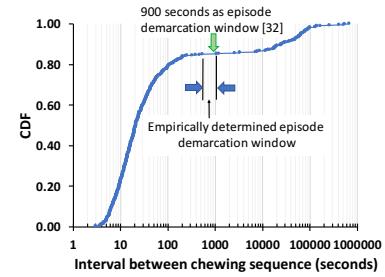


Fig. 3. Cumulative distribution function representing an empirical approach to determine eating episode boundaries based on the time between the end of one chewing sequence and the start of the next chewing sequence. Our empirically determined inter-episode gap is similar to the inter-episode gap as suggested by Leech et al. [32].

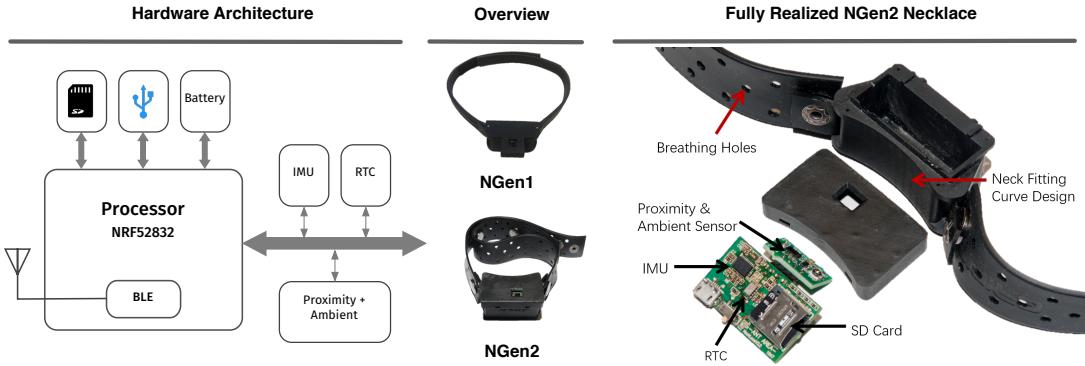


Fig. 4. Hardware architecture and overview of two generations of necklace. NGen1 was used in the Exploratory Study, while NGen2 was used in the Free-Living Study.

eating episodes. We decided to use $\delta = 900$ seconds in our evaluation. This choice of δ empirically validates the inter-episode interval reported or used by researchers previously [13, 32]. Applying this rule allowed us to turn chewing sequence labels to eating episode labels with exemption from ambiguity when evaluating eating episode prediction.

Our definition of eating episodes as well as the choice of δ is similar to the definition presented by Bi et al. [13]. However, Bi et al. do not provide a definition of the chewing sequence. On the other hand, Chun et al. propose the idea of sub-dividing the chewing action into both *chewing* and *chewing bouts*, where chewing was described as an uninterrupted sequence of chews that lasted for 5 seconds; chewing that occurred within 30 seconds of another chewing was part of the same chewing bout. Unlike prior works, our choice of grouping chewing sequences into episodes is from a data-driven approach.

3.2 Necklace Design

To evaluate the feasibility of using a neck-worn method for eating and chewing detection, we designed and developed a multi-sensor necklace prototype, *NGen1*, that was used in the Exploratory Study. Lessons learned from the study helped in designing and developing the *NGen2* prototype necklace that was used in the Free-Living Study (Section 5 presents the study details). Figure 4 presents the overview and hardware architecture of the necklace. The *NGen2* device is still a prototype, and the form-factor of a finished *NeckSense* product is beyond the scope of our research and will likely result in a much smaller prototype than the *NGen2*.

Signal and Sensors: Our proposed system design is based on observation and validation of eating activities in laboratory settings. We observed numerous eating activities and noticed that during most eating activity an individual leans forward to reach for the food; the hand feeds the food to the mouth; and the jaw continuously grinds to perform the repetitive chewing action. To capture these eating related signatures, we evaluated several sensors and finally chose an IMU sensor, a proximity sensor, and an ambient light sensor, all of which were embedded into a neck-worn device. The IMU sensor facilitates determining *leaning forward* movement (i.e., the Lean Forward Angle [LFA]). The proximity sensor on the necklace is directed towards the chin and allows monitoring the variation in signal while an individual is chewing. The on-necklace ambient light sensor's reading drops when the individual's hand moves towards the mouth during a feeding gesture, thus allowing detecting of the feeding gesture. We observed (described in Section 6) that although each sensor can individually detect the eating action, fusing the signals from these sensors improves the overall detection performance. We thus use all the sensors in the final *NeckSense*'s design.

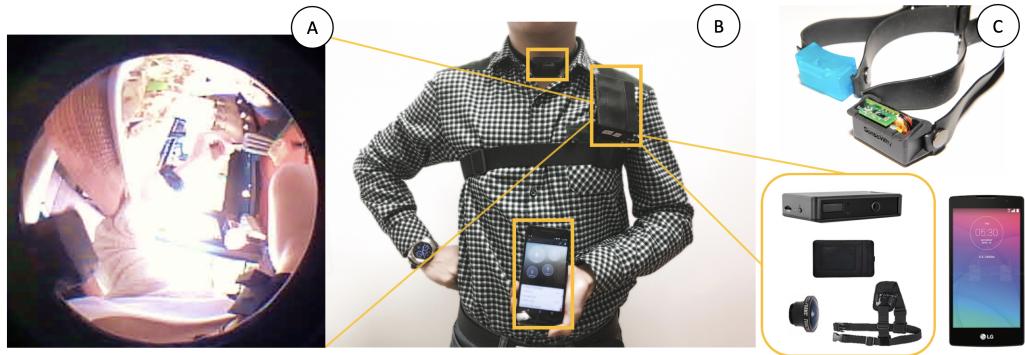


Fig. 5. Participant wearing the devices during the Free-Living Study . (A) Representative still image of the ground truth camera output showing the chin, hand, food, and utensil. (B) Participant wearing the necklace and Ground Truth camera. (C) All devices for Free-Living Study.

Hardware Design: Both NGen1 and NGen2 are centered around a Nordic NRF52832 SiP that has an ARM Cortex M4 CPU and a Bluetooth Low Energy (BLE) communication interface. Two sensors chips, VCNL4040 and BN0080, are embedded onto the necklace. The VCNL4040 is a combination of proximity and ambient light sensor. We intend to capture the chewing action using this sensor, and thus it faces upward, toward the user's chin. The BN0080 is an IMU sensor that is utilized for computing the Lean Forward Angle (LFA). All sensors were sampled at 20 Hz. The proximity, ambient light, and IMU sensors sequentially write to a memory buffer. When the buffer reaches its capacity, the data are written to a local micro SD (μ SD) card. The buffer helps in increasing the writing throughput and reducing the power consumption by writing periodically. The whole device is powered by a 350 mAh Li-ion battery, which guarantees more than 14 hours of battery life, making it possible to monitor eating activities during an entire waking day.

Time Keeping and Synchronization: Each chewing gesture is a short-lived action, and thus to ensure synchronization it is important to maintain a highly accurate time keeping system. The necklace has to ensure that it always records an accurate timestamp without any manual calibration and time setting. NGen1 did not have a real-time clock (RTC) module with a backup battery, and so when the main battery failed, the data had invalid timestamps, making synchronization of the necklace data and camera video challenging. As an improvement, we introduce an RTC in the NGen2 necklace so that it maintains time at a 2-ppm accuracy. With its rechargeable backup battery, it can keep time for more than 1 year after the main power source is dead. The accurate RTC timestamp ensures that the system's clock does not drift substantially. Additionally, whenever the necklace connects with a hub device (e.g., a smartphone), it automatically fetches the current time from a time server by calling the Current Time Service via BLE; the service helps to correct the time and keep it during a long period of study. We empirically observed that this feature eliminates the inevitable time drift of <180 msec per day.

Mechanical Design: We designed the necklace along with a case and a strap, as shown in Figure 4. This design has evolved over multiple iterations (with feedback from the Exploratory Study) to ensure that, in addition to being functional, it is comfortable to wear. The necklace's band is a sports-style porous silicone band that wraps around the neck. This was an improvement in NGen2 over NGen1 where we used a leather band. The case housing the printed circuit board (PCB) and sensors is connected to this band. The silicone material and carefully balanced weight of the case ensures that the sensor's position is stationary, even during vigorous activities. This property guarantees that we can collect quality data, even in noisy free-living environments. The resin material used for manufacturing the case (commonly used in wearable jewelry) is smooth and skin friendly. One face of the

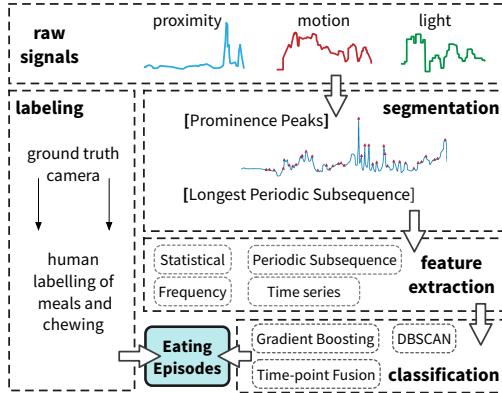


Fig. 6. Chewing and meal detection framework and validating it with the labeled ground truth.

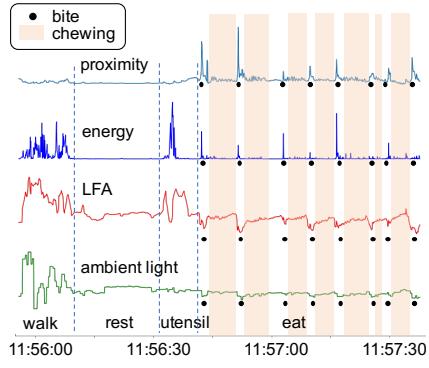


Fig. 7. Four signals (proximity, energy, LFA, and ambient light) captured while the wearer is walking, resting, utensiling, and eating.

case in NGen2 was curved to ensure that it was comfortable on the neck; this was an improvement based on feedback from NGen1. Based on user feedback, we also add breathable holes to the band to increase user comfort. Recent research indicates that a participant will adhere to wearing a device if the participant can decide on the device from a collection of devices with varied designs [1]. We thus manufactured devices with diverse designs and colors. We anecdotally observed that when participants chose their own device, they wore it for a longer duration.

