

# Advancing Food Nutrition Estimation via Visual-Ingredient Feature Fusion

Huiyan Qi  
Singapore Management University  
Singapore  
qhyan06040925@gmail.com

Bin Zhu †  
Singapore Management University  
Singapore  
binzhu@smu.edu.sg

Chong-Wah Ngo  
Singapore Management University  
Singapore  
cwngo@smu.edu.sg

Jingjing Chen  
Fudan University  
Shanghai, China  
chenjingjing@fudan.edu.cn

Ee-Peng Lim  
Singapore Management University  
Singapore  
eplim@smu.edu.sg

## ABSTRACT

Nutrition estimation is an important component of promoting healthy eating and mitigating diet-related health risks. Despite advances in tasks such as food classification and ingredient recognition, progress in nutrition estimation is limited due to the lack of datasets with nutritional annotations. To address this issue, we introduce FastFood, a dataset with 84,446 images across 908 fast food categories, featuring ingredient and nutritional annotations. In addition, we propose a new lampshade-agnostic Visual-Ingredient Feature Fusion (VIF<sup>2</sup>) method to enhance nutrition estimation by integrating visual and ingredient features. Ingredient robustness is improved through synonym replacement and resampling strategies during training. The ingredient-aware visual feature fusion module combines ingredient features and visual representation to achieve accurate nutritional prediction. During testing, ingredient predictions are refined using large multimodal models by data augmentation and majority voting. Our experiments on both FastFood and Nutrition5k datasets validate the effectiveness of our proposed method built in different backbones (e.g., Resnet, InceptionV3 and ViT), which demonstrates the importance of ingredient information in nutrition estimation. <https://huiyanqi.github.io/fastfood-nutrition-estimation/>.

## CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; *Neural networks*.

## KEYWORDS

Nutrition estimation, ingredient recognition, dataset

### ACM Reference Format:

Huiyan Qi, Bin Zhu †, Chong-Wah Ngo, Jingjing Chen, and Ee-Peng Lim. 2025. Advancing Food Nutrition Estimation via Visual-Ingredient Feature

† indicates corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

ICMR '25, June 30–July 3, 2025, Chicago, IL, USA

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1877-9/2025/06...\$15.00

<https://doi.org/10.1145/3731715.3733269>

Fusion. In *Proceedings of the 2025 International Conference on Multimedia Retrieval (ICMR '25)*, June 30–July 3, 2025, Chicago, IL, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3731715.3733269>

## 1 INTRODUCTION

Food plays a significant role in promoting healthy living, and related research topics attract increasing attention [30], such as food classification [1, 3, 11, 48], ingredient recognition [4, 5, 12, 25], cross-modal recipe retrieval [33, 36, 49], recipe generation [7, 22, 33, 38], and nutrition estimation [17, 26, 39]. A key objective of food research is nutrition estimation, which empowers individuals to make healthier dietary choices and mitigate the risk of diet-related diseases [2, 40]. Various datasets have been constructed in food domain, such as Food-101 [1], VIREO Food-172 and 251 [4, 6], ISIA Food-500 [31], and Food2K [32], these datasets provide food categories and ingredient labels. In contrast, datasets with nutrition labels are rare due to the much more challenging collection process. Nutrition5k [39] is a dataset comprising food from cafeterias and includes nutritional annotations, such as calorie and protein, but the data is still limited [14, 42]. Furthermore, current nutrition estimation methods heavily rely on precise portion size estimation, which is challenging to predict accurately. Additionally, methods based on depth images [27], while promising, often rely on datasets with limited variation, making them inadequate for handling complex and diverse dishes, thus limiting their generalization ability.

To address the limitations, we introduce a new dataset, FastFood, which focuses on fast food by collecting food images of popular fast food brands and their corresponding nutritional information. FastFood contains 908 different categories with 84,446 images in total, such as burgers, snacks, beverages, and desserts. In addition, we employ a semi-automatic approach to produce the ingredient labels for each image. We request GPT-4o to generate ingredients for the provided fast food first, following a manual refinement and correction to ensure the annotation quality.

In this paper, we further propose a novel model-agnostic Visual-Ingredient Feature Fusion (VIF<sup>2</sup>) method by integrating image and ingredient features for nutrition estimation. To aggregate ingredient features, we use the text encoder of the pre-trained CLIP model to extract embeddings for all ingredients and align them with the visual feature space through an ingredient projector. Additionally, we enhance robustness during training by leveraging ingredient synonyms replacement and sampling strategies. Ingredient-aware

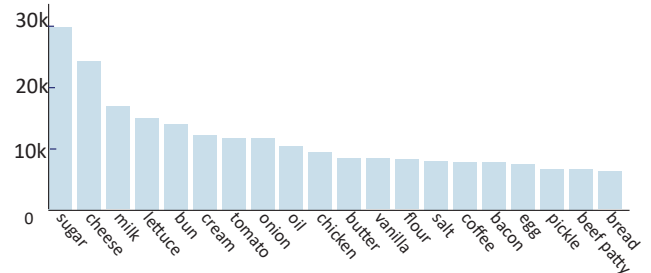
Visual Feature Fusion module is introduced to integrate ingredient features with visual representations extracted from convolutional neural networks (ResNets, InceptionV3) and Transformer-based (ViT) models. Moreover, we utilize separate task-specific heads to regress the nutrition values required for each prediction task. During the testing phase, we use large multimodal models (LMMs) to predict ingredients for test images, refining these predictions through data augmentation and a majority voting mechanism to reduce errors caused by hallucinations. Our proposed method demonstrates that integrating ingredient and visual features can significantly improve nutrition prediction accuracy in our FastFood and Nutrition5k datasets.

## 2 RELATED WORK

Food-related datasets play a crucial role in advancing research in this field. Food-101 [1] created the first Western food dataset, with 101,000 images categorized into 101 classes. VIREO Food-172 and 251 [4, 6] consist of Chinese food and provide ingredient information in addition to category labels. Recipe1M [33] is a more well-known large-scale food dataset offering images, recipes, and ingredient information. It is commonly used for cross-modal retrieval between recipes and images. Food2K [32] is another large-scale dataset in the food domain, providing category information for food images to advance scalable food visual feature learning. UEC Food256 [19] and FoodSeg103 [44] are two datasets designed for food segmentation tasks. DietDiary dataset [13] contains sequences of food images captured by users over multiple days, along with corresponding food weights, for the task of weight prediction. Nutrition5K [39] is widely used for nutrition estimation tasks. It provides food images, videos from different angles, and depth images from two cafeterias, as well as detailed ingredient and nutritional information. These datasets enable a wide range of research, from food image classification [21, 29] to ingredient recognition and recipe retrieval. Nutrition estimation presents greater challenges, requiring accurate recognition of ingredients and their quantities. Early studies, such as Im2Calories [28], use CNNs to estimate food mass and calories. Recent work highlights the effectiveness of multi-task learning, where models simultaneously predict multiple aspects of nutrition. For instance, [9] introduces a multi-task network for estimating calories, categories, and ingredients. Techniques like integrating RGB-Depth images [27] and combining object detection with calorie estimation [10] further enhance accuracy. Meanwhile, Nutrition5K[39] demonstrates the use of depth sensors to improve nutrition prediction. Recently, multimodal large language models [18, 46] have also been explored in the food domain to unify various food-related tasks within a single model.

