

# Navigating Weight Prediction with Diet Diary

Anonymous Author(s)

## ABSTRACT

Current research in food analysis primarily concentrates on tasks such as food recognition, recipe retrieval and nutrition estimation from a single image. Nevertheless, there is a significant gap in exploring the impact of food intake on physiological indicators (e.g., weight) over time. This paper addresses this gap by introducing the DietDiary dataset, which encompasses daily dietary diaries and corresponding weight measurements of real users. Furthermore, we propose a novel task of weight prediction with a dietary diary that aims to leverage historical food intake and weight to predict future weights. To tackle this task, we propose a model-agnostic time series forecasting framework. Specifically, we introduce a Unified Meal Representation Learning (UMRL) module to extract representations for each meal. Additionally, we design a diet-aware loss function to associate food intake with weight variations. By conducting experiments on the DietDiary dataset with two state-of-the-art time series forecasting models, NLinear and iTransformer, we demonstrate that our proposed framework achieves superior performance compared to the original models. We will make our dataset, code, and models publicly available.

## CCS CONCEPTS

• Information systems → Multimedia information systems.

## KEYWORDS

Weight prediction, food analysis, time series forecasting models

## 1 INTRODUCTION

Food plays a vital role in human existence, affecting the quality of life and various physiological indicators, with weight being one of the most fundamental and critical aspects. [24] indicates that the relative balance between energy expenditure and dietary intake determines weight gain or loss. Notably, the energy expenditure in non-exercise activities (such as sitting) usually accounts for a significantly larger portion of total energy expenditure than exercise [21]. Motivated by these insights, we aim to investigate the impact of food intake on weight prediction.

Numerous studies have focused on various tasks within the food domain, such as food classification [11, 25, 35, 37], ingredients recognition [1, 5, 6], recipe retrieval [19, 46, 48, 53, 54, 63, 64], food volume prediction [32, 34, 49] as well as calorie and nutritional information estimation [36, 50, 52]. Nevertheless, these existing methods typically analyze each food image independently and do

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.

ACM MM, 2024, Melbourne, Australia

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

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

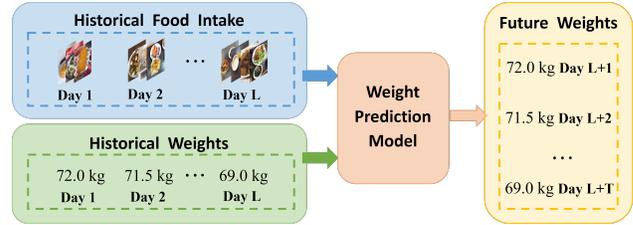


Figure 1: The overview of the proposed weight prediction with diet diary task. Given a historical food intake and corresponding weight measurement in the past  $L$  days, the task aims to predict the weights in the following  $T$  days.

not investigate the impact of food intake on physiological indicators (e.g., weight) over a period of time.

To address this gap, this paper navigates weight prediction with dietary diary. On the one hand, we construct a novel dataset named DietDiary, which includes daily dietary records and corresponding weight measurements for 611 participants from a health management system. Different from existing datasets in the food domain such as [2, 5, 38], DietDiary is the first dataset to provide both food intake and corresponding weight measurements over a period of time. On the other hand, as illustrated in Figure 1, we define a new task that aims to leverage historical weight and food intake data to predict future weights. This task poses two major challenges: (1) Meal representation learning. Given that food intake can be represented in various forms, such as textual ingredient labels or food images, developing an effective method for learning meal representations for weight prediction is crucial. (2) Modeling the correlation between food intake and weight changes. Since food intake is closely related to weight, weight prediction requires more than just exploring the temporal relationships between historical and future weights. Understanding the complex correlation and dependencies between dietary intake and weight fluctuations presents significant challenges to this task.

To address the aforementioned challenges, we propose a novel framework that integrates food intake information for weight prediction. Specifically, we introduce a Unified Meal Representation Learning (UMRL) module that leverages CLIP [43] text or image encoders to extract a unified feature representation of a meal from various forms of historical food intake. Additionally, we propose a diet-aware loss function to enable the model to capture the correlation and dependencies between food and weight changes. Importantly, our framework is designed to be compatible with any existing time series forecasting model for weight prediction that incorporates food intake information. We evaluate our framework on the DietDiary dataset using two representative advanced time series forecasting models, NLinear [59] and iTransformer [31]. The superior performance of our framework over these models demonstrates its effectiveness in leveraging food intake information for weight prediction.

The contributions of this paper can be summarized as follows:

- We construct a new DietDiary dataset, which is the first dataset providing food intake and corresponding weight measurements over a period of time.
- We introduce a novel task of weight prediction using dietary information. This task uniquely treats food intake as temporal data and investigates its potential to improve weight prediction.
- We propose a model-agnostic time series forecasting framework for weight prediction with food intake. Our framework demonstrates significant improvements over existing methods in weight prediction performance.

## 2 RELATED WORK

### 2.1 Time Series Forecasting

Time series forecasting is a fundamental task in various fields, including finance, traffic and meteorology. Over the years, to address the inherent challenges of predicting future values based on past observations, numerous models are proposed, such as the famous ARIMA[3]. With the rise of deep learning, neural network-based approaches have gained popularity in time series forecasting. Models such as Long Short-Term Memory (LSTM) networks[22], DeepAR[44] and Prophet[51] are used to capture long-term dependencies and nonlinear relationships in the data. As Transformer[31] demonstrates powerful sequence modeling capabilities in natural language processing[14], computer vision[4, 15] and other domains, models based on modifying the vanilla Transformer, especially the attention mechanism, tailored for time series forecasting tasks have been widely researched, such as Informer[60], Autoformer [56], Pyraformer [30], and FEDformer [61].

While various research efforts are ongoing to modify attention structures to achieve better Transformer-based solutions, [59] questions the effectiveness of the Transformers and proposes a set of simple and efficient linear models including Linear, NLinear, and DLinear, which have led more researchers to focus on linear-based model. TiDE [13] designs a Multi-layer Perceptron (MLP) based encoder-decoder model to capture covariates and non-linear dependencies. TSMixer [16] patches the time series data and enhances the learning capability of simple MLP structures. Chen et al. [10] propose a Time-mixing MLP to model temporal dependency and a Feature-mixing MLP to analyze the cross-variate information. Influenced by the outstanding performance of linear models, several works continue to employ transformer-based architectures but seek performance enhancement from alternative perspectives. PatchTST [41] uses patched data as input and adopts a channel-independence design that makes a token only contains information from one channel, instead of changing the structure of the transformer. iTransformer [31] inverts the duties of the attention mechanism and the feed-forward layer to capture temporal information to achieve promoted performance and generalization.

