Following Good Examples - Health Goal-Oriented Food Recommendation based on Behavior Data

Yabo Ling
Université de Montréal
Montréal, Québec, Canada
yabo.ling@umontreal.ca

Jian-Yun Nie
Université de Montréal
Montréal, Québec, Canada
nie@iro.umontreal.ca

Daiva Nielsen
McGill University
Montréal, Québec, Canada
daiva.nielsen@mcgill.ca

Bärbel Knäuper
McGill University
Montréal, Québec, Canada
baerbel.knaeuper@mcgill.ca

Nathan Yang
Cornell University
Ithaca, New York, USA
ncy6@cornell.edu

Laurette Dubé
McGill University
Montréal, Québec, Canada
laurette.dube@mcgill.ca

ABSTRACT

Typical recommender systems try to mimic the past behaviors of users to make future recommendations. For example, in food recommendations, they tend to recommend the foods the user prefers. While the recommended foods may be easily accepted by the user, it cannot improve the user's dietary habits for a specific goal such as weight control. In this paper, we build a food recommendation system that can be used on the web or in a mobile app to help users meet their goals on body weight, while also taking into account their health information (BMI) and the nutrition information of foods (calories). Instead of applying dietary guidelines as constraints, we build recommendation models from the successful behaviors of comparable users: the weight loss model is trained using the historical food consumption data of similar users who successfully lost weight. By combining such a goal-oriented recommendation model with a general model, the recommendations can be smoothly tuned toward the goal without disruptive food changes. We tested the approach on real data collected from a popular weight management app. It is shown that our recommendation approach can better predict the foods for test periods where the user truly meets the goal, than the typical existing approaches.

KEYWORDS

Food Recommendation, Health Factors, Weight Loss, Deep Learning

ACM Reference Format:

Yabo Ling, Jian-Yun Nie, Daiva Nielsen, Bärbel Knäuper, Nathan Yang, and Laurette Dubé. 2022. Following Good Examples - Health Goal-Oriented Food Recommendation based on Behavior Data . In *Proceedings of the ACM Web Conference 2022 (WWW '22), April 25–29, 2022, Virtual Event, Lyon, France.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3485447.3514193

WWW '22, April 25–29, 2022, Virtual Event, Lyon, France

ACM ISBN 978-1-4503-9096-5/22/04…$15.00
https://doi.org/10.1145/3485447.3514193

1 INTRODUCTION

Overweight and obesity are among the main causes to chronic diseases [29]. Body weight is a health aspect that many people care about today and try to manage. Many apps and web-based programs have been developed to help users manage their body weight [13, 21, 24, 43]. It is known that the main factors important to body weight are related to genetics, behavior, diet and environment. Among them, diet is one of the most important controllable factors that affect body weight and body health [3, 19]. However, according to the World Health Report 2020, the incidence rate of numerous diet-related diseases, such as diabetes, obesity, and malnutrition, is quickly growing over the world. [1] In particular, obesity has reached epidemic proportions, with at least 2.8 million people dying each year as a result of being overweighted or obese. Therefore, a well-balanced diet is crucial to maintain one's physical health.

Food recommendation has been an extensively studied topic because of its high potential impact on human physical health. Food recommendation can follow dietary guidelines developed by experts. [2] For example, it is recommended to limit the amount of sodium to less than 2,300 mg per day. While these guidelines are important, they can hardly be translated by the general population into their daily life [42]. Recommendation systems can offer more effective solutions by recommending concrete food items or ingredients to users on the fly. The popularity of such tools is growing with the widespread utilization of mobile devices or the web.

A typical type of food recommendation system attempts to capture users' previous food preferences and provide future recommendations accordingly, regardless of health factors [9, 12, 28, 36, 40]. For example, a recommender system may keep suggesting 'pizza' to a user who often consumed 'pizza' in the past. The recommendation may be easily adopted by the user, but this cannot lead the user to a healthy diet or eating pattern.

A second type of recommendation is to impose some diet constraints to users. For example, the caloric consumption is controlled within the recommended range. In a system developed to deal with malnutrition for elderly, Aberg [1] proposed a recommendation approach based on constraint satisfaction to meet a set of constraints

[1] https://www.who.int/data/gho/data/themes/topics/topic-details/GHO/world-health-statistics

[2] https://food-guide.canada.ca/en/guidelines/; https://www.who.int/news-room/fact-sheets/detail/healthy-diet

on several aspects. While such an approach can be applied in specific context where there is close monitoring, it is difficult to impose strict dietary constraints to general users.

