



MunchSonic: Tracking Fine-grained Dietary Actions through Active Acoustic Sensing on Eyeglasses

Saif Mahmud
Cornell University
Ithaca, NY, USA
sm2446@cornell.edu

Devansh Agarwal
Cornell University
Ithaca, NY, USA
da398@cornell.edu

Ashwin Ajit
Cornell University
Ithaca, NY, USA
aa794@cornell.edu

Qikang Liang
Cornell University
Ithaca, NY, USA
ql75@cornell.edu

Thalia Viranda
Cornell Tech
New York, NY, USA
tv74@cornell.edu

Francois Guimbretiere
Cornell University
Ithaca, NY, USA
francois@cs.cornell.edu

Cheng Zhang
Cornell University
Ithaca, NY, USA
chengzhang@cornell.edu

Abstract

We introduce MunchSonic, an AI-powered active acoustic sensing system integrated into eyeglasses to track fine-grained dietary actions. MunchSonic emits inaudible ultrasonic waves from the eyeglass frame, with the reflected signals capturing detailed positions and movements of body parts, including the mouth, jaw, arms, and hands involved in eating. These signals are processed by a deep learning pipeline to classify six actions: hand-to-mouth movements for food intake, chewing, drinking, talking, face-hand touching, and other activities (null). In an unconstrained study with 12 participants, MunchSonic achieved a 93.5% macro F1-score in a user-independent evaluation with a 2-second resolution in tracking these actions, also demonstrating its effectiveness in tracking eating episodes and food intake frequency within those episodes.

CCS Concepts

• **Human-centered computing** → **Ubiquitous and mobile computing systems and tools.**

Keywords

Eating Detection; Acoustic Sensing; Activity Recognition

ACM Reference Format:

Saif Mahmud, Devansh Agarwal, Ashwin Ajit, Qikang Liang, Thalia Viranda, Francois Guimbretiere, and Cheng Zhang. 2024. MunchSonic: Tracking Fine-grained Dietary Actions through Active Acoustic Sensing on Eyeglasses. In *Proceedings of the 2024 ACM International Symposium on Wearable Computers (ISWC '24)*, October 5–9, 2024, Melbourne, VIC, Australia. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3675095.3676619>

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.

ISWC '24, October 5–9, 2024, Melbourne, VIC, Australia

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1059-9/24/10

<https://doi.org/10.1145/3675095.3676619>

1 Introduction

Tracking dietary activity is crucial for measuring nutrition levels and managing chronic diseases. However, the lack of reliable automated dietary monitoring systems forces most people to rely on manual input-based applications, which are prone to self-reporting errors. To address this, researchers have developed wearable dietary activity recognition systems. Despite advancements in sensing modalities and machine learning algorithms to process the data, tracking fine-grained dietary actions such as food intake, chewing, and drinking, while distinguishing them from similar body movements like talking, or face touching, as well as from non-eating activities, remains a significant challenge for the wearable computing research community.

Dietary activities such as eating and drinking require coordinated hand, jaw, and mouth movements. Detecting fine-grained dietary actions necessitates the simultaneous tracking of both hand and mouth movements. Current state-of-the-art dietary monitoring systems, whether on eyeglasses [4, 42, 53], earables [10], or smartwatches [75], require sensing modalities at multiple body locations to achieve this, which can be costly, power-inefficient, and cumbersome for many users.

Recently, the active acoustic sensing system ActSonic [40] showed promising efficacy in tracking everyday activities, including identifying eating and drinking episodes, using a single point of instrumentation on eyeglasses. However, this system has not yet been able to predict fine-grained dietary actions, such as chewing or hand-to-mouth intake gestures, which are related to important biomarkers for longitudinal health monitoring [55], such as chewing patterns, drinking habits, or snacking behaviors. In this paper, we aim to address the following research question:

- *Can an active acoustic sensing system instrumented on smart glasses be used to further distinguish fine-grained dietary actions in unconstrained environments beyond only identifying eating moments?*

To address this challenge, we developed MunchSonic, a fine-grained dietary action recognition system in the form factor of

commodity eyeglasses using active acoustic sensing. Similar to ActSonic [40], our system employs two pairs of microphones and speakers attached to the hinges of the eyeglasses. MunchSonic transmits inaudible ultrasonic waves that reflect off various body parts, generating a two-dimensional echo profile, which serves as a range spectrogram of the scanned body parts. This profile is then analyzed by a lightweight deep learning framework to infer fine-grained dietary actions.

MunchSonic is based on the observation that fine-grained dietary actions involve movements of the mouth, jaw, and arms, which can be inferred from reflected ultrasound, as shown in prior work [28, 37, 39?]. However, the capacity of the active acoustic sensing-based general-purpose activity recognition system ActSonic [40] is limited to tracking moments of eating and drinking episodes. Since eating is a complex activity involving several fine-grained dietary actions such as hand-to-mouth movement for intake, chewing, swallowing, etc., we designed the MunchSonic system specifically for tracking body parts involved in eating, with modifications to ActSonic [40] hardware and inference pipeline. We evaluated our system in a user study with 12 participants in unconstrained environments of their choice. The results showed that MunchSonic can accurately recognize fine-grained dietary actions such as hand movements for intake, chewing, and drinking, as well as non-dietary actions like talking, face touching, and other activities (null), achieving an F1-score of 93.5%.

The contributions of the MunchSonic system are as follows:

- Advanced the knowledge of wearable-based active acoustic sensing by demonstrating the feasibility of using active acoustic sensing on glasses to recognize fine-grained dietary actions beyond only identifying eating moments.
- Evaluation of the proposed system through a 12-participant user study in unconstrained environments, with ground truth data annotated for each second.

2 Related Work

Various wearable devices with different sensing modalities have been proposed to track eating events. This section discusses non-eyeglass wearables, followed by eyeglasses. Table 1 summarizes prior studies on tracking eating activities.