In contrast, we propose a model-agnostic time series forecasting framework to predict weight with food intake. To validate the effectiveness of our framework, state-of-the-art linear-based model NLinear [59] and transformer-based model iTransformer [31] are chosen as our time series forecasting models respectively.

### 2.2 Food Analysis

With the development of computer vision and the emergence of various food datasets, research methods and tasks in the food domain have gradually become more diverse. Some traditional visual tasks have been extended to the food domain, such as food classification [11, 25, 27, 35, 37], ingredients recognition [1, 5, 6, 29], and food segmentation [17, 23, 28, 57]. With the release of Recipe1M [47], cross-modal food-related tasks have been extensively studied, such as the retrieval of finding the most relevant recipe for a food image or vice versa [8, 9, 19, 48, 63]. Some works focus on generating recipes from images [12, 20, 45] or corresponding images from recipes [42, 62]. These tasks all involve analyzing each image independently, whereas in this paper, we model temporal dietary data, which has not been explored in previous studies.

Another research branch involves predicting the nutritional components and calorie of food from food images, which aids in monitoring intake patterns. Some works estimate volume first from voxel [36], point cloud [18, 33] or 3D mesh [40], then mapping is achieved through data on the calories contained per unit volume. With the release of Nutrition5K [52], a dataset providing fine-grained nutritional attributes, food quality and food videos, some works predict nutrition directly from images by neural network [50, 52, 58]. However, these works only predict nutritional information without further linking food intake to weight. To the best of our knowledge, this paper is the first work to utilize dietary diary to predict weight.

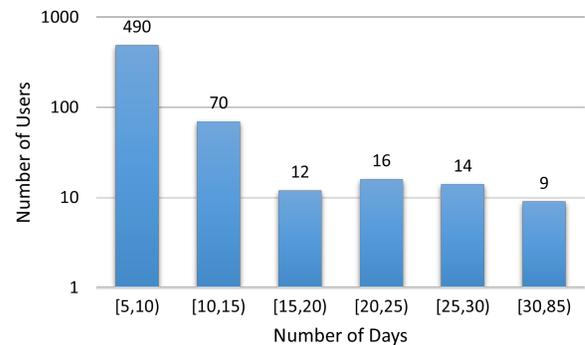


Figure 2: Distribution of the number of recording days for the participants in DietDiary dataset. The y-axis is in the log scale.

## 3 DATASET CONSTRUCTION

We introduce a novel dataset, DietDiary, specifically for analyzing weight in relation to food intake. In contrast to datasets such as Food-101 [2], VIREO Food-172 [5, 7], Food2K [39] and Recipe1M [47], DietDiary encompasses diet diary of three meals over a period of time, accompanied by daily weight measurement. To the best of our knowledge, DietDiary is the first dataset in the food domain to provide this kind of data, offering new opportunities for research in dietary pattern analysis and its impact on weight management.



Figure 3: Examples of food images from three meals in DietDiary.

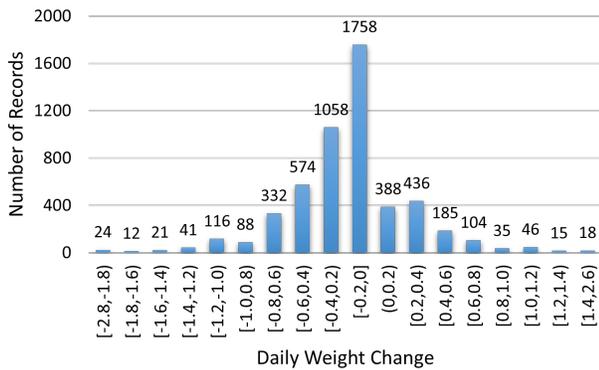


Figure 4: Distribution of Daily Weight Change of all records in DietDiary.

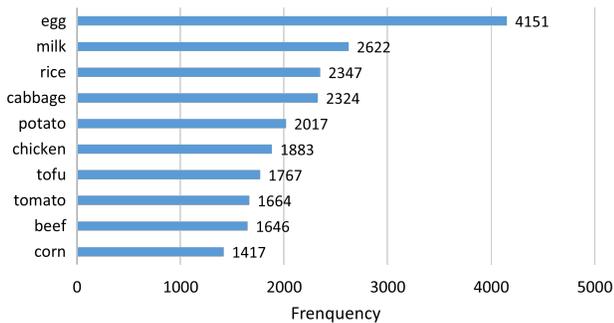


Figure 5: Top 10 ingredients by occurrence frequency in DietDiary.

### 3.1 DietDiary Dataset

The DietDiary dataset was collected within a health management system<sup>1</sup> wherein participants were required to meticulously log their daily dietary intake, along with their corresponding weight measurements. The dietary log consists of images of three meals (i.e., breakfast, lunch, and supper) along with manually labeled ingredients for each meal. The dataset encompasses a total of 611

<sup>1</sup><https://www.qiezilife.com/>

participants, with over 5k daily records, nearly 30k images, and over 15k ingredient annotations.

As depicted in Figure 2, the duration of record-keeping among participants varies, ranging from a minimum of one week to over one month. A significant majority (91.6%) of the participants maintained their records for a duration of one to two weeks. However, there were also 9 participants recorded their diet and weight for more than a month. Figure 4 illustrates the distribution of daily weight fluctuations across all participants. Most users experience weight changes within  $\pm 1$  kg per day, with weight loss records comprising the majority due to the data source. Additionally, a small portion of individuals experience changes exceeding 2 kg within a day.

### 3.2 Food Images and Ingredient Annotation

Figure 3 presents meal images from various participants, illustrating typical examples of breakfast, lunch, and supper. It is important to note that we have manually curated the dataset to exclude images that are not relevant to food.

The annotation of ingredients in the dataset was standardized through a series of preprocessing steps. First, we unified the separators used between phrases, as different participants employed varying symbols (e.g., “-”, “/”, “()”) to delineate ingredients. Second, the original annotations in Chinese were translated into English. Third, we adopted the approach as in inverse cooking [45] to aggregate and categorize ingredients based on common prefixes or suffixes, for example, “boiled eggs”, “steamed eggs”, and “egg slices” are all merged into the “egg” category. Fourth, ingredients that appeared less than five times were excluded from the dataset. After preprocessing, the total number of unique ingredients was 197. The ten most frequently occurring ingredients are depicted in Figure 5, such as “milk”, “egg” and “rice”, all of which are common in everyday meals. Furthermore, Figure 6 provides a comprehensive example from the dataset, showcasing food images, annotated ingredients, and corresponding weight measurements for two participants with different weight fluctuation trends. On the one hand, user A maintained a healthy dietary habit, and the weight measurements consistently show a decreasing trend. On the other hand, user B consumed high-calorie foods (such as fried shrimp and sugar) on the third and fifth days, and also consumed a large amount of staple food (dumplings and sushi) during supper, resulting in an increase in corresponding weight records.

