

LEARNING EATING ENVIRONMENTS THROUGH SCENE CLUSTERING

*Sri Kalyan Yarlagadda¹, Sriram Baireddy¹, David Güera¹,
Carol J. Boushey², Deborah A. Kerr³, Fengqing Zhu¹*

School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA¹
University of Hawaii Cancer Center, Honolulu, HI, USA²
School of Public Health, Curtin University, Perth, Western Australia³

ABSTRACT

It is well known that dietary habits have a significant influence on health. While many studies have been conducted to understand this relationship, little is known about the relationship between eating environments and health. Yet researchers and health agencies around the world have recognized the eating environment as a promising context for improving diet and health. In this paper, we propose an image clustering method to automatically extract the eating environments from eating occasion images captured during a community dwelling dietary study. Specifically, we are interested in learning how many different environments an individual consumes food in. Our method clusters images by extracting features at both global and local scales using a deep neural network. The variation in the number of clusters and images captured by different individual makes this a very challenging problem. Experimental results show that our method performs significantly better compared to several existing clustering approaches.

Index Terms— Image Clustering, Eating Environment, Deep Learning

1. INTRODUCTION

In 2016, \$7.5 trillion was spent on healthcare worldwide, which is approximately 10% of the world GDP [1]. While there are many factors that influence health, dietary habits have a significant impact [2, 3]. To understand the complex relationship between dietary habits and health, nutrition practitioners and researchers often conduct dietary studies to subjectively assess dietary intake of children and adults. Participants in these studies are asked to report the foods and drinks they consumed on a daily basis for a period of time. This data is then analyzed by researchers to understand the impact of certain dietary behaviors on health. For example, studies have shown that frequent consumption of fast food [4], skipping breakfast [5], and absence of home food [6] contribute to the increasing risk of obesity and overweight.

This work was sponsored by the National Science Foundation under grant 1657262. Address all correspondence to Fengqing Zhu, zhu0@ecn.purdue.edu

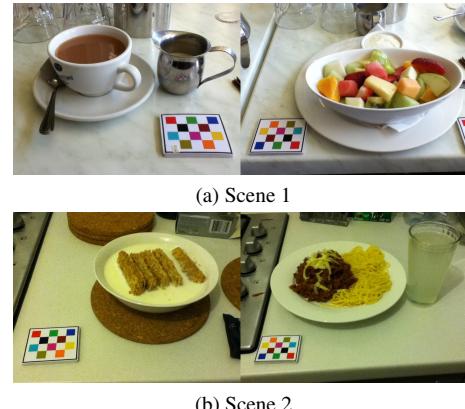


Fig. 1: Here are images of two different eating environments, captured by a single participant. The colored checkerboard in all the images is the FM.

While many studies have been conducted to understand how diet affects health [7], fewer work has been done to study the relationship between eating environments and health. However, researchers [8, 9] and organizations such as the World Health Organization and International Obesity Task Force have recognized the vital role of eating environments on health and diet. Studies have shown that some factors of the environment such as screen viewing during meals, family mealtime [10], and meal preparation time [11] influence health. For instance, family mealtime is shown to positively affect nutrient intake and meal preparation time is inversely related to Body Mass Index (BMI). While such patterns have been found, the relationship between eating environments and health is still poorly understood, partly due to the lack of valid, reliable measures of environmental factors. In this paper, we focus on understanding eating environments of an individual using digital images captured during eating occasions. In particular, we are interested in learning how many different environments an individual consumes food in.

Dietary Recall, 24-hr recall, and Food Frequency Questionnaire (FFQ) [12] are well-known dietary assessment methods used in most dietary studies. These methods require

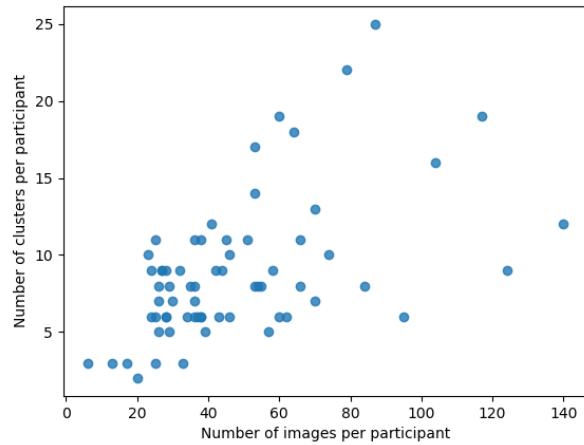


Fig. 2: Scatter plot of the number of images captured per participant versus the number of eating scene clusters per participant.

