

# Tiny Eats: Eating Detection on a Microcontroller

Maria T. Nyamukuru

*Thayer School of Engineering**Dartmouth College*

Hanover, NH, USA

Maria.T.Nyamukuru.th@dartmouth.edu

Kofi M. Odame

*Thayer School of Engineering**Dartmouth College*

Hanover, NH, USA

Kofi.M.Odame@dartmouth.edu

**Abstract**—There is a growing interest in low power highly efficient wearable devices for automatic dietary monitoring (ADM) [1]. The success of deep neural networks in audio event classification problems makes them ideal for this task. Deep neural networks are, however, not only computationally intensive and energy inefficient but also require a large amount of memory. To address these challenges, we propose a shallow gated recurrent unit (GRU) architecture suitable for resource-constrained applications. This paper describes the implementation of the Tiny Eats GRU, a shallow GRU neural network, on a low power microcontroller, Arm Cortex M0+, to classify eating episodes. Tiny Eats GRU is a hybrid of the traditional GRU [2] and eGRU [3] which makes it small and fast enough to fit on the Arm Cortex M0+ with comparable accuracy to the traditional GRU. The Tiny Eats GRU utilizes only 4% of the Arm Cortex M0+ memory and identifies eating or non-eating episodes with 6 ms latency and accuracy of 95.15%.

**Index Terms**—gated recurrent unit, neural networks, eating detection, automatic dietary monitoring, activity recognition, micro-controller, eating episodes, low power wearable devices

## I. INTRODUCTION

Poor eating habits have been linked to an increased risk of cardiovascular disease, cancer, obesity, diabetes, cataracts, and possibly dementia [4]. As such, there is growing interest in recording the food intake of at-risk individuals and exploring diet interventions to prevent long-term health problems [4]. Today's gold standard for recording food intake is the self-administered food diary. This approach depends on an individual's meticulous attention to their daily diet. As such, it is prone to human error and lacks an objective validation check for the self-reported data [5].

To address these problems, several automated alternatives to the food diary have been proposed. For instance, radio-frequency identification (RFID) equipped devices like the Dietary Intake Monitoring System (DIMS) were introduced in [6], but they cannot differentiate between users. Video-based systems are user-specific, but they are bulky, intrusive and require extensive memory and computational resources [7], [8]. Wearable devices [9], [10] are another promising approach, but the current wearable implementations have limited computational power, which requires them to transmit raw data for offline processing wirelessly; this continuous wireless streaming severely limits the wearable devices' battery life. In

Supported by National Science Foundation under award numbers CNS-1565269 and CNS-1565268. Special thanks to Shengjie Bi and Professor David Kotz at Dartmouth College for their support on the Auracle Project.

this paper, we introduce a power-efficient shallow GRU neural network for detecting eating episodes that can be implemented directly on a wearable device. Our eating detection model requires no offline processing or wireless communication, thus dramatically improving battery life and helping to make wearable devices a more effective alternative to the food diary.

## II. EATING DETECTION OVERVIEW

Our system for detecting eating episodes comprises a contact microphone that senses jaw movement [11] and a micro-controller that analyzes the microphone output for chewing sounds. The algorithm on the micro-controller is based on the gated recurrent unit (GRU) neural network [2], which has been applied successfully to various audio event detection problems [12], [13]. For our algorithm, we modified the GRU activation functions, quantized the weights, and implemented all computations as integer operations. These modifications allow our model to fit within the memory, computational, and latency constraints of low-power embedded processors.

We start by providing details on the neural network architecture in Section III. Next, Section IV shows the results of the model implementation and then a discussion in Section V followed by future work in Section VI. We then summarize the paper in Section VII.

## III. EMBEDDED NEURAL NETWORK DETAILS

### A. Arm Cortex M0+

We implemented our algorithm on the Arm Cortex M0+ because it is the most energy-efficient Arm processor available for constrained embedded applications [14], which makes it ideal for use in small-sized power conservative wearable devices. The challenge of using such a low-power processor is that it has no dedicated floating-point unit, operates at 48 MHz speed, and it has only 32 KB RAM and 256 KB flash memory. These limitations restrict the size of the neural network that can be implemented and the size of data that can be computed at a given time.