A third approach tries to balance the user preferences and the dietary needs. This approach may lead users to a healthier diet while taking into account the food preferences. The recommended foods may be more easily adopted by users. Our study falls into this category. The existing approaches in this category have been simplistic. Elsweiler and Harvey [10] make recommendations based on user preferences at first; then the recommended foods are filtered using the required dietary constraints (e.g. food with calories higher than a required value are filtered out). The recommendations made with such a simplistic solution may be hard to be adopted by users.

Different from the previous approaches, in this work, we propose an approach that makes food recommendations based on both user's preferences and the success stories of similar users. The approach we develop relies on a large amount of eating behavior data of real users from a weight management app, where we can observe the daily food consumption and the body weight of many users. The eating behaviors that lead to weight loss are used as good examples to guide food consumption of a similar user who wants to lose weight. This is done by employing a goal-oriented recommendation model trained with eating behavior data that truly led to the goal. This strategy differs from the one that relies on dietary guidelines in the sense that the guidelines are now hidden in the concrete successful behaviors. For a user who has a goal on body weight (e.g. losing weight), we make food recommendations by combining a general preference-based recommendation model with a goal-oriented model based on the successful behaviors.

As users differ in their dietary needs, it is important to rely on the goal-oriented model of similar users. To this end, we train several goal-oriented models for different groups of users. For this study, we group users into different BMI (Body Mass Index) groups by assuming that users in the same group would have (more) similar dietary needs. These groups can be refined in the future to take into account other factors such as age, sex, and so on. To train goal-oriented models, we segment the data collected from an app into subsets of periods, where we observe that user truly lose/maintain/gain body weights. The intuition is that these users would have adopted food patterns during these periods that lead to the observed weight loss/maintenance/gain. These patterns can be generalized to similar users with the same goal.

To evaluate the effectiveness, the ideal setting is to let the users use the recommender system and observe their final weight outcome. Unfortunately, this evaluation takes a long time and is not yet feasible. Instead, we use the recorded data as proxy in the following way: For a user with a weight goal (e.g. losing weight), we compare our recommendation approach to a general one for the test periods where the user reaches the goal. If the recommendation accuracy of our approach is higher for these periods, it is considered to be better suited to the the user and the goal. Experiments show that our combined recommendation system can produce higher recommendation accuracy for different weight goals. This suggests that the approach could effectively adapt the recommendations toward the user's goal, while also considering his/her preferences.

The main contributions of this paper are summarized as follows: (1) We propose goal-oriented recommendation models for different groups of users based on successful behavior data that meet the goals. (2) To lead user's food choices toward a goal, we combine a general model with the corresponding goal-oriented model. This model combination approach is different from the traditional one based on dietary guidelines. (3) Our experiments on a large set of real data demonstrate the effectiveness of this approach.

2 RELATED WORK

2.1 General Recommender Systems

Traditional recommender systems, such as Collaborative Filtering (CF) [20, 32], model user-item historical interactions in a static way to capture users' general preferences. However, these approaches ignore the sequential signals, such as the order of users' behaviors. So sequential recommendation is proposed to solve this problem. The problems can be further divided into Next-item and Next-basket recommendations depending on the output – only one item vs whole basket of items. Our task aims to recommend a set of food items for a day for the user. This falls in the Next-basket recommendation. There are two main methods to capture the sequential patterns from users historical interactions in Next-basket recommendation: (1) by Markov Chains (MCs). The representative system is FPMC proposed by Rendle et al. [28]. It combines both a common Markov Chain (MC) and the normal Matrix Factorization (MF) to model both sequential behaviors and general interests. (2) by Recurrent Neural Networks (RNNs). Many models are proposed based on RNNs in Next-basket recommendation, such as session-based GRU with ranking loss (GRU4Rec) [36], Dynamic REcurrent bAsket Model (DREAM) [49] and Attribute-aware Neural Attentive Model (ANAM) [4]. In these models, one considers the sequential purchase history, interactions between items within a basket, user's preferences and item properties. Recommendations are made by aggregating such information. Our basic food recommendation model is inspired by the next-basket recommendation models because recommending food for a day can be viewed as a similar problem to next-basket recommendation.