Form Factors Other than Eyeglasses. Smartwatch-based systems for eating [15, 32, 51, 60, 61, 68, 75] and drinking [25] use IMUs to track hand-to-mouth gestures but face challenges distinguishing similar motions, leading to high false positives [12, 41, 62, 63]. Ear-worn devices, or earables, track eating by monitoring jaw movement through in-ear proximity sensing [7, 8], passive acoustic sensing [10, 23], or both [5]. Neckbands [18, 46, 67] and necklaces [1, 13, 29, 54, 76] use passive acoustic and proximity sensing to track chewing and swallowing. Despite advancements and around 80% accuracy in free-living conditions, neckband and necklace form factors have not gained the same popularity as other wearables [3].

Eyeglasses Form Factor. Dietary monitoring using eyeglasses is promising due to their proximity to the mouth and jaw and higher social acceptance [31] (64% adoption in the US [21]). Researchers have used EMG electrodes [72], piezoelectric sensors [19, 20], inertial sensors [21], and load cells [14] on eyeglasses to track eating behaviors. Sensor fusion approaches [4, 6, 42, 47, 53], integrating

gyroscopes, accelerometers, proximity sensors, and microphones, have shown improved performance. Passive acoustic sensing using contact microphones [10] captures chewing sounds but requires close skin contact. Systems using acoustic sensing [40, 43] focus on eating episodes rather than fine-grained actions.

In summary, MunchSonic is the first to use active acoustic sensing on eyeglasses to infer fine-grained dietary actions by capturing jaw and hand movements. This allows for detailed dietary action extraction, such as intake and chewing, unlike other systems that primarily track moments of eating episodes.

3 System Implementation

The goal of MunchSonic is to capture the information related to movements on multiple body parts involved in dietary activities using an eyeglasses form factor. Previous research has shown active acoustic sensing's efficacy on wearables in tracking facial muscle movements [36, 37, 58, 73, 74], upper body limbs [39], hand poses [34, 69], gaze [35], respiration [65], sign language gestures [27, 28], physiological signals [17] and everyday activities [40]. Inspired by these prior works, especially [40], MunchSonic integrates active acoustic sensing into eyeglasses to track these fine-grained dietary actions. This section discusses the active acoustic sensing mechanism, the hardware implementation, and the deep learning framework for processing the captured data.

3.1 Sensing Mechanism

The design of MunchSonic's sensing system is inspired by the active acoustic sensing approach from [40]. In essence, MunchSonic uses a similar active acoustic sensing processing method, based on cross-correlation-based Frequency Modulated Continuous Wave (C-FMCW) [65]. These chirps are transmitted from eyeglasses' hinges, with frequency ranges of 18-21.5 KHz and 21.5-24.5 KHz for the left and right transmitters, respectively. The receiver samples at 50 KHz, with each chirp containing 600 samples and a sweep period of 0.012 seconds. The system can detect changes as small as 3.43 mm, given the slow speed of mouth, jaw, and hand movements compared to the speed of sound (343 m/s). The sensing range is up to 2.058 meters, though a shorter range is used for fine-grained dietary actions. MunchSonic uses a lower amplifier gain compared to [40] to focus on mouth, jaw, and hand movements near the face for tracking fine-grained dietary actions. Consequently, the microphones receive weaker reflections from lower body parts that are not involved in dietary actions.

MunchSonic's signal processing pipeline involves computing the cross-correlation of transmitted and received ultrasonic waves. The received signal is first filtered through a bandpass filter (18-21 KHz and 21.5-24.5 KHz) to eliminate audible frequencies. The cross-correlation output, called the *Echo Profile*, functions as a Range-FFT using acoustic C-FMCW waves. The *Echo Profile* is processed as a two-dimensional tensor (x -axis: time, y -axis: distance from MunchSonic eyeglasses form factor). To capture body part movements and eliminate static object reflections, the first derivative of the *Echo Profile*, termed the *Differential Echo Profile*, is used as input for the deep learning architecture to track fine-grained dietary actions.

Year	Study	Wearable	Sensing Modalitie(s)	Detected Event(s)	Evaluation Metric	Power Consumption
2017	EarBit [5]	Outer Ear Interface	IMU, IR Proximity Sensor, Microphone	Chewing	F1-score = 80.1%	-
2017	GlasSense [14]	Eyeglasses	Load Cell	Head Movement, Talking, Chewing	F1-score = 94.0%*	-
2017	Mirtchouk et al. [42]	Google Glass	IMU, Microphone, Motion Sensor	Meals	Precision = 31% Recall = 87%	-
2018	Auracle [10]	Earpiece	Contact Microphone	Eating Episodes	F1-score = 77.5%	14.47 mW
2018	Zhang et al. [72]	Eyeglasses	EMG Electrode	Eating Event	F1-score = 77%	81.96 mW
2020	FitByte [4]	Eyeglasses	1 Camera, 1 Proximity Sensor, 6 IMUs	Eating and Drinking Episodes	Precision = 82.8% (Eating), 56.7% (Drinking) Recall = 93.8% (Eating), 65.5% (Drinking)	105.08 mW
2022	MyDJ [53]	Eyeglasses	Accelerometer, Piezoelectric	Eating Episodes	F1-score = 92%	26.06 mW
2024	MunchSonic	Eyeglasses	Active Acoustic Sensing	Intake, Chewing, Drinking, Talking, Face Touch	F1-score = 93.50%	96.5 mW [†]

Table 1: Comparison of eating detection studies based on wearable form factors, sensing modalities, event granularity, evaluation metrics (F1-score or precision/recall), and power draw (mW). Asterisks (*) in the *Evaluation Metric* column indicate lab study evaluations; otherwise, evaluations were in free-living conditions. The dagger symbol (†) in the *Power Consumption* column indicates machine learning inference was performed on a cloud or mobile device, not the wearable.

3.2 Hardware and Form Factor

We replicated the sensing system from ActSonic [40] with minor changes to the sensors and controller unit, using the OWR-05049T-38D speaker and ICS-43434 microphone [59]. Our customized controller unit, shown in Figure 1(b), features the nRF52840 microcontroller [50], two MAX98357A audio amplifiers, a BLE SGW1110 [52] module, power management modules, and an SD card slot with a SanDisk Extreme Pro microSD card [2], optimized for power efficiency. The transceiver boards, each with a speaker and microphone, are attached to the eyeglass hinges, with placement refined for optimal tracking of mouth, jaw, and hand movements. The controller unit and LiPo battery are attached to one leg of the glasses frame, connected via Flexible Printed Circuit (FPC) cables and a JST connector. The sensing system on the eyeglasses draws 96.5 mW of power for collecting data, with a voltage of 4.02V and a current of 24.0 mA. The prototype used in the MunchSonic study, featuring a 290 mAh LiPo battery, lasts approximately 11.25 hours in SD card storage mode.




