

Automatic Recognition of Food Ingestion Environment from the AIM-2 Wearable Sensor

Yuning Huang,¹ M A Hassan,² Jiangpeng He,¹ J. Higgins,³ Megan McCrory,⁴
Heather Eicher-Miller,¹ J. Graham Thomas,⁵ Edward Sazonov,⁶ Fengqing Zhu¹

¹ Purdue University, West Lafayette, Indiana, U.S.A ² University of California, Davis, Davis, California, U.S.A

³ University of Colorado, Aurora, Colorado, U.S.A ⁴ Boston University, Boston, Massachusetts, U.S.A

⁵ Brown University, Providence, Rhode Island, U.S.A ⁶ The University of Alabama, Tuscaloosa, Alabama, U.S.A

Abstract

Detecting an ingestion environment is an important aspect of monitoring dietary intake. It provides insightful information for dietary assessment. However, it is a challenging problem where human-based reviewing can be tedious, and algorithm-based review suffers from data imbalance and perceptual aliasing problems. To address these issues, we propose a neural network-based method with a two-stage training framework that tactfully combines fine-tuning and transfer learning techniques. Our method is evaluated on a newly collected dataset called “UA Free Living Study”, which uses an egocentric wearable camera, AIM-2 sensor, to simulate food consumption in free-living conditions. The proposed training framework is applied to common neural network backbones, combined with approaches in the general imbalanced classification field. Experimental results on the collected dataset show that our proposed method for automatic ingestion environment recognition successfully addresses the challenging data imbalance problem in the dataset and achieves a promising overall classification accuracy of 96.63%.

1. Introduction

Most recent dietary assessment research mainly focuses on monitoring the intake of energy and nutrients in individuals’ diets [17, 22, 44, 47]. However, understanding ingestion behavior, including the impact of the environment and social context, is a relatively new and under-explored aspect of ingestion monitoring. The ingestion environment (see Fig. 1) impacts the dietary behavior of an individual through a range of sensory mechanisms related to food intake [50]. For instance, the environment could influence the quantity of food intake, an individual may consume more foods in one environment compared to another. Likewise, the physical posture during food intake may also vary depending on



Figure 1. Compilation of images showcasing various environments where food may be consumed. The montage is created using free-living data collected from AIM-2.

the ingestion environment. An individual is more likely to sit at a table in a restaurant but may sit at a table, lie on a bed, or sit on a sofa while consuming food at home [3].

The environment can play a vital role in food acceptance. Meiselman et al. [41] conduct a study demonstrating the impact of the ingestion environment on food acceptance. The participants were asked to rate their acceptance of the meal by completing a hedonic rating scale. Results show that the meal ratings varied with the environment, *e.g.*, meals are most liked in the restaurant and least liked in the cafeteria. A more recent study [18] examined how location and table settings affect people’s willingness to eat food revealed that participants were substantially more likely to consume food presented on a gourmet table (which is typical in restaurants) than on a home-style table and plastic tray.

The ingestion environment can also influence the type of food consumed. Bauer et al. [1] have shown that food choices made out-of-home are often less consistent with

people’s dietary plans, and eating out-of-home may be one of the main reasons for the failure of dietary goals. Claessens et al. [6] also indicate that people’s eating choices in restaurants are typically healthier and less sustainable than at home; their survey has shown that healthiness is the most important consideration for choosing home meals while being the third most important factor for choosing restaurant meals. Another recent study [3] investigated the distribution of eating locations for breakfast, lunch, dinner, and snacks and showed that the study participants tend to have more snacks in vehicles.

Compared to other aspects of measurement such as eating activity detection and dietary composition/energy estimation, ingestion environment recognition is still an under-explored area. A relevant study by Gemming et al. [19] used a wearable camera to capture and categorize environmental and social contexts to understand food intake behavior. The authors used the SenseCam wearable camera [53] to capture eating episodes by monitoring activity throughout the day. The images of the eating episode were manually annotated by two researchers, marking the eating location, external environment, physical position, social interaction, and viewing of media screens. The study reported that the duration of food intake varied with different food locations, such as home, workplace, and restaurant whereas the longest eating episodes tended to occur in restaurants and the shortest in the workplace.

An analysis of many previous studies is conducted by manually reviewing self-report questionnaires and images, however, this may not be an optimal solution. Self-reported data are subject to misreporting. A study from Thea and Mortel [53] reported that participants tend to present a favorable image of themselves when completing questionnaires that elicit an evaluative response. In addition, manual review of dietary images [45, 46] is time-consuming and subject to human error. Furthermore, egocentric wearable cameras such as eButton [51] and SensCam operate throughout the day, capturing 4,000 to 6,000 images, while only less than 20% of the images are relevant to eating events. For all of these egocentric cameras, an automatic method for ingestion environment recognition is needed to facilitate further studies on the influence of the environment on dietary activities.