### B. GRU Modifications for Arm Cortex M0+

In order to meet the resource constraints of the Arm Cortex M0+, we modify the traditional GRU cell [2] and design the Tiny Eats GRU described by:

$$\mathbf{r}_t = \frac{\varsigma(\mathbf{W}_r[\mathbf{x}, \mathbf{h}_{t-1}]) + 1}{2} \quad (1)$$

$$\mathbf{z}_t = \frac{\varsigma(\mathbf{W}_z[\mathbf{x}, \mathbf{h}_{t-1}]) + 1}{2} \quad (2)$$

$$\tilde{\mathbf{h}}_t = \varsigma(\mathbf{W}_h[\mathbf{x}, (\mathbf{r} \odot \mathbf{h}_{t-1})]) \quad (3)$$

$$\mathbf{h}_t = \mathbf{z}_t \mathbf{h}_{t-1} + (1 - \mathbf{z}_t) \tilde{\mathbf{h}}_t, \quad (4)$$

where  $\mathbf{x}$  represents the input features,  $\mathbf{W}$  represents the 8 bit quantized weights,  $\mathbf{r}$  represents the reset gate,  $\mathbf{z}$  represents the update gate,  $\tilde{\mathbf{h}}$  represents the candidate state,  $\mathbf{h}$  represents the new cell state, and  $\odot$  represents the element-wise multiplication of two vectors.

The Tiny Eats GRU cell diverges from the traditional GRU cell in that it uses a shifted soft-sign ( $\varsigma$ ) activation function in place of sigmoid activation function for the update and reset gates. A regular soft-sign is also used in place of tanh activation functions for candidate state as proposed by eGRU cell [3]. The sigmoid and tanh activation functions are memory- and resource-intensive since they require iterative accesses of look-up tables. The soft-sign is a more computationally-conservative alternative compared to the sigmoid and tanh. This computation is particularly crucial because the Arm Cortex M0+ lacks a floating-point unit (FPU) and Digital Signal Processing (DSP) instruction set. The Tiny Eats GRU cell diverges from the eGRU cell [3] in that it maintains the use of a reset gate as proposed by the traditional GRU cell. This helps extract the portion of the history that is relevant to the current computation, as highlighted in equation (3).

### C. Embedded Neural Network

The universal approximation theorem states that simple neural networks can represent a wide variety of interesting functions when given appropriate parameters. Instead of using a standard deep neural network, we use a shallow neural network with fewer layers and fewer neurons to represent the differences in STFT parameters of *eating* and *non-eating* data. The shallow neural network is computationally cheaper and conserves memory, which is necessary for implementation on the Arm Cortex M0+. In this paper, we propose the shallow neural network architecture shown in Fig. 1 that consists of 1 output layer and only 3 hidden layers; 2 gated recurrent units designed as explained in Section III-B, each with 16 neurons, and one fully connected layer with 8 neurons. The fully connected layer consists of a linear pre-activation function and a soft-sign activation. The output layer consists of a linear pre-activation function and a soft-max activation function. This design is very similar to the eGRU suggested by *Amoh and Odame* [3].

## IV. EXPERIMENTS AND RESULTS

### A. Eating Detection Dataset

The data used for this study was previously collected in a controlled laboratory setting from 20 participants using a contact microphone placed behind the participant's ear [15]. The contact microphone was connected to a data acquisition device (DAQ) that collected the activity data in real-time at a 20KHz sampling rate with 24-bit resolution. The contact microphone, used for acoustic sensing, is placed at the tip