2.2 Food Recommendation

As we mentioned earlier, approaches to food recommendation can be categorized into the following categories [39] that consider: (1) user's preferences [9]; (2) nutritional needs of users [1]; (3) both of them [10]. Our study is related to the third category, but we take a different approach than the existing work.

Current methods for food recommendation mainly consider the nutrition information via food analysis or external nutrition guidelines. For a weight goal, it may be recommended to consume some groups of foods instead of others, or consume different groups of food at some proportion. Most automatic methods have tried to incorporate healthiness into the recommendation process by substituting ingredients [6, 38], incorporating calorie counts [11], and generating food plans [10]. A critical problem is that the recommended foods or food plans may be very different from the user's habits and food preferences. In many cases, a user will start a diet program, but abandon it after a while [7, 33]. The key issue is that a too drastic change in food choices can be hardly adopted by a user. An alternative approach is to make gradual changes, by proposing some food alternatives that are acceptable to the user. Over time, the accumulated changes will eventually lead the user to the goal. This gradual approach has been found more successful, especially for users who are not monitored by a professional [2, 18]. This is the general approach we take in this work.

Many food guidelines are created for general population, without taking into account the specific health information of the user. While personalized food guides can be built with the help of a nutrition professional, this solution is not accessible to many people. The general food recommendation systems do not incorporate user's personal health information such as BMI. Some existing recommendation methods balance user's food preference and user's health via simple fusion inside the model. For example, Ge et al. [11] simply calculated the weighted mean between the preference component and health component. Elsweiler and Harvey [10] used a simple approach to increase or reduce the amount of calories by 500 kilocalories for users who want to gain or to lose weight. However, it may also be the case that the expected outcome is not produced, or the user will not adopt the recommended foods - a critical challenge in this domain [39]. We believe that making food recommendation by following real success examples observed in data can make a more acceptable plan for users toward the expected outcome. This motivates us to develop an approach combining goal-oriented recommendation with user preferences.

3 PROBLEM DESCRIPTION AND SOLUTION DESIGN

In this section, we explain the factors we consider in our food recommendation method, as well as the data we use.

3.1 Health Factors

Many health- and nutrition-related factors are important for human weight and health management. The first nutrition factor is food calories. To make a healthy and balanced diet, it is important to consume foods that provide enough but not too much calories [14, 25, 41]. Studies have shown that calorie intake less than the required amount may have impacts on one's weight management [45]. Caloric restriction is the most common method to control weight [47]. Weight loss occurs when energy intake is less than energy expenditure [5, 8]. Therefore, to successfully build a health recommendation system, we will incorporate calorie information of food in all the recommendation approaches in this paper.

Different users have different dietary needs. Several factors can influence the dietary needs of a user: body weight, height, etc. In this study, we take into account the Body Mass Index (BMI), which is a measure of body fat based on height and weight that applies to adult men and women. According to the BMI classification defined by Health Canada[3], users can be divided into 6 groups: 'Underweight', 'Normal Weight', 'Overweight', 'Obese class I', 'Obese class II' and 'Obese class III'). We simplify the classification by grouping all the classes of "Obese" together. Users are thus grouped into 4 categories.

In this study, we do not consider other factors such as food categories (meal, vegetable, etc.) [34], food energy density [30], etc. This is left to future work.

[3]https://www.canada.ca/en/health-canada/services/food-nutrition/healthy-eating/healthy-weights/canadian-guidelines-body-weight-classification-adults/body-mass-index-nomogram.html

Table 1: Average Calories Intake of Different Groups

User Segment	Item	Whole Day
All	137.56	1758.11
Under Weight	89.63	1358.09
Normal	124.12	1703.74
Overweight	142.81	1806.48
Obese	147.31	1768.51
Increase	140.36	1833.69
Decrease	132.12	1716.66
Maintenance	139.68	1754.48

3.2 Data and Data Segmentation

Our approach is built on anonymized real data collected from a popular weight management app. All personal information of the users are removed and we only keep an ID number for each user. The data is collected from January 2016 to December 2016. In the app, in addition to the general information about the user (age, sex, etc.), users record their daily food consumption, body weight, energy expenditure by exercises, etc. Food recorded are chosen from a list of 312 food items. The calories associated to each food are input by users and they can vary. To simplify, we take the average amount of calories to a food when making recommendation. To be able to train our models, we only keep users whose age falls in $[18, 65]$ and have continuous records (at least 2,000 food item records and 200 recorded days). Then duplicates within each day are removed, and days with fewer than two items are deleted. This provides us with 15,408,719 total food records, for 7,496 users. The users are partitioned into different BMI groups as follows: 39 underweight, 2,286 normal weight, 2,576 overweight and 2,595 obese.