Day	User	Breakfast	Lunch	Supper	Weight(kg)	User	Breakfast	Lunch	Supper	Weight(kg)
1	A				101.0 - 100.3	B				66.7 - 66.5
2					100.3 - 100.1					66.5 - 66.1
3					100.1 - 99.5					66.1 - 66.5
4					99.5 - 99.3					66.5 - 65.9
5					99.3 - 99.1					65.9 - 66.0
...	...	...	...	...	...	...	...	...	...	...

Figure 6: Data records for two participants with different weight fluctuation trends in DietDiary. The records leading to weight gain are highlighted in red.

## 4 METHOD

In this section, we provide detailed definition of the task in Section 4.1 and elaborate on the details of our proposed framework in Sections 4.2 and 4.3.

### 4.1 Task Definition

Given the historical weight sequence  $W = \{w_t\}_{t=1}^L$  and corresponding food intake  $F = \{Breakfast_t, Lunch_t, Supper_t\}_{t=1}^L$  for the past  $L$  days, the goal of our proposed task is to predict the future weight sequence  $\hat{W} = \{\hat{w}_t\}_{t=L+1}^{L+T}$  over the next  $T$  days, where  $w_t$  denotes the weight value of the  $t$ -th day,  $Breakfast_t$ ,  $Lunch_t$  and  $Supper_t$  represent sets of food images or ingredient labels for each meal on the  $t$ -th day. Note that the historical observation takes various data modalities, including textual ingredient labels or food images, and numerical weight.

### 4.2 Proposed Framework

Figure 7 depicts an overview of the proposed framework. The framework primarily consists of three key components: a unified Meal Representation Learning (UMRL) module, an agnostic time series forecasting model, and a diet-aware loss computation.

**Unified Meal Representation Learning (UMRL)** is a module to extract representation for each meal. Given the historical food intake (e.g., food images or ingredient labels), UMRL is designed to obtain representation for breakfast, lunch, and supper for each day respectively, denoted as  $f = \{b_t, l_t, s_t\}_{t=1}^L$ , where  $b_t$ ,  $l_t$  and  $s_t$  are representation for breakfast, lunch, and supper on the  $t$ -th day. More details of UMRL are presented in Section 4.3.

**Model-agnostic weight prediction.** After obtaining the representation of historical food intake from UMRL, we concatenate it with the historical weight  $W$ , denoted as  $X = \{b_t, l_t, s_t, w_t\}_{t=1}^L$ , and feed into a time series forecasting model  $M$ . The weight prediction result is obtained as follows:

$$\hat{Y} = \{\hat{b}_t, \hat{l}_t, \hat{s}_t, \hat{w}_t\}_{t=L+1}^{L+T} = M(X), \quad (1)$$

where  $\hat{b}_t$ ,  $\hat{l}_t$ ,  $\hat{s}_t$ ,  $\hat{w}_t$  are the predicted weights from breakfast, lunch, supper and weight respectively. Note that our proposed framework is agnostic to the time series forecasting model, which gives us the flexibility to utilize existing state-of-the-art models for our task. In this paper, we choose linear-based NLinear [59] and Transformer-based iTransformer [31] as  $M$  in Equation 1.

**Diet-aware loss computation.** We introduce a diet-aware loss to model the food intake and weight variations. Assume the weight variation  $\Delta_t$  of the  $t$ -th day is calculated by the following formula:

$$\Delta_t = w_t - w_{t-1}. \quad (2)$$

As each meal contributes to the weight variations, the diet loss is computed by combining all three meals as follows:

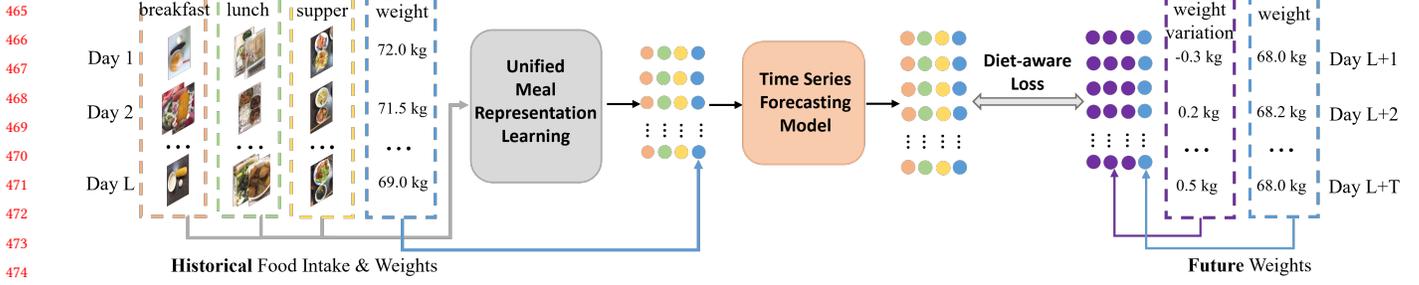
$$\mathcal{L}_{diet} = \frac{1}{3}((\Delta_t - \hat{b}_t)^2 + (\Delta_t - \hat{l}_t)^2 + (\Delta_t - \hat{s}_t)^2). \quad (3)$$

In addition, we compute a weight loss based on the weight prediction from historical weight as follows:

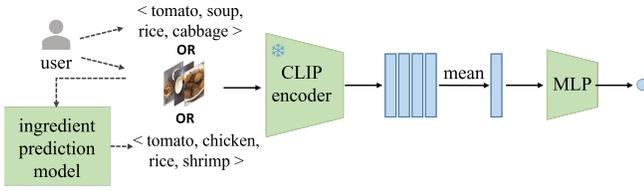
$$\mathcal{L}_{weight} = (w_t - \hat{w}_t)^2. \quad (4)$$

The overall diet-aware loss to train the proposed framework is to combine both diet and weight losses:

$$\mathcal{L} = \lambda \mathcal{L}_{weight} + (1 - \lambda) \mathcal{L}_{diet}, \quad (5)$$



**Figure 7: Framework Overview.** The “Unified Meal Representation Learning” module is proposed to map the historical food intake into a time series meal feature sequence. The features and historical weight sequence are combined and subsequently fed into an agnostic time series forecasting model to predict future weights.