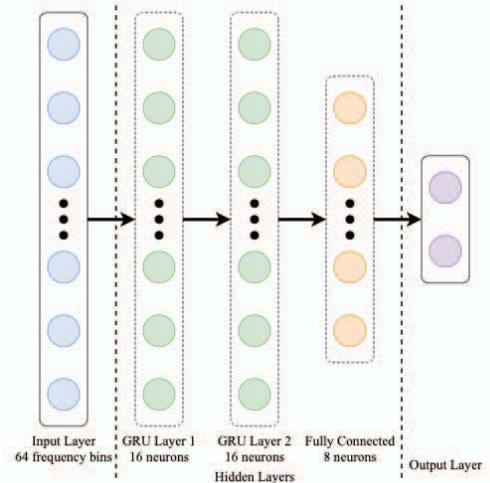


Fig. 1. Gated Recurrent Unit Neural Network Architecture

of the mastoid bone and secured using a headband to ensure contact with the body. The participants were instructed to perform an activity that involves both *eating* and *non-eating*. *Eating* activities included eating carrots, protein bars, crackers, canned fruit, instant food, and yogurt, sequentially for 2 minutes per food type. This resulted in a 4 hour total *eating* dataset. *Non-eating* activities included talking and silence for 5 minutes each and then coughing, laughing, drinking water, sniffling, and deep breathing for 24 seconds each. This resulted in a 7 hour total *non-eating* dataset. Each activity occurred separately and was classified based on activity type as *eating* or *non-eating*.

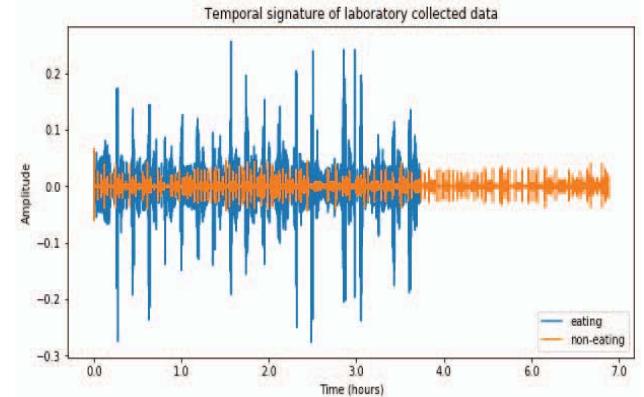


Fig. 2. Temporal signature of down-sampled and filtered laboratory collected raw data

### B. Data Pre-Processing

In this paper, we down-sampled the raw sound data to 500 Hz to reduce power consumption and memory usage and applied a high pass filter with a 20Hz cutoff frequency to attenuate noise [15] as shown in Fig. 2. We segmented the positive class data (*eating data*), and negative class data

(*non-eating data*) into 4-second windows, and computed a Short Time Fourier Transform (STFT), with 128 Fast Fourier Transform bins and no overlap, over the segments to view and extract features as shown in Fig. 3 and Fig. 4.

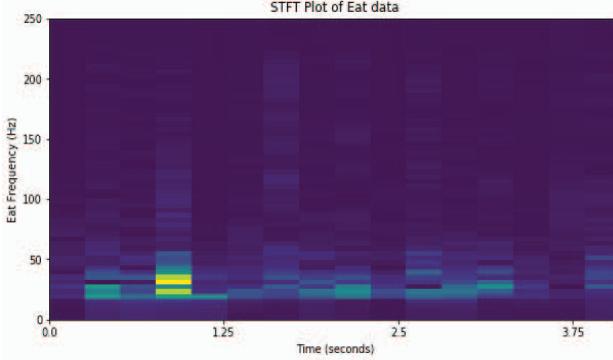


Fig. 3. STFT Results of Eating data

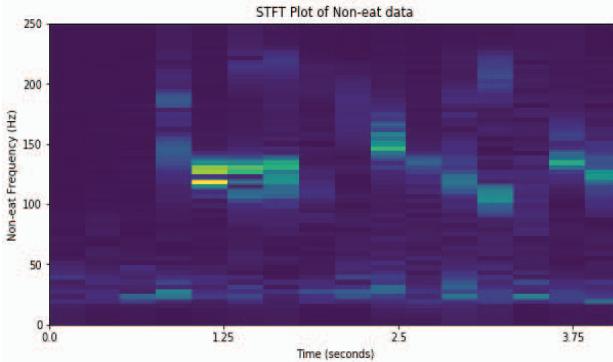


Fig. 4. STFT Results of Non-Eating data

The positive class data, as shown in Fig. 3 has a lower frequency and periodic envelope distribution in comparison to the negative class data shown in Fig. 4 that has a higher frequency and more compact distributions. These STFT features were used to train the neural network with each 4-second window being classified as an *eating* episode or a *non-eating* episode.