3.3 Ground Truth Video Camera

We used a fish-eye lens fitted wearable camera (that captured both audio and video) to continuously record participant behavior including fine-grained eating related activities. This recording provided the ground truth labels for learning models and validating free-living studies [2]. Once the neck-worn sensor is validated, we envision that the necklace's user will not use this camera. Figure 5 illustrates the various components of the wearable camera system. The system comprises of a Qcam QSD-722 camera, a fish-eye lens, a shoulder strap, and an ID card to conceal the device. The camera is positioned around the shoulder with the lens facing the participant's dominant hand rather than pointing outward (e.g., as in SenseCam [26]). This minimizes privacy concerns of bystanders, although bystanders to the side of the participant are noticeable. At this angle, the wearable camera can capture the mouth, the jaw, the neck, and the utensil and foods consumed (see Fig. 5A). As privacy is a significant concern for the participant, we provided participants with the option of deleting segments of video data they did not want study supervisors to see.

4 EATING DETECTION FRAMEWORK

We next present our method of predicting the chewing sequences and eating episodes from the signals generated by the sensors. The entire processing pipeline is presented in Figure 6. The pipeline consists of five steps that we next describe.

4.1 Signal Pre-processing

The signals extracted or calculated from the necklace's sensors include (i) proximity to chin, (ii) ambient light value, (iii) LFA, and (iv) energy signal defined as the sum of squares of tri-axial accelerometer values.

(i) Proximity to chin: We extract the proximity to the chin from the necklace's proximity sensor. The sensitivity range of the proximity sensor is set between 5 and 250 mm. Since the average neck length of a human being varies between 100 and 120 mm, this setting should sufficiently capture the chewing action. Figure 7 shows representative signals when the participant was walking, resting, utensiling (using utensils, but not eating), and eating. During the eating activity, smaller peaks in the proximity sensor's signal signify chews, while the larger peaks signify bites.

(ii) Ambient light value: The necklace's ambient light sensor provides the ambient light value. The ambient light value is highest when the users turn their heads right or left, since the sensor is not obstructed by the head. The value is lowest when the users lower their heads and move their hands toward their mouths during a feeding gesture. We can see the periodic drop in ambient light values during the eating activity in Figure 7.

(iii) LFA: We obtained the necklace's absolute orientation in the form of quaternions from the on-necklace IMU sensor. The quaternion is a 4D vector q representing the rotation axis and the angle of rotation around that axis. q can be projected into different planes to gain physical angles and infer activities, such as leaning forward and to the side, and to determine the orientation of the wearer. However, not all of these angles are related to the eating process. The most informative angle is the LFA, the angle between the IMU and the Earth's surface. When the wearer sits straight, the LFA is close to 90°. LFA is calculated by applying the dot product of the normal vectors of two planes:

$$LFA = \arccos \langle \mathbf{n}_1, \mathbf{n}_2 \rangle$$

where the normal vector of the Earth's surface is the z-axis, and the normal vector of the IMU is obtained through the quaternion transformation:

$$\mathbf{n}_1 = [0, 0, 1] \quad \mathbf{n}_2 = q\mathbf{n}_1q^{-1}$$

where q is a unit quaternion that rotates \mathbf{n}_1 to obtain the normal vector of the IMU. It is worthwhile to note that while LFA does not always occur in the field, particularly when snacking while sitting on a couch, features from LFA can enhance detection of bites.

(iv) Accelerometer's energy value: The on-necklace IMU also provides the tri-axial accelerometer data (a_x, a_y, a_z) capturing acceleration from the three axes. We calculate the energy signal as the sum of squares of the tri-axial acceleration components, $E = a_x^2 + a_y^2 + a_z^2$. Features computed from E help reduce false positive rates generated during physical activities. From Figure 7 we can see that during the eating gesture there are peaks in the energy signal. However, unlike the peaks observed in the signal during the walking activity, the peaks observed during eating are sparse.

4.2 Labeling

The process of labeling the data allows establishing the ground truth. Annotators labeled the start and end of every chewing sequence by visually and acoustically confirming the information using the video captured by the camera. If a participant continuously chewed during her entire meal, the entire episode was labeled as one chewing sequence. Annotators marked the end of a chewing sequence when the participant stopped chewing for at least 3 seconds. For example, if the participant took four breaks that were each at least 3 seconds long during the meal, there were five chewing sequence labels in that meal.

From the labeled chewing sequences, the annotators identified the eating episodes. As discussed in Section 3.1, we define an eating episode as a group of chewing sequences with chewing sequence breaks no longer than 900 seconds. Any two adjacent chewing sequences with a gap longer than 900 seconds were regarded as two separate eating episodes. Applying this rule allowed the annotators to establish eating episode labels from the chewing sequence labels.

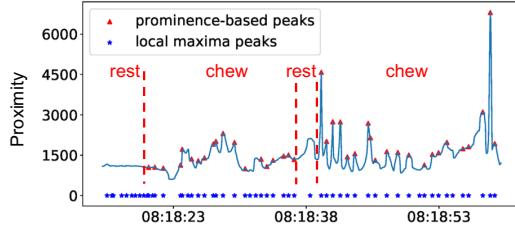


Fig. 8. Prominence-based peak finding algorithm (height=4.5) vs local maxima peaks (using 2 samples before and after the time point.)

4.3 Segmentation

We employed time-based methods on the necklace's proximity signals to detect periodic chewing patterns. The proximity sensor was selected for the segmentation step as it provides the right balance between identifying most periodic chewing patterns, which are considered eating segments (high recall), and not introducing many false positives (acceptable precision). As described in Section 4.3.1, to detect the candidate chewing subsequence, we used a peak-finding algorithm followed by a periodic subsequence algorithm in the segmentation step. Overall, we observed that both the proximity sensor and the IMU sensor could independently capture all eating episodes (recall = 1) for certain parameter value combinations. However, if we considered all sensor combinations and parameter choices that yielded a recall of over 85% in detecting eating episodes, the F1-score of detecting eating episodes using the proximity sensor was 17.2%, as compared with an F1-score of 8.8% when using the IMU sensor. This was almost a 2x improvement, and it indicated that more data could be safely discarded in the segmentation step if the proximity sensor data were used for the segmentation, as compared with using the IMU. Compared with proximity and IMU, the light sensor had a much lower recall during the segmentation step, indicating that we would miss several eating episodes if the ambient light sensor was used for the segmentation step. Thus it was not considered useful for the segmentation step.

4.3.1 Prominence-based Peak-finding Algorithm. The first step in detecting periodic sequences is to identify peaks in the proximity signal. We applied a typical peak-finding algorithm that returned both prominent and non-prominent peaks (non-prominent peaks may be associated with noise and other activities). Chewing peaks are often prominent peaks that stand out due to their intrinsic height and location relative to other nearby peaks. For example, in Figure 8 we observe several local maximum, yet non-prominent, peaks that are identified during the resting period. However, more prominent peaks are apparent during the eating (chewing action) period.

4.3.2 Longest Periodic Subsequence Algorithm. We adapted the longest periodic subsequence algorithm to identify chewing peaks that were ϵ -periodic [24]. The time points of the peaks from the prominence algorithm generated a sequence of timestamps for each peak. In this section we explain the significance of ϵ -periodic, define the periodic subsequence problem, and present a dynamic programming solution for the problem.

DEFINITION 1. ϵ -periodic: Given a sequence of increasing timestamps t_i , where $i \in \{1 \dots N\}$, the difference between consecutive numbers is $p_i = t_{i+1} - t_i, \forall i = \{1 \dots (N-1)\}$, if p_{\min} and p_{\max} are the smallest and largest values of these differences, respectively, then the sequence is defined to be ϵ -periodic if:

$$\frac{p_{\max}}{p_{\min}} < 1 + \epsilon$$

PROBLEM 1. Relative error periodic subsequence: Given a sequence of increasing numbers t_i , find all longest subsequences that are ϵ -periodic.

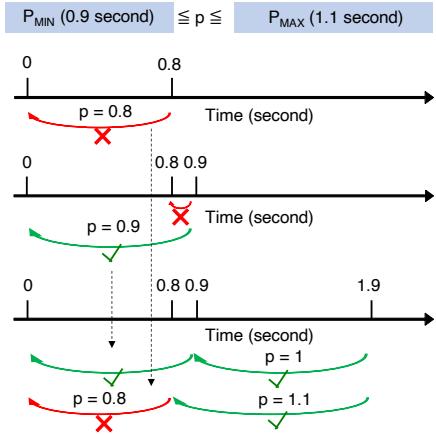


Fig. 9. Dynamic programming solution for absolute error periodic subsequence of proximity signal.

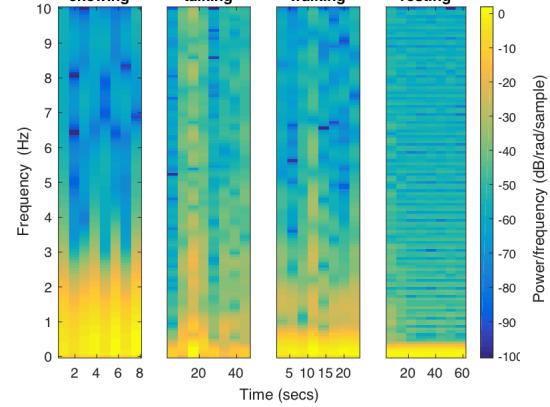


Fig. 10. Spectrogram of proximity signal for chewing, talking, walking, and resting.

PROBLEM 2. *Absolute error periodic subsequence:* Given a sequence of increasing numbers t_i , find all longest subsequences such that consecutive differences are bounded by p_{min} and p_{max} .