To successfully design a framework for the eating environment recognition, we need to first define the major challenges within the task. In comparison to the standard environment recognition problem, eating environment recognition has majorly suffered from limited publicly available datasets, perceptual aliasing, and most importantly, severe data imbalance. Although there are more public datasets for food-related tasks, they seldom provide the environment ground truth labels while food environment-related datasets are hard to get published due to privacy concerns with the



Figure 2. AIM-2, an egocentric wearable camera that monitors ingestion behavior.

egocentric view. In addition, images of different environments may look similar [15], such as the dining room of a house and the table setting of a restaurant. Also, the eating scene’s distribution may be imbalanced [3], with most consumption taking place at home and less consumption at other locations. Furthermore, the first two issues could exacerbate the impact of data imbalance because the limited dataset size and significant inter-class similarity impose a greater burden on the training of a classifier.

A recent study of automatic eating environment classification [40] has proposed a VGG network-based hierarchical method to automatically classify 15 food-related scenes with an overall accuracy of 56%. However, their method still requires a manual selection of eating episodes from recorded day photo streams and fails to address the challenges mentioned above.

In this paper, we propose to address the recognition of the ingestion environment by automatically detecting eating episodes through the use of a food intake sensor, AIM-2 [14], and the associated neural network-based classification framework that specifically targets the aforementioned challenges.

Our contributions are summarized as follows:

1. We identify major challenges in ingestion environment recognition and propose a simple, general, and effective deep-learning-based framework for addressing them.
2. We collect and analyze a dataset, *UA Free Living Study*, comprised of unrestricted ad-libitum food consumption for automatic recognition of ingestion environment using an egocentric wearable sensor.
3. We conduct various experiments on the collected dataset to verify the performance of the proposed method and show the advantage over generic techniques for imbalanced classification.

2. Related works

2.1. Environment Recognition

Multiple computer vision/early deep learning algorithms related to environment recognition were reviewed in [38]. Before the era of deep learning, researchers in traditional

computer vision and robotics attempted to address the issue of perceptual aliasing by using local and global feature descriptors. The local descriptors operate as a two-stage process of feature extraction and recognition. Feature descriptors such as SIFT [37] (scale-invariant feature transform) and SURF [2] (speeded-up robust features) are used to extract features and detect objects in an image. Similarly, global feature descriptor-based methods also operate as a two-stage process. Feature descriptors such as color histograms, descriptor-based PCA [29] (principal component analysis), and histogram of oriented gradients [11] methods are used to process the entire image and capture edges, corners, and color patches. The features are processed using machine learning methods such as SVM [7] (support vector machine). However, with the progress of network design and training, neural network-based environment recognition has achieved better accuracy than traditional methods since it is more powerful at effectively extracting classification-related features. Multiple network architectures such as VGG16 [48], ResNet [25], ConvNeXt [36] and Vision Transformer [12, 34] may be used to address environment recognition and classification problem. Among these architectures, the transformer-based methods have achieved the best overall performance. In this work, we select three representative network architectures as the backbone of the proposed framework, which are ResNet, ConvNeXt, and Swin transformer.

2.2. Imbalanced Classification

Imbalanced data is a challenging issue for classification where the number of samples belonging to different classes is largely different [4, 23, 54]. This problem, if not carefully addressed, often leads to unsatisfactory prediction accuracy on test data, especially for minority classes (classes that contain fewer samples). There are many general approaches proposed for training deep networks on an imbalanced dataset. Class re-balancing is a major paradigm in imbalanced learning and there are two widely-used effective categories of methods that belong to this paradigm, one is resampling [5, 21, 24, 33, 55] and another is weighted loss function [8, 32, 39, 52]. The resampling strategy is basically to re-assign the probability of the sample from each class to be presented in the training mini-batch. Samples from the minority class will appear with a higher probability while samples from the majority class will appear with a lower probability. The weighted loss function is used to adjust the training loss values for different classes by assigning them different weights corresponding to the number of samples in the class. The weighted loss has larger gradients from minority class samples and thus encourages the model to be more adaptive for learning features from minority classes. There are also other efforts for solving this challenging problem such as information augmentation [30]

and improvement on network module design [26]. In this work, we adopt two most widely used approaches, random resampling and weighted cross-entropy loss. We conduct experiments to verify our proposed framework has achieved a better performance boost while also maintaining a positive interaction with both approaches.

3. Dataset

In this section, we introduce how we collect and construct the dataset we need to perform automatic recognition of eating events, such as meals and snacks.

3.1. Sensor System

The sensor system used for the method development is the Automatic Ingestion Monitor, version 2 (AIM-2) [14], a second-generation egocentric wearable (Fig. 2) for monitoring dietary intake and eating behavior.

3.2. Data Collection

We collected experimental data from thirty volunteer participants (65% males and 35% females, aged 18 to 39 years old). The University of Alabama institutional review board approved the study, and participants were compensated for their participation. The subjects represented four races, non-Hispanic, African American, Asian, and Hispanic. The experiment was conducted in two parts: a controlled laboratory experiment and a free-living experiment.

For this study, we used the data from the free-living experiment. The participants were asked to wear the AIM-2 sensor for the entire day, follow their normal daily activities, and have at least a single meal at a place of their choice. The participants wore the device for 8.5 to 15.75 hours. The participants were not limited to any social/personal interaction, activities (except for those they considered private or water-based activities), consumption of particular food types, or how the food type was consumed.