Figure 1: MunchSonic Hardware and Form Factor: (a) User wearing eyeglasses form factor, (b) Customized controller unit with nRF52840 microcontroller, (c) Top view of the eyeglasses form factor, (d) MunchSonic transceiver for active acoustic sensing housing one speaker (top) and one microphone (bottom).

3.3 Deep Learning Framework

3.3.1 Data Processing and Model Architecture. We designed a lightweight deep learning framework to process the active acoustic data captured by the MunchSonic system. The differential echo profile described in Section 3.1 serves as the input to the model. We create overlapping sliding windows from the differential echo profile data. To augment the data, we apply Gaussian noise to 5% of the windows and mask 5% of the distance axis to simulate environmental noise and unrelated movements.

The preprocessed sliding windows are fed into a MobileNetV2 [48] convolutional neural network encoder, which generates a 256-dimensional embedding vector for each window. This vector is then fed into a classifier network with three feedforward layers (128, 64, and 6 neurons). Each layer, except the last, includes batch normalization [26], dropout [57] with a probability of 0.25, and Leaky ReLU [66] activation. Finally, a softmax operation provides the class probability distribution.

3.3.2 Model Training. The input to MunchSonic’s deep learning framework consists of sliding windows of differential echo profiles, optimally 2 seconds long with 50% overlap. This window length fits an eating intake gesture. The optimal sensing range is 150 pixels on the y -axis of the differential echo profile, corresponding to 51.45 cm from the MunchSonic device. With speakers and microphones on each hinge of the eyeglasses, transmitting acoustic waves in two frequency ranges, there are four channels in the differential echo profile. Thus, the shape of the input window is $(4 \times 150 \times 166)$, as each second contains 83 samples.

We use focal loss [38] for training to handle class imbalance in the MunchSonic dataset. The Adam [30] optimizer with a cosine annealing learning rate scheduler starts with an initial learning rate of 10^{-2} . The model, implemented using PyTorch and PyTorch Lightning, is trained for 30 epochs with a batch size of 128 on GeForce RTX 2080 Ti GPUs.

4 User Study

We conducted a 45-minute user study to evaluate MunchSonic in an unconstrained setting. Building on ActSonic’s [40] success in tracking eating and drinking episodes with high F1 scores, the MunchSonic study assumed active acoustic sensing on glasses can distinguish eating moments from other activities. This paper focuses on extracting fine-grained dietary actions, such as hand-to-mouth movements, chewing, and drinking. The study aimed to assess the system’s performance in tracking these actions and distinguishing them from non-eating activities in real-world conditions. We conducted a shorter-duration free-living study, leveraging ActSonic’s proven efficacy in longer studies. Ground truth annotation was done per second, unlike other systems [53] that rely on self-reporting, enabling high-resolution recognition of dietary actions.

Study Protocol We conducted a user study approved by our organization’s Institutional Review Board (IRB) with 13 participants, 6 identifying as male and 7 as female, with an average age of 26.5 years, ranging from 21 to 36 years. The participants wore the MunchSonic eyeglasses form factor and a chest-mounted GoPro HERO9 camera [24] facing upward to capture their eating activities around the face as the ground truth. The camera recorded 720p video at 30 fps with a diagonal field of view of 148°. Note that one participant (P09) accidentally turned off the chest-mounted camera, resulting in no ground truth data for that participant. Therefore, we discarded that participant from the study and evaluated the system on the remaining 12 participants.



Figure 2: User study in unconstrained conditions: (a) Foods consumed by user study participants, (b) Sample images of the chest-mounted camera view of MunchSonic data collection pipeline.

Participants were asked to bring a meal (breakfast, lunch, or dinner) from any source (home-cooked or restaurant) and come to the lab. After briefing them about the study procedure, we equipped them with the MunchSonic eyeglasses and the chest-mounted camera. We verified the data collection system with a short 5-minute session. Once verified, the data collection system was started, and participants were free to go anywhere while wearing the eyeglasses and camera. The only requirement was to finish the meal they brought to ensure sufficient samples of dietary movements. Apart from that, they could continue with their daily routine. The total duration of this unconstrained study was 45 minutes. Upon completion, participants returned to the lab to return the devices and were reimbursed with a \$20 USD gift card.

Dataset Statistics We collected 540 minutes of data from 12 participants. Two researchers labeled the reference video from the chest-mounted camera at each second using the ANU-CVML Video Annotation Tool (Vidat) [71]. The inter-rater reliability was measured through Cohen’s Kappa [33], κ , and its value is 0.82 ($\kappa > 0.80$ denotes near perfect agreement, and $0.6 < \kappa \leq 0.8$ denotes satisfactory agreement). The camera view that the annotators labeled is shown in Figure 2(b). The data was categorized into six classes: hand-to-mouth movement for food or drink intake, chewing, drinking, talking, face touch, and a null class for other activities. The annotators were instructed to segment the actions in two phases. First, they determined whether the action was eating or non-eating. If the action was identified as eating, they further categorized it into fine-grained actions such as hand-to-mouth movement, chewing, and drinking. Swallowing was integrated into the chewing category. For non-eating actions, they segmented the activities into talking, face touching, and other activities not included in the tracking set. A synchronization script was used to extrapolate ground truth labels for the differential echo profile sliding windows. The class distribution was 53.5% null, 5.9% hand-to-mouth for food intake, 22.0% chewing, 2.2% drinking, 15.1% talking, and 1.3% touching face with hand for non-eating activities.