During chewing activity, periodic peaks are generated in the proximity signal. In terms of ϵ -periodic subsequence, the value ϵ determines to what extent the frequency of these peaks is changing over time. A smaller ϵ indicates less change in the frequency during the periodic subsequence (i.e., candidate chewing activity). We empirically investigate and set $\epsilon = 0.1$ to remove periodic subsequences with high variability in the frequency, which are most likely generated by non-chewing activity. We empirically set the range of the number of peaks for a periodic subsequence to be between 4 and 100.

Problem 1 is not trivial when the lower p_{min} and upper bound p_{max} are not known. However, given that the chewing frequency range is known [44], these bounds can be estimated. The problem can then be solved by evoking multiple calls to a function that implements the absolute error periodic subsequence problem. Each time a function call is made by passing a new p_{min} and p_{max} starting from the smallest inter-chew distance min all the way until the largest inter-chew distance max (incrementing p_{min} by multiples of $[1 + \epsilon]$). Chewing activity has been shown to mainly occur in the range of 0.94 Hz (5th percentile) to 2.17 Hz (95th percentile); as a result we set $min = 0.4$ seconds and $max = 1.5$ seconds. We solve the absolute error periodic subsequence problem using dynamic programming, by defining the following recurrence:

$$OPT[i] = \max_L OPT[j], \forall j \ni p_{min} < t_i - t_j < p_{max}$$

The absolute error periodic subsequence algorithm is called by passing the entire array of timestamps $t[i]$, a value for p_{min} and p_{max} . It then iterates through the entire array, and for every index i it calculates the longest subsequence that ends at index i , searching only the previous indices in the array that satisfy the inter-chew distance p_{min} and p_{max} . Figure 9 shows an example where prominence peaks are detected at times 0, 0.8, 0.9, and 1.9 seconds. If p_{min} is 0.9 and p_{max} is 1.1 seconds, then the optimal subsequence is of length 2 at timestamps (0, 0.9, 1.9) seconds. Sequence (0, 0.8) is not valid because the difference between 0.8 and 0 is 0.8 which is $< p_{min}$, the same as sequence (0.8, 0.9).

The time complexity analysis of the absolute error periodic subsequence algorithm is $O(N)$, assuming the valid number of predecessors is constant. If the distance between p_{min} and p_{max} is a function of N , then this

Table 2. List of features extracted for classification

Category	Features
Statistics	Max, min, mean, median, std. dev., RMS, correlation, skewness, kurtosis, 1st & 3rd quartile, interquartile range
Frequency	Frequency amplitude of 0.25 Hz, 0.5 Hz, 0.75 Hz, 1 Hz, 1.25 Hz, 1.5 Hz, 1.75 Hz, 2 Hz, 2.25 Hz, 2.5 Hz
Statistics of frequency	Skewness and kurtosis of spectrum from frequency features
Time-series	Count below/above mean, first location of min/max, longest strike below/above mean, number of peaks
Periodic subsequence	$p_{min}, p_{max}, \epsilon$, length
Time	Hour of day

assumption does not hold. However, in practice, the difference between the smallest inter-chew distance *min*, and the longest inter-chew distance *max* is a small constant.

4.4 Feature Extraction

We extracted features from all the identified periodic subsequences to classify and validate whether the candidate subsequence identified in the *Segmentation* step (Section 4.3) was truly a chewing sequence. We extracted statistical-based features which are known to be useful in detecting physical activity [6] and eating [55], including maximum, minimum, mean, median, variance, root mean square (RMS), correlation, skewness, kurtosis, 1st and 3rd quartile values, and interquartile range. We plotted the spectrogram of the proximity signal for the dominant chewing frequencies (refer to Figure 10) and observed that the dominant frequencies during chewing occur between 1 and 2 Hz. As a comprehensive measure, we captured the amplitude of the dominant frequencies from 0.25 to 2.5 Hz.

Given a candidate sequence with start and end time $[c_1, c_2]$, the features listed in Table 2 are calculated from two window lengths. The first window is $CW = [c_1 - 2sec, c_2 + 2sec]$, referred to as the *ChewingWindow*. The second window is $BW = [c_1 - 2sec, c_1 + 2sec]$, referred to as the *BiteWindow*. Features from both these windows were concatenated into a single feature vector. CW is useful in capturing information related to the chewing segment, while BW captures bite-related features that occur at the beginning and end of the chewing sequence. Overall, we extracted 257 features for every sequence.

4.5 Classification

We used a gradient boosting classifier based on Friedman's Gradient Boosting Model (GBM), which is an ensemble classifier comprising multiple weak learners (high bias, low variance) [23] to train a model that classifies each subsequence as belonging to a chewing sequence or not. We employed an open source software package, XGBoost [14], for classification. Gradient boosting uses regularized model formalization that controls over-fitting the training data, providing more generalizable performance.

Gradient boosting has several parameters that can be tuned to optimize its performance including general, booster, and learning task parameters. For the general parameters, we used the gbtree model. For the booster parameters, we optimized the learning rate (*eta*), the maximum depth of the a tree (*max_depth*), the minimum loss reduction required to make a further partition on a leaf node of a the tree (*gamma*), the minimum sum of instance weight needed in a child (*min_child_weight*), and the subsample ratio of the training instance (*subsample*). We performed binary classification (chewing vs others) using a softmax objective.

Every positively detected chewing subsequence is then combined to generate a time-point distribution of predicted chewing sequences (at the per-second level). This distribution is a type of confidence weighting to estimate the likelihood of the duration belonging to a chewing sequence. Each second is then converted into a score according to the number of overlapping predicted chews. We then apply the DBSCAN algorithm to cluster the data. The sequences with low weights and sparse positioning are filtered during the clustering step. We then

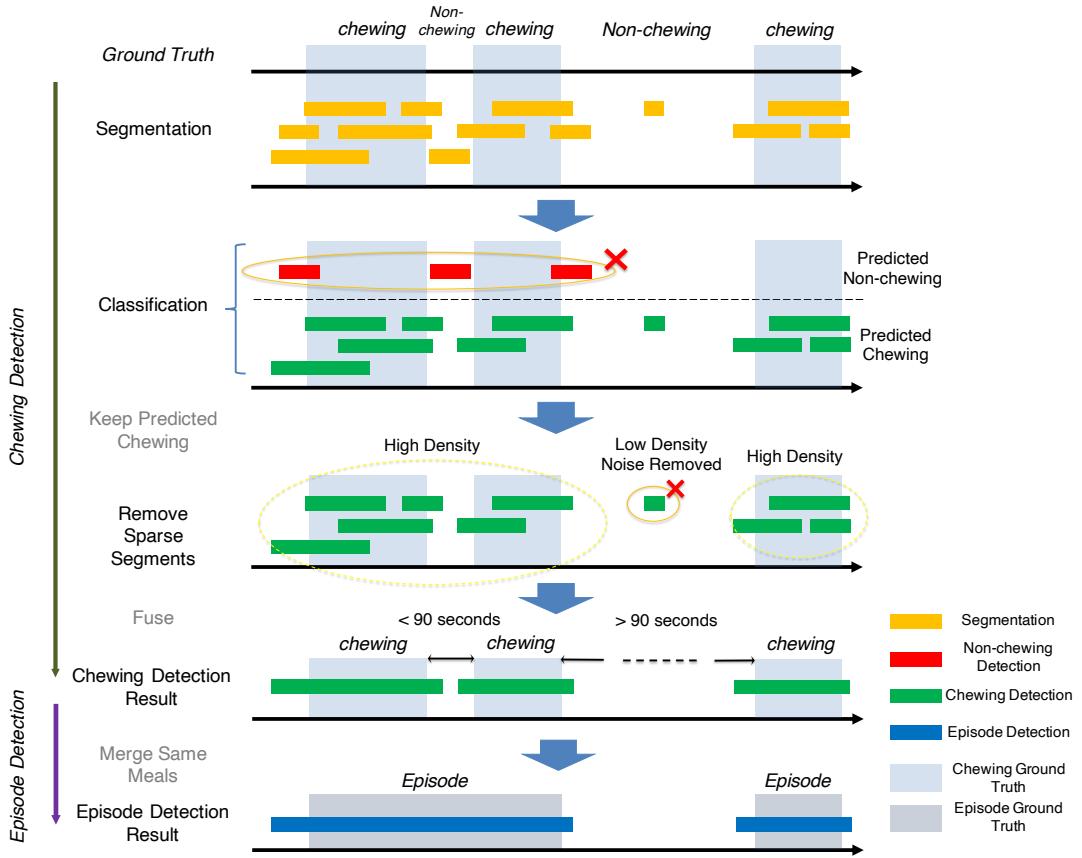


Fig. 11. After periodic subsequences are classified, DBSCAN filters out predicted single chewing subsequences and clusters the remainings into meals. If a ground truth eating episode overlaps with a predicted eating episode, then it is a true positive. If a predicted eating episode has no overlap with a ground truth eating episode, then it is a false positive. perform eating episode evaluation at the coarse-grained event level and fine-grained per-second level. Figure 11 provides an overview of the processing pipeline.

5 STUDY DESIGN & DATA COLLECTION

Demonstrating usefulness of our eating activity monitoring system necessitates that we test the system on a representative population, throughout the day while participants carry out their daily routine. Prior work has demonstrated that systems that are evaluated only in laboratory settings often perform poorly in naturalistic settings. This performance degradation is quite pronounced in eating- and behavior-tracking devices, as the behavior and habits of participants can easily be influenced by the in-lab setting, and the short duration of sessions rarely capture numerous real-life situations.

