



Diffusion Model with Clustering-based Conditioning for Food Image Generation

Yue Han
Purdue University
West Lafayette, IN, USA
han380@purdue.edu

Jiangpeng He
Purdue University
West Lafayette, IN, USA
he416@purdue.edu

Mridul Gupta
Purdue University
West Lafayette, IN, USA
gupta431@purdue.edu

Edward J. Delp
Purdue University
West Lafayette, IN, USA
ace@purdue.edu

Fengqing Zhu
Purdue University
West Lafayette, IN, USA
zhu0@purdue.edu

ABSTRACT

Image-based dietary assessment serves as an efficient and accurate solution for recording and analyzing nutrition intake using eating occasion images as input. Deep learning-based techniques are commonly used to perform image analysis such as food classification, segmentation, and portion size estimation, which rely on large amounts of food images with annotations for training. However, such data dependency poses significant barriers to real-world applications, because acquiring a substantial, diverse, and balanced set of food images can be challenging. One potential solution is to use synthetic food images for data augmentation. Although existing work has explored the use of generative adversarial networks (GAN) based structures for generation, the quality of synthetic food images still remains subpar. In addition, while diffusion-based generative models have shown promising results for general image generation tasks, the generation of food images can be challenging due to the substantial intra-class variance. In this paper, we investigate the generation of synthetic food images based on the conditional diffusion model and propose an effective clustering-based training framework, named ClusDiff, for generating high-quality and representative food images. The proposed method is evaluated on the Food-101 dataset and shows improved performance when compared with existing image generation works. We also demonstrate that the synthetic food images generated by ClusDiff can help address the severe class imbalance issue in long-tailed food classification using the VFN-LT dataset.

CCS CONCEPTS

• **Computing methodologies** → **Computer vision tasks; Machine learning**; • **Applied computing** → **Life and medical sciences**.

KEYWORDS

Food Image Generation, Diffusion Model, Food Classification, Image-based Dietary Assessment



This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

MADiMa '23, October 29, 2023, Ottawa, ON, Canada
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0284-6/23/10.
<https://doi.org/10.1145/3607828.3617796>

ACM Reference Format:

Yue Han, Jiangpeng He, Mridul Gupta, Edward J. Delp, and Fengqing Zhu. 2023. Diffusion Model with Clustering-based Conditioning for Food Image Generation. In *Proceedings of the 8th International Workshop on Multimedia Assisted Dietary Management (MADiMa '23)*, October 29, 2023, Ottawa, ON, Canada. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3607828.3617796>

1 INTRODUCTION

Dietary intake has a profound impact on personal health and disease. A healthy diet is an essential aspect of keeping lifelong well-being for current and future generations across the lifespan. Dietary factors are crucial in determining the risk of developing obesity, heart disease, diabetes, and even cancer [46]. Dietary assessment has been utilized as a fundamental approach for understanding diet's effects on human health. Traditional methods of dietary assessment are mainly based on self-reporting, which makes it difficult to obtain an accurate dietary assessment of individuals due to both random and systematic measurement errors [2]. Recently, image-based dietary assessment has been widely used to assist researchers and participants in recording their dietary intake with high accuracy and efficiency [4]. Methods such as FoodLog [1], DietCam [47], and FoodCam [44] allow participants to capture their food intake using a mobile device, then the captured images are analyzed by trained dietitians to estimate the nutrient composition. To further reduce the human effort in analyzing the captured eating occasion images, several recent approaches [69] automate this process by using deep learning-based techniques to recognize the food [27, 29, 53] and estimate the portion size [13, 14, 68] in the captured images. Besides using mobile devices to capture the food intake image, passive capture of eating occasion images utilizes wearable cameras assisted by on-device sensors to collect dietary intake data [24, 73].

The image-based dietary assessment methods reduce the time and labor required for nutrient analysis and allow real-time feedback to the participants. The major challenge with deep learning-based approaches is that the performance is heavily dependent on the quality and the quantity of the datasets used for training the models. It is difficult to obtain a large number of accurately annotated food images in real-world scenarios. Furthermore, food images collected in daily life usually suffer from data imbalance issues due to the nature that some foods are consumed more frequently than others [25, 32]. One way to increase the size of the

food image dataset is by using data augmentation methods to create more training samples. Traditional data augmentation methods utilize basic image manipulations including color space transformations, random erasing, geometric transformations, and kernel filtering [7, 57]. However, these augmentation methods do not work well when the training data is limited and can result in overfitting [71]. With the development of Generative Adversarial Networks (GANs), generating synthetic images from GANs for data augmentation is widely used for various deep learning methods [10, 23].

Generative adversarial networks (GANs) have shown great potential in generating realistic synthetic images [20]. Moreover, conditional generative models [56] can generate images for a specific food class. Previous works [17, 36, 37] for synthetic food image generation have typically focused on generating images of only one or a limited number of food classes, primarily due to the challenge of creating a dataset with a diverse variety of foods. Therefore, the limitation on scale of the current food image generation methods narrows the application of using them as data augmentation methods. Diffusion models have recently become the hottest research topic in the area of generative models, showing great capability in image generation [65], super-resolution [19], image modification [45], and image inpainting [51]. However, there is no existing work utilizing the diffusion model for food image generation. One of the major challenges of applying the diffusion model is that the food data usually suffers from high intra-class variance due to the differences in recipes or cooking methods with the same ingredient [38]. Therefore, it is relatively easy to have the mode collapse during the training of the image generation model. Mode collapse refers to the condition when the model only generates a limited set of examples instead of exploring the entire training data distribution [75].

In this paper, we first explore the performance of latent diffusion methods on food image generation by fine-tuning a stable diffusion model [65] on the Food-101 dataset [3] and compare results to the images generated using one of the latest GAN-based methods, StyleGAN3 [41]. We propose a novel clustering-based training framework, ClusDiff, to solve the high intra-class variance issue when training the food image generation model. For training ClusDiff, food images within each class are clustered as sub-classes, and then the model is trained with these sub-classes. When generating images for a specific class of food, ClusDiff is conditioned on its sub-class labels, where the occurrence of the labels follows the distribution of sub-classes within a class. We evaluate ClusDiff on the Food-101 dataset and demonstrate that our proposed method achieves better food image generation performance in terms of Frechet Inception Distance Metric (FID) [34] as compared to the baseline, which is a fine-tuned stable diffusion model. We also explore the benefit of using synthetic images for data augmentation by providing a case study that uses synthetic food images generated by ClusDiff to address the class-imbalance issue in the VFN-LT dataset [25] for long-tailed food image classification. Overall, the contribution of this work can be summarized into the following:

- This work represents the first investigation and application of the diffusion model to food image generation.
- We propose ClusDiff, a diffusion model with clustering-based conditioning, and show that the proposed method generates better food images than existing methods.