**Figure 8: The architecture of proposed Unified Meal Representation Learning Module.**

**Table 1: Statistics of training, testing, and validation sets for different settings. L-T refers to the setting of using weight and food intake in  $L$  history days to predict the weights of future  $T$  days.**

Data split	settings (L-T)				
	3-3	3-5	3-7/5-5/7-3	5-7	7-7
train	1,535	1,010	837	672	514
validation	221	149	123	74	57
test	440	296	241	216	164

$$E_n = \text{CLIPEncoder}(img_n),$$

$$\bar{E} = \frac{1}{N} \sum_{n=1}^N (E_n), \quad (6)$$

$$b_t = \text{MLP}(\bar{E}),$$

where  $\lambda$  is a hyper-parameter to balance the two losses.

### 4.3 Unified Meal Representation Learning

As shown in Figure 8, we propose Unified Meal Representation Learning (UMRL) Module to extract representation from historical food intake for each meal. As the food intake could be either food images or ingredient labels, we first employ CLIP [43] to encode both visual or textual inputs to unify the feature extraction. On the one hand, if the historical food intake is ingredient label, we adopt pretrained CLIP text encoder to extract the features. On the other hand, we adopt CLIP image encoder to extract the features from historical food images. The CLIP model is kept frozen during our training. Note that the users prefer taking pictures of their meals rather than typing the concrete ingredient annotations in practice. As a result, we also investigate employing a pre-trained ingredient recognition model to automatically obtain ingredient labels from images, which is illustrated in the bottom left of Figure 8.

The CLIP features of images or ingredient labels are subsequently averaged to get a feature vector for each meal. Finally, we feed this vector into a Multilayer Perceptron (MLP) to derive the final meal representation. Take breakfast food images as an example, denoted  $Breakfast_t = \{img_1, \dots, img_n\}_{n=1}^N$  as a set of breakfast food images on the  $t$ -th day, where  $N$  is the number of images, the process of UMRL can be formalized as follows:

where  $E_n$  is the CLIP image feature for  $img_n$ ,  $\bar{E}$  indicates the averaged CLIP image features and  $b_t$  is the breakfast representation.

## 5 EXPERIMENT

### 5.1 Experiment Settings

**Dataset.** We conduct experiments on our DietDiary dataset. Inspired by time series forecasting task [31, 59], we extensively explore different combinations of using weight and food intake history in  $L$  days to predict the weights in future  $T$  days, denoted as setting L-T. The settings examined include {3-3, 3-5, 3-7, 5-5, 5-7, 7-3 and 7-7}. For instance, in the 7-3 setting ( $L = 7, T = 3$ ), we utilize the historical weight sequence and corresponding food intake from the past 7 days to predict the weight for the subsequent 3 days during training. The dataset is partitioned into training, validation, and test sets in a 7:1:2 ratio for each setting. It is crucial to ensure that each participant appears exclusively in one of these splits, thereby preventing any overlap of participants across training, validation, and testing sets. The settings determine the minimum number of days required for training. To maximize the utility of the dataset, we employ different data partitions for different settings. For example, the settings 5-5 and 7-3 necessitate a minimum of 10 days of dietary intake records per participant, whereas the 7-7 setting

**Table 2: Performance comparison based on NLinear as time series forecasting model in terms of MSE and MAE. “image”, “ing-users” and “ing-LMM” represent the food intake is food image, ingredient labels provided by users, ingredient labels predicted by FoodLMM respectively.**

setting		3 - 3		3 - 5		3 - 7		5 - 5		5 - 7		7 - 3		7 - 7	
metric		MSE	MAE												
ours	NLinear [59]	3.729	1.107	5.390	1.448	6.949	1.710	5.991	1.619	5.331	1.589	5.464	1.600	6.791	1.958
	image	2.769	0.985	<b>4.765</b>	<b>1.371</b>	6.332	1.636	2.573	<b>1.330</b>	<b>3.219</b>	<b>1.341</b>	2.799	1.213	<b>2.454</b>	1.376
	ing-users	2.865	0.998	4.819	1.377	<b>6.307</b>	<b>1.633</b>	<b>2.539</b>	1.333	3.235	<b>1.341</b>	<b>2.742</b>	<b>1.211</b>	2.479	1.370
	ing-LMM	<b>2.048</b>	<b>0.962</b>	4.849	1.380	6.312	1.635	2.542	1.345	3.247	1.342	2.968	1.233	2.497	<b>1.365</b>

**Table 3: Performance comparison based on iTransformer as time series forecasting model. “image”, “ing-users” and “ing-LMM” represent the food intake is food image, ingredient labels provided by users, ingredient labels predicted by FoodLMM respectively.**

setting		3 - 3		3 - 5		3 - 7		5 - 5		5 - 7		7 - 3		7 - 7	
metric		MSE	MAE												
ours	iTransformer [31]	4.023	1.402	5.306	1.717	5.918	1.866	5.711	1.817	6.369	1.966	4.657	1.611	4.851	1.701
	image	<b>3.436</b>	<b>1.268</b>	5.268	1.710	5.791	1.835	5.478	1.783	4.299	1.616	4.524	1.584	<b>4.411</b>	<b>1.596</b>
	ing-users	3.544	1.293	5.202	1.696	5.922	1.861	<b>3.540</b>	<b>1.466</b>	3.746	1.541	4.322	1.545	4.645	1.641
	ing-LMM	4.133	1.424	<b>5.047</b>	<b>1.662</b>	<b>5.637</b>	<b>1.802</b>	4.626	1.611	<b>3.657</b>	<b>1.532</b>	<b>3.992</b>	<b>1.486</b>	4.452	1.610

requires at least 14 days. Table 1 presents the dataset statistics for various settings.

**Evaluation metrics.** We employ the Mean Squared Error (MSE) and Mean Absolute Error (MAE) as metrics for performance evaluation. Given that users are primarily concerned with long-term trends in weight prediction, we evaluate performance through autoregressive weight prediction. Specifically, during testing, given a history of  $L$  days, we initially forecast the weights for the subsequent  $T$  days. Subsequent predictions are then made autoregressively, using the outcomes of previous forecasts as input, until predictions have been generated for all recorded days of users in the test set.

**Implementation details.** Given the remarkable performance of FoodLMM [58], a Large Multi-modal Model tailored in the food domain, we choose FoodLMM as our ingredient prediction model. For convenience, we use “users” to refer to ingredients annotated by users and “LMM” to refer to ingredients predicted by the FoodLMM. We evaluate our framework on both the NLinear [59] model, representing linear time series forecasting models, and the iTransformer [31] model, representing transformer-based solutions. For iTransformer, the dimension of series representation  $D$  is set to 16 when  $L \in \{3, 5\}$ , 32 when  $L \in \{5, 7\}$ . The dimension of feed-forward layer is twice that of  $D$ . The number of inverted Transformer blocks and attention heads are set to 1 and 2 respectively. Both NLinear and iTransformer are trained using Adam optimizer [26] with batch size of 32 until early-stopping criteria is met (validation loss does not decrease after 7 epochs). The learning rate is initially set to 0.005 and decays with each epoch. The  $\lambda$  in equation 5 is set to  $\{0, 0.1, 0.25, 0.5, 0.75, 1\}$  to explore the influence of the proportion of diet loss, which will be shown in Section 5.4.