5 Performance Evaluation

We computed precision, recall, and macro F1-score to assess MunchSonic’s ability to track fine-grained dietary actions. Evaluations were conducted at two levels: frame-level inference for each two-second sliding window and episode-level inference for intake count and chewing time estimation. Employing a leave-one-participant-out evaluation strategy enabled user-independent assessment, using data from one participant as the test set and the remaining 11 for training and validation. This evaluation demonstrates the system’s performance without requiring training data from new users, which makes the system easier to deploy at scale.

5.1 Evaluation of Frame-Level Inference




Figure 3: Precision and recall of leave-one-participant-out evaluation of MunchSonic, where data from each participant on the x-axis serves as the test set.

We summarize the results of the leave-one-participant-out evaluation of sliding window prediction in Figure 3. According to the evaluation, the average macro F1-score across all 12 participants is 0.935 with a standard deviation of 0.031. It is evident that the MunchSonic system can track fine-grained dietary events with

precision and recall of more than 90% for most participants. P08 demonstrates the worst performance in terms of both precision and recall. Additionally, P04 and P13 yielded slightly lower recall compared to the average. Our analysis of the ground truth videos for those participants suggests that the slightly worse performance can be attributed to the participants touching their faces several times during the study for activities not related to eating. This phenomenon is also evident in the confusion matrix in Figure 4. The performance degradation of MunchSonic for face-touch actions can be related to the imbalanced nature of the dataset collected in a completely unconstrained manner, which contains very few samples of this particular activity. Nonetheless, MunchSonic exhibits robust performance across all activities, with a mean precision and recall of 0.941 ± 0.027 and 0.930 ± 0.041 respectively, across all 12 participants.




Figure 4: Normalized confusion matrix from the leave-one-participant-out evaluation across 12 participants in the user study of MunchSonic. The values in parentheses represent the total number of instances for each cell.

5.2 Evaluation of Episode-Level Inference




Figure 5: Episode-level evaluation of MunchSonic: (a) Segmentation of user study data into 4.5-minute-long episodes, (b) Detection of food intake within each episode.

5.2.1 Eating Episode Detection. To evaluate MunchSonic in detecting eating episodes and intake counts within episodes, we segmented the 45-minute data collected from each participant in the user study into 10 segments of equal length (4.5 minutes each). According to the definition provided in [60, 75], if a segment of length t_w contains $\frac{t_w+20}{5}$ intakes, which is 5 intakes for the aforementioned

segments, then that segment is defined as an eating episode. To evaluate MunchSonic’s efficacy in detecting these eating episodes, we deployed a majority voting strategy where a segment is labeled as an eating episode if 50% of the frames are labeled as either *food_intake* or *chewing*. The number of undetected eating episodes for the MunchSonic dataset is 1 out of 43 ground truth episodes, resulting in a False Negative Rate of 0.023. Additionally, there were 77 non-eating episodes in the dataset, of which 2 were incorrectly detected as eating by MunchSonic, leading to a False Positive Rate of 0.026.

5.2.2 Food Intake Counting. Furthermore, to count the number of intakes in each segment, we increment the intake count for that segment by 1 if we find one *food_intake* frame followed by at least two chewing frames or windows within the next 3.0 seconds. We compute the Mean Absolute Error (MAE) between the ground truth number of intakes and the predicted count as a metric to evaluate this. The average ground truth intake count in one 4.5-minute eating segment (containing more than 5 ground truth intakes within the timeframe) and non-eating segment (containing fewer than 5 ground truth intakes) is 17.65 and 1.19, respectively. The MAE for counting intakes in the eating segments is 2.01, leading to an average error of 11.39% in predicting the number of eating intakes. For the non-eating segments, the MAE for counting intakes is 0.047, leading to a mean error of 3.95% in predicting the intake counts during non-eating episodes.

5.2.3 Chewing Time Estimation. To evaluate the coverage of predicted chewing time, we calculate the Mean Absolute Error (MAE) between the total ground truth chewing time in a 4.5-minute episode and the predicted chewing time, measured in seconds, across all 12 user study participants. The average chewing time across eating episodes is 163.6 seconds (2.73 minutes) and 9.17 seconds for non-eating episodes in the MunchSonic dataset. Interpolating from the frame-level predictions of MunchSonic, the MAE of the estimated chewing time for eating episodes is 12.87 seconds, leading to a coverage of 92.13% of chewing time. For the non-eating episodes, the MAE of the estimated chewing time is 1.61 seconds, leading to a coverage of 82.5%. This higher coverage across non-eating episodes indicates MunchSonic’s potential in detecting snacking events, which was outside the scope of the user study.

6 Discussion

Analyzing the Passive Sensed Ultrasonic Range for Tracking Dietary Actions Theoretically, our system focuses on nearly inaudible sound frequencies above 18 KHz, which should not interfere with most daily activities, including eating. However, some eating sounds can have frequency components above 18 KHz. To evaluate the impact of these passively sensed ultrasonic signals on distinguishing dietary actions, we conducted a preliminary study with one researcher as a participant. In this study, we turned off MunchSonic’s speakers to prevent them from transmitting C-FMCW chirps and performed the same set of activities as in the user study. The microphones continued to record the surrounding acoustic signals. We then created differential echo profile windows and trained the MunchSonic deep learning model. The mean cross-session macro F1-score for this user-dependent model was 0.322. In contrast, with




Figure 6: Differential echo profiles of actions with active and passive acoustic sensing. The x -axis represents time, and the y -axis represents the echo distance from the MunchSonic form factor, ranging from 0 cm to 50 cm.

the speakers transmitting ultrasonic chirps, the mean cross-session macro F1-score was 0.976. As illustrated in Figure 6, active acoustic sensing plays a crucial role in distinguishing dietary activities by tracking the movements of body parts involved in dietary actions. While passive ultrasonic signals might have a minor impact on detecting fine-grained events, quantifying this impact numerically is challenging and beyond the scope of this paper. Our experiments indicate that chewing crunchy foods generated more ultrasonic components in the spectrogram. However, the signal strength of reflections from MunchSonic’s transmitted ultrasonic chirps is much higher, potentially overshadowing the passively sensed components’ feature importance. Further investigation into this will be left for future work.