- We demonstrate synthetic food images generated by ClusDiff can be used to address the data imbalance issue in food classification through a case study.

2 RELATED WORK

2.1 Image-based Dietary Assessment

With recent advances in machine learning and modern computer vision, image-based dietary assessment methods have been introduced to automatically record and analyze participants' food intake. The most important aspect of image-based dietary assessment is to detect and classify the food and beverages in the images captured by participants. DiaWear [72] is one of the pioneering works that uses a neural network to perform context-aware food recognition on food images captured by mobile devices. Deep learning-based detection methods, such as Faster R-CNN [64] and YOLO [62] are frequently used for food detection and have shown good results when trained on a large number of ground truth images [11, 24, 40]. The most recent work focuses on estimating food portion size or energy directly from the eating occasion images. In [15], a geometric model is introduced to estimate the food portion size. Conditional GAN [56] is used in [14] to generate the energy distribution map with an end-to-end regression framework for food energy estimation [13], which is further improved in [68] by integrating information from RGB images and in [70] by leveraging the 3D reconstruction. By combining these components, systems like goFOOD [50] and Technology-Assisted Dietary Assessment (TADA) [69] can provide an end-to-end automatic nutrient estimation for a meal.

Despite the promising advancements in image-based dietary assessment methods, the performance of deep learning-based methods relies heavily on the quality and quantity of the datasets. Training a model for food recognition or portion estimation requires a large, diverse, and accurately annotated dataset. Nevertheless, the food images we collect from real-world scenarios often exhibit stark differences from established food image datasets such as Food2K [55], Food-101 [3], and UEC-256 [43]. This discrepancy primarily stems from inherent limitations associated with real-world data, such as variability, inconsistent sizes, and potential inaccuracies in annotations. While various food recognition systems have been developed to target real-world scenarios, such as continual learning [26, 28, 30, 31, 61], few-shot learning [39], and long-tailed classification [18, 25, 32], they usually necessitate more sophisticated training regimes. This amplification in complexity invariably leads to a surge in computational requirements, which are typically constrained in real-world situations. A potential solution to mitigate the problem of data scarcity in the real world without requiring any changes to the existing image-based dietary assessment system is to use synthetic food images for the training process. However, generating synthetic food images could be challenging due to inter-class similarity and intra-class variability. In this work, we investigate the application of the stable diffusion model [65] for synthetic food image generation. We propose a novel framework for training a stable diffusion model to generate more diverse synthetic food images and demonstrate that these images can be used to address the data imbalance issue in tasks involving long-tailed food classification.

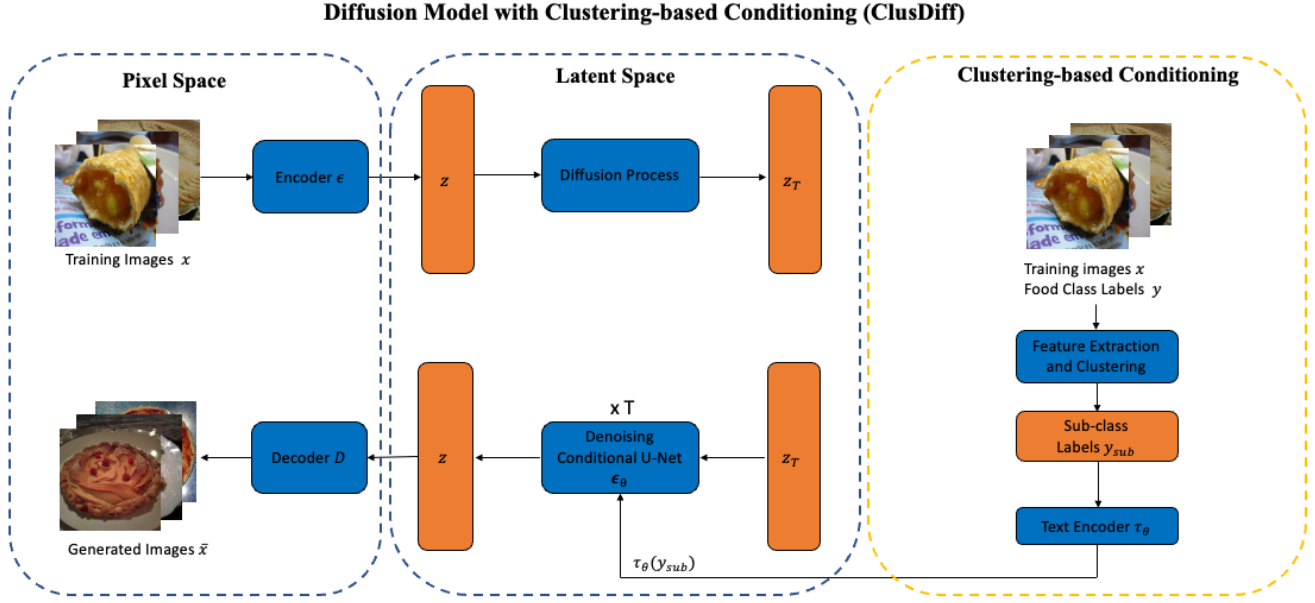


Figure 1: The block diagram of the proposed clustering-based conditional training strategy for latent diffusion model (ClusDiff). The blue blocks indicate the network structure, and the red blocks indicate data passing through the network.

2.2 Synthetic Image Generation

One of the main challenges of supervised machine learning or deep learning methods is the requirement for a large amount of annotated training samples, and the performance heavily relies on the quality of the dataset. One way to address this issue is by generating high-quality image data to augment the training dataset. For a long time, Generative Adversarial Networks (GANs) have been the most impactful methods and have shown remarkable results for synthetic image generation. These synthetic images have found applications in various fields, from video games to forensic scenarios. Many applications have utilized GANs to create synthetic images that can be used for generating ground truth information used in training machine learning networks [10, 23].