## 5.2 Performance Comparison

Tables 2 and 3 present the performance of our proposed method, which leverages state-of-the-art time series forecasting models NLinear [59] and iTransformer [31], respectively. Compared to models that do not incorporate food intake information, our method consistently achieves superior performance over NLinear [59] and iTransformer [31] across all evaluated settings, with a significant margin of improvement. Notably, the improvements are observed consistently with three different types of food intake inputs: food images, user-provided ingredient labels, and ingredient labels predicted by FoodLMM. These results clearly demonstrate the effectiveness of incorporating food intake information in enhancing the accuracy of weight prediction.

An inspiring observation is that the ingredient labels predicted by FoodLMM often perform comparably or even surpass manually provided ingredient labels by users in most settings. This finding is noteworthy considering that acquiring food images is generally easier than manually annotating ingredients for each meal. Figure 9 presents four examples comparing ingredient labels between users and FoodLMM. Typically, users are more prone to omitting ingredients rather than mislabeling them. For instance, “noodles”, “Chinese parsley/origanum”, and “green onion” in Figure 9 (a), and “corn” and “Shanghai cabbage” in Figure 9 (b) are missing from the user-provided labels. Remarkably, FoodLMM successfully identifies these omitted ingredients. However, ingredient recognition remains a challenging task [7], and the predictions from FoodLMM are still not perfect. For example, there are three incorrect ingredients predicted by FoodLMM in Figure 9 (b) and (c). In Figure 9 (d), although FoodLMM fails to accurately predict “tomato”, the remaining ingredients are correctly identified.

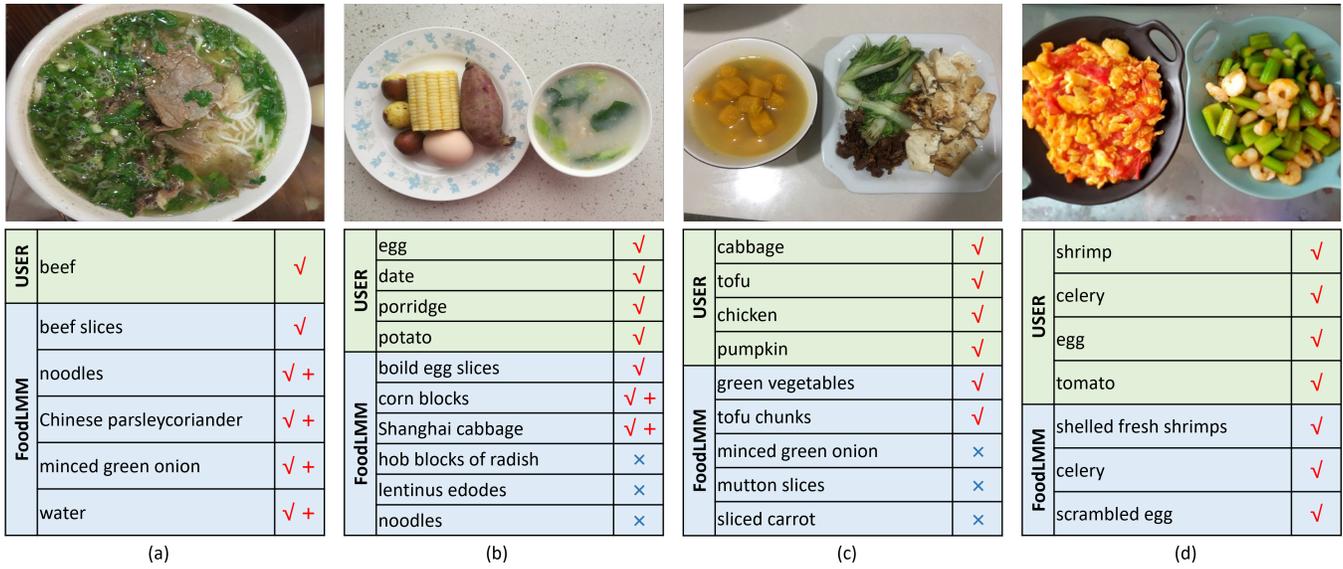


Figure 9: Examples of ingredient labels from user (green) and LMM (blue) for the same food image. The "✓" sign means the correct ingredients, while the "x" sign means the incorrect. The "✓+" sign indicates the ingredients from LMM supplement those missed by users.

Table 4: The multi-modal fusion results of images and ingredient annotations from users based on NLinear.

setting	5 - 5	5 - 7	7 - 3	7 - 7					
metric	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	
NLinear [59]	5.991	1.619	5.331	1.589	5.464	1.600	6.791	1.958	
ours	image	2.573	1.330	2.751	1.469	3.722	1.340	4.390	1.576
	ing-users	<b>2.539</b>	1.333	2.773	1.480	4.327	1.430	4.340	1.567
	fusion	2.567	<b>1.319</b>	<b>2.748</b>	<b>1.464</b>	<b>3.623</b>	<b>1.323</b>	<b>4.334</b>	<b>1.566</b>

### 5.3 Multi-modal Fusion of Food Intake

We further investigate multi-modal fusions of food intake from both images and ingredient labels by averaging the UMRs of the same meal across different modalities. Tables 4 and 5 present the fusion results of images with ingredient annotations from users and FoodLMM respectively. Except for the MSE of setting 5-5, the fusion results consistently outperform those of single modality, whether the ingredient labels are manually provided or predicted by FoodLMM. This demonstrates that food information from different modalities can complement each other, which leads to incorporating both of them achieving better performance. Additionally, the effectiveness of multi-modal fusion also validates that our proposed UMRL module can embed food images and ingredient annotations into a common space and learn unified representations of meals from different modalities.

### 5.4 Ablation study

**Impact of number of meals.** Table 6 presents the impact of incorporating different numbers of daily meals on weight prediction performance. Generally, the inclusion of food intake information, regardless of the number of meals, manages to improve performance. When only one meal is considered, the inclusion of lunch

Table 5: The multi-modal fusion results of images and ingredient annotations from FoodLMM based on NLinear.