The average food calories consumed by different groups of users are shown in Table 1. As we can see, the general trend is that the "underweight" group consumes less calories than "normal weight" group, which is lower than "overweight" and "obese" groups, both on calorie per food item and daily calorie intake. This is consistent with our intuition. However, we do not observe a large difference between the "overweight" and "obese" groups. The numbers shown in the table confirm that calorie consumption may be strongly associated to the body weight (BMI). So any food recommendation system should incorporate food calorie information.

Exercise calorie expenditure is not considered in this work. It would be analyzed further in the future work.

3.3 Aligning User Intent and Weight Outcome

The app recorded the global intent of users, indicating their goal when registering into the app. A large proportion of the users intended to lose weight. However, the real outcomes do not match the global intent in many cases. During the course of utilization of the app, the user's body weight fluctuates and the final weight change may, or may not, correspond to the global intent. This can be seen in Figure 1 where we plot the real weight changes of three representative users. The weights of all users fluctuate, but the general trend of user #450 increases, that of user #28403 decreases, while the weight of user #10378 generally maintains. Looking at

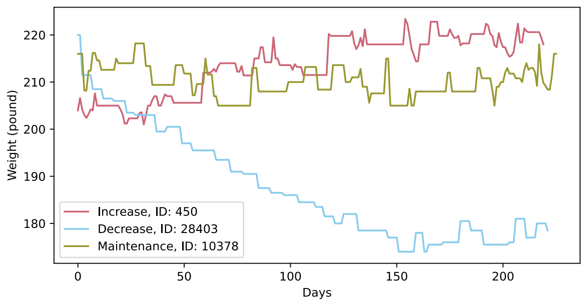

Figure 1: Weight Trends for Randomly selected Users with Three Different Intents

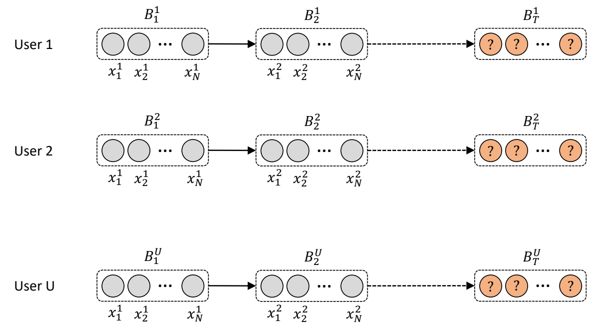

Figure 2: Next basket food recommendation task

the global intent, user #10378 wanted to lose weight, but did not succeed.

The data we collected suggest that the food choices made by users are not always consistent with their global intent. Imposing food selections according to the global intent would make the food choices very different from users' eating habits, thereby jeopardizing their adoption by the users. On the other hand, the real weight changes during a period of time may better reflect the true impact of food choices during the period, i.e. the food choices truly led to the weight outcome. Therefore, we can consider such a period as a good example that accomplishes an implicit goal corresponding to the weight outcome. Yet the foods chosen by the user during the period can be considered as an acceptable food pattern by the user. An intent model trained on such examples may be a smoother adaptation of the user's normal eating pattern toward the implicit goal, than the traditional food recommendation approaches based on dietary guidelines [1]. Based on the above observation, we propose an approach to model intents based on behavior data of the periods (i.e. good examples) that produced the intended weight change. The philosophy of our approach to food recommendation can be stated as "*following good examples*".

In order to obtain example data to train intent models, we segment the behavior data into different periods of time, each corresponding to a specific type of weight outcome: increase/maintenance /decrease. We consider a week as an appropriate period for this purpose, as this has been found to be a good time granularity to observe weight changes [27, 31, 45]. Our intent models will be trained using data of the corresponding weeks to reflect the eating behaviors of weight increase/maintenance/decrease.