## 3 FASTFOOD DATASET CONSTRUCTION

The FastFood dataset is constructed from ten highly renowned fast-food brands worldwide, such as McDonald’s and Burger King. We begin by collecting images and nutritional information for each food item directly from the official websites of these fast-food brands. These websites provided low-quality images with clean backgrounds and precise nutritional information, such as calories, fats, carbohydrates, and proteins. In total, we gathered 908 types of food items, such as bacon cheeseburger and French fries.



**Figure 1: Distribution of Top20 ingredients in FastFood dataset.**

Figure 2 provides an overview of the fast-food brands included and showcases the image examples with ingredients and nutritional information.

**Semi-automatic Ingredient Annotation.** We use GPT-4o following a manual correction to derive ingredient annotations in each category of our dataset. The generation process involves providing food images to GPT-4o, accompanied by prompt instructing it to output ingredients relevant to food, while filtering out unrelated terms such as “straw” or “fork”. While the generated ingredients are generally good, they exhibit issues such as duplicates, vagueness, and context mismatch. Specifically, the generated ingredients often include duplication, such as plural forms. To ensure consistency, these variations are standardized according to a predefined vocabulary. Additionally, we also remove irrelevant phrases from the ingredients, such as the extra notes within parentheses, such as “lettuce (200g)”, which is revised to “lettuce”. Vagueness is another issue, where broad terms like “meat” are used instead of more specific ones like “beef patty”. We manually checked each ingredient to make sure the ingredient annotations were accurate and consistent in our dataset.

Figure 1 showcases the top 20 ingredients in the FastFood dataset, with high-frequency ingredients not only serving as foundational ingredients of fast food but also closely tied to their nutrition characteristics. For example, sugar reflects the high-calorie nature, while cheese and milk are rich in fat and protein. This distribution reflects the impact of high-frequency ingredients on the nutritional composition and health attributes of fast food.

**Food Image Expansion.** The images obtained from the official websites are still very limited by numbers and diversity. Thus We expand the dataset via Google image search in two ways: (i) Text-Based Image Retrieval: Searching for related images using each food item’s name as a keyword. (ii) Image-Based Image Retrieval: Searching for related images using the image of each food item collected from the official website. These two methods balance the quality and diversity of the collected images. When searching by food names, the results include images with more complex backgrounds, such as “a person in a car holding a beef burger”. In contrast, image-based image retrieval typically returns higher-quality images that feature less background and prominently display the food item. After collecting the images via Google search, we manually check each crawled image to ensure the quality of the dataset.

Figure 3 shows some examples of images collected using text-based image retrieval and image-based image retrieval. After collection, we exclude images that exhibit blurring, contain multiple

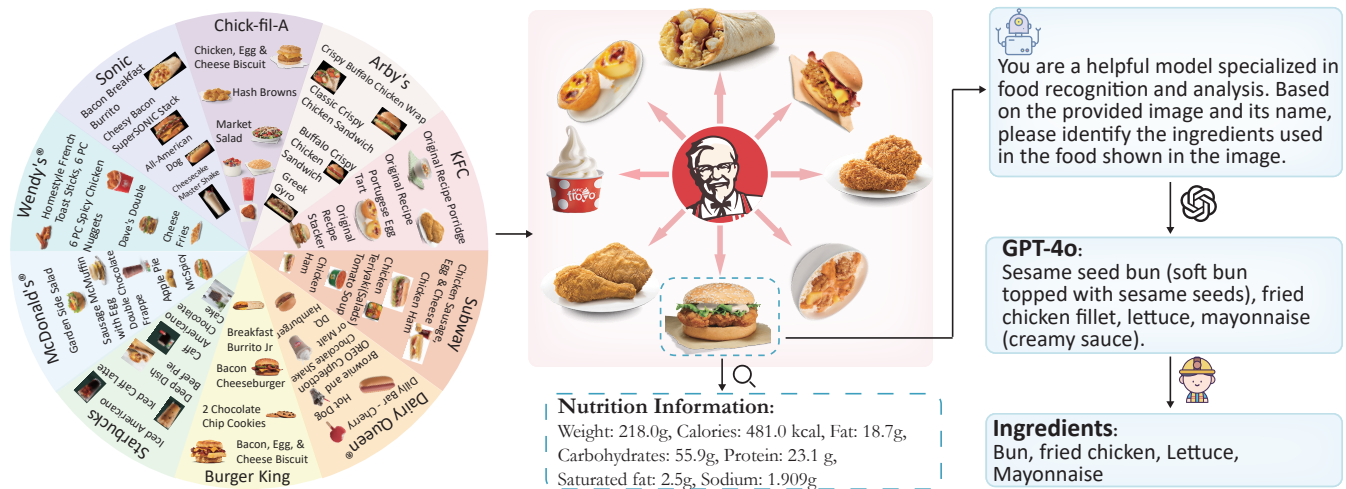


Figure 2: Overview of our FastFood dataset, highlighting brands, sample images, and annotated ingredients and nutrition.



Figure 3: Data Expansion for Food Items Using Different Retrieval Methods. The first row shows official food item images, while the second and third rows illustrate data expansion achieved through Image-Based Image Retrieval and Text-Based Image Retrieval, respectively.

food items (except for combo meals), or mismatch the query. Additionally, we exclude search results where the food size does not match the official image to ensure the accuracy of the nutritional information. For instance, the official website might label a small bottle of milk as fresh milk, but search results using this keyword might include images of large containers of milk. We review the collected nutritional information for each food item and standardize the units of the same nutritional attributes across different fast food brands.

Figure 4 and Figure 5 illustrate the distribution of samples in our dataset across different calorie, fat, carbohydrate, and protein ranges, where fat, carbohydrates, and proteins are predominantly in the 0–10 g, 0–20 g, and 0–8 g ranges, respectively, with fewer samples in higher ranges due to the prevalence of single-item food categories in the FastFood dataset, while combo items contribute to the upper ranges.

We eventually obtain 84,446 images, including images obtained from the official websites. We divide these images into training, validation, and test sets in a 70%, 20%, and 10% split, respectively.

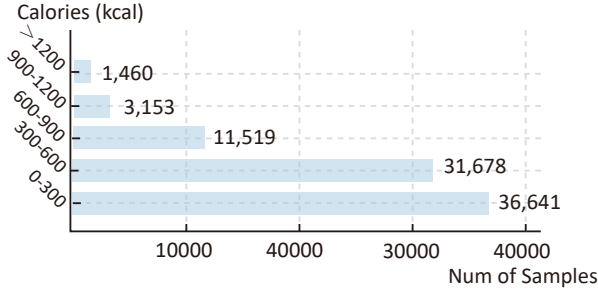


Figure 4: Distribution of calorie in FastFood dataset.