With this context, we conducted an *Exploratory Study* for optimizing various system parameters. Using the learning outcome from the Exploratory Study, we conducted a *Free-Living Study* to determine the system's performance in a completely free-living condition. Both studies were conducted in naturalistic settings, while the participants performed their everyday activities. We recruited 20 participants (10 for the Exploratory Study and 10 for the Free-Living Study) from an urban Midwestern city in the United States using paper flyers and via

Table 3. Number of hours of video and necklace data in the Exploratory Study and the Free-Living Study. Necklace valid hours exclude data with incorrect timestamp resulting from RTC going out of battery or when the necklace was not worn by the participant. Participants with obesity are highlighted in orange. Days are unique days in the data. Only data captured by both the camera and necklace were used in validation. From the 271.3 hours analyzed in the Exploratory Study and the Free-Living Study, 14.3 hours correspond to eating, while 257 hours are non-eating.

Exploratory Study						Free-Living Study					
Participant	Necklace		Overlap w/ Videos			Participant	Necklace		Overlap w/ Videos		
	Total hours	Ave. Per Day	Total hours	Ave. Per Day	Meal		Total hours	Ave. Per Day	Total hours	Ave. Per Day	Meal
P1	35.7	5.9	27.8	5.6	7	P11	23.9	11.9	17.2	8.6	4
P2	8.7	1.7	3.5	0.9	7	P12	14.6	7.3	15.2	7.6	1
P3	74.6	6.8	31.6	2.9	15	P13	12.4	6.2	12.4	6.2	5
P4	20.3	3.4	10.1	3.3	5	P14	15.1	7.6	12.8	6.4	3
P5	40.4	10.1	33.3	8.3	11	P15	19.2	9.6	17.4	8.7	4
P6	19.3	3.9	4.9	1.6	9	P16	20.2	10.1	16.9	8.5	5
P7	23.0	7.7	14.0	7.0	3	P17	23.2	10.6	15.1	7.6	6
P8	20.1	2.5	5.2	0.7	11	P18	14.9	7.5	11.0	5.5	4
P9	20.0	10.0	0.7	0.4	2	P19	20.2	10.1	13.6	6.8	3
P10	15.0	2.5	3.1	0.8	6	P20	27.0	13.5	21.8	10.9	5
Total	277.1	-	134.2	-	76	Total	193.0	-	137.1	-	40

ResearchMatch. The inclusion criteria were: 18-63 years of age and BMI above 18. The exclusion criteria included anyone unwilling to wear the study devices (due to a history of skin irritations or device sizing limitations) and anyone who did not own a smartphone. None of the participants were members of our research team. Overall, we used 134.2 hours of data from the Exploratory Study to fine-tune our system and 137.1 hours of data from the Free-Living Study to evaluate our system. Specific details about data collection are presented in Table 3. We plan to anonymize and release this dataset for use by clinicians and researchers for evaluating their own devices and approaches to eating detection.

During their first laboratory visit (for both the Exploratory Study and Free-Living Study), we trained the participants about how to wear and charge their devices. After the final day of the study, participants returned the devices and completed a post-study survey. During this visit, participants were given the option to review captured video and remove segments they felt uneasy about sharing.

5.1 Exploratory Study

To determine the eating sensing system’s feasibility, we recruited 10 participants (4 males, 6 females; between 19 and 54 years old) and instructed them to wear the prototype device, including the NGen1 necklace for 2 weeks. Participants were free to wear the device for as many or as few hours as they wanted during this study. However, we instructed them to wear the prototype device during as many meals as possible. Since we are interested in ensuring that our system performs reliably across a varied BMI range, we allocated participants so that 50% of the participants in this study were categorized as obese ($BMI > 30 \text{ kg/m}^2$). The BMI of the 10 participants ranged between 21 and 46 kg/m^2 . Participants were compensated monetarily for their time (smartwatch and \$100; total value of \$300). Overall, we collected 277.1 hours of data during the Exploratory Study. However, after removing

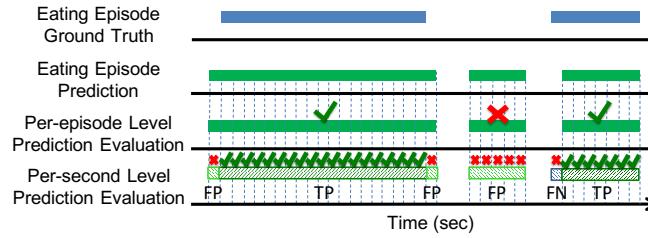


Fig. 12. Evaluation criteria for the eating activity detection at two levels: a commonly used per-episode level evaluation approach and a challenging per-second level evaluation approach. If there is a 50% overlap at the per-episode level, we infer that the episode has been detected correctly.

data that had synchronization issues or raised privacy concerns, we acquired 134.2 hours of usable data from the Exploratory Study for our analysis.

5.2 Free-Living Study

After optimizing reliability and usability of the neck-worn sensor, we designed the newer necklace, NGen2. To test this device, we recruited 10 participants (5 obese, 5 non-obese) to participate in a 2-day Free-Living Study. None of these participants had participated in the previous Exploratory Study. Participants were between 24 and 60 years old, and their BMI ranged from 20.1 to 38.1 kg/m². Table 3 summarizes the device usage for each participant. Unlike the Exploratory Study, for this study we instructed the participants to wear the device during the entire waking day, removing the device only when it was completely discharged or when they had privacy or other concerns. Participants did not delete any data in this study. Overall, after removing data segments with synchronization issues, we extracted 137.1 hours of usable data. We provided \$50 compensation to participants for their time participating in the study.

6 EVALUATION AND RESULTS

As stated in Section 1, the overall goal of *NeckSense* is to detect eating activity, while ensuring that the device has acceptable battery life. We thus evaluated *NeckSense* using the Exploratory Study and Free-Living Study data, while answering the following questions:

- Q1: Can *NeckSense* effectively detect eating activity? How does the system perform as compared with other similar techniques?
- Q2: How do factors such as sensor choice, device position, and classification features affect the detection performance?
- Q3: Can *NeckSense*'s battery support monitoring an individual's eating activities throughout the waking day?

Before answering the questions, we describe the evaluation metric used to evaluate *NeckSense*.

6.1 Evaluation Criteria and Metric

Eating Activity detection: We evaluate the eating activity detection at two levels: (1) the possibility of detecting eating at a *per-second level* and (2) the possibility of detecting the overall eating at a *per-episode level*. Figure 12 pictorially describes the two levels. For each level we compute the *precision*, *recall*, and *F1-score*. A high precision value indicates that the seconds (or episodes) identified by *NeckSense* as eating were actually seconds (or episodes) when eating occurred, whereas a high recall indicates that *NeckSense* could find most of the moments (either per-second level or per-episode level) when a participant performed an eating action. The *F1-score* presents the

Table 4. Comparison of performance (F1-score) of various sensor combinations for eating episode detection. The proximity sensor is used for segmenting the chewing sequences, and various sensor combinations subsequently detect the eating episodes. One-way repeated measures ANOVA determined that the mean F1-score differed significantly between the four sensor combinations.

Sensor(s) used	Exploratory Study	Free-Living Study
Proximity only (ref)	73.4%	66.4%
Proximity + IMU*	81.5%	78.7%
Proximity + ambient light	72.7%	70.3%
All Sensors*	81.6%	77.1%

*Post hoc analyses with Bonferroni correction show statistically significant improvement of Proximity+IMU and All Sensors over Proximity only at the P<.05 level.

harmonic mean between the precision and recall. We evaluated *NeckSense* by performing a leave-one-subject-out cross validation (LOSOCV) and reporting the average performance of *NeckSense*, which is the average of every participant's precision, recall, or F1-score. The output of the GBM classifier, after sparse segment removal, is utilized for the per-second level evaluation, while the output after merging adjacent chewing segments is utilized for the per-episode level evaluation. If there was 50% overlap between the time of the predicted episode and the actual episode, we considered the episode as a true positive episode.

Battery Lifetime: Since the necklace operates in several modes (e.g., time-keeping mode, data collection mode, data logging mode), it is necessary to monitor the power consumption in each mode. Additionally, since our goal is to ensure that the necklace can operate without frequent charging, it is necessary to understand the average battery lifetime. We measure the power consumption of each mode of operation in milliWatts and the average battery lifetime in terms of number of hours for which the device can operate after it is fully charged, until it is completely discharged.

Now that we have established the evaluation criteria and metric, we next evaluate the performance of *NeckSense*.

6.2 Q1 : Eating Activity Detection

Although several previous studies have demonstrated the feasibility of automatically detecting eating in laboratory settings, very few researchers have explored the possibility of detecting eating over multiple days in free-living conditions and evaluating their system at a per-second level. We thus conducted the Exploratory Study to evaluate this possibility. We performed an LOSOCV to generate the classification model, as well as to determine the DBSCAN clustering parameters.

Analyzing the performance of NeckSense in the Exploratory Study: At a per-second level analysis, the system attained an F1-score of 76.2%. At per-episode level analysis, we observed that among the 76 meals that participants consumed in the Exploratory Study, we could correctly detect 63 meals. The system's average precision, recall, and F1-score across the participants during this Exploratory Study were 80.8%, 86.8%, and 81.6%, respectively. These results indicated that it was indeed possible to detect most meals that the individual consumes, even in naturalistic settings.