### C. Neural Network Training

The neural network shown in Fig. 1 is first programmed and trained in Python using TensorFlow Keras v2.0. This model uses floating-point. The training dataset contains more representation of *non-eating* data than it does *eating* data as described in Section IV-A. This imbalance inherently creates a bias in the training algorithm. In order to debias the algorithm, we normalize the binary cross-entropy loss function with each class's representative weight during training. This ensures that both classes are represented fairly.

In order to guarantee that the best performing model is used, a model checkpoint is implemented to save and retrieve the model with the highest validation accuracy and minimal

validation loss. The results of this FPU model are as shown in Table I and Fig. 5.

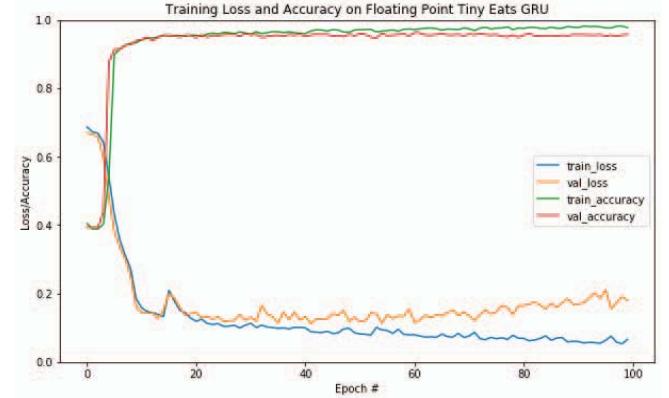


Fig. 5. Loss and Accuracy Training Graphs for Floating Point Tiny Eats GRU

An integer quantization based model [3] that uses 8-bit quantization for all weights is then implemented in Pytorch and trained, as shown in Fig. 6. The weight normalization applied to the cross-entropy loss function, and the model checkpointing used in the FPU model is implemented here as well.

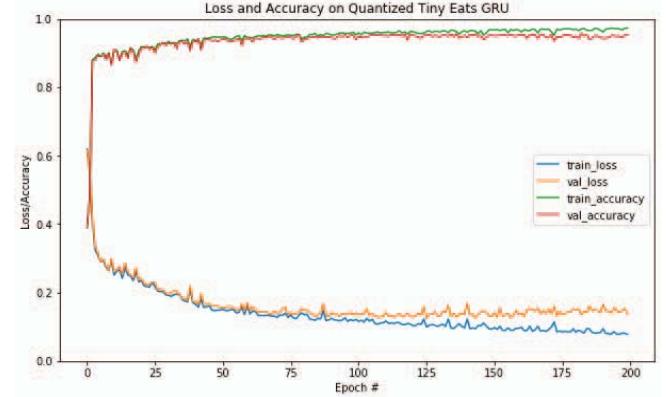


Fig. 6. Quantized (8 bit) Tiny Eats GRU Loss and Accuracy Training Graphs

The cross-validation accuracy and loss from both models are compared, as shown in Table I.