The participants were asked to self-report all eating events (both solids and liquids) using the ASA24 [31] (Automated Self-Administered 24) in food diary mode after completing the day of AIM-2 monitoring during free-living.

In total, the participants consumed 89 meals in four different environments: vehicle, home, restaurant, and workplace. We corrected the falsely reported self-assessment data by using the self-reporting data correction approach described by Giacchi et al. [20]. We performed the expert review for the entire dataset as our population was significantly smaller compared to the population reported in [20]. During the expert review, we reviewed each image sequence including images before, after, and during the eating episode to determine the actual ingestion environment. After this, we made corrections to the self-reported ingestion environment, see more details of the construction of the dataset in supplementary materials.




Figure 3. Description of the proposed method for automatic ingestion environment recognition.

Table 1. Sample distribution in the dataset

	Vehicle	Home	Restaurant	Workplace
Percentage at seq-level	2.2%	71.9%	11.3%	14.6%
Percentage at img-level	2.0%	67.0%	13.6%	17.4%

3.3. Statistics and Challenges

After labeling the eating environment, we constructed the dataset for training and testing the proposed method. However, as we discussed in Section 1, the ingestion environment recognition task contains several inherent issues that require specific considerations in designing methods for it.

First, the number of samples in each class is very imbalanced (see Table 1) with the ratio between the major class and minor class being 35. This severe data-imbalanced problem is an inherent issue as most people consume more meals at home and consume fewer meals in other places. The imbalance issue of the dataset introduces unfair favors of the classifier towards the majority class and ignorance toward the minority class.

Second, the total number of sequences and images collected in our dataset is limited, with only 89 sequences and 5,351 images. The size of the dataset makes the imbalance issue more difficult to handle since normal methods may introduce overfitting in training toward the minority class.

In addition, the dataset is egocentric, which is different than general classification datasets (e.g. ImageNet [10] and Place365 [56]). The egocentric sequence may bear great inter-class similarity due to their viewpoint. For example, when the participant is looking directly at the food, it may be hard to distinguish between “Home” and “Restaurant” since the scene captured may only contain a table with food on it with limited background information.

The aforementioned problems make our dataset even more challenging compared to the normal imbalanced datasets. In the experiments, we show that chosen general techniques for addressing data imbalanced classification may not directly work well for our dataset.

In Section 4, we introduce an automatic scene recognition method with a two-stage training framework that addresses the challenging issue in the proposed dataset.

4. Method

In this paper, we choose to perform neural network-based scene recognition. We select ResNet [25], ConvNeXt [35] and Swin transformer [34] as our representative network backbones. Note that our method is not restricted to the special design of any general classification architectures and thus can be easily adapted to other network backbones.

4.1. Overall Design

The overall architecture of the proposed automatic scene recognition method is shown in Fig. 3. We use sensor-captured signals to determine the start time and end time of eating episodes [13, 16] (more details in supplementary materials). Note that the network used in the figure is already trained and finetuned with the proposed framework.

As mentioned in Section 3.3, the collected dataset only contains a limited number of samples and has severe data imbalance issues. To address the challenges introduced by the dataset, we propose a two-stage drop-then-maintain training framework that tactfully adopts the techniques from the finetuning and transfer-learning field to achieve balanced training with more training samples.

We propose to drop the feature classifier (see definition in the next sub-section) in the first stage while finetuning the model on Places365 database [56] and maintain the feature classifier while finetuning the model on our dataset, which we referred to as the drop-then-maintain strategy. To enable the second stage of training, we need to additionally process the Place365 database using a semantic-based class filtering and merging.

It is worth noting that the drop strategy is well established and used in the transfer-learning field on the classification task while the maintain strategy is commonly used in other tasks like image restoration when finetuning is useful. However, our proposed drop-then-maintain strategy is less explored, and to the best of our knowledge, the first one to be proposed and applied to the environment recognition field which successfully addresses the inherent challenges mentioned before.

The strategy is designed based on our previous analysis of challenges. In addition, experiments verified its effective-




Figure 4. Two-stage drop-then-maintain training framework. After the ImageNet pretraining, the feature classifier g_1 is dropped and replaced by g_2 (since the number of classes has changed from 1,000 to 4), and g_2 can remain from stage 1 to stage 2 because of the semantic-based class filtering and merging we performed to Place365 dataset.

tiveness against the simple drop strategy or maintain strategy and showed that this tactful combination of the two is helpful and necessary.