Importance of each sensor in NeckSense: To understand how each sensor assisted in the detection of meals in the Exploratory Study, we segmented the signal from all the participants' data using the proximity sensor data and the longest subsequence algorithm. For each segment, we used each sensor either independently or in combination with other sensors. Overall we observed that we could identify eating episodes with an F1-score of 73.4% when we used only the proximity sensor, whereas including other sensors with the proximity sensor helped in improving the overall F1-score by more than 8%, demonstrating that the other sensors were useful, as compared with using the proximity sensor in isolation. Table 4 shows the performance of various sensors in

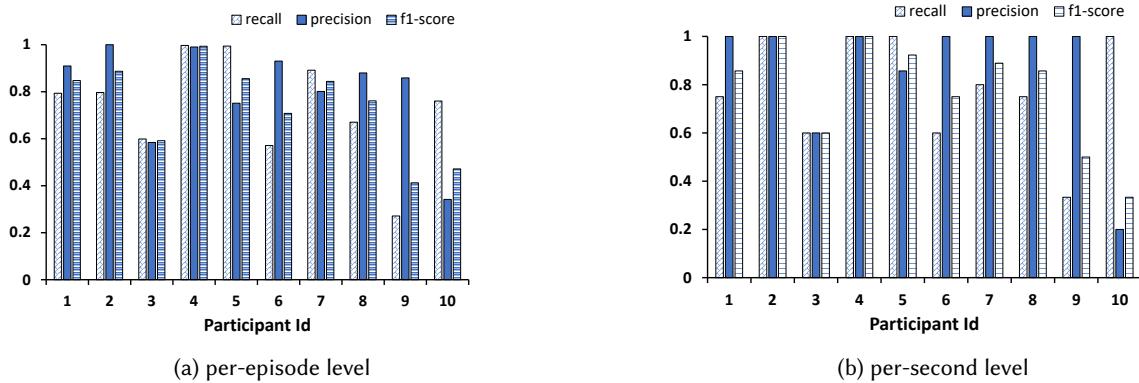


Fig. 13. Eating episode prediction evaluation for the Free-Living Study using LOSOCV method for *per-second level* and *per-episode level* analysis.

detecting the eating episodes in the Exploratory Study and Free-Living Study. To determine significant difference between sensor combinations, we combined participants' results from both the Exploratory and Free-Living study (both studies used the same algorithm and sensor combinations). A one-way repeated measures ANOVA (used to determine whether three or more group means are different when the participants are the same in each group) determined that the mean F1-score differed significantly between the four groups: Proximity only, Proximity+IMU, Proximity+ambient light, and All Sensors ($F[3,57]=8.555$, $P<.0001$). Post hoc tests using a paired sample t-test with Bonferroni correction revealed that Proximity+IMU and All Sensors exhibited significant improvement in mean F1-Score ($P<.05$) compared to Proximity only, while Proximity only was not significantly different from Proximity+ambient light ($P>.05$). Therefore, we conclude that there is statistically significant improvement when the proximity sensor is used either in conjunction with the IMU, or when using all signals, compared to using the proximity sensor only. This indicates that the eating activity generates some specific movement pattern, which the IMU sensor can capture. Orthogonally, it is interesting to note that in both studies, the system performs poorly when only the light and proximity sensor data are used for the classification. However, light sensor data augmented with the IMU data helps improve the system's overall performance, indicating that additional features from the light sensor might contribute toward improving performance. While annotating the dataset we observed that several participants consumed their meal while watching television and sitting in a dark location. Understandably, the light sensor was not useful for these episodes.

Analyzing the performance of NeckSense in the Free-Living Study: Through the Exploratory Study we demonstrated that the necklace could indeed determine eating in a semi-free-living condition. To understand the system's performance in a completely uncontrolled setting, we analyzed the data from the Free-Living Study. Overall, we found that at a per-second level, our system could detect eating with an average F1-score of 73.7% (average precision = 80.5%, average recall = 73.4%), while at a per-episode level, the F1-score was 77.1% (average precision = 86.6%, average recall = 78.3%). This was an improvement of over 10% as compared with using only data from the proximity sensor. Figure 13 presents the per-participant performance of NeckSense. It is interesting to note here that every sensor combination provided a better F1-score as compared with using only the proximity data. Overall, the system could identify 35 of the 40 meals when using either the proximity and IMU data or when using all sensor data. Although our system performed well in semi-free-living settings, the performance degrades while evaluating in an uncontrolled free-living setting. This observation should motivate researchers to evaluate their systems not only in semi-controlled free-living conditions, but also in truly free-living conditions to identify their system's actual performance.

Table 5. Performance of *NeckSense* for different training and testing groups

		Test	
		Obese	Non-obese
Train	Obese	71.21%	75.33%
	Non-obese	66.75%	79.88%

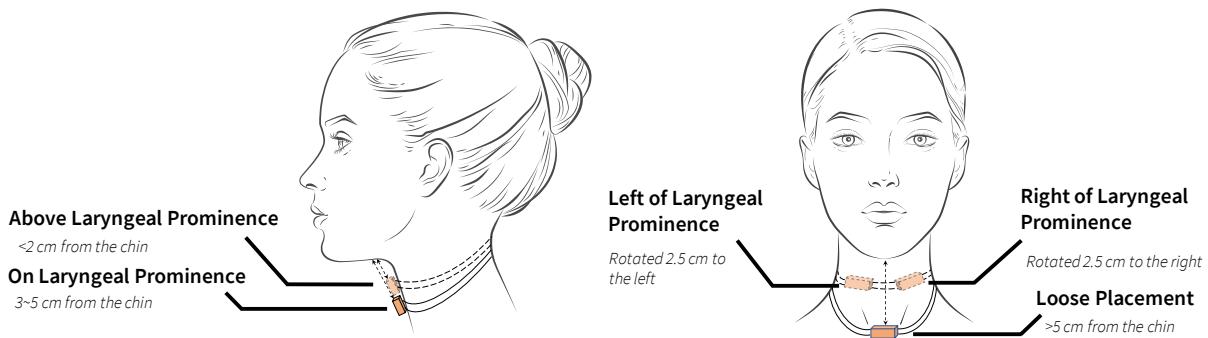


Fig. 14. Five different positions for the necklace on the neck: above the laryngeal prominence, on the laryngeal prominence, loose placement, left of the laryngeal prominence, and right of the laryngeal prominence.

Exploring BMI-based models: To evaluate whether models trained on participants with obesity perform reliably when tested on participants without obesity and vice-versa, we generated two models, one trained with data only from participants without obesity (M_{N-O}) and the other with data from only the participants with obesity (M_O). When we tested M_{N-O} with the data from people with obesity, we observed that the F1-score for detecting eating episodes dropped to 66.75%. This is substantially lower than the generalized model trained on multiple body types and tested in participants with obesity. Table 5 presents the results for different combinations of the training and test sets. The system performs fairly consistently when tested in people without obesity, independent of the BMI of the training set. However, the performance is lower when tested in participants with obesity. Several factors can be attributed to this discrepancy. These included factors such as differences in movement patterns while eating, change in the distance of the proximity sensor from the neck, and difference in posture during the eating activity. In the future, to ensure generalizability, researchers developing automatic dietary monitoring systems should consider testing their wearable system in participants that are recruited from a population with different BMI profiles and refrain from solely testing in a single, homogeneous population to prove efficacy of their system.

6.3 Q2 : Effect of Various Factors

Performance of various sensors: To understand the usefulness of each sensor in determining the chewing sequence and the eating activity, we analyzed the Exploratory Study's data and observed that for the per-episode level evaluation, we could achieve an average precision of 77.2%, recall of 74.0%, and F1-score of 73.4% across participants when we use only the proximity sensor's signal, as compared with an average precision, recall, and F1-score of 80.8%, 86.8%, and 81.6%, respectively, when we use all sensors. This improvement in performance validates

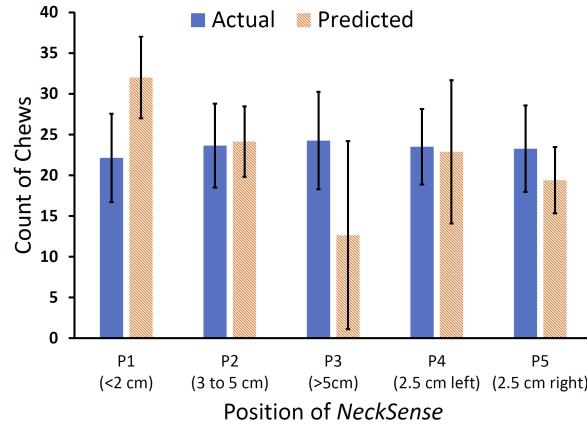


Fig. 15. Performance in detecting the chewing activity when *NeckSense* is placed at different locations from the chin. It shows the actual and predicted average and standard deviation of the number of chews at various on-neck positions of *NeckSense*. the usefulness of employing multi-sensor modalities over using a single sensing modality, proximity, for eating detection.

Feature importance: Since the number of instances for each feature in the gradient boosting tree is proportional to the feature's contribution among all the features, we ranked the features accordingly. In Section 4.4 we described the features were extracted from either the bite window (*BW*) or the chewing window (*CW*). We observed that the top five features that aided in classification, as selected by the gradient boosting algorithm, were: (1) Frequency: FFT 2.5 Hz of energy signal (*BW*), (2) Time-series: first location of minimum of energy signal (*CW*), (3) Time-series: first location of maximum of energy signal (*CW*), (4) Frequency: FFT 0.5 Hz of ambient signal (*CW*), and (5) Time-series: count above mean of ambient signal (*CW*). All top five features were extracted from both the ambient light and energy signals. Two of the five features were FFT-based features, while three were time-series-based features. Four features were extracted from the *CW*, and one was from the *BW*.