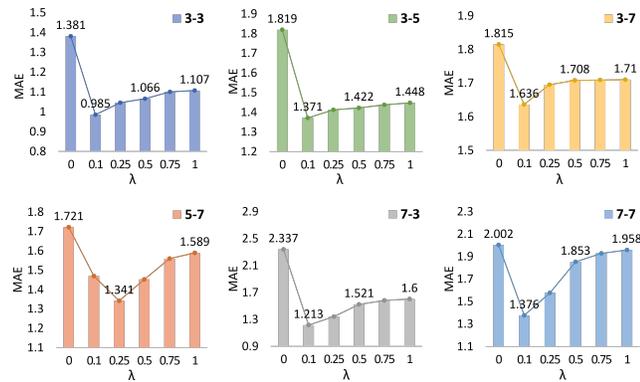
setting	5 - 5	5 - 7	7 - 3	7 - 7					
metric	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	
NLinear [59]	5.991	1.619	5.331	1.589	5.464	1.600	6.791	1.958	
ours	image	2.573	1.330	2.751	1.469	3.722	1.340	4.390	1.576
	ing-LMM	<b>2.542</b>	1.345	2.765	1.476	4.523	1.461	4.367	1.572
	fusion	2.563	<b>1.325</b>	<b>2.741</b>	<b>1.460</b>	<b>3.653</b>	<b>1.328</b>	<b>4.361</b>	<b>1.571</b>

Table 6: Ablation study of number of meals as food intake based on NLinear as time series forecasting model. Food images are used as food intake and MAE is reported.

setting	3-3	3-5	3-7	5-7
NLinear [59]	1.107	1.448	1.710	1.589
+ breakfast	0.990	1.407	1.701	1.452
+ lunch	0.986	1.408	1.692	1.450
+ supper	1.032	1.413	1.697	1.450
+ breakfast + lunch	1.004	1.404	1.695	1.375
+ breakfast + supper	0.987	1.424	1.702	1.373
+ lunch + supper	1.012	1.410	1.686	1.374
+ all three meals	<b>0.985</b>	<b>1.371</b>	<b>1.636</b>	<b>1.341</b>

data yields the most significant improvement under most settings. This is attributed to lunch meals providing the most diverse and distinct food information within the dataset. Breakfast and supper, often featuring less distinctive items like congee-like foods (e.g., meal replacement powders), do not contribute as significantly as lunch to the encoder's effectiveness. For example, the images in the last row of both breakfast and supper in Figure 3. However,

the combination of all three meals offers complementary benefits, leading to the best overall results for weight prediction.

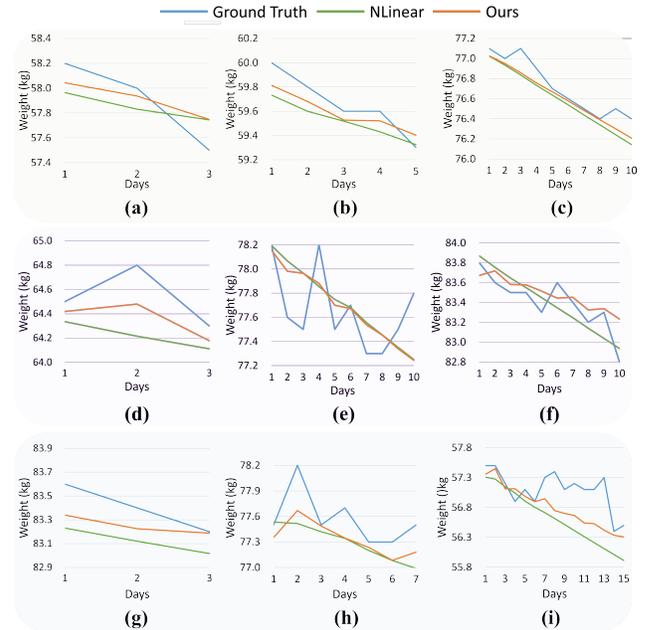


**Figure 10: Ablation study of hyper-parameter of  $\lambda$  based on NLinear model in terms of MAE.**

**Hyper-parameter  $\lambda$ .** We further conduct an ablation study to investigate the impact of the hyper-parameter  $\lambda$  in Equation 5, which balances the contributions of the weight loss  $\mathcal{L}_{weight}$  and the diet loss  $\mathcal{L}_{diet}$ . Figure 10 illustrates the performance of weight prediction under different settings with  $\lambda$  values ranging from 0 to 1, specifically  $\lambda \in \{0, 0.1, 0.25, 0.5, 0.75, 1\}$ . As  $\lambda$  increases from 0 to 1, which corresponds to a decreasing emphasis on  $\mathcal{L}_{diet}$ , we observe an initial improvement in performance followed by a gradual decline. The best results are achieved at  $\lambda = 0.1$  in most of the settings, demonstrating the robustness of the hyper-parameters across settings. Notably, when  $\lambda = 1$ , only  $\mathcal{L}_{weight}$  is used to optimize the model, which is equivalent to NLinear model by excluding food intake. On the contrary, when  $\lambda = 0$  and the model is solely optimized using  $\mathcal{L}_{diet}$ , resulting in inferior performance by combining the weight loss. The results show the significance of integrating both weight and diet losses for effective weight prediction.

## 5.5 Weight Prediction Visualization

As illustrated in Figure 11, we qualitatively compare the weight prediction visualization among ground-truth weight (blue), NLinear model (green), and our framework based on NLinear (orange) using images as dietary information for different users. It is evident that both the trend and the exact predicted values of our framework are closer to the ground truth than those of the NLinear model. For instance, in subplot (b), the ground truth weight trend shows an initial decrease, followed by a slight increase, and then another decrease. The orange line, representing our prediction results, exhibits the same trend and is numerically closer to the ground truth than the green line of the NLinear model. Similar trends can be observed in subplots (a), (d), (f), and (g) as well. However, in subplots (c) and (i), even though our prediction results are superior to the NLinear model, our framework's predictions do not align well with the ground truth trend. This discrepancy can be attributed to variable changes in weight and the accumulation of prediction errors in long-term predictions. This issue remains an unresolved challenge in time series forecasting tasks [55, 59].