participants to manually enter details of their diet information through a web interface, or an in-person or phone interview. These methods are known to be time-consuming and prone to errors because participants may not recall all foods and beverages they consumed or cannot accurately estimate the food portion [12]. To overcome these limitations, researchers have leveraged advances in mobile technology to develop image-based dietary assessment methods to record daily food intake. Some examples of image based dietary assessment tools are TADA™ [13] and Food-Log [14]. In these approaches, participants record their diet by capturing images of foods and beverages consumed using mobile cameras. These images can then be analyzed by trained analysts [15] or using image analysis methods [16, 17, 18] to extract nutrient information. In this paper, we leverage eating occasion images captured using the TADA™ system to cluster eating scenes based on their visual appearance to help understand the relationship between a person's eating environment and dietary quality.

2. DATASET AND RELATED WORK

We used a dataset \mathcal{D} that consists of 3137 images from 66 participants collected in a community dwelling dietary study [19]. In this study, participants were asked to take pictures of their foods and to place a colored checkerboard with known dimensions, called the Fiducial Marker (FM), in the scene. An example of two pairs of images belonging to two different eating environments are shown in Fig. 1. The FM serves two purposes: to aid in color correction, and more importantly, to provide a reference scale for food portion estimation, so the nutrient content of the foods can be computed. One of the challenges of this dataset is the large variance in

the number of images and the number of eating environments for different participant, as shown in Fig. 2.

Our goal is to cluster food images of a participant based on eating environments. Clustering is not a new concept. Several classical [20] and deep learning-based [21] methods have been proposed. Classical clustering approaches are applied on relevant features extracted from the data [22]. However, the problem of clustering food images based on the eating environments captured has not been studied before, so the relevant features are not defined. On the other hand, deep learning-based methods, given sufficient data, are capable of simultaneously learning to extract relevant features and cluster images [21]. A common assumption of such approaches is that the number of clusters is known *a priori* [21]. However, in our case, we do not know *a priori* the number of eating environments for each participant. In addition, the number of images collected by each participant is not sufficient to apply deep learning-based methods directly, as they usually require hundreds of images per cluster for training a good model. As both classical and deep learning-based approaches have their shortcomings for clustering eating environments from images, new techniques need to be developed to solve this problem.

3. METHOD

In this section, we describe the details of the proposed method. First, we introduce the notation used throughout the paper. P_k denotes the k^{th} participant in the dataset. The i^{th} image captured by participant P_k is denoted by $x_i^{P_k}$. Our goal is to cluster food images based on eating environment. We do this by extracting features at the local and global level from relevant pixels of $x_i^{P_k}$ and then applying a clustering method on said features.

3.1. Global Feature Extraction

The image $x_i^{P_k}$ contains many salient objects such as the food, drinks, FM, and silverware. However, pixels belonging to the salient objects do not contain any credible information regarding the eating environment because a person can eat different food with different plates in the same eating environment. Pixels belonging to the FM are also not relevant because the FM is common to all food images, irrespective of the eating environment. Instead, we are interested in pixels belonging to the non-salient regions. To extract these relevant pixels in $x_i^{P_k}$, we first extract its binary saliency map denoted by $s_i^{P_k}$ using a state-of-the-art saliency estimator R3NET [23]. The salient pixels of $x_i^{P_k}$ have a value of 1 in $s_i^{P_k}$ and the rest have a value of 0. The relevant pixels of $x_i^{P_k}$ are captured by $\tilde{x}_i^{P_k}$, and is defined as

$$\tilde{x}_i^{P_k} = (1 - s_i^{P_k}) * x_i^{P_k} \quad (1)$$

Features are extracted from $\tilde{x}_i^{P_k}$ using a Convolutional Neural Network (CNN) VGG16 [24] pre-trained on the ImageNet