Effect of various necklace positions: The necklace form factor, worn loosely around the neck, lets participants adjust and move the necklace closer or farther from their mouths. To investigate the effect of positioning of the necklace on the body, we conducted a small in-laboratory investigatory study where two participants (one male, one female) wore the necklace in multiple positions and orientations and consumed a meal. We present the five different positions in Figure 14. Overall, each participant wore the necklace at five positions: (i) above the laryngeal prominence, where the necklace was less than 2 cm from the chin, (ii) on the laryngeal prominence, where the necklace was between 3 and 5 cm from the chin, (iii) loose placement, where the necklace was more than 5 cm from the chin, (iv) left of the laryngeal prominence, where the necklace was rotated 2.5 cm to the left of the laryngeal prominence, and (v) right of the laryngeal prominence, where the necklace was rotated 2.5 cm to the right of the laryngeal prominence. Figure 15 presents the result for the average (with standard deviation) of actual and predicted number of chews for every mouthful during these episodes. The periodic subsequence algorithm could detect chews with an average difference of less than 4 chews when the necklace is placed on the laryngeal prominence, and the prediction was not affected when the necklace rotated a couple of centimeters toward the left or right of the laryngeal prominence. Our algorithm is designed to detect a chewing sequence even when some of the periodic peaks (or chews) are detected during a chewing event. Thus, even for loose or high placement, we can still detect the occurrence of a chewing sequence. This indicates that the necklace's eating activity detection performance is robust to various on-neck positions. Figure 16 presents a signal trace

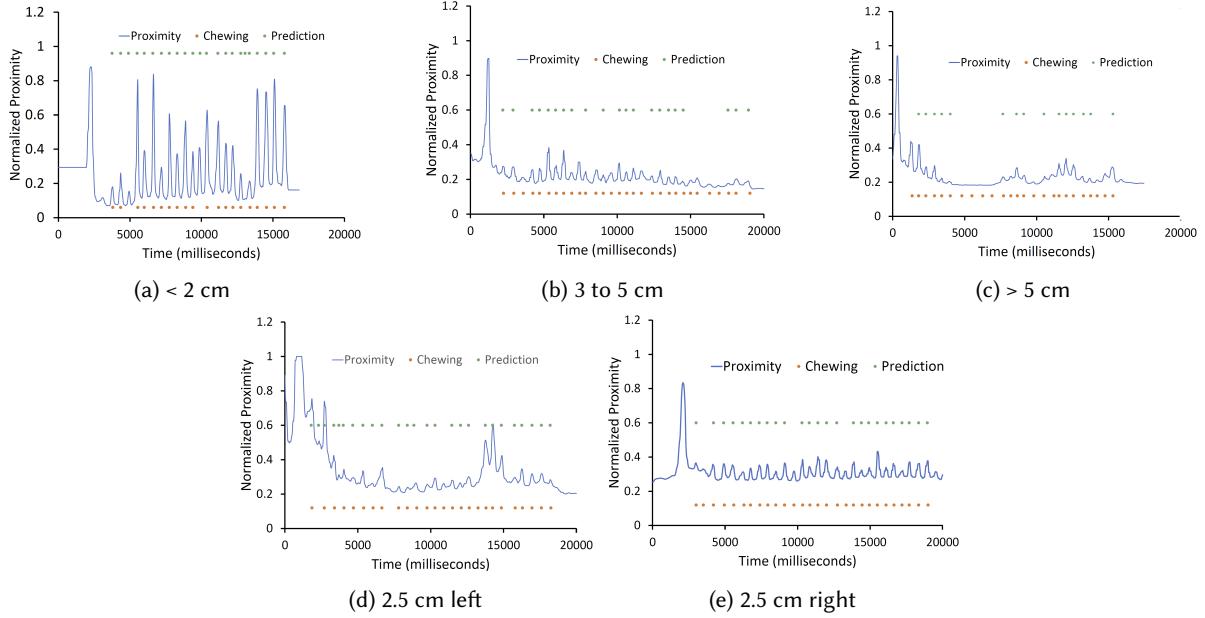


Fig. 16. Sample proximity sensor signal captured from different positions. The orange dot indicates the time frame when a chewing action occurred, while the green dot indicates whether the time point was identified as a chew by *NeckSense*. collected from various on-neck positions during the study. We have also indicated the time points when the chewing action occurred, along with our periodic subsequence algorithm's output. Most chewing gestures were captured by the periodic subsequence algorithm.

6.4 Q3 : Battery Lifetime

Now that we established that *NeckSense* can indeed detect eating activities in a naturalistic setting, we next analyse whether the device can collect sensor data continuously during an individual's waking hours. The necklace is powered by a 350-mAh battery to ensure a relatively small size. To conserve power we implemented several power saving schemes. First, we selected the Nordic NRF52832 SiP, which includes a low-powered microcontroller and a BLE interface. Overall, this choice helped in reducing power consumption and size. Second, we identified that the SD card writes consumed substantial power. We thus implemented a batch-and-write approach for data logging. Third, we designed the super-low-power time keeping mode with an extra-small battery so we maintain valid system time even if the main battery is off. The necklace can operate in four modes: timekeeping mode (only running the RTC), time syncing mode (communicating over BLE to synchronize time), data collecting mode (recording sensor reading), and data logging mode (writing data to SD card). We measured the average power for each mode using the STM32 Nucleo expansion board [52] for power measurement and tested the battery life time for a single charge. The power consumption for each mode, as shown in Figure 17, was 3.8 μ W, 36.8 mW, 56.6 mW, and 64.9 mW for timekeeping, time syncing, data collecting, and data logging mode, respectively. Overall, in the Exploratory Study the average battery life was 13 hours, and it improved to 15.8 hours in the Free-Living Study, which is sufficient to record data for most of the day and cover all meals that occur during a waking day. This lifetime allows for simplified deployments and little, if any, study coordinator oversight. The participants can charge the device once a day before sleeping.

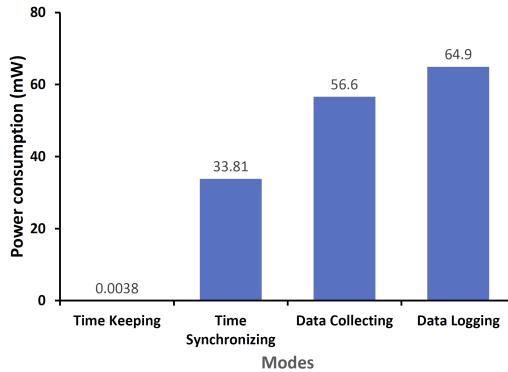


Fig. 17. Power consumption for each active mode of *NeckSense*.

6.5 Comfort and Usability

At the end of the Exploratory Study, we asked participants to indicate their feedback about *NeckSense*. Of the participants who completed the survey, 90% (9/10) claimed the device did not change how they performed their daily activities, how they socialized, or how they felt at home; 80% (8/10) claimed the device did not change how they ate or felt in public places. Two participants suggested that the velcro in the back of the necklace bothered them, suggesting a clasp or button may be more comfortable, and two other participants suggested that the device should be miniaturized to appear less obvious.

All the participants responded with agreement to the question, "I am willing to wear the necklace for 2 weeks continuously with compensation (if I am paid in cash)". The compensation amount preference varied largely among participants. In general, a greater number of participants reported willingness to wear with increasing monetary value. Interestingly, on average, people with obesity were willing to wear the device for less monetary compensation, compared with people with a normal BMI, suggesting we may have improved adherence in future studies in a population group that feels they will benefit most from such a device or system. When we asked participants to select all the factors that motivated them to wear the necklace for 2 weeks; 9 participants selected monetary compensation, 5 selected health, and 8 selected contribution to research.

6.6 Summary of Results

From these results we observe that: (i) It is indeed possible to detect the eating activity in a *completely naturalistic setting* at an *extremely challenging per-second granularity* with a F1-score of 73.7% (76.2% in semi-free living conditions), thus presenting the possibility of performing real-time interventions. This is an improvement over similar existing techniques. (ii) It is possible to determine the eating episodes, even as individuals perform their everyday activities, with an F1-score of 77.1% (81.6% in semi-free living conditions). (iii) A combination of multi-sensor eating detection outperforms the eating detection by a single sensing modality by 8% to 12%. (iv) It is necessary to build eating activity detection models while using data from various demographic groups because we identified that models trained on participants without obesity fail to accurately detect eating episodes for participants with obesity. (v) *NeckSense* is robust to the on-neck position and can capture chewing sequences during an eating episode. (vi) Once *NeckSense*'s battery is completely charged, it can continuously monitor a participant's eating activity in free-living conditions for over 15 hours, thus making it possible to monitor individuals through an entire waking day. These results demonstrate that the necklace can be a promising avenue

for identifying eating-related activities through an individual's entire waking day and for deploying and testing eating-related interventions in real time.

Overall, our results are promising. However, given that both the studies were conducted using a prototype device and in a truly free-living population setting, challenging eating postures and scenarios confounded some of our chewing and eating episode detection. For example, lying down while eating confounded our ability to capture the lean forward motion during eating and also chewing from the proximity sensor. Similarly, eating in total darkness confounded the classification model because the ambient light sensor gave readings far below the normal range, while eating during exercising resulted in several false chewing detection, although our aggregate filter successfully filtered out these isolated chewing sequences. Meals that included no chewing also presented a challenge; this was seen in one of our participants who had 50% of the labeled eating episodes as ice cream, yogurt, or other non-chewing foods. This issue has also been reported by other researchers exploring chewing-based activity detection [13]. In the future we will explore techniques to overcome these challenges.

7 CONCLUSION

In this paper we present the design, implementation, and evaluation of a necklace suite for detection and validation of chewing sequences and eating episodes in free-living conditions. We utilize sensor data from a proximity sensor, an ambient light sensor, and an IMU sensor to detect chewing activity. We performed two free-living studies using the necklace. In our exploratory study, where selective meals were captured and non-continuous wear time was observed, we were able to outperform other methods with an F1-score of 76.2% at the per-second level and 81.6% at the per-episode level. Overall, our system achieved an F1-score of 73.7% in detecting the chewing sequences at an extremely challenging per-second-level granularity in a truly free-living study. Additionally, our system could detect eating episodes with an F1-score of 77.1%. This is an improvement of over 10% as compared with using only the neck-worn proximity sensor for eating detection. Moreover, our work showed that models trained on participants without obesity underperformed when tested on participants with obesity. To ensure generalizability, future automatic dietary monitoring researchers should ensure BMI diversity in their training set while developing the models. In terms of hours of battery life, the necklace could monitor the participants for 15.8 hours during the waking day, making it possible to monitor an individual's eating activity occurring during an entire day.