There are few studies on applying GANs in the field of food image synthesis, such as RamenGAN [37], Multi-ingredient Pizza Image Generator (MPG) [22], and CookGAN [21]. RamenGAN [37] and its related work [36] generate synthetic food images of noodles and rice based on a conditional GAN network. MPG [22] uses conditional StyleGAN2 [42] to generate realistic pizza images with various ingredients. CookGAN [21] proposed a method to generate a meal image conditioned on a list of ingredients, with the types of foods limited to salad, a cookie, and a muffin.

In recent years, diffusion models have shown the ability to produce high-quality images while maintaining ease of use. The diffusion model defines a Markov chain to gradually add random noise to the data, and the model learns to reverse the diffusion process. As a result, the trained model is capable of constructing the desired data samples from the noise [35]. In [9], the authors demonstrate that the diffusion model outperforms GANs and achieves better image generation results on multiple datasets. Additionally, in [65], the authors propose latent diffusion models, which apply the diffusion

process in the latent space of the data, resulting in a computational reduction and a boost in the visual fidelity of the generated images. Despite diffusion models being applied to many downstream tasks [8, 52, 60], the study of applying them to food image generation is still lacking. In this paper, our work represents the first investigation and application of the diffusion model to food image generation.

3 METHOD

A block diagram of our proposed clustering-based conditional training strategy for the latent diffusion model, ClusDiff, is shown in Fig. 1. The system includes three parts: latent diffusion model training and inference, feature space food image clustering using affinity propagation, and clustering-based conditioning. The motivation for this work is that food datasets usually exhibit high intra-class variance, with different cooking methods leading to variations within the same food class. As a result, training a generative model with simple class labels may not sufficiently cover the distribution in the training dataset. In the training step of the generative model, we employ a clustering method to divide food images from a class label into several sub-classes based on their visual representations. Subsequently, the generative model is trained conditionally on these sub-classes to further capture the variations within a food class. When we use this trained generative model to generate food images, a food class label for conditional generation is represented as one of its sub-classes, determined by the frequency of occurrence in the training dataset. By doing this, the food images generated using ClusDiff better represent the distribution of the training dataset.

3.1 Latent Diffusion Model

There are two processes in diffusion models [35]: (i) a forward diffusion process to diffuse an image into random noise, and (ii) a reverse diffusion process that converts the noise into the image from a learned data distribution during the iterative denoising process.

In the forward diffusion process, a Markov chain is employed to gradually add noise to transform the training data into Gaussian noise. This discrete Markov chain is constructed with a predefined sequence of variance values $0 < \beta_1, \dots, \beta_T < 1$. Given the input data distribution $q(x_0)$, a forward diffusing process at time t is defined as Equation 1:

$$\begin{aligned} q(x_t|x_{t-1}) &\sim \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I}), \\ q(x_t|x_0) &\sim \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)\epsilon), \end{aligned} \quad (1)$$

where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ denotes the injected noise, and the forward process can be interpreted as sampling x_t at any time step t in a closed form by using the notation: $\alpha_t := 1 - \beta_t$ and $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$, which represent the noise level at each time step. At the final time step T , x_T approximately becomes an isotropic Gaussian noise.

In the reverse diffusion process, we can recover the input data distribution $q(x_0)$ from $p(x_T) \sim \mathcal{N}(x_T; 0, \mathbf{I})$. The equations for this process are defined as follows:

$$\begin{aligned} p_\theta(x_{0:T}) &\sim p(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t), \\ p_\theta(x_{t-1}|x_t) &\sim \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \sum_{\theta} (x_t, t)). \end{aligned} \quad (2)$$

The training object of the diffusion model is to find a denoising autoencoder $\epsilon_\theta(x_t, t)$ to predict the injected noise ϵ . The corresponding objective can be simplified as follows:

$$L_D = \mathbb{E}_{t, x_0, \epsilon} [\|\epsilon - \epsilon_\theta(x_t, t)\|^2], \quad (3)$$

where t is uniformly sampled from $\{1, \dots, T\}$.

The latent diffusion[65] follows similar processes as the diffusion model, except that it uses the latent space of images as input samples to the diffusion model. It utilizes a perceptual compression model[12] consisting of an encoder \mathcal{E} and a decoder \mathcal{D} . The encoder is used to map the image x to the latent space $z = \epsilon(x)$ and the decoder is used to reconstruct the image $\hat{x} = \mathcal{D}(z)$ from the latent space. The latent space representation reduces the dimension of the input data, therefore focusing on the critical features of the input data and reducing the computational complexity of the diffusion model.

3.2 Image Clustering using Affinity Propagation

We utilize Affinity Propagation [16] to cluster the training images from one class into sub-classes. This helps in addressing the high intra-class variance problem of the food dataset. A pre-trained ResNet-18 [33], fine-tuned on the Food-101 dataset for food classification, is used to map the images into feature vectors. These feature vectors represent the visual features of the input food images and are used for clustering. Affinity propagation creates clusters by sending messages between samples until convergence, and it does not require a pre-defined number of clusters. A small number of




Figure 2: The block diagram of the denoising conditional U-Net architecture with cross-attention mechanism. Noted the attention layers (shown in green blocks) are mapped to each intermediate layer of U-Net.

most representative samples are referred to as exemplars to describe the dataset. Messages sent between pairs of samples indicate how well one sample can serve as the exemplar of the other, and these messages are iteratively updated in response to the values from other pairs. The updating process continues until convergence, and the final exemplars are used to form the clusters.

Suppose z_i and z_j are features vectors of two food images, we define $s(i, j) = 1 - \frac{z_i \cdot z_j}{\|z_i\| \cdot \|z_j\|}$ as the similarity between z_i and z_j , which is the cosine distance between two vectors. In Affinity Propagation, two messages sent between the sample are responsibility r and availability a . Responsibility $r(i, j)$ describes how well z_j to be the exemplar for x_i , relative to other candidate exemplars for z_i , and it is defined as:

$$r(i, j) = s(i, j) - \max_{j' \neq j} [a(i, j') + s(i, j')]. \quad (4)$$

Availability $a(i, j)$ describes suitability for z_i to pick z_j as its exemplar, considering other samples' preference for z_j as an exemplar, and it is defined as:

$$a(i, j) = \min[0, a(j, j) + \sum_{i' \notin \{i, k\}} r(i', j)]. \quad (5)$$

Both responsibility r and availability a are set to zero at the beginning and updated during each iteration until convergence. The damping factor λ is used during the process to prevent numerical oscillations. The updated functions for both responsibility and availability during each iteration are shown in Equation 6:

$$\begin{aligned} r_{n+1}(i, j) &= \lambda \cdot r_n(i, j) + (1 - \lambda) \cdot r_{n+1}(i, j), \\ a_{n+1}(i, j) &= \lambda \cdot a_n(i, j) + (1 - \lambda) \cdot a_{n+1}(i, j), \end{aligned} \quad (6)$$

where n indicates the iteration step.