**Figure 11: Weight prediction visualization. In each subplot, the x-axis represents the number of predicted days, and the y-axis refers to the weight(kg). The settings for 3-3, 5-5, and 7-3 are displayed in the first (i.e., (a), (b) and (c)), second (i.e., (d), (e) and (f)), and third (i.e., (g), (h) and (i)) rows respectively.**

## 6 CONCLUSION

We have investigated weight prediction using diet diary. A new dataset named DietDiary is constructed which comprises dietary intake and corresponding daily weight measurements over a period. We introduce a new task of predicting weights by leveraging historical food intake. To address this task, we propose a model-agnostic time series forecasting framework that achieves significant improvements in weight prediction. Our experiments highlight the effectiveness of the proposed UMRL module and diet-aware loss in learning unified intake representation and establishing a correlation between food intake and weight change. Furthermore, we examine three forms of food intake for weight prediction, including images, ingredient annotations provided by users, and ingredient labels predicted from a pre-trained ingredient prediction model. Experimental results on two representative time series prediction models, NLinear [59] and iTransformer [31], demonstrate that incorporating food intake leads to improved accuracy in weight prediction across all three forms. We also investigate multi-modal fusions of dietary information from different modalities (image and text) and achieve better performance, which demonstrates that multi-modal food information can complement each other and provide more effective food information. Through ablation studies, we further demonstrate the necessity of using all three meals for accurate weight prediction. While encouraging, this paper primarily focuses on weight prediction from food intake, exploring the impact of food to other physiological indicator such as blood glucose, will be our future work.

## REFERENCES

- [1] Marc Bolaños, Aina Ferrà, and Petia Radeva. 2017. Food ingredients recognition through multi-label learning. In *New Trends in Image Analysis and Processing-ICIAP 2017: ICIAP International Workshops, WBICV, SSPandBE, 3AS, RGBD, NIVAR, IWBAAS, and MADiMa 2017, Catania, Italy, September 11-15, 2017, Revised Selected Papers 19*. Springer, 394–402.
- [2] 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, 446–461.
- [3] George EP Box, Gwilym M Jenkins, Gregory C Reinsel, and Greta M Ljung. 2015. *Time series analysis: forecasting and control*. John Wiley & Sons.
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *European conference on computer vision*. Springer, 213–229.
- [5] 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*. 32–41.
- [6] 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*, Vol. 34. 10542–10550.
- [7] 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* 30 (2020), 1514–1526.
- [8] Jing-jing Chen, Chong-Wah Ngo, and Tat-Seng Chua. 2017. Cross-modal recipe retrieval with rich food attributes. In *Proceedings of the 25th ACM international conference on Multimedia*. 1771–1779.
- [9] Jing-Jing Chen, Chong-Wah Ngo, Fu-Li Feng, and Tat-Seng Chua. 2018. Deep understanding of cooking procedure for cross-modal recipe retrieval. In *Proceedings of the 26th ACM international conference on Multimedia*. 1020–1028.
- [10] Si-An Chen, Chun-Liang Li, Nate Yoder, Sercan O Arik, and Tomas Pfister. 2023. Tsmixer: An all-mlp architecture for time series forecasting. *arXiv preprint arXiv:2303.06053* (2023).
- [11] Xin Chen, Yu Zhu, Hua Zhou, Liang Diao, and Dongyan Wang. 2017. Chinese-foodnet: A large-scale image dataset for chinese food recognition. *arXiv preprint arXiv:1705.02743* (2017).
- [12] Prateek Chhikara, Dhiraj Chaurasia, Yifan Jiang, Omkar Masur, and Filip Ilievski. 2024. Fire: Food image to recipe generation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 8184–8194.
- [13] Abhimanyu Das, Weihao Kong, Andrew Leach, Rajat Sen, and Rose Yu. 2023. Long-term forecasting with tide: Time-series dense encoder. *arXiv preprint arXiv:2304.08424* (2023).
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [16] Vijay Ekambaram, Arindam Jati, Nam Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. 2023. Tsmixer: Lightweight mlp-mixer model for multivariate time series forecasting. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 459–469.
- [17] Charles NC Freitas, Filipe R Cordeiro, and Valmir Macario. 2020. Myfood: A food segmentation and classification system to aid nutritional monitoring. In *2020 33rd SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*. IEEE, 234–239.
- [18] Anqi Gao, Frank P-W Lo, and Benny Lo. 2018. Food volume estimation for quantifying dietary intake with a wearable camera. In *2018 IEEE 15th International Conference on Wearable and Implantable Body Sensor Networks (BSN)*. IEEE, 110–113.
- [19] Ricardo Guerrero, Hai X Pham, and Vladimir Pavlovic. 2021. Cross-modal retrieval and synthesis (x-mrs): Closing the modality gap in shared subspace learning. In *Proceedings of the 29th ACM International Conference on Multimedia*. 3192–3201.
- [20] Helena H. Lee, Ke Shu, Palakorn Achananuparp, Philips Kokoh Prasetyo, Yue Liu, Ee-Peng Lim, and Lav R Varshney. 2020. RecipeGPT: Generative pre-training based cooking recipe generation and evaluation system. In *Companion Proceedings of the Web Conference 2020*. 181–184.
- [21] Marc T Hamilton, Deborah G Hamilton, and Theodore W Zderic. 2007. Role of low energy expenditure and sitting in obesity, metabolic syndrome, type 2 diabetes, and cardiovascular disease. *Diabetes* 56, 11 (2007), 2655–2667.
- [22] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [23] Yuma Honbu and Keiji Yanai. 2022. Unseen food segmentation. In *Proceedings of the 2022 International Conference on Multimedia Retrieval*. 19–23.
- [24] Eric Jéquier and Luc Tappy. 1999. Regulation of body weight in humans. *Physiological reviews* 79, 2 (1999), 451–480.
- [25] Shuqiang Jiang, Weiqing Min, Linhu Liu, and Zhengdong Luo. 2019. Multi-scale multi-view deep feature aggregation for food recognition. *IEEE Transactions on Image Processing* 29 (2019), 265–276.
- [26] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [27] Chairi Kiourt, George Pavlidis, and Stella Markantonatou. 2020. Deep learning approaches in food recognition. *Machine Learning Paradigms: Advances in Deep Learning-based Technological Applications* (2020), 83–108.
- [28] Xing Lan, Jiayi Lyu, Hanyu Jiang, Kun Dong, Zehai Niu, Yi Zhang, and Jian Xue. 2023. Foodsam: Any food segmentation. *IEEE Transactions on Multimedia* (2023).
- [29] Chengxu Liu, Yuanzhi Liang, Yao Xue, Xueming Qian, and Jianlong Fu. 2020. Food and ingredient joint learning for fine-grained recognition. *IEEE transactions on circuits and systems for video technology* 31, 6 (2020), 2480–2493.
- [30] Shizhan Liu, Hang Yu, Cong Liao, Jianguo Li, Weiyao Lin, Alex X Liu, and Shahram Dustdar. 2021. Pyraformer: Low-complexity pyramidal attention for long-range time series modeling and forecasting. In *International conference on learning representations*.
- [31] Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. 2023. iTransformer: Inverted Transformers Are Effective for Time Series Forecasting. In *The Twelfth International Conference on Learning Representations*.
- [32] Frank P-W Lo, Yingnan Sun, Jianing Qiu, and Benny PL Lo. 2019. Point2volume: A vision-based dietary assessment approach using view synthesis. *IEEE Transactions on Industrial Informatics* 16, 1 (2019), 577–586.
- [33] Frank P-W Lo, Yingnan Sun, Jianing Qiu, and Benny PL Lo. 2019. Point2volume: A vision-based dietary assessment approach using view synthesis. *IEEE Transactions on Industrial Informatics* 16, 1 (2019), 577–586.
- [34] Sepehr Makhsous, Hashem M Mohammad, Jeannette M Schenk, Alexander V Mamishev, and Alan R Kristal. 2019. A novel mobile structured light system in food 3D reconstruction and volume estimation. *Sensors* 19, 3 (2019), 564.
- [35] Niki Martinel, Gian Luca Foresti, and Christian Micheloni. 2018. Wide-slice residual networks for food recognition. In *2018 IEEE Winter conference on applications of computer vision (WACV)*. IEEE, 567–576.
- [36] 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*. 1233–1241.
- [37] 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).
- [38] 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).
- [39] 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).
- [40] Shu Naritomi and Keiji Yanai. 2021. 3D Mesh Reconstruction of Foods from a Single Image. In *Proceedings of the 3rd Workshop on AIxFood*. 7–11.
- [41] Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. 2023. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *International Conference on Learning Representations*.
- [42] Siyuan Pan, Ling Dai, Xuhong Hou, Huating Li, and Bin Sheng. 2020. ChefGAN: Food image generation from recipes. In *Proceedings of the 28th ACM International Conference on Multimedia*. 4244–4252.
- [43] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PMLR, 8748–8763.
- [44] David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. 2020. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International journal of forecasting* 36, 3 (2020), 1181–1191.
- [45] Amaia Salvador, Michal Drozdal, Xavier Giró-i Nieto, and Adriana Romero. 2019. Inverse cooking: Recipe generation from food images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10453–10462.
- [46] Amaia Salvador, Erhan Gundogdu, Loris Bazzani, and Michael Donoser. 2021. Revamping cross-modal recipe retrieval with hierarchical transformers and self-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15475–15484.
- [47] 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*. 3020–3028.
- [48] Mustafa Shukor, Guillaume Couairon, Asya Grechka, and Matthieu Cord. 2022. Transformer decoders with multimodal regularization for cross-modal food retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*