**Generating Diet Suggestion dialogue with GPT-4o.** We selected 5,800 images from our original dataset to generate multi-turn dialogues focused on dietary planning, nutritional analysis, and tailored suggestions using GPT-4o to further show the potential of the dataset. Each image is associated with 2 to 5 rounds of dialogue, resulting in a total of 20,204 dialogue turns, which provide a rich source for training and evaluation of large language models.

The template for generating dietary planning and suggestions data using the GPT-4o API is shown in Figure 6. We provide GPT-4o with the nutritional content of the food and include an example to ensure the model understands our format. In the template, *cal*, *fat*, *carb*, *pro*, and *N* are placeholders that are replaced with the nutritional information and the number of dialogue turns, respectively.

## 4 METHOD

In this section, we introduce Visual-Ingredient Feature Fusion (VIF<sup>2</sup>), a model-agnostic method that integrates visual and ingredient features to enable accurate prediction of food nutrition. Figure 7 illustrates the overall framework of our proposed method.

### 4.1 Robust Ingredient Feature Aggregation

Given ingredient list  $\{\text{ing}_1, \text{ing}_2, \dots, \text{ing}_N\}$  for a food image, we employ CLIP’s text encoder to each ingredient to generate the corresponding text embedding. Then, all the generated embeddings are averaged to obtain the final feature representation of the ingredients  $\mathbf{T}_{\text{ingredients}}$  as follows:

$$\mathbf{T}_{\text{ingredients}} = \frac{1}{N} \sum_{i=1}^N \text{CLIP}_{\text{text}}(\text{ing}_i). \quad (1)$$

To ensure that the ingredient features and image features are mapped to the same feature space, as shown in Eq.2, we apply a linear layer and an activation function to  $\mathbf{T}_{\text{ingredients}}$ , where  $\mathbf{W}$  and  $\mathbf{b}$  are learnable parameters.

$$\mathbf{T}_{\text{transformed}} = \text{ReLU}(\mathbf{W} \cdot \mathbf{T}_{\text{ingredients}} + \mathbf{b}). \quad (2)$$

To make our model more robust for ingredient feature aggregation, we employ GPT-4o to generate a set of synonyms for each ingredient. We ensure that the generated synonyms are semantically similar to the original ingredient name but do not expand into overly general concepts. For example, synonyms for “lettuce” could be “romaine lettuce” but not a more generalized term like “vegetables”. This approach aims to preserve the semantic specificity

of ingredients while avoiding the introduction of overly vague or unrelated variants. In the training phase, we apply the following two strategies to the ingredient list for each sample: (1) Individual Ingredient Replacement: For each ingredient, we randomly select one of its synonyms with a 50% probability to replace the original word. This operation enhanced the model’s robustness to variations in word forms. (2) Ingredient Sampling: For each image’s ingredient list, we sample the entire list with a 50% probability, randomly selecting a subset as input. These two strategies aim to simulate realistic noise in user-provided inputs. In practice, ingredients are often imprecise, incomplete, or expressed in free-form natural language, resulting in inherent ambiguity. By introducing controlled noise during training, we encourage the model to generalize better to such uncertain and variable real-world conditions, where even the “ground-truth” ingredient list may contain ambiguities that are difficult to resolve definitively.

### 4.2 Ingredient-aware Visual Feature Fusion

We propose a model-agnostic ingredient-aware visual feature fusion method to integrate image and ingredient features. We select two well-known convolutional models, ResNet [15] and InceptionV3 [37], as well Transformer-based model ViT [8] as our baselines. ResNet and Inception models generate multi-scale feature maps for food images through convolution and pooling operations, transforming the image representation from low to high level [15, 37], thus we select intermediate feature maps  $\mathbf{X}_{\text{img}}$  to fuse ingredient features. Specifically, for ResNet50 and ResNet101, we choose the convolutional output of block2 for feature fusion, while for InceptionV3, we select the feature map after its second pooling layer. To incorporate ingredient information, as shown in Eq.3, the linearly transformed text embeddings  $\mathbf{T}_{\text{transformed}}$  are added channel-wise to the feature map  $\mathbf{X}_{\text{img}}$ , where  $\mathbf{T}_{\text{transformed}} \in \mathbb{R}^C$  is expanded to match the dimensions of the feature map through a broadcasting mechanism.

$$\mathbf{X}_{\text{fused}} = \mathbf{X}_{\text{img}} + \mathbf{T}_{\text{transformed}}. \quad (3)$$

In contrast, images are processed through the patch embedding module to generate a series of patch tokens in ViT model. To facilitate better interaction between ingredient information and image features, we introduce linearly transformed ingredient feature as a separate token, which is input into the transformer encoder alongside the class token and patch tokens. Denote the input sequence  $\mathbf{X}_{\text{tokens}} = [\mathbf{X}_{\text{cls}}, \mathbf{T}_{\text{transformed}}, \mathbf{X}_{\text{patches}}]$ , where  $\mathbf{X}_{\text{cls}}$  represents the classification token used to capture global information, and  $\mathbf{X}_{\text{patches}}$  refers to the patch embeddings. We combine the input sequence with positional encoding  $\mathbf{P}_{\text{pos}}$  and pass it through the Transformer layers:

$$\mathbf{X}_{\text{encoded}} = \text{Transformer}(\mathbf{X}_{\text{tokens}} + \mathbf{P}_{\text{pos}}). \quad (4)$$

Finally, we use the representation of the class token for subsequent nutrition prediction tasks.

### 4.3 Multi-task Nutrition Estimation

Following Nutrition5k [39], we train a separate head for each nutrition (i.e., calories, fats, carbohydrates, and proteins), which is a two-layer fully connected neural network. The first layer projects

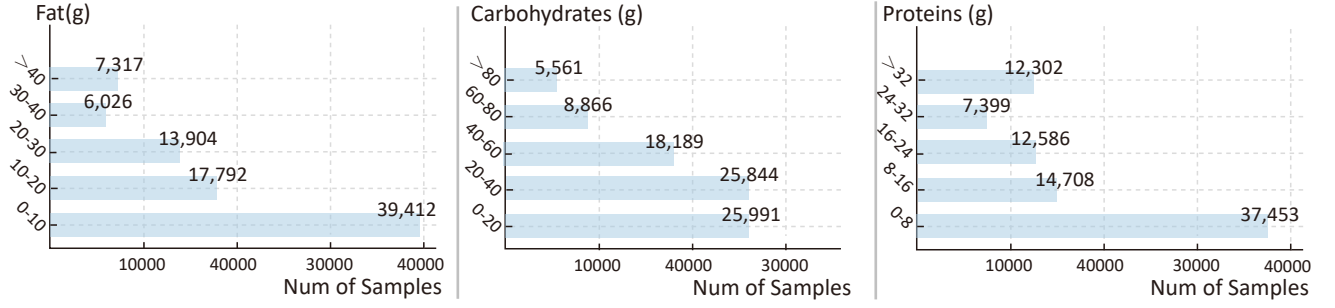


Figure 5: Distribution of samples in the dataset across different ranges of fat, carbohydrates, and protein content.