3.3 Clustering-based Conditioning

After the input food images are clustered into sub-classes, we denote the new class label of each food image by its food name followed by the cluster ID (e.g., burger_1 for an image of a burger in cluster 1 of its class). A diffusion model is then trained to model the conditional distributions $p(x|y)$, where x is the input image and y is the class label of the input image. To achieve this, we train a conditional denoising autoencoder $\epsilon_\theta(x_t, t, y)$, as mentioned in section 3.1, which allows us to control the image generation process through the class label y . By conditioning the diffusion model on the new class label,

we can generate food images that belong to a specific cluster of a food class, and this helps in generating more diverse food images within each food class.

Similar to the latent diffusion model[65], we use a U-Net with an attention mechanism as our conditional denoising autoencoder. Figure 2 shows how the sub-class labels y_{sub} control the denoising process through U-Net. To incorporate the sub-class information into the denoising process, we use a pre-trained CLIP text encoder τ_θ to project y_{sub} into an embedding vector $\tau_\theta(y_{sub})$, which is then added to the intermediate layers of the U-Net using a cross-attention mechanism. The implementation of the cross-attention mechanism is shown in the following equation:

$$q = \pi_l(z_t) \cdot W_q^l, \quad k = \tau_\theta(y_{sub}) \cdot W_k^l, \quad v = \tau_\theta(y_{sub}) \cdot W_v^l, \quad (7)$$

$$A = \text{softmax}\left(\frac{qk^T}{\sqrt{d}}\right),$$

where q , k , and v are the query, key, and value for the attention mechanism, respectively. l denotes the l -th intermediate layer of U-Net, and $\pi_l(z_t)$ is the intermediate vector representation of U-Net. W_q , W_k , and W_v are learnable projection matrices, A is the attention map of cross-attention mechanism.

The new training objective of the conditional latent diffusion model is shown below:

$$L_{LDM} = \mathbb{E}_{t, \mathcal{E}(z), y_{sub}, \epsilon} [\|\epsilon - \epsilon_\theta(z_t, t, \tau_\theta(y_{sub}))\|^2], \quad (8)$$

where both denoising autoencoder ϵ_θ and text encoder τ_θ are jointly optimized together.

When generating images for a specific class of food, the latent diffusion model is conditioned with its sub-class labels to control the generation process, where the occurrence of the labels will follow the distribution of sub-classes within the class.

4 EXPERIMENTS

In this section, we first train a conditional StyleGAN3[41] on the Food-101[3] dataset as the GAN-based method. We then fine-tune a pre-trained stable diffusion model[65] on the Food-101 dataset as the baseline diffusion method. Finally, we fine-tune a pre-trained stable diffusion with clustering-based conditioning as our proposed ClusDiff. We evaluate ClusDiff and compare it to the GAN-based method and the baseline diffusion method using FID score [34]. Finally, we demonstrate the benefit of using generated synthetic food images as data augmentation by providing a case study to address the class-imbalance issue on the VFN-LT[25] dataset for the long-tailed food image classification.

4.1 Dataset

4.1.1 Food-101. Food-101[3] dataset consists of 101 food categories and each category contains 1000 images. The dataset contains a small amount of noisy data which are images with intense colors or wrong labels. All the images in Food-101 were rescaled to have a maximum side dimension of 512 pixels. We further resize the images to 512 by 512 pixels for training generative models.

4.1.2 VFN-LT. VFN-LT[25] dataset is the long-tailed version of the VFN dataset[54] constructed by removing data from each food class based on the food consumption frequency[48] in the real world.

Datasets	Food-101
Methods	FID
StyleGAN3[41]	39.05
Finetuned Latent Diffusion [65]	30.39
ClusDiff	27.73

Table 1: Food image generation results in terms of FID. The lower FID score indicates better diversity and fidelity of generated images and the best result is marked in bold.

Overall, the training set of VFN-LT is highly imbalanced containing 2.5K images from 74 common food classes with a maximum of 288 images per class and a minimum of 1 image per class. The test set of VFN-LT is kept balanced as 25 images per class for the long-tailed classification problem.

4.2 Evaluation Metric

4.2.1 Frechet Inception Distance Metric (FID). FID[34] is a metric that calculates the distance between the feature vectors of training and generated images, captured by a pre-trained Inception-v3[74] model. The FID measures the similarity between the features extracted from the training images and the generated images. A lower FID score indicates the feature distributions of generated images are close to real images, suggesting that the generator well captures the data distribution of the training images. A higher FID score indicates the generated images deviate from the real images, meaning the generator does not have a good performance.

4.2.2 Long-tailed classification. We use Top-1 image classification accuracy as the evaluation metric. In addition, we provide the performance on both head (instance-rich) and tail (instance-rare) classes. Following the benchmark protocol [25], there are 22 head classes and 52 tail classes in the VFN-LT dataset.

4.3 Implementation Details

4.3.1 Food image generation. The StyleGAN3 model is trained on the Food-101 dataset for 150 epochs with a batch size of 8, and the fixed learning rate is set to 0.002. The reason we did not use pre-trained weights for StyleGAN3 is that the pre-trained models of StyleGAN3 have only been trained on face images, which does not contribute effectively to food image generation. Both the baseline diffusion method and the clustering-based conditional diffusion method (ClusDiff) are fine-tuned on the pre-trained weight 'stable-diffusion-v1-4' for 100 epochs with a batch size of 4, and the fixed learning rate is set to $1e-5$. The pre-trained weight of the stable diffusion model is trained on a subset of the LAION-5B dataset [67], which consists of 5.85 billion general image-text pairs, including some food images. It is a common practice to finetune this pre-trained weight for downstream applications.

4.3.2 Long-tailed classification. We use ResNet-18 [33] pre-trained on ImageNet [66] as the backbone and train the model for 150 epochs with batch size 128 using the SGD optimizer. The initial learning rate is set as 0.1 and decreased to 0.0001 using a cosine learning rate scheduler.