4.2. Method Formulation

For any general classification network, we can partition them into two major functional parts: feature extractor and feature classifier, where the classifier is the last fully connected layer in the neural network that utilizes highly condensed information extracted by the feature extractor (all previous layers) to perform the classification. The feature extractor can be viewed as a function that maps from image space to feature space, and the feature classifier can be viewed as a function that maps from feature space to label space, which can be formulated as follows:

$$f : \mathbb{R}^{H \times W \times 3} \mapsto \mathbb{R}^C \quad (1)$$

$$g : \mathbb{R}^C \mapsto \mathbb{R}^N \quad (2)$$

where f denotes the feature extractor, g denotes the feature classifier that predicts class probabilities, H and W are the height and width of the input image, C is the feature dimension, and N is the number of classes. Using f and g , the classification network can be abstracted as a composed function that maps from image space to label space:

$$g \circ f : \mathbb{R}^{H \times W \times 3} \mapsto \mathbb{R}^N \quad (3)$$

In normal transfer learning for classification, the feature classifier g is often dropped and replaced by a new classifier that has the same output dimension as the number of target classes while only the feature extractor is transferred. This is a common practice in the field and has been used in [56] and [40].

However, we argue that this is sub-optimal for food environment classification since training new feature classifiers on these datasets suffers a strong bias towards the majority class (*i.e.*, “Home”). To train a powerful feature extractor and a robust feature classifier, we believe our proposed drop-then-maintain strategy is necessary.

In the next sub-section, we will go into the details of our proposed two-stage drop-then-maintain training framework that utilizes semantic-based class filtering and merging.

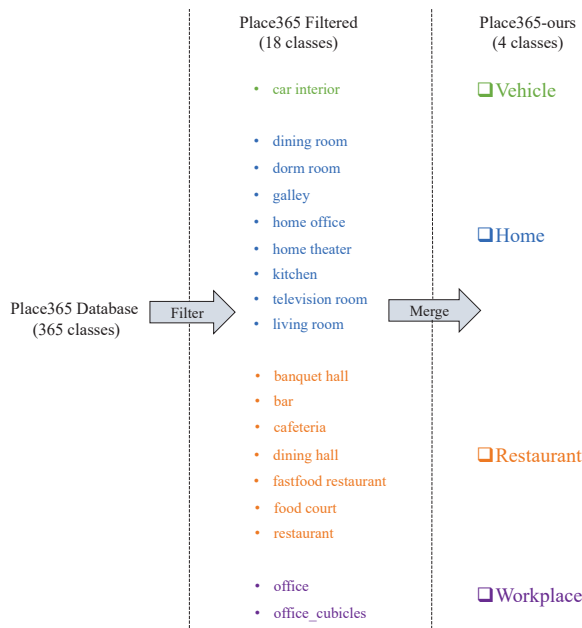


Figure 5. Semantic-based dataset filtering and merging.

4.3. Training Scheme

To address the data imbalance problem, a straightforward but effective approach is to increase the number of training data [28]. We adopt a two-stage training strategy to maximize the utilization of the public dataset to increase the numbers of data in the whole training process, see Fig. 4.

Dataset selection: We identify two public datasets that may benefit the performance of the model, one is ImageNet [9], and another is Place365 [56].

ImageNet contains 1,000 classes that are different from the four classes we are interested in. Due to its large volume, we consider using the pre-trained model on it since it helps to improve the ability of the feature extractor f . It has been shown that pre-training on ImageNet can generally fasten the converging speed of the classification model as well as improve its performance [27].

Besides, it is also worthwhile to utilize the Place365 dataset, which is targeted for general scene recognition

tasks. Since Place365 contains relevant classes to our target classes, we expect it helpful for improving the performance of feature classifier g . As mentioned before, we want to train g on Place365 and finetune it on our dataset without dropping and retraining a new classifier. However, the maintaining strategy requires the class match between two datasets and the Place365 dataset has many more classes than our dataset. This is where semantic-based class filtering and merging come into play, which ensures both a semantic and dimensional match between the processed Place365 dataset and our collected dataset.

Semantic-based class filtering and merging: To fulfill the semantic match and keep the output dimension consistent for the feature classifier g , we propose a semantic-based class filtering and merging strategy to pre-process the Place365 dataset. First, we identify the relevant classes in the dataset and then discard all the irrelevant classes. After dataset filtering, we merge the remaining classes into four groups in a semantic manner as described in Fig. 5. After filtering and merging, the Place365 dataset has the same number and semantic meaning of classes as our dataset, which we refer to as *Place365-ours*. The Place365-ours is not perfectly balanced but significantly better than our dataset since the ratio between major class and minor class drops from 35 to 8. Besides, the semantic-matched training samples can further increase the intra-class diversity to improve the generalizability of the model.

Two-stage drop-then-maintain training: To utilize the Imagenet and Place365-ours datasets to improve both the feature extractor f and feature classifier g . We first train the model on Place365-ours with an ImageNet pre-trained feature extractor f_1 along with a randomly initialized feature classifier g_2 , note that the ImageNet pre-trained feature classifier g_1 is dropped here. Then, we train the model on our dataset with both the pre-trained feature extractor and feature classifier obtained from the first stage training (see Fig. 4). Note that f_1 is transferred through both stages while g_1 is dropped out at the first stage and g_2 is maintained at the second stage. This setting maximizes the utilization of ImageNet and Place365-ours to train a powerful and representative feature extractor and a less biased robust feature classifier.

In both stages of training, we do not freeze the weights of any layers so that the update of the gradient is passed through all layers. We want to finetune both the feature extractor and feature classifier to address the potential domain shift issue between Imagenet, Place365-ours, and our dataset.