TABLE I  
COMPARISON OF FPU AND QUANTIZED TINY EATS GRU CROSS  
VALIDATION ACCURACY

Model Metrics	Model Designs	
	FPU	Quantized
Cross-Validation Accuracy	96.13%	94.41%
Cross-Validation Loss	0.12	0.13
Epochs Used	100	200

The quantized model is evaluated on an independent test set, and the test accuracy, loss, precision, recall, and F1-score are reported, as shown in Table II. The precision score represents

TABLE II  
EVALUATION OF TINY EATS GRU ON INDEPENDENT TEST SET VS OTHER MODELS

Model Measure	Models Used				
	Tiny Eats	Model in [15]	Ear Bit [16]	Splendid [17]	EGG [18]
Model Test Acc.	GRU	LR	RF	SVM	FFNN
95.15%	91.5%	90.1%	93.8%	90.1%	
F1-Score	94.89%	NA	90.9%	76.1%	NA
Precision	94.68%	95.1%	86.2%	79.4%	88.4%
Recall	95.12%	87.4%	96.1%	80.7%	91.8%

the ratio of correctly classified positive observations of the total number of observations classified as positive. The recall score (sensitivity) represents the ratio of correctly classified positive observations to the total number of true positive observations, as shown in equation (5) and equation (6) where  $TP$  is True Positive,  $FN$  is false negative, and  $FP$  is false positive.

$$Precision = \frac{TP}{TP + FP} \quad (5)$$

$$Recall = \frac{TP}{TP + FN} \quad (6)$$

The trained integer quantized model is then implemented on the Arm Cortex M0+.

#### D. Arm Cortex M0+ Implementation

We perform the feature extraction on the Arm Cortex M0+ by implementing a Fast Fourier Transform which utilizes only 9% (23 KB) of the memory. The features generated from the Arm FFT are then transmitted to the implemented Tiny Eats GRU. The implementation of the GRU on the Arm Cortex M0+ takes 6 ms to execute one sample and utilizes only 4% (12 KB) of the memory.

The accuracy of the GRU implementation is verified by transmitting test data via serial communication and reading the output label classification. This was done with a subset of the test data set, which recorded a label match between the Pytorch quantized model and the Arm Cortex M0+ implementation, as shown in Fig. 7.

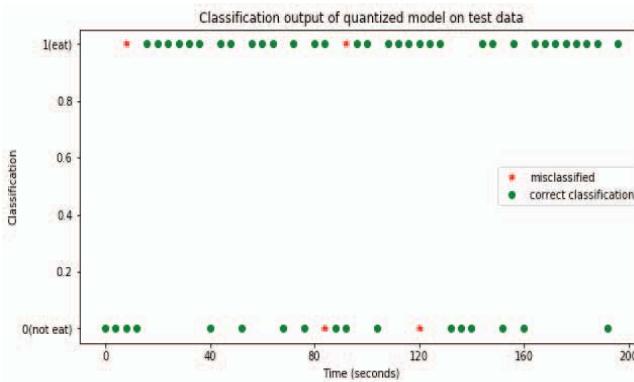


Fig. 7. Quantized (8bit) Tiny Eats GRU Predicted label vs true label

## V. DISCUSSION

In order to account for the limitation in the size of the neural network used, advanced feature engineering is paramount. The Tiny Eats GRU utilizes STFT features shown in Section IV-A as the only class of input features. This is crucial because it can be implemented on the Arm Cortex M0+ at a low cost, as shown in Section IV-D, while ensuring significant accuracy is achieved. The efficient implementation of the STFT on the Arm Cortex M0+ reduces the burden on the neural network. It consequentially allows for a reduction in the number of layers and the number of neurons required for classification. This minimization of neural network parameters is resource-efficient and allows for implementation on resource-constrained micro-controllers while achieving high accuracy.

In order to reduce memory usage and run time execution performance, we utilize quantized weights instead of floating-point weights and fixed-point arithmetic. Using 8-bit quantized weights reduces the memory footprint of the network, and when coupled with fixed-point arithmetic, reduce the power consumption and increases computational speed, with minimal impact on the accuracy of the overall network, making it efficient for deployment on embedded hardware [19], [20]. As shown in Table I, there is only a 2% decrease in accuracy by switching to 8-bit integer quantization from floating-point. 8-bit weight quantizations also ensure that the Tiny Eats GRU is small enough to fit on the Arm Cortex M0+ and run efficiently, as seen in Section IV-D.