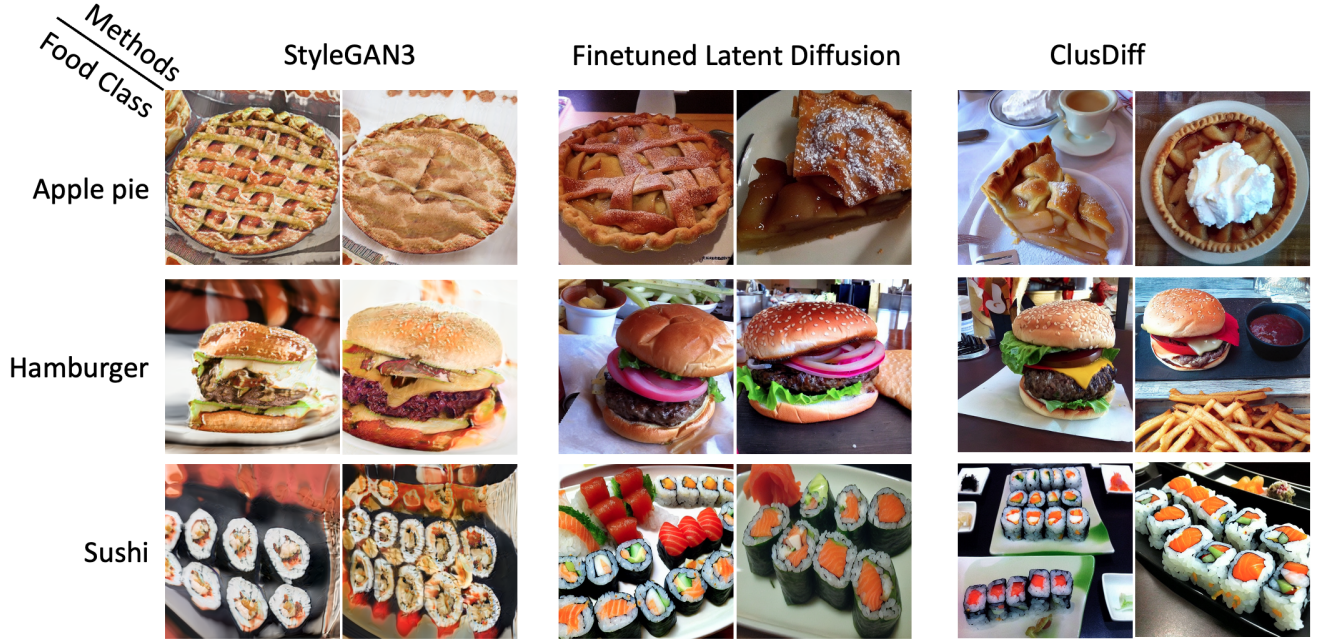


Figure 3: Examples of synthetic food images generated by each model. We manually selected three classes for demonstration, and the synthetic food images were randomly chosen. We can observe significant artifacts in the background of the food images generated by StyleGAN3. Furthermore, both diffusion-based methods, finetuned latent diffusion and ClusDiff, generate more distinct and refined details of the food images, such as the ingredients of the burger and the filling of the sushi.

4.4 Image Generation Results on Food-101

Compared methods. We compare our proposed ClusDiff with two other techniques: one of the latest GAN-based methods, StyleGAN3, and the baseline latent diffusion method. StyleGAN architectures have shown an ability to generate highly realistic images across multiple applications, including generating food images [22]. Therefore, we choose the latest variant, StyleGAN3 [41], as the GAN-based method used for comparison. As for the baseline diffusion method, we select the latent diffusion model[65] to showcase the effectiveness of our proposed clustering-based training framework. This will allow us to demonstrate the improvements achieved by our modifications compared to the baseline.

We generate 100 synthetic images for each type of food in the Food-101 dataset using each image generation method and compare the FID scores. Fig. 3 shows the sample results of synthetic food images generated by each method. We manually selected three classes of food for demonstration, and the generated food images were randomly chosen. Table 1 shows the comparison results for different food image generation methods.

From Fig. 3, we can see that compared to the synthetic food images generated by StyleGAN3, the images generated by the latent diffusion method show significant improvement in the background of food images. Furthermore, the latent diffusion method generates more refined details in the food region compared to StyleGAN3, as evident by the intricate textures of the sushi fillings and the distinct presentation of hamburger ingredients. The FID scores from Table 1 further confirm our visual inspections, showing that the baseline

diffusion method outperforms StyleGAN3 by approximately 10 FID points.

As shown in Fig. 3, the visual difference between the synthetic images generated by the fine-tuned latent diffusion method and our proposed ClusDiff, is minimal. This is because ClusDiff maintains the original network structure of the latent diffusion model.

The comparison between ClusDiff and the baseline latent diffusion model is shown in Table 1. Since both methods are fine-tuned on the same pre-trained weight, this comparison demonstrates that ClusDiff significantly enhances the FID results of the latent diffusion model. This improvement indicates that our proposed modifications in ClusDiff effectively assist the latent diffusion model in capturing the underlying distribution from the training images, leading to better diversity and fidelity in the generated images.

4.5 Long-tailed Food Classification Results on VFN-LT

We demonstrate the effective use of synthetic food images in addressing the challenges of class imbalance associated with the long-tailed classification problem. Specifically, we generate food images for each of the 74 food classes in the VFN-LT dataset to construct a balanced training set. The food images are generated by using a pre-trained stable diffusion model with our clustering-based conditioning on the Food-101 dataset for a fair evaluation.

Compared methods. We compare with existing long-tailed image classification methods, including random over-sampling

Datasets	VFN-LT		
	Head (%)	Tail (%)	Overall (%)
Baseline	62.3	24.4	35.8
ROS[76]	61.7	24.9	35.9
RUS[5]	54.6	26.3	34.8
CMO[58]	60.8	33.6	42.1
LDAM[6]	60.4	29.7	38.9
BS[63]	61.3	32.9	41.9
IB[59]	60.2	30.8	39.6
Focal[49]	60.1	28.3	37.8
Food2stage[25]	61.9	37.8	45.1
Baseline + ClusDiff	68.7	42.4	49.5

Table 2: Top-1 accuracy on VFN-LT with tail class (Tail) and head class (Head) accuracy. The best results are marked in bold.

(ROS)[76], random under-sampling (RUS) [5], context-rich oversampling (CMO)[58] and visual-aware hybrid sampling (Food2stage)[25] as data augmentation-based approaches. In addition, several loss functions have been commonly used to address class imbalance including label-distribution-aware margin loss (LDAM) [6], balanced Softmax (BS) [63], influence-balanced loss (IB) [59] and Focal loss [49]. We use vanilla training as the baseline and our method as Baseline + ClusDiff by using a balanced training set containing original and synthetic food images generated by our proposed ClusDiff.