4.4. Approaches for Imbalanced Classification

We additionally include two approaches for imbalanced classification in our training stage and compare the performance boost in the experiment section.

Weighted loss function: We utilized the weighted cross-entropy loss in our model, where the weights are based on the inverse of the sample sizes of each class.

Resampling strategy: We apply the resampling strategy using both over-sampling and under-sampling to ensure the four classes have the same probability of appearing in one training mini-batch.

5. Experiment and Analysis

In this section, we report the results of our proposed method and compare it to three Baseline methods:

Baseline 0 (direct, or DR): Use ImageNet pre-trained model, directly finetune on our dataset.

Baseline 1 (weighted loss, or WL): Baseline 0 with the weighted loss used in finetuning.

Baseline 2 (resampling, or RS): Baseline 0 with the resampling technique used in finetuning.

Ours: Use the proposed two-stage training scheme, resampling and weighted loss are not applied.

Ours (WL): Use the proposed two-stage training scheme together with the weighted loss in the second stage.

Ours (RS): Use the proposed two-stage training scheme together with the resampling in the second stage.

By comparing our proposed method with Baselines 0, 1, and 2, we can verify if the proposed method can mitigate the data imbalance problem in the collected dataset and improve the overall performance of the classifier.

Due to the limited number of data samples in the dataset, we adopt the k-fold cross-validation [49] technique to utilize all the data for more reliable performance estimation. In our experiment, we empirically set k to 5 to balance the computation cost for each setting. See more details in the supplementary material.

5.1. Evaluation Metrics

Since the goal of our method is to recognize the environment associated with eating activities, we use sequence-level accuracy instead of frame-level accuracy. The top-1 prediction result for each frame in a meal sequence is aggregated using majority voting to select the class label for this sequence (See Fig. 6).




Figure 6. Process of generating the class label for each meal sequence, $f_n \in \mathbb{R}^{F \times H \times W \times 3}$ is the frames within one meal, $L_n \in \mathbb{R}^{F \times 1}$ is the top-1 prediction for each frame and $L_s^{Pred} \in \mathbb{R}^1$ is the final label for this sequence

Table 2. Main Result: Overall Performance

Architectures	Model Paras.	Method	Seq-level Acc	Macro Precision	Macro Recall	Macro F1
ResNet	42.51M	Baseline 0 (DR)	88.76	67.64	59.42	62.67
		Baseline 1 (WL)	88.76	60.90	64.01	62.04
		Baseline 2 (RS)	86.52	61.15	55.95	57.45
		Ours	89.89	90.06	72.30	78.00
		Ours (WL)	91.01	90.99	74.80	80.20
		Ours (RS)	92.13	91.58	76.72	81.37
ConvNeXt	87.58M	Baseline 0 (DR)	89.89	64.59	63.45	63.95
		Baseline 1 (WL)	87.64	61.65	60.56	60.76
		Baseline 2 (RS)	88.76	62.34	63.06	62.69
		Ours	92.13	93.69	73.08	80.17
		Ours (WL)	94.38	94.28	81.72	85.82
		Ours (RS)	94.38	92.14	81.72	84.54
Swin transformer	86.75M	Baseline 0 (DR)	91.01	65.35	65.37	65.27
		Baseline 1 (WL)	91.01	64.40	65.37	64.73
		Baseline 2 (RS)	91.01	68.66	63.84	65.91
		Ours	94.38	95.29	80.19	85.34
		Ours (WL)	96.63	95.31	84.61	87.57
		Ours (RS)	96.63	95.31	84.61	87.57

The **sequence-level accuracy (SLA)** is defined as (4)

$$SLA = \frac{S}{N} \quad (4)$$

S denotes the number of correctly classified sequences (meals), and N denotes the total number of meals.

For the evaluation of the proposed ingestion environment recognition method, we use imbalance accuracy metrics [42] such as macro-average precision, recall, and F1-score [43]. See definitions in the supplementary material.

5.2. Main Result

Overall performance: The overall performance of the baseline methods and our proposed methods are summarized in Table 2. We report the sequence-level accuracy (after majority voting in each sequence) and the macro-average metrics for three different network backbones. The best results for each network backbone are bolded.

As reported in Table 2, comparing the proposed two-stage training scheme (Ours) to Baselines 0, 1, and 2, it is noted that applying weight loss functions or resampling (Baseline 1 or Baseline 2) does not help improve the performance for normal finetuning (direct finetune on our dataset). In contrast, our method achieves better overall performance with better sequence-level accuracy and much better macro-average metrics. As we see later in Table 3, Baselines 1 and 2 cannot predict the “Vehicle” class correctly.

Due to severe data imbalance and limited training sample, simple re-weighting (WL/RS) can cause overfitting of limited “Vehicle” class data for the feature classifier g . However, by effectively incorporating more training

samples from the Place365-ours dataset utilizing our proposed drop-then-maintain strategy, the severe data imbalance problem is alleviated with more training samples and the improvement of intra-class diversity in minority classes. Results (Ours) on all network backbones get a significant performance boost in all evaluation metrics. It is worth noting that the macro-average metrics of the model show even more improvement.