Impact of Sensing Range and Temporal Context We evaluated the impact of sensing range and temporal context for the MunchSonic system. Figure 7 shows that a sliding window size of 2.00 seconds with a 50% overlap yields the best tracking performance. We also found that a sensing range of 50 to 80 cm is optimal for tracking fine-grained dietary actions. This is because most movements related to dietary actions occur in the jaw and mouth region, within 30 cm of the eyeglasses form factor. Additionally, MunchSonic tracks hand movements for food intake, so extending the sensing range to 50 cm yields the best performance. Poorer performance at longer ranges is likely due to various unrelated activities occurring in that region.




Figure 7: Ablation study to measure the impact of sensing range and sliding window size on the performance of MunchSonic. Here, all the sliding window sizes mentioned have a 50% overlap.

Comfort and Safety of Wearing MunchSonic Prototype

After participating in the user study described in Section 4, participants completed an IRB-approved questionnaire on the comfort of wearing the MunchSonic prototype. They rated their comfort from 0 (most uncomfortable) to 5 (most comfortable), with a mean rating of 3.462 and a standard deviation of 1.050. Since MunchSonic operates using the ultrasonic range, participants were not expected to hear anything from the system, and indeed, none reported hearing any noise. Three participants noted that the prototype did not fit their head sizes and prescription lenses well, which could be addressed in future iterations by creating an attachable sensing module. Additionally, one participant mentioned that the chest-mounted camera used for ground truth acquisition occasionally touched their dining table, interrupting their eating.

We measured the transmitted signal intensity using a CDC-provided app [11], finding it to be 68 dB(A), below the NIOSH limit of 85 dB [45]. While MHz range ultrasonic exposure can cause discomfort [44], MunchSonic operates in the KHz range with no reported issues. Future studies will explore potential audibility among animals and children.

Potential Application Scenarios

We envision MunchSonic being useful for eating behavior assessment, chronic disease management, and overall well-being. Our system passively and objectively measures eating episodes and micro-level actions (e.g., pauses, drinking, chewing) in real-world settings. By accurately classifying these actions, we can derive meal micro-structure metrics [9] like eating duration, speed, and chewing rate, traditionally assessed through interviews or self-report questionnaires [49], which are prone to biases. Our automated approach provides precise, naturalistic data. Clinical studies link these metrics to health outcomes like higher Body Mass Index [56, 64], eating disorders [16, 22], and metabolic diseases [70]. Accurate assessment enhances monitoring and management of these conditions, facilitating real-time interventions. A day-long evaluation of MunchSonic will provide crucial insights into snacking behavior and bulimia or binge eating disorder detection.

Acknowledgements

This project was supported by the National Science Foundation Grant No. 2239569 and partially by the Cornell University IGNITE Innovation Acceleration Program.