Our dataset is unique in that it comprises a 2-week and 2-day study of eating detection collected from 20 participants in free-living conditions. These data are accompanied by videos that were validated and professionally labeled, marking the ground truth. This dataset includes many unique eating events and eating situations, including eating in cars, slouching, talking and eating, and eating a variety of soft foods. We hope this dataset, the associated code base, and the hardware design will provide useful information from which researchers can glean insights from the data and inform future studies. Understanding the effects of chewing speed and duration on overeating can help interventionists design and test improved treatments for monitoring eating behaviors. Our user survey showed that we may have improved adherence in future studies in a population that feels they will benefit most from such a system. *NeckSense* is the next step in enabling behaviorists to imagine, design, deploy, and test for effective interventions during and immediately after an eating episode.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Institute of Diabetes and Digestive and Kidney Diseases under award number K25DK113242 (NIDDK). We would also like to acknowledge support by the National Science Foundation under award number CNS1915847. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Institutes of Health or the National Science Foundation.

REFERENCES

- [1] Rawan Alharbi, Angela Pfammatter, Bonnie Spring, and Nabil Alshurafa. 2017. WillSense: Adherence Barriers for Passive Sensing Systems That Track Eating Behavior. In *Proceedings of the CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA)*. <https://doi.org/10.1145/3027063.3053271>
- [2] Rawan Alharbi, Tammy Stump, Nilofar Vafaie, Angela Pfammatter, Bonnie Spring, and Nabil Alshurafa. 2018. I Can't Be Myself: Effects of Wearable Cameras on the Capture of Authentic Behavior in the Wild. *ACM Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)* 2, 3, Article 90 (2018), 40 pages. <https://doi.org/10.1145/3264900>
- [3] Rawan Alharbi, Mariam Tolba, Lucia C Petito, Josiah Hester, and Nabil Alshurafa. 2019. To Mask or Not to Mask? Balancing Privacy with Visual Confirmation Utility in Activity-Oriented Wearable Cameras. *ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)* 3, 3, Article 72 (2019), 29 pages. <https://doi.org/10.1145/3351230>
- [4] Rawan Alharbi, Nilofar Vafaie, Kitty Liu, Kevin Moran, Gwendolyn Ledford, Angela Pfammatter, Bonnie Spring, and Nabil Alshurafa. 2017. Investigating barriers and facilitators to wearable adherence in fine-grained eating detection. In *IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. 407–412. <https://doi.org/10.1109/PERCOMW.2017.7917597>
- [5] Nabil Alshurafa, Annie Wen Lin, Fengqing Zhu, Roozbeh Ghaffari, Josiah Hester, Edward Delp, John Rogers, and Bonnie Spring. 2019. Counting Bites With Bits: Expert Workshop Addressing Calorie and Macronutrient Intake Monitoring. *Journal of Medical Internet Research (JMIR)* 21, 12 (2019). <https://doi.org/10.2196/14904>
- [6] Nabil Alshurafa, Wenyao Xu, Jason J. Liu, Ming-Chun Huang, Bobak Mortazavi, Majid Sarrafzadeh, and Christian Roberts. 2013. Robust human intensity-varying activity recognition using Stochastic Approximation in wearable sensors. In *IEEE International Conference on Body Sensor Networks*. <https://doi.org/10.1109/BSN.2013.6575515>
- [7] Oliver Amft and Gerhard Tröster. 2006. Methods for detection and classification of normal swallowing from muscle activation and sound. In *IEEE Pervasive Health Conference and Workshops*. 1–10. <https://doi.org/10.1109/PCTHEALTH.2006.361624>
- [8] Oliver Amft and Gerhard Tröster. 2009. On-body sensing solutions for automatic dietary monitoring. *IEEE Pervasive Computing* 8, 2 (2009), 62–70. <https://doi.org/10.1109/MPRV.2009.32>
- [9] Abdelkareem Bedri, Richard Li, Malcolm Haynes, Raj Prateek Kosaraju, Ishaan Grover, Temiloluwa Prioleau, Min Yan Beh, Mayank Goel, Thad Starner, and Gregory Abowd. 2017. EarBit: Using Wearable Sensors to Detect Eating Episodes in Unconstrained Environments. *ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)* 1, 3, Article 37 (2017). <https://doi.org/10.1145/3130902>
- [10] Abdelkareem Bedri, Apoorva Verlekar, Edison Thomaz, Valerie Avva, and Thad Starner. 2015. Detecting mastication: A wearable approach. In *ACM International Conference on Multimodal Interaction*. 247–250. <https://doi.org/10.1145/2818346.2820767>
- [11] Brooke Bell, Ridwan Alam, Nabil Alshurafa, Edison Thomaz, Abu Mondol, Kayla de la Haye, John Stankovic, John Lach, and Donna Spruijt-Metz. 2020. Automatic, wearable-based, in-field eating detection approaches for public health research: a scoping review. *npj Digital Medicine* 3, Article 38 (2020). <https://doi.org/10.1038/s41746-020-0246-2>
- [12] France Bellisle. 2004. Impact of the daily meal pattern on energy balance. *Scandinavian Journal of Nutrition* 48, 3 (2004), 114–118. <https://doi.org/10.1080/11026480410000454>
- [13] Shengjie Bi, Tao Wang, Nicole Tobias, Josephine Nordrum, Shang Wang, George Halvorsen, Sougata Sen, Ronald Peterson, Kofi Odame, Kelly Caine, Ryan Halter, Jacob Sorber, and David Kotz. 2018. Auracle: Detecting Eating Episodes with an Ear-Mounted Sensor. *ACM Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)* 2, 3, Article 92 (2018). <https://doi.org/10.1145/3264902>
- [14] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. In *ACM International Conference on Knowledge Discovery and Data Mining (KDD)*. 785–794. <https://doi.org/10.1145/2939672.2939785>
- [15] Keum San Chun, Sarnab Bhattacharya, and Edison Thomaz. 2018. Detecting Eating Episodes by Tracking Jawbone Movements with a Non-Contact Wearable Sensor. *ACM Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)* 2, 1, Article 4 (2018), 21 pages. <https://doi.org/10.1145/3191736>
- [16] CM Doak, TLS Visscher, CM Renders, and JC Seidell. 2006. The prevention of overweight and obesity in children and adolescents: a review of interventions and programmes. *Obesity Reviews* 7, 1 (2006), 111–136. <https://doi.org/10.1111/j.1467-789X.2006.00234.x>
- [17] Yujie Dong, Adam Hoover, Jenna Scisco, and Eric Muth. 2012. A new method for measuring meal intake in humans via automated wrist motion tracking. *Applied Psychophysiology and Biofeedback* 37, 3 (2012). <https://doi.org/10.1007/s10484-012-9194-1>
- [18] Yujie Dong, Jenna Scisco, Wilson Mike, Muth Eric, and Hoover Adam. 2014. Detecting periods of eating during free-living by tracking wrist motion. *IEEE Journal of Biomedical and Health Informatics (JBHI)* 18, 4 (2014). <https://doi.org/10.1109/JBHI.2013.2282471>
- [19] Muhammad Farooq, Megan A McCrory, and Edward Sazonov. 2017. Reduction of energy intake using just-in-time feedback from a wearable sensor system. *Obesity* 25, 4 (2017), 676–681. <https://doi.org/10.1002/oby.21788>
- [20] Muhammad Farooq and Edward Sazonov. 2016. Detection of chewing from piezoelectric film sensor signals using ensemble classifiers. In *Annual International Conference of the Engineering in Medicine and Biology Society (EMBC)*. <https://doi.org/10.1109/EMBC.2016.7591833>
- [21] Muhammad Farooq and Edward Sazonov. 2016. Segmentation and characterization of chewing bouts by monitoring temporalis muscle using smart glasses with piezoelectric sensor. *IEEE Journal of Biomedical and Health Informatics (JBHI)* 21, 6 (2016), 1495–1503. <https://doi.org/10.1109/JBHI.2016.2640142>