Table 2 summarizes the results on the VFN-LT dataset. Our method achieves the best performance in terms of head, tail, and overall classification accuracy. Compared with existing data augmentation based approaches, using synthetic images is more effective, which improves the generalization ability and avoids the over-fitting issue for both instance-rich and instance-rare classes. Moreover, the use of synthetic data streamlines the training process, eliminating the need for any new loss functions, thus offering an advantage over existing methods and showing great potential for real-life applications.

5 CONCLUSION AND FUTURE WORK

In this paper, we initially demonstrate the promising performance of latent diffusion methods in food image generation by comparing their results with one of the latest GAN-based methods, StyleGAN3. We then propose a clustering-based conditional training framework, ClusDiff, to address the challenge of high intra-class variance within the food dataset. This framework enables us to generate more diverse and representative food images that better reflect the underlying distribution of the training data. Experimental results show that our proposed method outperforms the baseline diffusion model, showcasing its effectiveness in enhancing food image generation. Additionally, we conduct a thorough study on the utilization of synthetic food images to address the class-imbalance issue in long-tailed food classification.

In the future, we plan to investigate integrating cluster information into the loss function or network structure of the diffusion model to better capture the variance of food images. This approach

could potentially lead to further improvements in the quality and diversity of the generated food images.

REFERENCES

- [1] Kiyoharu Aizawa and Makoto Ogawa. 2015. FoodLog: Multimedia Tool for Healthcare Applications. *IEEE MultiMedia* 22, 2 (April 2015), 4–8.
- [2] Regan L Bailey. 2021. Overview of dietary assessment methods for measuring intakes of foods, beverages, and dietary supplements in research studies. *Current Opinion in Biotechnology* 70 (August 2021), 91–96.
- [3] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. 2014. Food-101—mining discriminative components with random forests. *Proceedings of the European Conference on Computer Vision* (September 2014), 446–461. Zurich, Switzerland.
- [4] CJ Boushey, M Spoden, FM Zhu, EJ Delp, and DA Kerr. 2017. New mobile methods for dietary assessment: review of image-assisted and image-based dietary assessment methods. *Proceedings of the Nutrition Society* 76, 3 (August 2017), 283–294. London, UK.
- [5] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. 2018. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks* (October 2018), 249–259.
- [6] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Archiga, and Tengyu Ma. 2019. Learning imbalanced datasets with label-distribution-aware margin loss. *Advances in Neural Information Processing Systems* 32 (December 2019), 1567–1578.
- [7] Eduardo Castro, Jaime S Cardoso, and Jose Costa Pereira. 2018. Elastic Deformations for Data Augmentation in Breast Cancer Mass Detection. *Proceedings of the IEEE EMBS International Conference on Biomedical & Health Informatics* (March 2018), 230–234. Las Vegas, NV.
- [8] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. 2023. Executing your Commands via Motion Diffusion in Latent Space. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (June 2023), 18000–18010. Vancouver, Canada.
- [9] Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems* 34 (December 2021), 8780–8794.
- [10] K. W. Dunn, C. Fu, D. J. Ho, S. Lee, S. Han, P. Salama, and E. J. Delp. 2019. DeepSynth: Three-dimensional nuclear segmentation of biological images using neural networks trained with synthetic data. *Scientific Reports* 9, 1 (December 2019), 18295–18309.
- [11] Takumi Ege and Keiji Yanai. 2018. Multi-task learning of dish detection and calorie estimation. *Proceedings of the Joint Workshop on Multimedia for Cooking and Eating Activities and Multimedia Assisted Dietary Management* (July 2018), 53–58. Stockholm, Sweden.
- [12] Patrick Esser, Robin Rombach, and Bjorn Ommer. 2021. Taming transformers for high-resolution image synthesis. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (June 2021), 12873–12883. Virtual.
- [13] S. Fang, Z. Shao, D. A. Kerr, C. J. Boushey, and F. Zhu. 2019. An End-to-End Image-Based Automatic Food Energy Estimation Technique Based on Learned Energy Distribution Images: Protocol and Methodology. *Nutrients* 11, 4 (April 2019), 877.
- [14] Shaobo Fang, Zeman Shao, Runyu Mao, Chichen Fu, Edward J Delp, Fengqing Zhu, Deborah A Kerr, and Carol J Boushey. 2018. Single-View Food Portion Estimation: Learning Image-to-Energy Mappings Using Generative Adversarial Networks. *Proceedings of the IEEE International Conference on Image Processing* (October 2018), 251–255. Athens, Greece.
- [15] Shaobo Fang, Fengqing Zhu, Chufan Jiang, Song Zhang, Carol J Boushey, and Edward J Delp. 2016. A comparison of food portion size estimation using geometric models and depth images. *Proceedings of the IEEE International Conference on Image Processing (ICIP)* (September 2016), 26–30. Phoenix, AZ.
- [16] Brendan J. Frey and Delbert Dueck. 2007. Clustering by Passing Messages Between Data Points. *Science* 315, 5814 (February 2007), 972–976.
- [17] Wenjin Fu, Yue Han, Jiangpeng He, Sriram Baireddy, Mridul Gupta, and Fengqing Zhu. 2023. Conditional synthetic food image generation. *Proceedings of the IS&T International Symposium on Electronic Imaging* 35 (January 2023), 1–6. San Francisco, CA.
- [18] Jixiang Gao, Jingjing Chen, Huazhu Fu, and Yu-Gang Jiang. 2022. Dynamic Mixup for Multi-Label Long-Tailed Food Ingredient Recognition. *IEEE Transactions on Multimedia* (June 2022), 1–10.
- [19] Sicheng Gao, Xuhui Liu, Bohan Zeng, Sheng Xu, Yanjing Li, Xiaoyan Luo, Jianzhuang Liu, Xiantong Zhen, and Baochang Zhang. 2023. Implicit diffusion models for continuous super-resolution. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (June 2023), 10021–10030. Vancouver, Canada.
- [20] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. *Proceedings of the International Conference on Neural Information Processing Systems* (December 2014), 2672–2680. Montreal, Canada.