References

- [1] Nabil Alshurafa, Haik Kalantarian, Mohammad Pourhomayoun, Jason J Liu, Shruti Sarin, Behnam Shahbazi, and Majid Sarrafzadeh. 2015. Recognition of nutrition intake using time-frequency decomposition in a wearable necklace using a piezoelectric sensor. *IEEE sensors journal* 15, 7 (2015), 3909–3916.
- [2] Amazon.com. [n.d.]. SanDisk Extreme Pro 32GB SDHC UHS-I Card (SDSDXXG-032G-GN4IN). <https://www.amazon.com/SanDisk-Extreme-32GB-UHS-I-SDSDXXG-032G-GN4IN/dp/B01J5RHBQ4>. [Online; accessed 21-May-2024].
- [3] Judith Amores, Mae Dotan, and Pattie Maes. 2019. An exploration of form factors for sleep-olfactory interfaces. In *2019 41st annual international conference of the IEEE engineering in medicine and biology society (EMBC)*. IEEE, 1456–1460.
- [4] Abdelkareem Bedri, Diana Li, Rushil Khurana, Kunal Bhuwanka, and Mayank Goel. 2020. FitByte: Automatic Diet Monitoring in Unconstrained Situations Using Multimodal Sensing on Eyeglasses. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (<conf-loc>, <city>Honolulu</city>, <state>HI</state>, <country>USA</country>, </conf-loc>) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376869>
- [5] 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. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 1, 3 (2017), 1–20.
- [6] Abdelkareem Bedri, Yuchen Liang, Sudershan Boovaraghavan, Geoff Kaufman, and Mayank Goel. 2022. FitNibble: A Field Study to Evaluate the Utility and Usability of Automatic Diet Monitoring in Food Journaling Using an Eyeglasses-based Wearable. In *Proceedings of the 27th International Conference on Intelligent User Interfaces* (<conf-loc>, <city>Helsinki</city>, <country>Finland</country>, </conf-loc>) (IUI '22). Association for Computing Machinery, New York, NY, USA, 79–92. <https://doi.org/10.1145/3490099.3511154>
- [7] Abdelkareem Bedri, Apoorva Verlekar, Edison Thomaz, Valerie Avva, and Thad Starner. 2015. Detecting Mastication: A Wearable Approach. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction* (Seattle, Washington, USA) (ICMI '15). Association for Computing Machinery, New York, NY, USA, 247–250. <https://doi.org/10.1145/2818346.2820767>
- [8] Abdelkareem Bedri, Apoorva Verlekar, Edison Thomaz, Valerie Avva, and Thad Starner. 2015. A wearable system for detecting eating activities with proximity sensors in the outer ear. In *Proceedings of the 2015 ACM International Symposium on Wearable Computers*. 91–92.
- [9] France Bellisle. 2020. Edograms: recording the microstructure of meal intake in humans—a window on appetite mechanisms. *International Journal of Obesity* 44, 12 (Dec. 2020), 2347–2357. <https://doi.org/10.1038/s41366-020-00653-w> Publisher: Nature Publishing Group.
- [10] 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. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 3, Article 92 (sep 2018), 27 pages. <https://doi.org/10.1145/3264902>
- [11] PE Churci A. Kardous, MS and Ph.D. Peter B. Shaw. [n.d.]. CDC: So How Accurate Are These Smartphone Sound Measurement Apps? <https://blogs.cdc.gov/niosh-science-blog/2014/04/09/sound-apps/>. [Online; accessed 29-Nov-2023].
- [12] Keum San Chun, Sarnab Bhattacharya, and Edison Thomaz. 2018. Detecting eating episodes by tracking jawbone movements with a non-contact wearable sensor. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 2, 1 (2018), 1–21.
- [13] Keum San Chun, Sarnab Bhattacharya, and Edison Thomaz. 2018. Detecting Eating Episodes by Tracking Jawbone Movements with a Non-Contact Wearable Sensor. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 1, Article 4 (mar 2018), 21 pages. <https://doi.org/10.1145/3191736>
- [14] Jungman Chung, Jungmin Chung, Wonjun Oh, Yongkyu Yoo, Won Gu Lee, and Hyunwoo Bang. 2017. A glasses-type wearable device for monitoring the patterns of food intake and facial activity. *Scientific reports* 7, 1 (2017), 41690.
- [15] Yujie Dong, Jenna Scisco, Mike Wilson, Eric Muth, and Adam Hoover. 2013. Detecting periods of eating during free-living by tracking wrist motion. *IEEE journal of biomedical and health informatics* 18, 4 (2013), 1253–1260.
- [16] Christopher Fairburn and G. Wilson. 1993. Binge Eating: Nature, Assessment, and Treatment. *Journal of Nervous and Mental Disease - J NERV MENT DIS* 183 (Jan. 1993).
- [17] Xiaoran Fan, David Pearl, Richard Howard, Longfei Shangguan, and Trausti Thormundsson. 2023. Apg: Audioplethysmography for cardiac monitoring in hearables. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*. 1–15.
- [18] Muhammad Farooq, Juan M Fontana, and Edward Sazonov. 2014. A novel approach for food intake detection using electroglottography. *Physiological measurement* 35, 5 (2014), 739.
- [19] Muhammad Farooq and Edward Sazonov. 2016. A novel wearable device for food intake and physical activity recognition. *Sensors* 16, 7 (2016), 1067.
- [20] 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* 21, 6 (2016), 1495–1503.
- [21] Muhammad Farooq and Edward Sazonov. 2018. Accelerometer-based detection of food intake in free-living individuals. *IEEE sensors journal* 18, 9 (2018), 3752–3758.
- [22] Centre for Research on Eating Disorders at Oxford (CREDO). [n.d.]. Eating Disorder Examination (Edition 17.0D). <https://www.cbte.co/site/download/ede-17-0d/?wpdmml=615&masterkey=5c644ef9b6149>. Accessed: 2024-05-26.
- [23] Yang Gao, Ning Zhang, Honghao Wang, Xiang Ding, Xu Ye, Guanling Chen, and Yu Cao. 2016. iHear food: eating detection using commodity bluetooth headsets. In *2016 IEEE First International Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE)*. IEEE, 163–172.
- [24] GoPro. 2020. HERO9 Black. <https://gopro.com/en/us/shop/cameras/hero9-black/CHDXH-901-master.html>. [Online; accessed 22-May-2024].
- [25] Takashi Hamatani, Moustafa Elhamshary, Akira Uchiyama, and Teruo Higashino. 2018. FluidMeter: Gauging the human daily fluid intake using smartwatches. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 3 (2018), 1–25.
- [26] Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. 448–456. <http://jmlr.org/proceedings/papers/v37/loffe15.pdf>
- [27] Yincheng Jin, Seokmin Choi, Yang Gao, Jiyang Li, Zhengxiong Li, and Zhanpeng Jin. 2023. TransASL: A Smart Glass based Comprehensive ASL Recognizer in Daily Life. In *Proceedings of the 28th International Conference on Intelligent User Interfaces* (<conf-loc>, <city>Sydney</city>, <state>NSW</state>, <country>Australia</country>, </conf-loc>) (IUI '23). Association for Computing Machinery, New York, NY, USA, 802–818. <https://doi.org/10.1145/3581641.3584071>
- [28] Yincheng Jin, Yang Gao, Yanjun Zhu, Wei Wang, Jiyang Li, Seokmin Choi, Zhangyu Li, Jagmohan Chauhan, Anind K. Dey, and Zhanpeng Jin. 2021. Soni-cASL: An Acoustic-based Sign Language Gesture Recognizer Using Earphones. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 2, Article 67 (jun 2021), 30 pages. <https://doi.org/10.1145/3463519>
- [29] 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.
- [30] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [31] Nataliya Kosmyna, Caitlin Morris, Utkarsh Sarawgi, and Pattie Maes. 2019. AttentivU: A Biofeedback System for Real-time Monitoring and Improvement of Engagement. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI EA '19). Association for Computing Machinery, New York, NY, USA, 1–2. <https://doi.org/10.1145/3290607.3311768>
- [32] Konstantinos Kyritsis, Christos Diou, and Anastasios Delopoulos. 2019. Detecting meals in the wild using the inertial data of a typical smartwatch. In *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 4229–4232.
- [33] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. 2017. *Research methods in human-computer interaction*. Morgan Kaufmann.
- [34] Chi-Jung Lee, Ruidong Zhang, Devansh Agarwal, Tianhong Catherine Yu, Vipin Gunda, Oliver Lopez, James Kim, Sicheng Yin, Boao Dong, Ke Li, Mose Sakashita, Francois Guimbretiere, and Cheng Zhang. 2024. EchoWrist: Continuous Hand Pose Tracking and Hand-Object Interaction Recognition Using Low-Power Active Acoustic Sensing On a Wristband. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 403, 21 pages. <https://doi.org/10.1145/3613904.3642910>
- [35] Ke Li, Ruidong Zhang, Boao Chen, Siyuan Chen, Sicheng Yin, Saif Mahmud, Qikang Liang, Francois Guimbretiere, and Cheng Zhang. 2024. GazeTrak: Exploring Acoustic-based Eye Tracking on a Glass Frame. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking* (Washington D.C., DC, USA) (ACM MobiCom '24). Association for Computing Machinery, New York, NY, USA, 497–512. <https://doi.org/10.1145/3636534.3649376>
- [36] Ke Li, Ruidong Zhang, Siyuan Chen, Boao Chen, Mose Sakashita, François Guimbretière, and Cheng Zhang. 2024. EyeEcho: Continuous and Low-power Facial Expression Tracking on Glasses. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–24.
- [37] Ke Li, Ruidong Zhang, Bo Liang, François Guimbretière, and Cheng Zhang. 2022. EarLo: A Low-Power Acoustic Sensing Earable for Continuously Tracking Detailed Facial Movements. 6, 2, Article 62 (jul 2022), 24 pages. <https://doi.org/10.1145/3534621>
- [38] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*. 2980–2988.
- [39] Saif Mahmud, Ke Li, Guilin Hu, Hao Chen, Richard Jin, Ruidong Zhang, François Guimbretière, and Cheng Zhang. 2023. PoseSonic: 3D Upper Body Pose Estimation Through Egocentric Acoustic Sensing on Smartglasses. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 7, 3, Article 111 (sep 2023), 28 pages.