In order to increase computational speed, the Tiny Eats GRU uses soft-sign activation functions, as explained in Section III-B. The results shown in Section IV-C affirm that the proposed Tiny Eats GRU architecture is a viable replacement for the traditional GRU [2] in eating episode classification. The GRU network with fast soft-sign activation functions achieves a cross-validation accuracy of 96.13% and 94.41% for floating-point and quantized weight implementation, respectively, as shown in Table I. This soft-sign implementation achieves high accuracy while minimizing computation time.

The proposed Tiny Eats GRU's performance outperforms some of the recently proposed approaches to eating detection, as shown in Table II. The accuracy, precision, recall, and F1-score of the Ear Bit Sensor [16] that uses a random forest (RF) model, the SPLENDID device [17] that uses support vector machine (SVM) classification, and the electroglottography (EGG) device [18] that utilizes feedforward neural networks (FFNN), are compared to that of the Tiny Eats GRU. The Tiny Eats GRU has better overall performance on all evaluation metrics, recording the highest F1-score of 94.89%, as shown in Table II. These results are comparable to the model design in [15] that uses linear regression (LR).

The Tiny Eats GRU is extremely useful in resource-constrained environments like the Arm Cortex M0+, which can be used in low power wearable devices. It operates with only 6 ms latency, which is significantly smaller than the 250 ms FFT feature extraction time and uses only 4% of the M0+ memory, which allows for the execution of other processes

that could include implementing more features.

## VI. FUTURE WORK

In the future, we can further improve the accuracy by increasing the number of feature classes to represent more relationships besides the STFT. Feature classes like the zero-crossing rate and number of peaks have defining characteristics that are essential in classifying *eating* and *non-eating* data. This would also imply that we could reduce the number of neurons required in the Tiny Eats GRU, and consequentially reducing its computation time and memory requirements, making it even more efficient.

Furthermore, we will implement and train this algorithm on data collected outside the lab to infer its performance in an unconstrained, free-living environment. We will also explore classifying different food types and inferring the volume of food eaten during an eating session to gain more insight into the user's diet.

## VII. CONCLUSION

This paper proposed the Tiny Eats GRU architecture, an implementation on Arm Cortex M0+ for the detection of eating episodes. The Tiny Eats GRU was demonstrated to have significantly low classification costs by being memory-conservative, computationally efficient with significant accuracy results while consuming minimal power. This allows for the implementation of Tiny Eats GRU on severely resource-constrained devices incapable of handling traditional deep neural networks, that can be used in wearable devices.

## ACKNOWLEDGMENT

This work was supported by the National Science Foundation under award numbers CNS-1565269 and CNS-1565268. We would like to thank Shengjie Bi and Professor David Kotz at Dartmouth College for their support on the Auracle Project. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the sponsors.