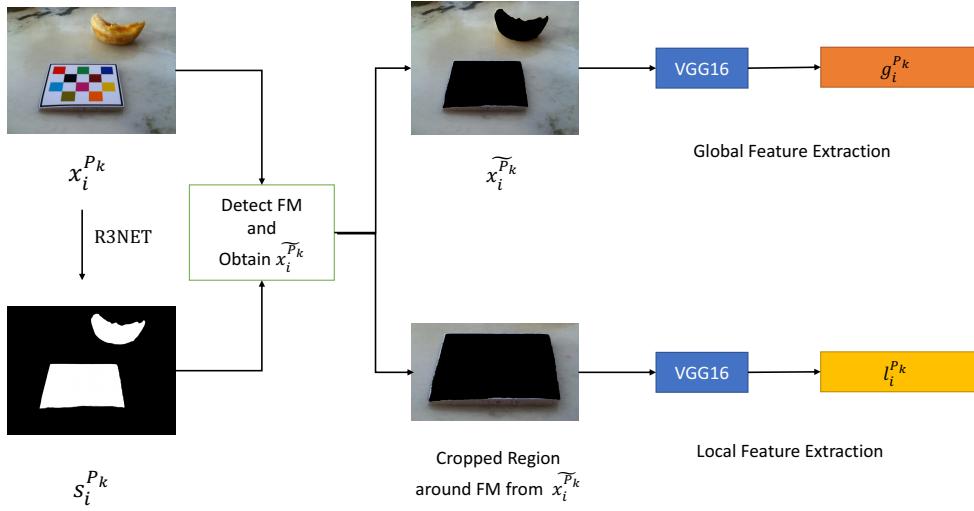


Fig. 3: Overview of feature extraction

dataset [25]. We use a pre-trained VGG16 for feature extraction because these features are robust to artifacts such as noise, lighting changes, and differing viewpoints.

The 2nd convolutional layer is chosen for feature extraction and the reasoning behind this is explained in Section 4.1. Global Average Pooling (GAP) is applied to the output of the 2nd layer to spatially summarize the information into a 64-dimensional vector. This vector is denoted as $g_i^{P_k}$ and we refer to it as the global feature.

3.2. Local Feature Extraction

Since the FM is always placed in close proximity to the foods, we assume it is on the same surface as the foods. Examples of such surfaces may include desk, dining table, and kitchen counter, to name a few. These surfaces give us a lot of information about the eating environment because it is very likely that a person uses the same surface during an eating event in a particular eating environment. Therefore, we extract features from this surface and they are denoted by $l_i^{P_k}$.

By identifying the FM in the image, we can extract information about the eating surface. The FM is detected by finding the salient object with the highest number of interest points. An interest point is a pixel that ORB (Oriented FAST and Rotated BRIEF) [26] finds useful for image registration. The FM has a lot of interest points because of the colored checkerboard pattern and it is unlikely for other salient objects such as foods to have as many distinct interest points. The salient objects of the image can be found by performing connected component analysis on $s_i^{P_k}$ and treating each connected component as a salient object. Once we have located the FM, we extract a region around it from $\tilde{x}_i^{P_k}$. $l_i^{P_k}$ is obtained from this region using the pre-trained VGG16 in the

same way as done in 3.1. Fig. 3 illustrates the the global and local feature extraction process.

3.3. Feature Fusion and Clustering

We fuse the local and global features using their distance matrices. The distance matrices G^{P_k} and L^{P_k} are defined as follows

$$\begin{aligned} L_{(i,j)}^{P_k} &= \|l_i^{P_k} - l_j^{P_k}\|_2 \\ G_{(i,j)}^{P_k} &= \|g_i^{P_k} - g_j^{P_k}\|_2 \end{aligned} \quad (2)$$

The fused distance matrix D^{P_k} is defined as follows

$$D^{P_k} = \alpha * L^{P_k} + (1 - \alpha) * G^{P_k} \quad (3)$$

Here $\alpha \in [0, 1]$ and controls the relative importance of the local and global features. A higher value of α indicates local features are more important and vice versa. In our case , $\alpha = 0.44$. The reason behind this is explained in Section 4.1. Affinity Propagation (AP) [27] is applied to D^{P_k} to obtain the final clusters.

4. EXPERIMENTS

In this section we describe all the experiments conducted and compare our method to four existing clustering methods, namely: DBSCAN [28], MeanShift [29], OPTICS [30], and AP [27]. We use the Adjusted Rand Index (ARI) and Normalized Mutual Information (NMI) to measure the accuracy of our clusters. ARI ranges from -1.0 to 1.0 while NMI ranges from 0 to 1.0 . Higher values indicate better clustering for both ARI and NMI. Accuracy over a dataset is reported as the average accuracy among all participants.