- <https://doi.org/10.1145/3610895>
- [40] Saif Mahmud, Vineet Parikh, Qikang Liang, Ke Li, Ruidong Zhang, Ashwin Ajit, Vipin Gunda, Devansh Agarwal, François Guimbretière, and Cheng Zhang. 2024. ActSonic: Recognizing Everyday Activities from Inaudible Acoustic Waves Around the Body. *arXiv:2404.13924* [cs.HC]
 - [41] Saif Mahmud, M. T. H. Tonmoy, Kishor Kumar Bhaumik, A. M. Rahman, M. A. Amin, M. Shoyab, Muhammad Asif Hossain Khan, and A. Ali. 2020. Human Activity Recognition from Wearable Sensor Data Using Self-Attention. In *ECAI 2020 - 24th European Conference on Artificial Intelligence, 29 August-8 September 2020, Santiago de Compostela, Spain*.
 - [42] 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. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 85 (sep 2017), 20 pages. <https://doi.org/10.1145/3131894>
 - [43] Vimal Mollyn, Karan Ahuja, Dhruv Verma, Chris Harrison, and Mayank Goel. 2022. SAMoSA: Sensing Activities with Motion and Subsampled Audio. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 3 (2022), 1–19.
 - [44] David Baeza Moyano, Daniel Arranz Paraiso, and Roberto Alonso González-Lezcano. 2022. Possible effects on health of ultrasound exposure, risk factors in the work environment and occupational safety review. In *Healthcare*, Vol. 10. MDPI, 423.
 - [45] William J Murphy and John R Franks. 2002. Revisiting the NIOSH criteria for a recommended standard: Occupational noise exposure. *The Journal of the Acoustical Society of America* 111, 5 (2002), 2397.
 - [46] Tauhidur Rahman, Alexander T. Adams, Mi Zhang, Erin Cherry, Bobby Zhou, Huaishu Peng, and Tanzeem Choudhury. 2014. BodyBeat: a mobile system for sensing non-speech body sounds. In *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services* (Bretton Woods, New Hampshire, USA) (*MobiSys '14*). Association for Computing Machinery, New York, NY, USA, 2–13. <https://doi.org/10.1145/2594368.2594386>
 - [47] Tauhidur Rahman, Mary Czerwinski, Ran Gilad-Bachrach, and Paul Johns. 2016. Predicting "About-to-Eat" Moments for Just-in-Time Eating Intervention. In *Proceedings of the 6th International Conference on Digital Health Conference* (Montréal, Québec, Canada) (*DH '16*). Association for Computing Machinery, New York, NY, USA, 141–150. <https://doi.org/10.1145/2896338.2896359>
 - [48] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4510–4520.
 - [49] S. Sasaki, A. Katagiri, T. Tsuji, T. Shimoda, and K. Amano. 2003. Self-reported rate of eating correlates with body mass index in 18-y-old Japanese women. *International Journal of Obesity* 27, 11 (Nov. 2003), 1405–1410. <https://doi.org/10.1038/sj.ijo.0802425> Publisher: Nature Publishing Group.
 - [50] Nordic Semiconductors. [n.d.]. nRF52840 Multiprotocol Bluetooth 5.4 SoC supporting Bluetooth Low Energy, Bluetooth mesh, NFC, Thread and Zigbee. <https://www.nordicsemi.com/products/nrf52840>. [Online; accessed 29-Nov-2023].
 - [51] Sougata Sen, Vigneshwaran Subbaraju, Archan Misra, Rajesh Balan, and Youngki Lee. 2018. Annapurna: building a real-world smartwatch-based automated food journal. In *2018 IEEE 19th International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)*. IEEE, 1–6.
 - [52] SGWireless. [n.d.]. SGW111X BLE Modules. <https://www.sgwireless.com/product/SGW111X>. [Online; accessed 29-Nov-2023].
 - [53] Jaemin Shin, Seungjoo Lee, Taesik Gong, Hyungjun Yoon, Hyunchul Roh, Andrea Bianchi, and Sung-Ju Lee. 2022. MyDJ: Sensing Food Intakes with an Attachable on Your Eyeglass Frame. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI '22*). Association for Computing Machinery, New York, NY, USA, Article 341, 17 pages. <https://doi.org/10.1145/3491102.3502041>
 - [54] Jaemin Shin, Seungjoo Lee, and Sung-Ju Lee. 2019. Accurate eating detection on a daily wearable necklace. In *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services*. 649–650.
 - [55] Arnold Slyper. 2021. Oral Processing, Satiation and Obesity: Overview and Hypotheses. *Diabetes, Metabolic Syndrome and Obesity: Targets and Therapy* 14 (July 2021), 3399–3415. <https://doi.org/10.2147/DMSO.S314379>
 - [56] Chikanobu Sonoda, Hideki Fukuda, Masayasu Kitamura, Hideaki Hayashida, Yumiko Kawashita, Reiko Furugen, Zenya Koyama, and Toshiyuki Saito. 2018. Associations among Obesity, Eating Speed, and Oral Health. *Obesity Facts* 11, 2 (April 2018), 165–175. <https://doi.org/10.1159/000488533>
 - [57] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
 - [58] Rujia Sun, Xiaohu Zhou, Benjamin Steeper, Ruidong Zhang, Sicheng Yin, Ke Li, Shengzhang Wu, Sam Tilsen, François Guimbretière, and Cheng Zhang. 2023. EchoNose: Sensing Mouth, Breathing and Tongue Gestures inside Oral Cavity using a Non-contact Nose Interface. In *Proceedings of the 2023 ACM International Symposium on Wearable Computers* (Cancun, Quintana Roo, Mexico) (*ISWC '23*). Association for Computing Machinery, New York, NY, USA, 22–26. <https://doi.org/10.1145/3594738.3611358>