- [22] Juan M Fontana, Muhammad Farooq, and Edward Sazonov. 2014. Automatic ingestion monitor: A novel wearable device for monitoring of ingestive behavior. *IEEE Transactions on Biomedical Engineering (TBE)* 61, 6 (2014), 1772–1779. <https://doi.org/10.1109/TBME.2014.2306773>
- [23] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* (2001), 1189–1232. <http://www.jstor.org/stable/2699986>
- [24] Beat Gfeller. 2011. Finding longest approximate periodic patterns. In *Workshop on Algorithms and Data Structures (WADS)*. Springer, 463–474. https://doi.org/10.1007/978-3-642-22300-6_39
- [25] Berit Lilenthal Heitmann and Lauren Lissner. 1995. Dietary underreporting by obese individuals—is it specific or non-specific? *British Medical Journal (BMJ)* 311, 7011 (1995), 986–989. <https://doi.org/10.1136/bmj.311.7011.986>
- [26] Steve Hodges, Lyndsay Williams, Emma Berry, Shahram Izadi, James Srinivasan, Alex Butler, Gavin Smyth, Narinder Kapur, and Ken Wood. 2006. SenseCam: A retrospective memory aid. In *International Conference on Ubiquitous Computing (UbiComp)*. Springer, 177–193. https://doi.org/10.1007/11853565_11
- [27] AK Illner, H Freisling, H Boeing, Inge Huybrechts, SP Crispim, and N Slimani. 2012. Review and evaluation of innovative technologies for measuring diet in nutritional epidemiology. *International Journal of Epidemiology (IJE)* 41, 4 (2012), 1187–1203. <https://doi.org/10.1093/ije/dys105>
- [28] Luis G James, Juan Calderon, Juan Lopez, and Andrew Raji. 2015. Trends in mobile cyber-physical systems for health just-in time interventions. In *IEEE SoutheastCon*. 1–6. <https://doi.org/10.1109/SECON.2015.7132887>
- [29] Adrienne S Juarascio, Megan N Parker, Madeline A Lagace, and Kathryn M Godfrey. 2018. Just-in-time adaptive interventions: A novel approach for enhancing skill utilization and acquisition in cognitive behavioral therapy for eating disorders. *International Journal of Eating Disorders* 51, 8 (2018), 826–830. <https://doi.org/10.1002/eat.22924>
- [30] Haik Kalantarian, Nabil Alshurafa, Tuan Le, and Majid Sarrafzadeh. 2015. Monitoring eating habits using a piezoelectric sensor-based necklace. *Computers in Biology and Medicine* 58 (2015), 46–55. <https://doi.org/10.1016/j.combiom.2015.01.005>
- [31] Haik Kalantarian, Nabil Alshurafa, and Majid Sarrafzadeh. 2014. A wearable nutrition monitoring system. In *IEEE International Conference on Wearable and Implantable Body Sensor Networks (BSN)*. 75–80. <https://doi.org/10.1109/BSN.2014.26>
- [32] Rebecca M Leech, Anthony Worsley, Anna Timperio, and Sarah A McNaughton. 2015. Characterizing eating patterns: a comparison of eating occasion definitions. *The American Journal of Clinical Nutrition (AJCN)* 102, 5 (2015). <https://doi.org/10.3945/ajcn.115.114660>
- [33] Jindong Liu, Edward Johns, Louis Atallah, Claire Pettitt, Benny Lo, Gary Frost, and Guang-Zhong Yang. 2012. An intelligent food-intake monitoring system using wearable sensors. In *International Conference on Wearable and Implantable Body Sensor Networks (BSN)*. <https://doi.org/10.1109/BSN.2012.11>
- [34] BessH Marcus, Neville Owen, LeighAnnH Forsyth, NickA Cavill, and Fred Fridinger. 1998. Physical activity interventions using mass media, print media, and information technology. *American Journal of Preventive Medicine (AJPM)* 15, 4 (1998), 362–378. [https://doi.org/10.1016/S0749-3797\(98\)00079-8](https://doi.org/10.1016/S0749-3797(98)00079-8)
- [35] Mark Mirtchouk, Drew Lustig, Alexandra Smith, Ivan Ching, Min Zheng, and Samantha Kleinberg. 2017. Recognizing Eating from Body-Worn Sensors: Combining Free-living and Laboratory Data. *ACM Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)* 1, 3, Article 85 (2017), 20 pages. <https://doi.org/10.1145/3131894>
- [36] Inbal Nahum-Shani, Eric B Hekler, and Donna Spruijt-Metz. 2015. Building health behavior models to guide the development of just-in-time adaptive interventions: A pragmatic framework. *Health Psychology* 34, S (2015), 1209. <https://doi.org/10.1037/he0000306>
- [37] Inbal Nahum-Shani, Shawna N Smith, Ambuj Tewari, Katie Witkiewitz, Linda M Collins, Bonnie Spring, and Susan Murphy. 2014. Just in time adaptive interventions (jitaits): An organizing framework for ongoing health behavior support. *Methodology Center Technical Report* 2014 (2014), 14–126. <https://doi.org/10.1007/s12160-016-9830-8>
- [38] Linda Neuhauser and Gary L Kreps. 2003. Rethinking communication in the e-health era. *Journal of Health Psychology* 8, 1 (2003), 7–23. <https://doi.org/10.1177/1359105303008001426>
- [39] Linda Neuhauser and Gary L Kreps. 2010. eHealth communication and behavior change: promise and performance. *Social Semiotics* 20, 1 (2010), 9–27. <https://doi.org/10.1080/10350330903438386>
- [40] Jun Nishimura and Tadahiro Kuroda. 2008. Eating habits monitoring using wireless wearable in-ear microphone. In *IEEE International Symposium on Wireless Pervasive Computing (ISWPC)*. <https://doi.org/10.1109/ISWPC.2008.4556181>
- [41] Gillian O'Loughlin, Sarah Jane Cullen, Adrian McGoldrick, Siobhan O'Connor, Richard Blain, Shane O'Malley, and Giles D Warrington. 2013. Using a wearable camera to increase the accuracy of dietary analysis. *American Journal of Preventive Medicine* 44, 3 (2013), 297–301. <https://doi.org/10.1016/j.amepre.2012.11.007>
- [42] Vasileios Papapanagiotou, Christos Diou, Lingchuan Zhou, Janet van den Boer, Monica Mars, and Anastasios Delopoulos. 2017. A Novel Chewing Detection System Based on PPG, Audio, and Accelerometry. *IEEE Journal of Biomedical and Health Informatics (JBHI)* 21, 3 (2017), 607–618. <https://doi.org/10.1109/JBHI.2016.2625271>
- [43] Sarah E Pember and Adam P Knowlden. 2017. Dietary change interventions for undergraduate populations: Systematic review and recommendations. *American Journal of Health Education* 48, 1 (2017), 48–57. <https://doi.org/10.1080/19325037.2016.1250018>
- [44] J.M.C. Po, Jules Kieser, Luigi M Gallo, A.J. Tésenyi, P Herbison, and Mauro Farella. 2011. Time-frequency analysis of chewing activity in the natural environment. *Journal of Dental Research* 90, 10 (2011), 1206–1210. <https://doi.org/10.1177/0022034511416669>

- [45] Tauhidur Rahman, Alexander Travis Adams, Mi Zhang, Erin Cherry, Bobby Zhou, Huaishu Peng, and Tanzeem Choudhury. 2014. BodyBeat: a mobile system for sensing non-speech body sounds.. In *Annual International Conference on Mobile Systems, Applications, and Services (MobiSys)*. <https://doi.org/10.1145/2594368.2594386>
- [46] Sasank Reddy, Andrew Parker, Josh Hyman, Jeff Burke, Deborah Estrin, and Mark Hansen. 2007. Image Browsing, Processing, and Clustering for Participatory Sensing: Lessons from a DietSense Prototype. In *Workshop on Embedded Networked Sensors (EmNets)*. 5. <https://doi.org/10.1145/1278972.1278975>
- [47] Polly Ryan and Diane Ruth Lauver. 2002. The efficacy of tailored interventions. *Journal of Nursing Scholarship* 34, 4 (2002), 331–337. <https://doi.org/10.1111/j.1547-5069.2002.00331.x>
- [48] Jenna Scisco, Eric Muth, and Adam Hoover. 2014. Examining the utility of a bite-count-based measure of eating activity in free-living human beings. *Journal of the Academy of Nutrition and Dietetics (JAND)* 114, 3 (2014), 464–469. <https://doi.org/10.1016/j.jand.2013.09.017>
- [49] Sougata Sen, Vigneshwaran Subbaraju, Archan Misra, Rajesh Balan, and Youngki Lee. 2018. Annapurna: Building a Real-World Smartwatch-Based Automated Food Journal. In *IEEE International Symposium on “A World of Wireless, Mobile and Multimedia Networks” (WoWMoM)*. <https://doi.org/10.1109/WoWMoM.2018.8449755>
- [50] Sougata Sen, Vigneshwaran Subbaraju, Archan Misra, Rajesh Krishna Balan, and Youngki Lee. 2015. The case for smartwatch-based diet monitoring. In *IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom workshops)*. <https://doi.org/10.1109/PERCOMW.2015.7134103>
- [51] Donna Spruijt-Metz and Wendy Nilsen. 2014. Dynamic models of behavior for just-in-time adaptive interventions. *IEEE Pervasive Computing* 13, 3 (2014), 13–17. <https://doi.org/10.1109/MPRV.2014.46>
- [52] STMicroelectronics. 2019. *STM32 Nucleo expansion board*. Date accessed: 11/11/2019.
- [53] Akio Tada and Hiroko Miura. 2018. Association of mastication and factors affecting masticatory function with obesity in adults: a systematic review. *BMC oral health* 18, 1 (2018), 76. <https://doi.org/10.1186/s12903-018-0525-3>
- [54] Jason Tang, Charles Abraham, Elena Stamp, and Colin Greaves. 2015. How can weight-loss app designers' best engage and support users? A qualitative investigation. *British Journal of Health Psychology (BJHP)* 20, 1 (2015), 151–171. <https://doi.org/10.1111/bjhp.12114>
- [55] Edison Thomaz, Irfan Essa, and Gregory D. Abowd. 2015. A Practical Approach for Recognizing Eating Moments with Wrist-mounted Inertial Sensing. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. <https://doi.org/10.1145/2750858.2807545>
- [56] Edison Thomaz, Aman Parnami, Irfan Essa, and Gregory D. Abowd. 2013. Feasibility of Identifying Eating Moments from First-Person Images Leveraging Human Computation. In *Proceedings of International SenseCam & Pervasive Imaging Conference (SenseCam)*. <https://doi.org/10.1145/2526667.2526672>
- [57] Edison Thomaz, Cheng Zhang, Irfan Essa, and Gregory D. Abowd. 2015. Inferring Meal Eating Activities in Real World Settings from Ambient Sounds: A Feasibility Study. In *International Conference on Intelligent User Interfaces (IUI)*. <https://doi.org/10.1145/2678025.2701405>
- [58] James T Tufano and Bryant T Karras. 2005. Mobile eHealth interventions for obesity: a timely opportunity to leverage convergence trends. *Journal of Medical Internet Research (JMIR)* 7, 5 (2005). <https://doi.org/10.2196/jmir.7.5.e58>
- [59] Koji Yatani and Khai N. Truong. 2012. BodyScope: A Wearable Acoustic Sensor for Activity Recognition. In *ACM Conference on Ubiquitous Computing (UbiComp)*. <https://doi.org/10.1145/2370216.2370269>
- [60] Rui Zhang and Oliver Amft. 2017. Monitoring chewing and eating in free-living using smart eyeglasses. *IEEE Journal of Biomedical and Health Informatics (JBHI)* (2017). <https://doi.org/10.1109/JBHI.2017.2698523>
- [61] Shibo Zhang, Rawan Alharbi, Matthew Nicholson, and Nabil Alshurafa. 2017. When Generalized Eating Detection Machine Learning Models Fail in the Field. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing and ACM International Symposium on Wearable Computers (UbiComp)*. <https://doi.org/10.1145/3123024.3124409>