- [21] Fangda Han, Ricardo Guerrero, and Vladimir Pavlovic. 2020. CookGAN: Meal Image Synthesis from Ingredients. *Proceedings of the IEEE Winter Conference on Applications of Computer Vision* (March 2020), 1450–1458. Colorado, USA.
- [22] Fangda Han, Guoyao Hao, Ricardo Guerrero, and Vladimir Pavlovic. 2020. MPG: A Multi-ingredient Pizza Image Generator with Conditional StyleGANs. *arXiv preprint arXiv:2012.02821* (December 2020).
- [23] Yue Han, Yang Lei, Viktor Shkolnikov, Daisy Xin, Alicia Auduong, Steven Barcelo, Jan Allebach, and Edward J Delp. 2023. An Ensemble Method With Edge Awareness for Abnormally Shaped Nuclei Segmentation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops* (June 2023), 4314–4324. Vancouver, Canada.
- [24] Yue Han, Sri Kalyan Yarlagadda, Tonmoy Ghosh, Fengqing Zhu, Edward Sazonov, and Edward J Delp. 2021. Improving Food Detection For Images From a Wearable Egocentric Camera. *Proceedings of the IS&T International Symposium on Electronic Imaging 33* (January 2021), 1–7. Burlingame, CA.
- [25] Jiangpeng He, Luotao Lin, Heather A Eicher-Miller, and Fengqing Zhu. 2023. Long-Tailed Food Classification. *Nutrients* 15, 12 (June 2023), 2751.
- [26] Jiangpeng He, Luotao Lin, Jack Ma, Heather A. Eicher-Miller, and Fengqing Zhu. 2023. Long-Tailed Continual Learning For Visual Food Recognition. *arXiv preprint arXiv:2307.00183* (June 2023).
- [27] Jiangpeng He, Runyu Mao, Zeman Shao, Janine L Wright, Deborah A Kerr, Carol J Boushey, and Fengqing Zhu. 2021. An end-to-end food image analysis system. *Proceedings of the IS&T International Symposium on Electronic Imaging 2021*, 8 (January 2021), 285–1. Burlingame, CA.
- [28] Jiangpeng He, Runyu Mao, Zeman Shao, and Fengqing Zhu. 2020. Incremental Learning In Online Scenario. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (June 2020), 13926–13935. Virtual.
- [29] Jiangpeng He, Zeman Shao, Janine Wright, Deborah Kerr, Carol Boushey, and Fengqing Zhu. 2020. Multi-task Image-Based Dietary Assessment for Food Recognition and Portion Size Estimation. *Proceedings of the IEEE Conference on Multimedia Information Processing and Retrieval* (August 2020), 49–54. Virtual.
- [30] Jiangpeng He and Fengqing Zhu. 2021. Online Continual Learning for Visual Food Classification. *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops* (October 2021), 2337–2346. Virtual.
- [31] Jiangpeng He and Fengqing Zhu. 2022. Exemplar-Free Online Continual Learning. *2022 IEEE International Conference on Image Processing* (October 2022), 541–545.
- [32] Jiangpeng He and Fengqing Zhu. 2023. Single-Stage Heavy-Tailed Food Classification. *arXiv preprint arXiv:2307.00182* (July 2023).
- [33] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (June 2016), 770–778. Las Vegas, NV.
- [34] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems* 30 (December 2017), 6629–6640.
- [35] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising Diffusion Probabilistic Models. *Advances in Neural Information Processing Systems* 33 (December 2020), 6840–6851.
- [36] Daichi Horita, Wataru Shimoda, and Keiji Yanai. 2019. Unseen Food Creation by Mixing Existing Food Images with Conditional StyleGAN. *Proceedings of the International Workshop on Multimedia Assisted Dietary Management* (October 2019), 19–24. Nice, France.
- [37] Yoshifumi Ito, Wataru Shimoda, and Keiji Yanai. 2018. Food Image Generation Using a Large Amount of Food Images with Conditional GAN: RamenGAN and RecipeGAN. *Proceedings of the Joint Workshop on Multimedia for Cooking and Eating Activities and Multimedia Assisted Dietary Management* (July 2018), 71–74. Stockholm, Sweden.
- [38] Landu Jiang, Bojia Qiu, Xue Liu, Chenxi Huang, and Kunhui Lin. 2020. DeepFood: food image analysis and dietary assessment via deep model. *IEEE Access* 8 (February 2020), 47477–47489.
- [39] Shuqiang Jiang, Weiqing Min, Yongqiang Lyu, and Linhu Liu. 2020. Few-Shot Food Recognition via Multi-View Representation Learning. *ACM Transactions on Multimedia Computing, Communications, and Applications* 16, 3 (July 2020), 1–20.
- [40] Fahad Jubayer, Janibul Alam Soeb, Abu Naser Mojumder, Mitun Kanti Paul, Pranta Barua, Shahidullah Kayshar, Syeda Sabrina Akter, Mizanur Rahman, and Amirul Islam. 2021. Detection of mold on the food surface using YOLOv5. *Current Research in Food Science* 4 (October 2021), 724–728.
- [41] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2021. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems* 34 (December 2021), 852–863.
- [42] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020. Analyzing and Improving the Image Quality of StyleGAN. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (March 2020), 8110–8119. Colorado, USA.
- [43] Yoshiyuki Kawano and Keiji Yanai. 2014. Automatic expansion of a food image dataset leveraging existing categories with domain adaptation. *Proceedings of the European Conference on Computer Vision Workshop* (September 2014), 3–17. Zurich, Switzerland.
- [44] Yoshiyuki Kawano and Keiji Yanai. 2014. FoodCam: A Real-Time Mobile Food Recognition System Employing Fisher Vector. *Proceedings of the International Conference on Multimedia Modeling* (January 2014), 369–373. Dublin, Ireland.
- [45] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. 2023. Imagic: Text-based real image editing with diffusion models. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (June 2023), 6007–6017. Vancouver, Canada.
- [46] Timothy J Key, Kathryn E Bradbury, Aurora Perez-Cornago, Rashmi Sinha, Konstantinos K Tsilidis, and Shoichiro Tsugane. 2020. Diet, nutrition, and cancer risk: what do we know and what is the way forward? *BMJ* 368 (March 2020), m511.
- [47] Fanyu Kong and Jindong Tan. 2012. DietCam: Automatic dietary assessment with mobile camera phones. *Pervasive and Mobile Computing* 8, 1 (February 2012), 147 – 163.
- [48] Luotao Lin, Fengqing Zhu, Edward Delp, and Heather Eicher-Miller. 2021. The Most Frequently Consumed and the Largest Energy Contributing Foods of US Insulin Takers Using NHANES 2009–2016. *Current Developments in Nutrition* 5 (June 2021), 426–426.
- [49] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. *Proceedings of the IEEE International Conference on Computer Vision* (July 2017), 2980–2988. Honolulu, HI.
- [50] Ya Lu, Thomai Stathopoulou, Maria F Vasiloglou, Lillian F Pinault, Colleen Kiley, Elias K Spanakis, and Stavroula Mougiakakou. 2020. goFOODTM: an artificial intelligence system for dietary assessment. *Sensors* 20, 15 (July 2020), 4283.
- [51] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. 2022. Repaint: Inpainting using denoising diffusion probabilistic models. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (June 2022), 11461–11471. New Orleans, LA.
- [52] Shitong Luo and Wei Hu. 2021. Diffusion probabilistic models for 3d point cloud generation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (June 2021), 2837–2845. Virtual.
- [53] Runyu Mao, Jiangpeng He, Luotao Lin, Zeman Shao, Heather A. Eicher-Miller, and Fengqing Maggie Zhu. 2021. Improving Dietary Assessment Via Integrated Hierarchy Food Classification. *Proceedings of the IEEE International Workshop on Multimedia Signal Processing* (June 2021), 1–6. Virtual.
- [54] Runyu Mao, Jiangpeng He, Zeman Shao, Sri Kalyan Yarlagadda, and Fengqing Zhu. 2021. Visual aware hierarchy based food recognition. *Proceedings of the International Conference on Pattern Recognition Workshop* (February 2021), 571–598. Virtual.
- [55] Weiqing Min, Zhiling Wang, Yuxin Liu, Mengjiang Luo, Liping Kang, Xiaoming Wei, Xiaolin Wei, and Shuqiang Jiang. 2023. Large scale visual food recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45 (August 2023), 9932–9949.
- [56] Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (November 2014).
- [57] Daniel Mas Montserrat, Qian Lin, Jan Allebach, and Edward J Delp. 2017. Training Object Detection and Recognition CNN Models Using Data Augmentation. *Proceedings of the IS&T International Symposium on Electronic Imaging 2017*, 10 (January 2017), 27–36. Burlingame, CA.
- [58] Seulki Park, Youngkyu Hong, Byeongho Heo, Sangdoo Yun, and Jin Young Choi. 2022. The Majority Can Help The Minority: Context-rich Minority Oversampling for Long-tailed Classification. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (June 2022), 6887–6896. New Orleans, LA.
- [59] Seulki Park, Jongin Lim, Younghun Jeon, and Jin Young Choi. 2021. Influence-balanced loss for imbalanced visual classification. *Proceedings of the IEEE/CVF International Conference on Computer Vision* (June 2021), 735–744. Virtual.
- [60] Walter HL Pinaya, Petru-Daniel Tudosiu, Jessica Dafflon, Pedro F Da Costa, Virginia Fernandez, Parashkev Nachev, Sebastien Ourselin, and M Jorge Cardoso. 2022. Brain imaging generation with latent diffusion models. *Proceedings of the MICCAI Workshop on Deep Generative Models* (September 2022), 117–126. Singapore, Singapore.
- [61] Siddeshwar Raghavan, Jiangpeng He, and Fengqing Zhu. 2023. Online Class-Incremental Learning For Real-World Food Classification. *arXiv preprint arXiv: 2301.05246* (January 2023).
- [62] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. You only look once: Unified, real-time object detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (June 2016), 779–788. Las Vegas, NV.
- [63] Jiawei Ren, Cunjun Yu, Xiao Ma, Haiyu Zhao, Shuai Yi, et al. 2020. Balanced meta-softmax for long-tailed visual recognition. *Advances in Neural Information Processing Systems* 33 (December 2020), 4175–4186.
- [64] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2017. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39, 6 (June 2017), 1137–1149.
- [65] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (June 2022), 10684–10695. New Orleans, LA.