Ours (WL) and Ours (RS) achieve better performance in all metrics compared to Baselines 1, 2, and Ours, which shows a positive interaction between both approaches and our proposed framework. The aforementioned overfitting problem is alleviated here since more relevant training samples are included for training a less biased feature classifier using Place365-ours, which is made possible by the proposed semantic-based class filtering and merging. This result further verifies our previous argument of the necessity of the two-stage drop-then-maintain training strategy.

Performance on the minority class: Since our dataset has a severe data imbalance problem where the least represented class “Vehicle” only contains 2 sequences, we want to further verify if our proposed method can help to improve the classification accuracy for it. We report the class accuracy for “Vehicle” in Tabel 3. We can see that all of the Baseline methods 0, 1, and 2 fail to predict the minority class “Vehicle”. Even though baseline methods 1 and 2 have utilized weighted loss and resampling techniques to mitigate the negative influence of lacking training samples in the minority class, they still suffer from the overfitting problem since the number of training samples is too small for the model to capture the general characteristic of data

Table 3. Main Result: Minority Class Accuracy

Architecture	Sequence-level Accuracy (%) for Vehicle					
	Baseline Method			Our Method		
	DR	WL	RS	Ours	+WL	+RS
ResNet	0.00	0.00	0.00	50.00	50.00	50.00
ConvNeXt	0.00	0.00	0.00	50.00	50.00	50.00
Swin transformer	0.00	0.00	0.00	50.00	50.00	50.00

Table 4. Dataset Utilization of Training Strategies

	Dataset Utilization		
	ImageNet	Place365-Ours	Our dataset
Strategy 1			✓
Strategy 2	✓		✓
Strategy 3	✓	✓	
Strategy 4	✓	✓	✓

Table 5. Different Training Strategies: Overall Performance

	ResNet		ConvNeXt		Swin transformer	
	Acc	F1	Acc	F1	Acc	F1
Strategy 1	77.53	39.34	73.03	28.74	78.65	42.28
Strategy 2	85.39	56.87	88.76	61.15	88.76	77.21
Strategy 3	82.02	69.42	89.89	79.20	89.89	81.02
Strategy 4	89.89	78.00	92.13	80.17	94.38	85.34

F1 score is the macro-average F1 score calculated from the confusion matrix for each result.

Table 6. Maintaining Strategy for Feature Classifier

	ResNet		ConvNeXt		Swin transformer	
	Acc	F1	Acc	F1	Acc	F1
Ours w/o M	88.76	61.67	85.39	54.69	87.64	58.96
Ours w/ M	89.89	78.00	92.13	80.17	94.38	85.34

F1 score is the macro-average F1 score calculated from the confusion matrix for each result. M denotes the maintaining of the feature classifier.

distribution for “Vehicle” class.

However, by incorporating more relevant training samples from Place365-ours, the proposed method can classify the eating scene for the “Vehicle” class. This is consistent through all network architectures. The large intra-class difference in our dataset is the main reason for the inability to correctly classify both two sequences labeled as “Vehicle”. In our dataset, one “Vehicle” sequence is captured on the front seat of a family car while another sequence is captured on a bus. Since the second stage of training (on our dataset) uses one sequence for training and another for testing, the model learning to classify the environment on a family car may fail to classify the environment on a bus and vice versa.

The significant improvement in classification accuracy in the minority class can further verify that our proposed method helps handle the inherent data imbalance problem existing in eating scene recognition tasks.

5.3. Ablation Studies

Additional datasets: Since we have two additional datasets (ImageNet and Place365-Ours) that are used for training our model, there are four training strategies presented in the experiments to explore the effectiveness of including the additional datasets. Table 4 summarizes the differences between the four training strategies. See the description of strategies in the supplementary material.

From Table 5, we observe that strategy 4, the proposed method, is the best training strategy since it achieves the highest accuracy and Macro-Average F1 score for all network architectures. This is expected since the proposed two-stage training strategy optimally utilizes two additional datasets as well as our dataset for improving performance.

Maintaining of feature classifier: As mentioned in Section 4.2, we maintain the feature classifier g from stage 1 to stage 2 instead of dropping it (enabled by semantic-based class filtering and merging). We verify the strategy and see results in Table 6.

We observe a significant improvement in sequence-level accuracy and even more improvement in macro-average F1 score by maintaining the feature classifier g for the second stage of training. The classifier is much more robust to data imbalance issues, demonstrating the effectiveness of our proposed semantic-based label filtering and merging that enables the drop-then-maintain strategy.

6. Conclusion

In this paper, we present an automatic ingestion environment recognition method to aid nutritionists and dietitians in overcoming the limitations of self-report and manual review of eating scene images. We use the data from the accelerometer and flexible sensor to detect eating episodes and explore the proposed two-stage drop-then-maintain training framework on several neural network architectures to perform scene classification. The experimental results indicate our proposed method outperforms the baseline methods in both sequence-level accuracy and minority-class accuracy, demonstrating its effectiveness in addressing the data imbalance problem.