...I will provide a fast food image and its nutrition information. Please help me generate a multi-round dialogue about the image on dietary planning and suggestions as the model's training data.

For example, nutrition information description like this: TOTAL(194, 15.12, 0.5, 13.3) means this food item has 194 kcal calories and 15.12g fat, 0.5g carbohydrate and 13.3 protein.

**One 1-round example is:** \Image Nutrition Information: TOTAL (500, 23, 49, 24) **USER:** Can you provide some dietary suggestions based on this image? **ASSISTANT:** Sure! To balance this meal, consider pairing it with low-calorie, nutrient-rich side such as a green salad...

**Nutrition Information:** TOTAL ( {cal}, {fat}, {carb}, {pro} ). Please help me generate a {N}-round dialogue closely related to dietary planning and suggestions. The questions are expected to be professional or general, the answers are expected to be explanatory.

Figure 6: Prompt template for generating dietary planning and suggestion data.

the input features to a 4096-dimensional hidden layer, and the second layer maps the hidden layer's output to a single scalar value. A ReLU activation function is applied after the first layer to enhance the model's nonlinear representational capacity.

We employ the Mean Absolute Error (MAE) to train the model as in Eq. 5,  $n_i$  represents the number of samples for task  $i$ ,  $\hat{y}_{i,j}$  and  $y_{i,j}$  are the predicted value and the true value of the  $j_{th}$  sample in task  $i$ , respectively.

$$\text{MAE}(\hat{y}_i, y_i) = \frac{1}{n_i} \sum_{j=1}^{n_i} |\hat{y}_{i,j} - y_{i,j}| \quad (5)$$

The total loss of the model is represented as follows:

$$\mathcal{L} = \sum_{i=1}^N \text{MAE}(\hat{y}_i, y_i). \quad (6)$$

#### 4.4 Augmented Ingredient Prediction

Acquiring ingredient labels is often challenging in practice. Therefore, during the testing phase, we assume only images are available. Consequently, we utilize off-the-shelf LMMs (e.g., LLaVa [20]) to generate the ingredients for these test images. Since hallucination is a common problem with LLMs [16, 24, 41, 45], we notice that the results from current LLMs frequently produce ingredients that do not exist in the image. To address this issue, we introduce augmented ingredient prediction. Specifically, for each test image  $I$ , a set of data augmentation methods  $\{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_K\}$  (e.g. Random Rotation, Horizontal Flip, Random Crop with maintaining at least 70% of the original size, and Convert to Grayscale) are applied to produce  $K$  augmented images:

$$I_k = \mathcal{T}_k(I), \quad k \in \{1, 2, \dots, K\}, \quad (7)$$

$\mathcal{T}_k$  represents the  $k$ -th augmentation method. Subsequently, each augmented image  $I_k$  is fed into the LMM to predict the ingredients  $S_k$ :

$$S_k = \text{LMM}(I_k). \quad (8)$$

Finally, a majority voting mechanism is employed to filter the ingredients from all the predicted results  $S_{\text{all}}$ . We count the frequency of occurrence for each ingredient, and any ingredient with a count exceeding the predefined threshold  $\tau$  is included in the final prediction result:

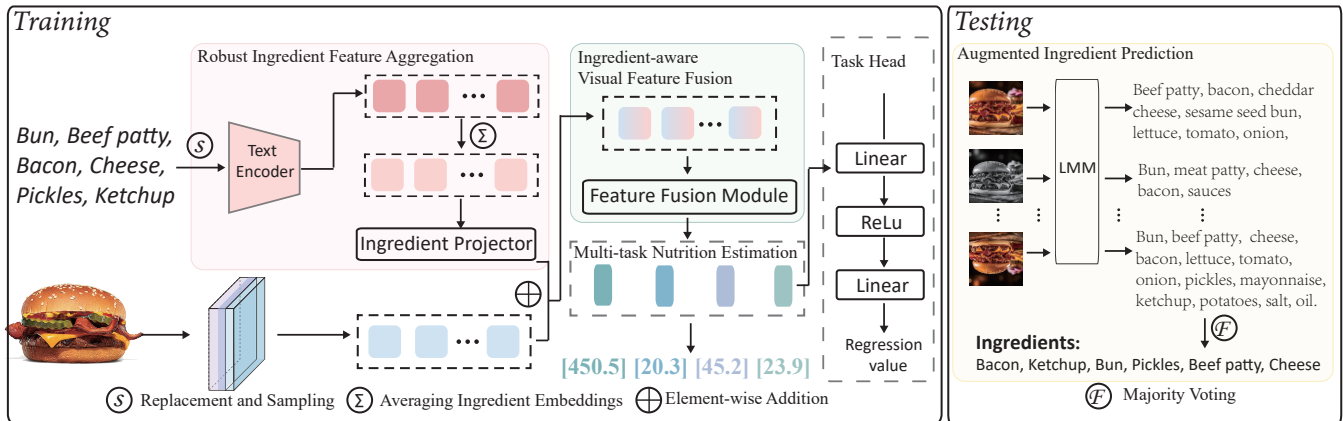
$$S_{\text{final}} = \{ \text{ing} \mid n(\text{ing}) \geq \tau, \text{ing} \in S_{\text{all}} \}. \quad (9)$$

The final predicted ingredients and image are then used as input for our model to predict nutrition values.

## 5 EXPERIMENTS

### 5.1 Experimental Setup

**Datasets and Evaluation.** Our experiments are conducted on our FastFood and Nutrition5k [39] datasets. Nutrition5k provides RGB images, multi-angle videos, and depth images of dishes captured in two cafeterias using sensor devices. It is worth noting that publicly available large-scale datasets with rich nutritional annotations and multimodal data (e.g., RGB, depth, video) are extremely scarce. Nutrition5k is one of the few such datasets, making it a valuable benchmark for evaluating nutrition estimation models. During training, we follow [39] by sampling every five frames from the video data in the Nutrition5k dataset. After extracting video



**Figure 7: Overview of the proposed method for nutrition prediction. During training, the ingredient projector aligns robust ingredient features extracted from the ingredient list with visual features. The transformed ingredient features are fused with visual features through the fusion module to produce a comprehensive representation for multi-task nutrition estimation. During testing, a large language model (LMM) is employed to predict ingredients via data augmentation for nutrition estimation.**

frames, Nutrition5k contains 230,000 images for training, 5,700 for validation, and 38,000 for testing. In all experiments, we evaluate model performance using MAE and also report the Relative Error Percentage, which expresses MAE as a percentage of the corresponding field’s mean value. Additionally, we evaluate Nutrition5k under two testing protocols. First, we sample one frame for each frame in test videos. Second, based on the first protocol, we select the optimal frame per video to mitigate the potential redundancy and noise, such as blurriness or distortions.