 - [59] TDK. [n.d.]. ICS-43434 Multi-Mode Microphone with I²S Digital Output. <https://invensense.tdk.com/products/ics-43434/>. [Online; accessed 29-Nov-2023].
 - [60] Edison Thomaz, Irfan Essa, and Gregory D Abowd. 2015. A practical approach for recognizing eating moments with wrist-mounted inertial sensing. In *Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing*. 1029–1040.
 - [61] 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 *Proceedings of the 20th International Conference on Intelligent User Interfaces* (Atlanta, Georgia, USA) (*IUI '15*). Association for Computing Machinery, New York, NY, USA, 427–431. <https://doi.org/10.1145/2678025.2701405>
 - [62] Catherine Tong, Shyam A. Tailor, and Nicholas D. Lane. 2020. Are Accelerometers for Activity Recognition a Dead-End?. In *Proceedings of the 21st International Workshop on Mobile Computing Systems and Applications* (Austin, TX, USA) (*HotMobile '20*). Association for Computing Machinery, New York, NY, USA, 39–44. <https://doi.org/10.1145/3376897.3377867>
 - [63] M. Tanjid Hasan Tonmoy, Saif Mahmud, A. K. M. Mahbubur Rahman, M. Ashraf Amin, and Amin Ahsan Ali. 2021. Hierarchical Self Attention Based Autoencoder for Open-Set Human Activity Recognition. In *Advances in Knowledge Discovery and Data Mining*, Kamal Karlapalem, Hong Cheng, Naren Ramakrishnan, R. K. Agrawal, P. Krishna Reddy, Jaideep Srivastava, and Tanmoy Chakraborty (Eds.). Springer International Publishing, Cham, 351–363.
 - [64] B. Timothy Walsh, Harry R. Kissileff, Susan M. Cassidy, and Sondra Dantzie. 1989. Eating Behavior of Women With Bulimia. *Archives of General Psychiatry* 46, 1 (Jan. 1989), 54–58. <https://doi.org/10.1001/archpsyc.1989.01810010056008>
 - [65] Tianben Wang, Daqing Zhang, Yuanqing Zheng, Tao Gu, Xingshe Zhou, and Bernadette Dorizzi. 2018. C-FMCW Based Contactless Respiration Detection Using Acoustic Signal. 1, 4, Article 170 (jan 2018), 20 pages. <https://doi.org/10.1145/3161188>
 - [66] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. 2015. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853* (2015).
 - [67] Koji Yatani and Khai N Truong. 2012. Bodyscope: a wearable acoustic sensor for activity recognition. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. 341–350.
 - [68] Xu Ye, Guanling Chen, Yang Gao, Honghao Wang, and Yu Cao. 2016. Assisting food journaling with automatic eating detection. In *Proceedings of the 2016 CHI conference extended abstracts on human factors in computing systems*. 3255–3262.
 - [69] Tianhong Catherine Yu, Guilin Hu, Ruidong Zhang, Hyunchul Lim, Saif Mahmud, Chi-Jung Lee, Ke Li, Devansh Agarwal, Shuyang Nie, Jinseok Oh, et al. 2024. Ring-a-Pose: A Ring for Continuous Hand Pose Tracking. *arXiv preprint arXiv:2404.12980* (2024).
 - [70] Shu-qian Yuan, Ying-ming Liu, Wei Liang, Fei-fei Li, Yuan Zeng, Yin-yue Liu, Shu-zhen Huang, Quan-yuan He, Binh Quach, Jiao Jiao, Julien S. Baker, and Yi-de Yang. 2021. Association Between Eating Speed and Metabolic Syndrome: A Systematic Review and Meta-Analysis. *Frontiers in Nutrition* 8 (Oct. 2021). <https://doi.org/10.3389/fnut.2021.700936> Publisher: Frontiers.
 - [71] Jiahao Zhang, Stephen Gould, and Itzik Ben-Shabat. 2020. Vidat—ANU CVML Video Annotation Tool. <https://github.com/anucvml/vidat>.
 - [72] Rui Zhang and Oliver Amft. 2017. Monitoring chewing and eating in free-living using smart eyeglasses. *IEEE journal of biomedical and health informatics* 22, 1 (2017), 23–32.
 - [73] Ruidong Zhang, Hao Chen, Devansh Agarwal, Richard Jin, Ke Li, François Guimbretière, and Cheng Zhang. 2023. HPSpeech: Silent Speech Interface for Commodity Headphones. In *Proceedings of the 2023 ACM International Symposium on Wearable Computers*. 60–65.
 - [74] Ruidong Zhang, Ke Li, Yihong Hao, Yufan Wang, Zhengnan Lai, François Guimbretière, and Cheng Zhang. 2023. EchoSpeech: Continuous Silent Speech Recognition on Minimally-Obtrusive Eyewear Powered by Acoustic Sensing. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (*CHI '23*). Association for Computing Machinery, New York, NY, USA, Article 852, 18 pages. <https://doi.org/10.1145/3544548.3580801>
 - [75] Ruidong Zhang, Jihai Zhang, Nitish Gade, Peng Cao, Seyun Kim, Junchi Yan, and Cheng Zhang. 2022. EatingTrak: Detecting Fine-Grained Eating Moments in the Wild Using a Wrist-Mounted IMU. *Proc. ACM Hum.-Comput. Interact.* 6, MHC1, Article 214 (sep 2022), 22 pages. <https://doi.org/10.1145/3546749>
 - [76] Shibo Zhang, Yuqi Zhao, Dzong 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 (jun 2020), 26 pages. <https://doi.org/10.1145/3397313>