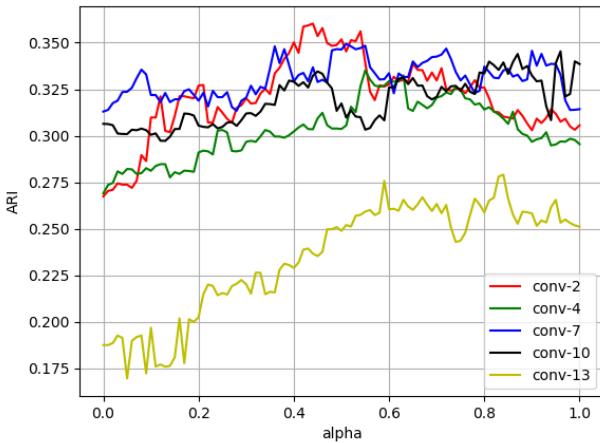


Fig. 4: ARI scores for different α and convolutional layer m on \mathcal{D}_{val}

4.1. Hyperparameter Tuning

Our method has two hyperparameters: the weighting factor for feature fusion α and the convolutional layer of VGG16 for feature extraction m . Our dataset \mathcal{D} is split into \mathcal{D}_{val} and $\mathcal{D}_{\text{test}}$. \mathcal{D}_{val} consists of images from 10 participants and $\mathcal{D}_{\text{test}}$ contain images from 56 participants. Our method is evaluated on \mathcal{D}_{val} by varying α and m . α ranges from 0 to 1 in steps of 0.01. To find the optimal m for our dataset, we extract features from convolutional layers preceding a max-pooling layer. There are five such layers in VGG16. ARI is used to select the optimal values. Fig. 4 shows the ARI scores for different values of α and convolutional layer m , and the optimal value for α is 0.44 and for m is 2.

From equation 3, we can infer that as the value of α increases, more weight is given to local features and vice-versa. An optimal value of 0.44 for α suggests that our method performs best when approximately equal weight is given to both local and global features. This shows that using only one of these features is less optimal. In Fig. 4, we can see that the performance of our method degrades once the chosen feature extraction layer gets very deep. Later layers, like conv-13, extract abstract features and completely lose information about edges, colors, and textures. We suspect this loss of information is the reason for decrease in performance.

4.2. Testing

We evaluated the performance of our method on $\mathcal{D}_{\text{test}}$ after selecting the optimal values for α and m using the validation set \mathcal{D}_{val} . We choose four well-known clustering methods for comparison, namely DBSCAN [28], MeanShift [29], OPTICS [30], and AP [27] using the eating scene image as the input. ARI and NMI for the five methods are reported in Ta-

ble 1, where our method achieves the best performance. It is worth noting that although we use AP [27] for clustering after feature fusion, our method performs significantly better than AP [27] when meaningful features are not known. This indicates that our feature extraction strategy is highly relevant and very important to our problem.

Table 1: ARI and NMI scores for methods tested on $\mathcal{D}_{\text{test}}$. The best results are reported in bold.

Methods	ARI	NMI
Ours	0.39	0.68
DBSCAN [28]	0.24	0.47
MeanShift [29]	0.12	0.35
OPTICS [30]	0.08	0.39
AP [27]	0.2	0.49

5. CONCLUSION

We proposed a method to cluster food images based on their eating environment. Our method extracts features from a pre-trained CNN at multiple levels. These features are fused using their distance matrices and a clustering algorithm is applied after feature fusion. Our method is evaluated on a dataset containing 3137 eating scene images collected from a dietary study with a total of 585 clusters. We compared our method to state-of-the-art clustering methods and showed improved performance.

6. REFERENCES

- [1] World Health Organization et al., “Public spending on health: a closer look at global trends,” Tech. Rep., World Health Organization, 2018.
- [2] A. E. Mesas and et al., “Selected eating behaviours and excess body weight: a systematic review,” *Obes Rev*, vol. 13, no. 2, pp. 106–135, Feb 2012, 21955734[pmid].
- [3] Karin Nordström, Christian Coff, Håkan Jönsson, Lennart Nordenfelt, and Ulf Görman, “Food and health: individual, cultural, or scientific matters?,” *Genes Nutr*, vol. 8, no. 4, pp. 357–363, Jul 2013, 23494484[pmid].
- [4] Garcia Ginny and et al., “The fast food and obesity link: consumption patterns and severity of obesity,” *Obes Surg*, vol. 22, no. 5, pp. 810–818, May 2012.
- [5] Deshmukh-Taskar PR1 and et al., “The relationship of breakfast skipping and type of breakfast consumption with nutrient intake and weight status in children and adolescents: the National Health and Nutrition Examination Survey 1999–2006,” *J Am Diet Assoc*, vol. 110, no. 6, pp. 869–878, Jun 2010.

[6] Amber J. Hammons and Barbara H. Fiese, "Is frequency of shared family meals related to the nutritional health of children and adolescents?," *Pediatrics*, vol. 127, no. 6, pp. e1565–1574, Jun 2011.