- [66] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* 115, 3 (2015), 211–252.
- [67] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems* 35 (November 2022), 25278–25294.
- [68] Zeman Shao, Shaobo Fang, Runyu Mao, Jiangpeng He, Janine L Wright, Deborah A Kerr, Carol J Boushey, and Fengqing Zhu. 2021. Towards learning food portion from monocular images with cross-domain feature adaptation. *Proceedings of the IEEE 23rd International Workshop on Multimedia Signal Processing (MMSP)* (October 2021), 1–6. Tampere, Finland.
- [69] Zeman Shao, Yue Han, Jiangpeng He, Runyu Mao, Janine Wright, Deborah Kerr, Carol Jo Boushey, and Fengqing Zhu. 2021. An Integrated System for Mobile Image-Based Dietary Assessment. *Proceedings of the 3rd Workshop on AIxFood* (October 2021), 19–23. Virtual.
- [70] Zeman Shao, Gautham Vinod, Jiangpeng He, and Fengqing Zhu. 2023. An End-to-end Food Portion Estimation Framework Based on Shape Reconstruction from Monocular Image. *arXiv preprint arXiv:2308.01810* (July 2023).
- [71] Connor Shorten and Taghi M Khoshgoftaar. 2019. A survey on image data augmentation for deep learning. *Journal of Big Data* 6, 1 (July 2019), 1–48.
- [72] Geeta Shroff, Asim Smailagic, and Daniel P Siewiorek. 2008. Wearable context-aware food recognition for calorie monitoring. *Proceedings of the IEEE International Symposium on Wearable Computers* (September 2008), 119–120. Pittsburgh, PA.
- [73] Mingui Sun, Wenyan Jia, Guangzong Chen, Mingke Hou, Jiacheng Chen, and Zhi-Hong Mao. 2022. Improved wearable devices for dietary assessment using a new camera system. *Sensors* 22, 20 (September 2022), 8006.
- [74] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (June 2016), 2818–2826. Las Vegas, NV.
- [75] Hoang Thanh-Tung and Truyen Tran. 2020. Catastrophic forgetting and mode collapse in GANs. *Proceedings of the International Joint Conference on Neural Networks (IJCNN)* (July 2020), 1–10. Glasgow, UK.
- [76] Jason Van Hulse, Taghi M Khoshgoftaar, and Amri Napolitano. 2007. Experimental perspectives on learning from imbalanced data. *Proceedings of the 24th International Conference on Machine Learning* (June 2007), 935–942. Corvallis, OR.