A larger study is planned for the future to extend the work to more environment categories for a longer duration of device wear. At the same time, a major advantage of this study is that the image database is fully natural and contains no staged images or artificially created environments. Therefore, the performance of the proposed method is representative of what is expected in everyday living.

References

- [1] Jan M Bauer, Kristian S Nielsen, Wilhelm Hofmann, and Lucia A Reisch. Healthy eating in the wild: An experience-sampling study of how food environments and situational factors shape out-of-home dietary success. *Social Science & Medicine*, 299:114869, 2022.
- [2] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. *Lecture notes in computer science*, 3951:404–417, 2006.
- [3] Matthew Breit, Jonathan Padia, Tyson Marden, Dan Forjan, Pan Zhaoxing, Wenru Zhou, Tonmoy Ghosh, Graham Thomas, Megan A McCrory, Edward Sazonov, et al. The spectrum of eating environments encountered in free living adults documented using a passive capture food intake wearable device. *Frontiers in Nutrition*, 10:1119542, 2023.
- [4] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arachiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. *Advances in neural information processing systems*, 32, 2019.
- [5] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- [6] Iris WH Claessens, Marleen Gillebaart, and Denise TD de Ridder. Personal values, motives, and healthy and sustainable food choices: Examining differences between home meals and restaurant meals. *Appetite*, 182:106432, 2023.
- [7] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.
- [8] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9268–9277, 2019.
- [9] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and F. Li. Imagenet: A large-scale hierarchical image database. *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.
- [11] Oscar Déniz, Gloria Bueno, Jesús Salido, and Fernando De la Torre. Face recognition using histograms of oriented gradients. *Pattern recognition letters*, 32(12):1598–1603, 2011.
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations*, 2021.
- [13] Abul Doulah, Xin Yang, Jason Parton, Janine A Higgins, Megan A McCrory, and Edward Sazonov. The importance of field experiments in testing of sensors for dietary assessment and eating behavior monitoring. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 5759–5762, 2018.
- [14] Abul Doulah, Tonmoy Ghosh, Delwar Hossain, Masudul H Imtiaz, and Edward Sazonov. “automatic ingestion monitor version 2”—a novel wearable device for automatic food intake detection and passive capture of food images. *IEEE journal of biomedical and health informatics*, 25(2):568–576, 2020.
- [15] Lydia Dubourg, Ana Rita Silva, Christophe Fitamen, Chris JA Moulin, and Céline Souchay. Sensecam: A new tool for memory rehabilitation? *Revue Neurologique*, 172(12):735–747, 2016.
- [16] Muhammad Farooq and Edward Sazonov. Accelerometer-based detection of food intake in free-living individuals. *IEEE sensors journal*, 18(9):3752–3758, 2018.
- [17] Muhammad Farooq, Abul Doulah, Jason Parton, Megan A McCrory, Janine A Higgins, and Edward Sazonov. Validation of sensor-based food intake detection by multicamera video observation in an unconstrained environment. *Nutrients*, 11(3):609, 2019.
- [18] Purificación García-Segovia, Robert J Harrington, and Han-Seok Seo. Influences of table setting and eating location on food acceptance and intake. *Food Quality and Preference*, 39:1–7, 2015.
- [19] Luke Gemming, Aiden Doherty, Jennifer Utter, Emma Shields, and Cliona Ni Mhurchu. The use of a wearable camera to capture and categorise the environmental and social context of self-identified eating episodes. *Appetite*, 92: 118–125, 2015.
- [20] M Giacchi, R Mattei, and S Rossi. Correction of the self-reported bmi in a teenage population. *International journal of obesity*, 22(7):673–677, 1998.
- [21] Jiangpeng He and Fengqing Zhu. Single-stage heavy-tailed food classification. *2023 IEEE International Conference on Image Processing*, pages 1115–1119, 2023.
- [22] Jiangpeng He, Zeman Shao, Janine Wright, Deborah Kerr, Carol Boushey, and Fengqing Zhu. Multi-task image-based dietary assessment for food recognition and portion size estimation. *2020 IEEE Conference on Multimedia Information Processing and Retrieval*, pages 49–54, 2020.
- [23] Jiangpeng He, Luotao Lin, Heather A Eicher-Miller, and Fengqing Zhu. Long-tailed food classification. *Nutrients*, 15(12):2751, 2023.
- [24] Jiangpeng He, Luotao Lin, Jack Ma, Heather A Eicher-Miller, and Fengqing Zhu. Long-tailed continual learning for visual food recognition. *arXiv preprint arXiv:2307.00183*, 2023.
- [25] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [26] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5375–5384, 2016.
- [27] Minyoung Huh, Pulkit Agrawal, and Alexei A Efros. What makes imagenet good for transfer learning? *arXiv preprint arXiv:1608.08614*, 2016.