## REFERENCES

- [1] T. Prioleau, E. Moore, and M. Ghovanloo, "Unobtrusive and wearable systems for automatic dietary monitoring," *IEEE Transactions on Biomedical Engineering*, vol. PP, pp. 1–1, 01 2017.
- [2] K. Cho, B. van Merriënboer, C. Gulcehre, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," in *Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*, 2014.
- [3] J. Amoh and K. M. Odame, "An optimized recurrent unit for ultra-low-power keyword spotting," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 3, no. 2, Jun. 2019. [Online]. Available: <https://doi.org/10.1145/3328907>
- [4] I. H. Rosenberg, "Nutrition research: An investment in the nation's health," *Nutrition Reviews*, vol. 54, no. 1, pp. S5–S6, Jan. 2009. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pubmed/8919670>
- [5] L. Burke and J. Wang, "Self-monitoring in weight loss: A systematic review of the literature," *Journal of the American Dietetic Association*, vol. 111, pp. 92–102, 01 2011.
- [6] K. Ofei, M. Dobroczynski, M. Holst, H. Rasmussen, and B. Mikkelsen, *The Dietary Intake Monitoring System (DIMS): an Innovative Device for Capturing Patient's Food Choice, Food Intake and Plate Waste in a Hospital Setting*, ser. Proceedings of Measuring Behavior 2014. Proceedings of Measuring Behavior 2014, Wageningen, 8 2014, vol. 60.
- [7] Wen Wu and Jie Yang, "Fast food recognition from videos of eating for calorie estimation," in *2009 IEEE International Conference on Multimedia and Expo*, June 2009, pp. 1210–1213.
- [8] S. Cadavid, M. Abdel-Mottaleb, and A. Helal, "Exploiting visual quasi-periodicity for real-time chewing event detection using active appearance models and support vector machines," *Personal and Ubiquitous Computing*, vol. 16, no. 6, pp. 729–739, Aug 2012. [Online]. Available: <https://doi.org/10.1007/s00779-011-0425-x>
- [9] Y. Zheng, X. Ding, C. Poon, B. Lo, H. Zhang, X. Zhou, G.-Z. Yang, N. Zhao, and Y.-T. Zhang, "Unobtrusive sensing and wearable devices for health informatics," *IEEE Transactions on Biomedical Engineering*, vol. 61, pp. 1538–1554, 03 2014.
- [10] O. D. Lara and M. A. Labrador, "A survey on human activity recognition using wearable sensors," *IEEE Communications Surveys Tutorials*, vol. 15, no. 3, pp. 1192–1209, Third 2013.
- [11] S. Bi, T. Wang, N. Tobias, J. Nordrum, S. Wang, G. Halvorsen, S. Sen, R. Peterson, K. Odame, K. Caine, and et al., "Auracle: Detecting eating episodes with an ear-mounted sensor," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 2, no. 3, Sep. 2018. [Online]. Available: <https://doi.org/10.1145/3264902>
- [12] H.-G. Kim and J. Y. Kim, "Acoustic event detection in multichannel audio using gated recurrent neural networks with high-resolution spectral features," *ETRI Journal*, vol. 39, no. 6, pp. 832–840, 2017. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.4218/etrij.17.0117.0157>
- [13] A. Kusupati, M. Singh, K. Bhatia, A. Kumar, P. Jain, and M. Varma, "Fastgrnn: A fast, accurate, stable and tiny kilobyte sized gated recurrent neural network," in *NIPS 2018*, December 2018.
- [14] H. Thurnburg. (2020, Jan.) Cortex-m0+. [Online]. Available: <https://developer.arm.com/ip-products/processors/cortex-m/cortex-m0-plus>
- [15] S. Bi, T. Wang, E. Davenport, R. Peterson, R. Halter, J. Sorber, and D. Kotz, "Toward a wearable sensor for eating detection," in *Proceedings of the 2017 Workshop on Wearable Systems and Applications*, ser. WearSys '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 17–22. [Online]. Available: <https://doi.org/10.1145/3089351.3089355>
- [16] A. Bedri, R. Li, M. Haynes, R. P. Kosaraju, I. Grover, T. Prioleau, M. Y. Beh, M. Goel, T. Starner, and G. Abowd, "Earbit: Using wearable sensors to detect eating episodes in unconstrained environments," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 1, no. 3, Sep. 2017. [Online]. Available: <https://doi.org/10.1145/3130902>
- [17] V. Papapanagiotou, C. Diou, L. Zhou, J. van den Boer, M. Mars, and A. Delopoulos, "A novel chewing detection system based on ppg, audio, and accelerometry," *IEEE Journal of Biomedical and Health Informatics*, vol. 21, no. 3, pp. 607–618, May 2017.
- [18] M. Farooq, J. Fontana, and E. Sazonov, "A novel approach for food intake detection using electroglottography," *Physiological measurement*, vol. 35, pp. 739–751, 03 2014.
- [19] S. Han, H. Mao, and W. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding," in *ICLR*, 10 2016.
- [20] V. Vanhoucke, A. Senior, and M. Mao, "Improving the speed of neural networks on cpus," in *Computer Science*, 01 2011.