**Implementation details.** We adopt the widely used architectures, including ResNet50, ResNet101, InceptionV3 and ViT (base16) as backbones for experiments. When InceptionV3 serves as the backbone network, the input image resolution is set to 299x299, which is the standard input size required for InceptionV3. For the other five models, the input image resolution is 224x224. During training, we optimize the models using the RMSProp algorithm with an initial learning rate of  $1e-4$ , momentum of 0.9, decay of 0.9, and epsilon set to 1.0. All models are initialized with the weights pre-trained on ImageNet. Each model is trained for 100 epochs, with a batch size of 64 for the FastFood dataset and 256 for the Nutrition5k dataset.

## 5.2 Performance Comparison

**Table 1: Performance comparison between baselines and their counterparts enhanced with our proposed VIF<sup>2</sup> on the FastFood dataset.**

Method	Caloric MAE	Fat MAE	Carb MAE	Protein MAE
Resnet101 [15]	118.04 / 29.84%	7.06 / 43.98%	15.57 / 39.14%	7.18 / 40.30%
Resnet50 [15]	118.07 / 29.85%	6.97 / 43.40%	15.32 / 38.52%	7.03 / 39.43%
ViT [8]	141.39 / 35.76%	8.20 / 51.07%	18.76 / 47.17%	8.58 / 48.17%
InceptionV3 [37]	142.57 / 36.04%	10.51 / 65.49%	22.74 / 57.18%	11.62 / 65.17%
Resnet101 + VIF <sup>2</sup>	<b>61.26 / 15.49%</b>	<b>3.60 / 22.40%</b>	<b>8.76 / 22.02%</b>	<b>3.87 / 21.73%</b>
Resnet50 + VIF <sup>2</sup>	63.86 / 16.15%	3.99 / 24.85%	9.67 / 24.32%	4.18 / 23.44%
ViT + VIF <sup>2</sup>	76.62 / 19.38%	4.75 / 29.60%	12.01 / 30.20%	5.29 / 29.71%
InceptionV3 + VIF <sup>2</sup>	80.86 / 20.44%	4.93 / 30.70%	11.02 / 27.72%	4.94 / 27.71%

**Table 2: Performance comparison between baselines and their counterparts with our VIF<sup>2</sup> on the Nutrition5k dataset.**

Method	Caloric MAE	Fat MAE	Carb MAE	Protein MAE
Testing Protocol 1				
Resnet101 [15]	103.03 / 42.81%	9.94 / 69.35%	12.91 / 72.86%	7.56 / 56.56%
Resnet50 [15]	104.68 / 43.49%	9.97 / 69.58%	12.92 / 72.91%	7.64 / 57.09%
ViT [8]	148.47 / 61.69%	11.57 / 80.78%	12.67 / 71.48%	9.49 / 70.98%
InceptionV3 [37]	106.50 / 44.25%	9.93 / 69.33%	13.06 / 73.70%	7.72 / 57.73%
Testing Protocol 1				
Resnet101 + VIF <sup>2</sup>	<b>83.75 / 34.80%</b>	<b>6.94 / 48.45%</b>	5.64 / 31.85%	3.27 / 24.46%
Resnet50 + VIF <sup>2</sup>	84.49 / 35.11%	7.01 / 48.95%	<b>5.63 / 31.77%</b>	<b>3.15 / 23.55%</b>
ViT + VIF <sup>2</sup>	87.31 / 36.28%	7.26 / 50.71%	5.58 / 31.49%	3.24 / 24.19%
InceptionV3 + VIF <sup>2</sup>	90.65 / 37.67%	7.63 / 53.27%	6.57 / 37.05%	3.75 / 28.02%
Testing Protocol 2				
Resnet101 + VIF <sup>2</sup>	<b>23.32 / 9.69%</b>	<b>1.93 / 13.50%</b>	<b>1.57 / 8.87%</b>	<b>0.91 / 6.82%</b>
Resnet50 + VIF <sup>2</sup>	28.84 / 11.98%	2.39 / 16.71%	1.92 / 10.84%	1.08 / 8.04%
ViT + VIF <sup>2</sup>	31.07 / 12.91%	2.58 / 18.05%	1.99 / 11.21%	1.15 / 8.61%
InceptionV3 + VIF <sup>2</sup>	31.81 / 13.22%	2.68 / 18.69%	2.31 / 13.00%	1.32 / 9.83%

Tables 1 and 2 list the results on the FastFood and Nutrition5k dataset respectively. After incorporating our proposed VIF<sup>2</sup>, the performances are significantly improved compared with baseline counterparts for both datasets. For instance, the best-performing model, ResNet101+VIF<sup>2</sup>, achieves a Caloric MAE of 61.26—a remarkable 48.1% reduction from the baseline ResNet101’s 118.04. Additionally, the relative error decreases from 29.84% to 15.49%. Similarly, on the Nutrition5k dataset, ResNet101+VIF<sup>2</sup> achieves the best results, reducing the Caloric MAE to 83.75, an 18.7% improvement over the baseline ResNet101. It is not surprising that the performance of testing protocol 2 shows much better performance than that of testing protocol 1 as we select the best-performed frame for each testing video. In addition, our proposed method Resnet101+VIF<sup>2</sup> with testing protocol 2 demonstrates a significant improvement over state-of-the-art (SOTA) methods, including Google-Nutrition-depth [39], Google-Nutrition-monocular [39], RGB-D Nutrition [35], and Swin-Nutrition [34]. Among these, DPF-Nutrition [14] achieves the best caloric estimation performance on



**Figure 8: Examples of typical samples with low and high prediction errors.**

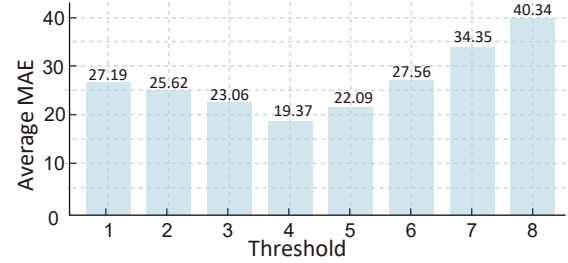
Nutrition5k using monocular images. Compared to this method, our approach reduces caloric MAE by 38.47% and protein MAE by 74.44%, highlighting its effectiveness in food nutrition estimation. Since the evaluation protocol for these methods is unclear, we do not include them in the comparison tables. Under a fair comparison, our proposed method consistently outperforms all baselines.