- [28] Nathalie Japkowicz and Shaju Stephen. The class imbalance problem: A systematic study. *Intelligent data analysis*, 6(5): 429–449, 2002.
- [29] Yan Ke and Rahul Sukthankar. Pca-sift: A more distinctive representation for local image descriptors. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2:II–II, 2004.
- [30] Jaehyung Kim, Jongheon Jeong, and Jinwoo Shin. M2m: Imbalanced classification via major-to-minor translation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13896–13905, 2020.
- [31] Julia Kupis, Sydney Johnson, Gregory Hallihan, and Dana Lee Olstad. Assessing the usability of the automated self-administered dietary assessment tool (asa24) among low-income adults. *Nutrients*, 11(1):132, 2019.
- [32] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *Proceedings of the IEEE International Conference on Computer Vision*, pages 2980–2988, 2017.
- [33] Xu-Ying Liu, Jianxin Wu, and Zhi-Hua Zhou. Exploratory undersampling for class-imbalance learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(2):539–550, 2008.
- [34] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *IEEE/CVF International Conference on Computer Vision*, pages 9992–10002, 2021.
- [35] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022.
- [36] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11966–11976, 2022.
- [37] David G Lowe. Object recognition from local scale-invariant features. *Proceedings of the IEEE International Conference on Computer Vision*, 2:1150–1157, 1999.
- [38] Stephanie Lowry, Niko Sünderhauf, Paul Newman, John J Leonard, David Cox, Peter Corke, and Michael J Milford. Visual place recognition: A survey. *IEEE Transactions on Robotics*, 32(1):1–19, 2015.
- [39] Runyu Mao, Jiangpeng He, Zeman Shao, Sri Kalyan Yarlagadda, and Fengqing Zhu. Visual aware hierarchy based food recognition. *Proceedings of the International Conference on Pattern Recognition Workshop*, pages 571–598, 2021.
- [40] Estefania Talavera Martinez, Maria Leyva-Vallina, Md Mostafa Kamal Sarker, Domenec Puig, Nicolai Petkov, and Petia Radeva. Hierarchical approach to classify food scenes in egocentric photo-streams. *IEEE journal of biomedical and health informatics*, 24(3):866–877, 2019.
- [41] Herbert L Meiselman, JL Johnson, W Reeve, and JE Crouch. Demonstrations of the influence of the eating environment on food acceptance. *Appetite*, 35(3):231–237, 2000.
- [42] Ebrahim Mortaz. Imbalance accuracy metric for model selection in multi-class imbalance classification problems. *Knowledge-Based Systems*, 210:106490, 2020.
- [43] Lawrence Mosley. *A balanced approach to the multi-class imbalance problem*. PhD thesis, Iowa State University, 2013.
- [44] ES Sazonov, SAC Schuckers, P Lopez-Meyer, O Makeyev, EL Melanson, MR Neuman, JO Hill, et al. Toward objective monitoring of ingestive behavior in free-living population. *Obesity*, 17(10):1971–1975, 2009.
- [45] Zeman Shao, Yue Han, Jiangpeng He, Runyu Mao, Janine Wright, Deborah Kerr, Carol Jo Boushey, and Fengqing Zhu. An integrated system for mobile image-based dietary assessment. *Proceedings of the 3rd Workshop on AIXFood*, page 19–23, 2021.
- [46] Zeman Shao, Jiangpeng He, Ya-Yuan Yu, Luotao Lin, Alexandra Cowan, Heather Eicher-Miller, and Fengqing Zhu. Towards the creation of a nutrition and food group based image database. *arXiv preprint arXiv:2206.02086*, 2022.
- [47] Zeman Shao, Gautham Vinod, Jiangpeng He, and Fengqing Zhu. An end-to-end food portion estimation framework based on shape reconstruction from monocular image. *2023 IEEE International Conference on Multimedia and Expo*, pages 942–947, 2023.
- [48] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations*, 2015.
- [49] Mervyn Stone. Cross-validators: choice and assessment of statistical predictions. *Journal of the royal statistical society: Series B (Methodological)*, 36(2):111–133, 1974.
- [50] Nanette Stroebele and John M de Castro. Influence of physiological and subjective arousal on food intake in humans. *Nutrition*, 22(10):996–1004, 2006.
- [51] Mingui Sun, Lora E Burke, Zhi-Hong Mao, Yiran Chen, Hsin-Chen Chen, Yicheng Bai, Yuecheng Li, Chengliu Li, and Wenyan Jia. ebutton: a wearable computer for health monitoring and personal assistance. *Proceedings of the 51st annual design automation conference*, pages 1–6, 2014.
- [52] Jingru Tan, Changbao Wang, Buyu Li, Quanquan Li, Wanli Ouyang, Changqing Yin, and Junjie Yan. Equalization loss for long-tailed object recognition. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11662–11671, 2020.
- [53] Thea F Van de Mortel. Faking it: social desirability response bias in self-report research. *Australian Journal of Advanced Nursing*, 25(4):40–48, 2008.
- [54] Xudong Wang, Long Lian, Zhongqi Miao, Ziwei Liu, and Stella Yu. Long-tailed recognition by routing diverse distribution-aware experts. *International Conference on Learning Representations*, 2021.
- [55] Zizhao Zhang and Tomas Pfister. Learning fast sample re-weighting without reward data. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 725–734, 2021.
- [56] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464, 2017.