- 1045 *Pattern Recognition*. 4567–4578.
- 1046 [49] Takuo Suzuki, Kana Futatsuishi, Kana Yokoyama, and Nobuko Amaki. 2020. Point cloud processing method for food volume estimation based on dish space. In *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 5665–5668.
- 1047
- 1048 [50] Chi-en Amy Tai, Matthew Keller, Saejith Nair, Yuhao Chen, Yifan Wu, Olivia Markham, Krish Parmar, Pengcheng Xi, Heather Keller, Sharon Kirkpatrick, et al. 2023. NutritionVerse: Empirical Study of Various Dietary Intake Estimation Approaches. In *Proceedings of the 8th International Workshop on Multimedia Assisted Dietary Management*. 11–19.
- 1050 [51] Sean J Taylor and Benjamin Letham. 2018. Forecasting at scale. *The American Statistician* 72, 1 (2018), 37–45.
- 1052 [52] 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*. 8903–8911.
- 1054 [53] Hao Wang, Doyen Sahoo, Chenghao Liu, Ee-peng Lim, and Steven CH Hoi. 2019. Learning cross-modal embeddings with adversarial networks for cooking recipes and food images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 11572–11581.
- 1055 [54] Hao Wang, Doyen Sahoo, Chenghao Liu, Ke Shu, Palakorn Achananuparp, Ee-peng Lim, and Steven CH Hoi. 2021. Cross-modal food retrieval: learning a joint embedding of food images and recipes with semantic consistency and attention mechanism. *IEEE Transactions on Multimedia* 24 (2021), 2515–2525.
- 1056 [55] XiaoFeng Wang and Ying Zhang. 2020. Multi-step-ahead time series prediction method with stacking LSTM neural network. In *2020 3rd International Conference on Artificial Intelligence and Big Data (ICAIBD)*. IEEE, 51–55.
- 1057
- 1058
- 1059
- 1060
- 1061
- 1062
- 1063
- 1064
- 1065
- 1066
- 1067
- 1068
- 1069
- 1070
- 1071
- 1072
- 1073
- 1074
- 1075
- 1076
- 1077
- 1078
- 1079
- 1080
- 1081
- 1082
- 1083
- 1084
- 1085
- 1086
- 1087
- 1088
- 1089
- 1090
- 1091
- 1092
- 1093
- 1094
- 1095
- 1096
- 1097
- 1098
- 1099
- 1100
- 1101
- 1102
- [56] Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. 2021. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in neural information processing systems* 34 (2021), 22419–22430.
- [57] Xiongwei Wu, Xin Fu, Ying Liu, Ee-Peng Lim, Steven CH Hoi, and Qianru Sun. 2021. A large-scale benchmark for food image segmentation. In *Proceedings of the 29th ACM international conference on multimedia*. 506–515.
- [58] 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).
- [59] Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. 2023. Are transformers effective for time series forecasting?. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 37. 11121–11128.
- [60] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 11106–11115.
- [61] Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. 2022. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International conference on machine learning*. PMLR, 27268–27286.
- [62] Bin Zhu and Chong-Wah Ngo. 2020. CookGAN: Causality based text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5519–5527.
- [63] Bin Zhu, Chong-Wah Ngo, Jingjing Chen, and Yanbin Hao. 2019. R2gan: Cross-modal recipe retrieval with generative adversarial network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11477–11486.
- [64] Bin Zhu, Chong-Wah Ngo, and Jing-jing Chen. 2020. Cross-domain cross-modal food transfer. In *Proceedings of the 28th ACM International Conference on Multimedia*. 3762–3770.
- 1103
- 1104
- 1105
- 1106
- 1107
- 1108
- 1109
- 1110
- 1111
- 1112
- 1113
- 1114
- 1115
- 1116
- 1117
- 1118
- 1119
- 1120
- 1121
- 1122
- 1123
- 1124
- 1125
- 1126
- 1127
- 1128
- 1129
- 1130
- 1131
- 1132
- 1133
- 1134
- 1135
- 1136
- 1137
- 1138
- 1139
- 1140
- 1141
- 1142
- 1143
- 1144
- 1145
- 1146
- 1147
- 1148
- 1149
- 1150
- 1151
- 1152
- 1153
- 1154
- 1155
- 1156
- 1157
- 1158
- 1159
- 1160