Table 3 provides the performance of different models combined with VIF<sup>2</sup> when testing with the ground truth ingredients of Fast-Food dataset. Compared to the baselines presented in the main text, all models achieved significant reductions in nutrition prediction. Among them, ResNet101 + VIF<sup>2</sup> performs the best, with a caloric MAE reduced to 44.71 and a relative error of only 11.30%. Similar to the experiment conducted with ingredients generated by other LMMs, ResNet101 remains the best-performing model, while the other models perform slightly worse than ResNet101. This result can be regarded as the upper bound of VIF<sup>2</sup> under ideal conditions, providing a reference for evaluating the impact of generated ingredients on model performance and demonstrating the general applicability of VIF<sup>2</sup> across different architectures.

**Table 3: Performance comparison of different architectures combined with VIF<sup>2</sup> using ground truth ingredients from the FastFood dataset during testing.**

Method	Caloric MAE	Fat MAE	Carb MAE	Protein MAE
Resnet101 + VIF <sup>2</sup>	44.71 / 11.30%	2.59 / 16.13%	7.06 / 17.75%	2.85 / 15.99%
Resnet50 + VIF <sup>2</sup>	47.39 / 11.98%	3.08 / 19.21%	8.07 / 20.28%	3.21 / 18.02%
ViT + VIF <sup>2</sup>	48.29 / 12.21%	3.42 / 21.29%	9.19 / 23.09%	3.41 / 19.12%
InceptionV3 + VIF <sup>2</sup>	48.53 / 12.27%	3.07 / 19.15%	7.77 / 19.53%	3.06 / 17.18%

Figure 8 illustrates typical samples with low and high prediction errors. The three images on the left, with more accurate predictions, demonstrate the model’s strong performance in handling foods with clear shapes and distinct ingredients. For instance, the ‘Bacon, Egg & Swiss Croissant’ features prominent visual characteristics and well-defined ingredients, making it easy for the model to recognize. In contrast, the image on the right, which shows higher prediction errors (‘Mac & Cheese Tray’), has more complex visual features, highly mixed ingredients, and lacks clear boundaries, making it difficult for the model to extract key features effectively. This highlights the challenges the model faces when dealing with foods that exhibit complex visual characteristics and diverse ingredient compositions. These results clearly demonstrate the effectiveness of visual-ingredient feature fusion in enhancing nutrition estimation performance. Furthermore, our findings indicate that convolutional networks, such as ResNets and InceptionV3, outperform transformer-based models like ViT in nutrition-related tasks [39].



**Figure 9: The average MAE across four tasks as ResNet101 integrates ingredient features generated by LLaVA1.6 under different thresholds.**

### 5.3 Ablation Study

**Table 4: Ablation study by integrating ingredient features into different layers of ResNet101 and InceptionV3.**

model	Injection layer	Average MAE	model	Injection layer	Average MAE
Resnet101 [15]	block1	22.20/24.76%	InceptionV3 [37]	PostMaxPool2	27.10/28.26%
	block2	20.47/22.22%		Mixed6e (w/ AUX)	25.44/26.64%
	block3	21.57/22.20%		Mixed6e (w/o AUX)	25.66/26.77%
	block4	33.09/32.94%		PostMixed7c	32.75/33.55%

**Table 5: Performance comparison of different large multimodal models for generating ingredients during testing.**

LMM	Caloric MAE	Fat MAE	Carb MAE	Protein MAE
LLaVA1.6(7B) [23]	61.26 / 15.49%	3.60 / 22.40%	8.76 / 22.02%	3.87 / 21.73%
Qwen2-VL(7B) [43]	69.56 / 17.59%	4.30 / 26.82%	10.14 / 25.49%	4.54 / 25.45%
LLaVA-NeXT(LLama-3 8B) [47]	73.10 / 18.48%	4.50 / 28.03%	10.54 / 26.49%	4.63 / 26.00%

**Ingredient feature injection layer.** In Table 4, we investigate the impact of integrating ingredient features at different layers of ResNet101 and InceptionV3. For ResNet101, integrating ingredient features after block2 achieves the best performance, while doing so after block4 significantly increases the average MAE. For InceptionV3, the best performance is observed when ingredient features are incorporated after the Mixed6e module and also fused in the auxiliary branch. The results show that injecting ingredient features in the intermediate layers achieves the best performance as intermediate layers retain semantic information while still preserving local details.

### 5.4 Zero-Shot Evaluation of Diet Analysis and Nutritional Reasoning

**Threshold  $\tau$  for augmented ingredient prediction.** Figure 9 illustrates the effect of threshold for augmented ingredient prediction. As the threshold increases from 1 to 4, the average MAE decreases, reaching its lowest value of 61.26 at a threshold of 4. This indicates that an appropriate threshold effectively filters out irrelevant or noisy ingredients, enhancing nutrition prediction performance. However, when the threshold increases beyond 5, the average MAE begins to rise, suggesting that overly high thresholds may result in insufficient ingredient information, which compromises the model’s generalization capability.

**Large multimodal models.** To evaluate the impact of different large multimodal models (LMMs) on ingredient generation during

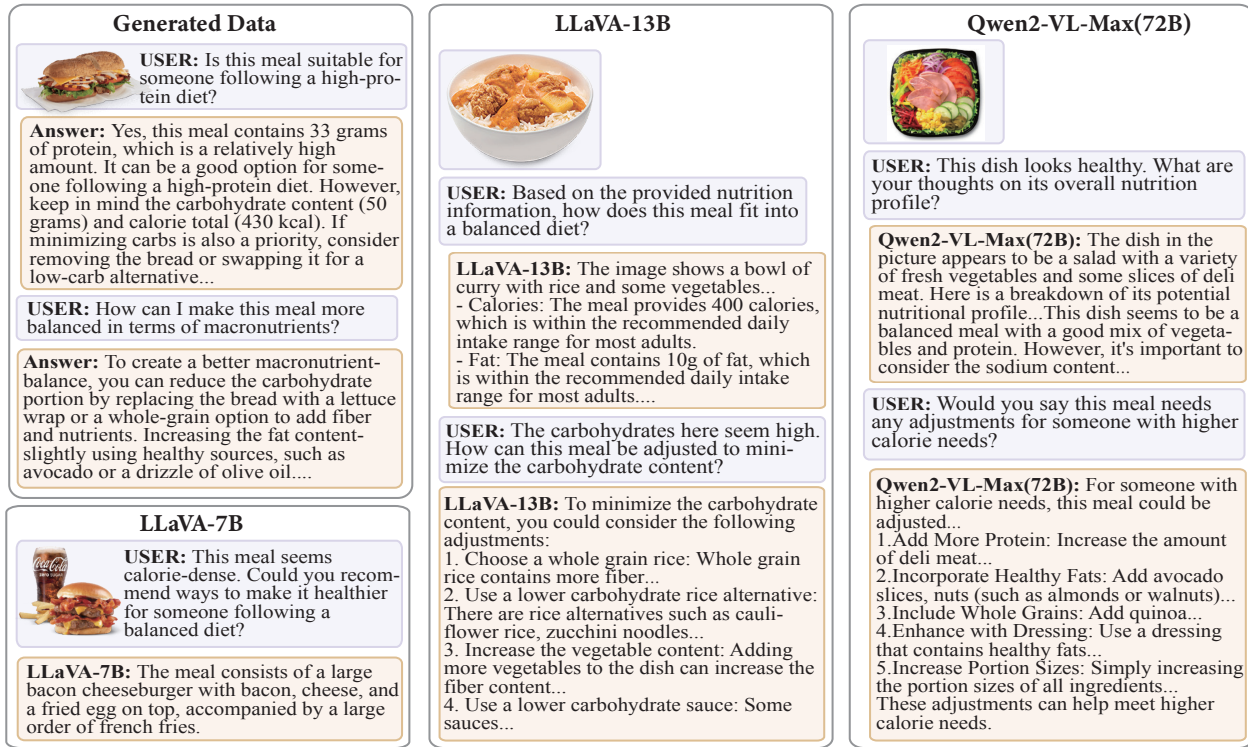


Figure 10: Qualitative examples showcasing zero-shot diet analysis of multimodal large language models.