[7] Weiqing Min and et al., "A survey on food computing," *ACM Comput. Surv.*, vol. 52, no. 5, pp. 92:1–92:36, Sept. 2019.

[8] Sarah L. Booth and et al., "Environmental and societal factors affect food choice and physical activity: Rationale, influences, and leverage points," *Nutrition Reviews*, vol. 59, no. 3, pp. S21–S36, 2001.

[9] James O. Hill and et al., "Obesity and the environment: Where do we go from here?," *Science*, vol. 299, no. 5608, pp. 853–855, 2003.

[10] Brian Wansink and Ellen Van Kleef, "Dinner rituals that correlate with child and adult bmi," *Obesity (Silver Spring, Md.)*, vol. 22, 05 2014.

[11] Brad Appelhans and et al., "Meal preparation and cleanup time and cardiometabolic risk over 14years in the study of women's health across the nation (swan)," *Preventive medicine*, vol. 71, 12 2014.

[12] Jee-Seon Shim and et al., "Dietary assessment methods in epidemiologic studies," *Epidemiol Health*, vol. 36, pp. e2014009–e2014009, Jul 2014, 25078382[pmid].

[13] Fengqing Zhu and et al., "The use of mobile devices in aiding dietary assessment and evaluation," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 4, pp. 756–766, Aug 2010.

[14] Kiyoharu Aizawa and Makoto Ogawa, "Foodlog: Multimedia tool for healthcare applications," *IEEE MultiMedia*, vol. 22, no. 2, pp. 4–8, Apr 2015.

[15] Erica Howes and et al., "Image-Based Dietary Assessment Ability of Dietetics Students and Interns," *Nutrients*, vol. 9, no. 2, pp. 114, feb 2017.

[16] Fengqing Zhu and et al., "Multiple Hypotheses Image Segmentation and Classification With Application to Dietary Assessment," *IEEE Journal of Biomedical and Health Informatics*, vol. 19, no. 1, pp. 377–388, jan 2015.

[17] Fang, Shao, Kerr, Boushey, and Zhu, "An End-to-End Image-Based Automatic Food Energy Estimation Technique Based on Learned Energy Distribution Images: Protocol and Methodology," *Nutrients*, vol. 11, no. 4, pp. 877, apr 2019.

[18] Yu Wang and et al., "Context based image analysis with application in dietary assessment and evaluation," *Multimedia Tools and Applications*, vol. 77, no. 15, pp. 19769–19794, aug 2018.

[19] Deborah A Kerr and et al., "The connecting health and technology study: a 6-month randomized controlled trial to improve nutrition behaviours using a mobile food record and text messaging support in young adults," *The international journal of behavioral nutrition and physical activity*, vol. 13, no. 1, 2016.

[20] Rui Xu and D. Wunsch, "Survey of clustering algorithms," *IEEE Transactions on Neural Networks*, vol. 16, no. 3, pp. 645–678, May 2005.

[21] E. Min and et al., "A survey of clustering with deep learning: From the perspective of network architecture," *IEEE Access*, vol. 6, pp. 39501–39514, 2018.

[22] Seema Wazarkar and Bettahally N Keshavamurthy, "A survey on image data analysis through clustering techniques for real world applications," *Journal of Visual Communication and Image Representation*, vol. 55, pp. 596–626, 2018.

[23] Zijun Deng and et al., "R3net: Recurrent residual refinement network for saliency detection," *International Joint Conference on Artificial Intelligence*, pp. 684–690, July 2018.

[24] Karen Simonyan and Andrew Zisserman, "Very deep convolutional networks for large-scale image recognition," *International Conference on Learning Representations*, May 2015, San Diego, CA.

[25] Jia Deng and et al., "Imagenet: A large-scale hierarchical image database," *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, June 2009.

[26] Ethan Rublee and et al., "Orb: An efficient alternative to sift or surf.," in *ICCV*. Citeseer, 2011, vol. 11, p. 2.

[27] Brendan J Frey and Delbert Dueck, "Clustering by passing messages between data points," *science*, vol. 315, no. 5814, pp. 972–976, 2007.

[28] Martin Ester and et al., "A density-based algorithm for discovering clusters in large spatial databases with noise.," in *Kdd*, 1996, vol. 96, pp. 226–231.

[29] D. Comaniciu and P. Meer, "Mean shift: a robust approach toward feature space analysis," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 5, pp. 603–619, May 2002.

[30] Mihael Ankerst and et al., "Optics: ordering points to identify the clustering structure," in *ACM Sigmod record*. ACM, 1999, vol. 28, pp. 49–60.