inference, we utilize three LMMs during the testing phase with augmented ingredient prediction: LLaVA1.6 (7B) [23], Qwen2-VL (7B), and LLaVA-NeXT (Llama-3-8B) [47]. Table 5 shows that LLaVA1.6 achieves the best performance across all metrics. Qwen2-VL follows closely, showing slightly lower performance than LLaVA1.6 in Caloric and Fat but demonstrating improvements in Carb and Protein. In contrast, LLaVA-NeXT exhibited weaker overall performance, falling behind the other two models in all metrics.

Figure 10 illustrates our generated data and evaluates the reasoning of diet-related questions answered by LLaVA(7B/13B) and Qwen2-VL-Max(72B) under zero-shot setting. In our generated data, the question-answer pairs are capable of providing accurate nutritional analysis and formulating reasonable diet suggestions based on the provided image. While clearly addressing users' needs, our QA pairs also consider the balance of other macronutrients (such as carbohydrates and fats) and offer specific suggestions for improving nutritional balance accordingly.

During testing, LLaVA-7B demonstrates a correct understanding of both the images and user queries, such as analyzing the high-calorie characteristics of the food. However, its responses lack specific support for nutrition content and do not adequately. For instance, when a query mentions that "high protein intake should be balanced with carbohydrates and fats," it fails to provide actionable suggestions. LLaVA-13B shows significant improvements over LLaVA-7B in terms of the accuracy and interpretability of nutrition information. Its responses encompass a comprehensive analysis from total calorie content to macronutrient breakdown and offer

specific recommendations to address user queries, reflecting greater practicality. Qwen2-VL-Max(72B) is capable of conducting a relatively comprehensive analysis too. Moreover, this model identifies potential issues in its responses, such as high sodium and calorie sources, and offers highly feasible and diverse solutions tailored to user requirements.

## 6 CONCLUSION

In this paper, we have presented a new dataset named FastFood with ingredient and nutrition annotations. Furthermore, we propose a model-agnostic method by integrating the ingredient and visual features for nutrition estimation. The results on both FastFood and Nutrition-5K show that our proposed method significantly outperforms the existing methods. In the future, we aim to explore the development of large multimodal models for advanced nutrition estimation and personalized diet suggestion generation.

## ACKNOWLEDGMENT

This research/project is supported by the Ministry of Education, Singapore, under Academic Research Fund (AcRF) Tier 1 grant (No. MSS23C018) and Tier 2 (Proposal ID: T2EP20222-0046). Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of the Ministry of Education, Singapore.

## REFERENCES

- [1] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. 2014. Food-101—mining discriminative components with random forests. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part VI 13*. Springer.
- [2] Rebecca G Boswell, Wendy Sun, Shosuke Suzuki, and Hedy Kober. 2018. Training in cognitive strategies reduces eating and improves food choice. *Proceedings of the National Academy of Sciences* (2018).
- [3] Chi-Sheng Chen, Guan-Ying Chen, Dong Zhou, Di Jiang, and Dai-Shi Chen. 2024. Res-vmamba: Fine-grained food category visual classification using selective state space models with deep residual learning. *arXiv preprint arXiv:2402.15761* (2024).
- [4] Jingjing Chen and Chong-Wah Ngo. 2016. Deep-based ingredient recognition for cooking recipe retrieval. In *Proceedings of the 24th ACM international conference on Multimedia*.
- [5] Jingjing Chen, Liangming Pan, Zhipeng Wei, Xiang Wang, Chong-Wah Ngo, and Tat-Seng Chua. 2020. Zero-shot ingredient recognition by multi-relational graph convolutional network. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [6] Jingjing Chen, Bin Zhu, Chong-Wah Ngo, Tat-Seng Chua, and Yu-Gang Jiang. 2020. A study of multi-task and region-wise deep learning for food ingredient recognition. *IEEE Transactions on Image Processing* (2020).
- [7] Prateek Chhikara, Dhiraj Chaurasia, and Yifan Jiang. 2024. Fire: Food image to recipe generation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*.
- [8] Alexey Dosovitskiy. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [9] Takumi Ege and Keiji Yanai. [n. d.]. Image-based food calorie estimation using knowledge on food categories, ingredients and cooking directions. In *Proceedings of the on Thematic Workshops of ACM Multimedia 2017*.
- [10] Takumi Ege and Keiji Yanai. 2019. Simultaneous estimation of dish locations and calories with multi-task learning. *IEICE TRANSACTIONS on Information and Systems* (2019).
- [11] Xinle Gao, Zhiyong Xiao, and Zhaohong Deng. 2024. High accuracy food image classification via vision transformer with data augmentation and feature augmentation. *Journal of Food Engineering* (2024).
- [12] Yinxuan Gui, Bin Zhu, Jingjing Chen, and Chong-Wah Ngo. 2025. Efficient Prompt Tuning for Hierarchical Ingredient Recognition. In *2025 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE.
- [13] Yinxuan Gui, Bin Zhu, Jingjing Chen, Chong Wah Ngo, and Yu-Gang Jiang. 2024. Navigating weight prediction with diet diary. In *Proceedings of the 32nd ACM International Conference on Multimedia*. 127–136.
- [14] Yuzhe Han, Qimin Cheng, Wenjin Wu, and Ziyang Huang. 2023. Dpf-nutrition: Food nutrition estimation via depth prediction and fusion. *Foods* (2023).
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [16] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *Comput. Surveys* (2023).
- [17] Shuqiang Jiang. 2024. Food Computing for Nutrition and Health. In *2024 IEEE 40th International Conference on Data Engineering Workshops (ICDEW)*. IEEE.
- [18] Pengkun Jiao, Xinlan Wu, Bin Zhu, Jingjing Chen, Chong-Wah Ngo, and Yugang Jiang. 2024. Rode: Linear rectified mixture of diverse experts for food large multi-modal models. *arXiv preprint arXiv:2407.12730* (2024).
- [19] Yoshiyuki Kawano and Keiji Yanai. 2015. Automatic expansion of a food image dataset leveraging existing categories with domain adaptation. In *Computer Vision—ECCV 2014 Workshops: Zurich, Switzerland, September 6–7 and 12, 2014, Proceedings, Part III 13*. Springer.
- [20] Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. 2023. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *arXiv preprint arXiv:2306.00890* (2023).
- [21] Chang Liu, Yu Cao, Yan Luo, Guanling Chen, Vinod Vokkarane, and Yunsheng Ma. 2016. Deepfood: Deep learning-based food image recognition for computer-aided dietary assessment. In *Inclusive Smart Cities and Digital Health: 14th International Conference on Smart Homes and Health Telematics, ICOST 2016, Wuhan, China, May 25–27, 2016, Proceedings 14*. Springer.
- [22] Guoshan Liu, Hailong Yin, Bin Zhu, Jingjing Chen, Chong-Wah Ngo, and Yu-Gang Jiang. 2025. Retrieval augmented recipe generation. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2453–2463.
- [23] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. [n. d.]. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [24] Hanchao Liu, Wenyan Xue, Yifei Chen, Dapeng Chen, Xiutian Zhao, Ke Wang, Liping Hou, Rongjun Li, and Wei Peng. 2024. A survey on hallucination in large vision-language models. *arXiv preprint arXiv:2402.00253* (2024).
- [25] Yuxin Liu, Weiqing Min, Shuqiang Jiang, and Yong Rui. 2024. Convolution-Enhanced Bi-Branch Adaptive Transformer with Cross-Task Interaction for Food Category and Ingredient Recognition. *IEEE Transactions on Image Processing* (2024).
- [26] Frank Po Wen Lo, Yao Guo, Yingnan Sun, Jianing Qiu, and Benny Lo. 2022. An intelligent vision-based nutritional assessment method for handheld food items. *IEEE Transactions on Multimedia* (2022).
- [27] Ya Lu, Thomai Stathopoulou, Maria F Vasiloglou, Stergios Christodoulidis, Zeno Stanga, and Stavroula Mouggiakakou. 2020. An artificial intelligence-based system to assess nutrient intake for hospitalised patients. *IEEE transactions on multimedia* (2020).
- [28] Austin Meyers, Nick Johnston, Vivek Rathod, Anoop Korattikara, Alex Gorban, Nathan Silberman, Sergio Guadarrama, George Papandreou, Jonathan Huang, and Kevin P Murphy. 2015. Im2Calories: towards an automated mobile vision food diary. In *Proceedings of the IEEE international conference on computer vision*.
- [29] Simon Metzger and Barbara Koroušić Seljak. 2017. NutriNet: a deep learning food and drink image recognition system for dietary assessment. *Nutrients* (2017).
- [30] Weiqing Min, Shuqiang Jiang, Linhu Liu, Yong Rui, and Ramesh Jain. 2019. A survey on food computing. *ACM Computing Surveys (CSUR)* (2019).
- [31] Weiqing Min, Linhu Liu, Zhiling Wang, Zhengdong Luo, Xiaoming Wei, Xiaolin Wei, and Shuqiang Jiang. [n. d.]. Isia food-500: A dataset for large-scale food recognition via stacked global-local attention network. In *Proceedings of the 28th ACM International Conference on Multimedia*.
- [32] 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* (2023).
- [33] Amaia Salvador, Nicholas Hynes, Yusuf Aytar, Javier Marin, Ferda Ofli, Ingmar Weber, and Antonio Torralba. 2017. Learning cross-modal embeddings for cooking recipes and food images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [34] Wenjing Shao, Sujuan Hou, Weikuan Jia, and Yuanjie Zheng. 2022. Rapid non-destructive analysis of food nutrient content using swin-nutrition. *Foods* 11, 21 (2022), 3429.
- [35] Wenjing Shao, Weiqing Min, Sujuan Hou, Mengjiang Luo, Tianhao Li, Yuanjie Zheng, and Shuqiang Jiang. 2023. Vision-based food nutrition estimation via RGB-D fusion network. *Food Chemistry* 424 (2023), 136309.
- [36] Fangzhou Song, Bin Zhu, Yanbin Hao, and Shuo Wang. 2024. Enhancing recipe retrieval with foundation models: A data augmentation perspective. In *European Conference on Computer Vision*. Springer, 111–127.
- [37] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [38] Karan Taneja, Richard Segal, and Richard Goodwin. 2024. Monte Carlo tree search for recipe generation using GPT-2. *arXiv preprint arXiv:2401.05199* (2024).
- [39] Quin Thames, Arjun Karpur, Wade Norris, Fangting Xia, Liviu Panait, Tobias Weyand, and Jack Sim. 2021. Nutrition5k: Towards automatic nutritional understanding of generic food. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- [40] David Tilman and Michael Clark. 2014. Global diets link environmental sustainability and human health. *Nature* (2014).
- [41] SM Tomnroy, SM Zaman, Vinija Jain, Anku Rani, Vipula Rawte, Aman Chadha, and Amitava Das. 2024. A comprehensive survey of hallucination mitigation techniques in large language models. *arXiv preprint arXiv:2401.01313* (2024).
- [42] Gautham Vinod, Zeman Shao, and Fengqing Zhu. 2022. Image based food energy estimation with depth domain adaptation. In *2022 IEEE 5th International Conference on Multimedia Information Processing and Retrieval (MIPR)*. IEEE.
- [43] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. 2024. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191* (2024).
- [44] Xiongwei Wu, Xin Fu, Ying Liu, Ee-Peng Lim, Steven CH Hoi, and Qianru Sun. [n. d.]. A large-scale benchmark for food image segmentation. In *Proceedings of the 29th ACM International Conference on Multimedia*.
- [45] Ziwei Xu, Sanjay Jain, and Mohan Kankanhalli. 2024. Hallucination is inevitable: An innate limitation of large language models. *arXiv preprint arXiv:2401.11817* (2024).
- [46] Yuehao Yin, Huiyan Qi, Bin Zhu, Jingjing Chen, Yu-Gang Jiang, and Chong-Wah Ngo. 2023. Foodlmm: A versatile food assistant using large multi-modal model. *arXiv preprint arXiv:2312.14991* (2023).
- [47] Yuanhan Zhang, Jiming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. 2024. Video instruction tuning with synthetic data. *arXiv preprint arXiv:2410.02713* (2024).
- [48] Bin Zhu, Chong-Wah Ngo, and Wing-Kwong Chan. 2021. Learning from web recipe-image pairs for food recognition: Problem, baselines and performance. *IEEE Transactions on Multimedia* (2021).
- [49] Bin Zhu, Chong-Wah Ngo, Jingjing Chen, and Wing-Kwong Chan. 2022. Cross-lingual adaptation for recipe retrieval with mixup. In *Proceedings of the 2022 International Conference on Multimedia Retrieval*. 258–267.