As we mentioned earlier, users of different BMI groups should have different eating behaviors. Therefore, we segment the three intent sub-groups within each BMI group and use them to train goal-oriented models specific for the BMI groups. However, for the underweight group, we only have 39 users. Indeed, the underweight group are atypical users of the app. So, we do not deal with this group of users specifically. The statistics on the number of days and food records in different intent sub-groups within each BMI group are shown in Table 2.

4 METHODOLOGY

In this section, we formulate the task of food recommendation for the next day as next-basket recommendation and then introduce the proposed model in detail. The model is trained on a set of training data. Depending on the training data, we will obtain different models to reflect the eating behaviors of different groups of users and with different intents, which are then combined.

4.1 Problem Formulation

Next-basket recommendation aims at recommending a set of items to put in a basket [4, 28, 46, 48, 49]. This is different from recommending a single item at a time and the interactions between items in a basket could be captured. Our recommendation task is similar to next-basket recommendation: we aim to recommend foods of a whole day to the user. The relative positions of each item within the day/basket is not important – **[*coffee, egg, chocolate*]** and **[*chocolate, egg, coffee*]** are considered the same basket. In the following description, we will keep using basket to denote a day in order to correspond to the notions used in the previous studies on next-basket recommendation.

The next-basket recommendation task is as follows (see Figure 2): Given a sequential basket records $B^u = \{B_1^u, B_2^u, \cdots, B_T^u\}$ of user u, we predict the last basket B_T^u based on the previous basket sequence $\{B_1^u, B_2^u, \cdots, B_{T-1}^u\}$.

4.2 Model Architecture

In this paper, we propose a **H**ealth-factor-aware **H**ierarchical **F**ood **R**ecommender (HHFR) architecture to incorporate health-related information, which is illustrated in Figure 3. This architecture is an adaptation of those proposed for next-basket recommendation to our problem. HHFR utilizes a hierarchical RNN to model the user's sequential behavior over time, and incorporates health information (calorie) of each item into a basket.

In the following, we first model the information of a basket: encoding information of items and health factors (calorie) in each basket; and learning the intra-basket representation. Then we model user's sequential inter-basket behavior by the basket-level RNN.

4.3 Item and Calorie Encoding

The input instances of HHFR model are the representations of a sequence of baskets $\{B_1, B_2, \cdots B_T\}$ where T denotes the total number of basket. Every basket representation is of the form $B_t = \{C_j^t | j = 1, 2, \cdots, N\}$, where N is the number of items in basket B_t.

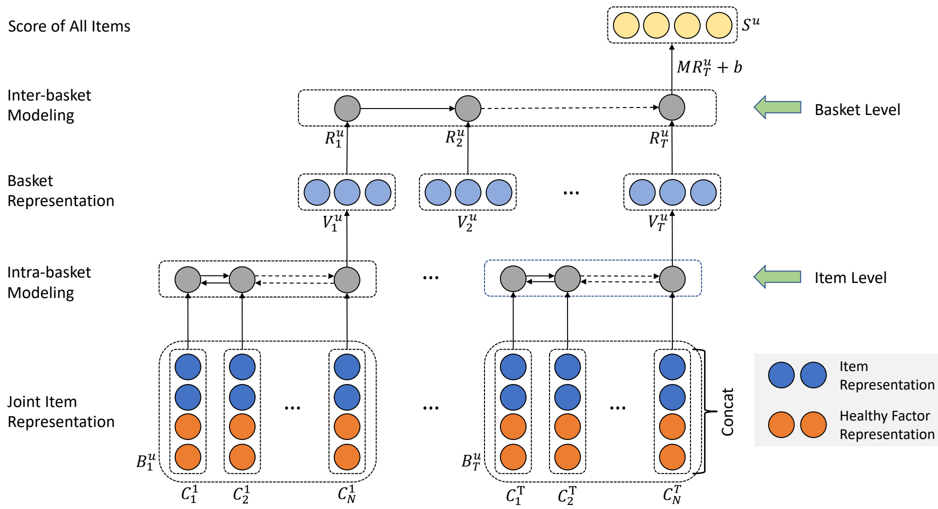

Figure 3: The framework of HHFR

Each C_j^t is a concatenated representation of a food item and its amount of calories.

For initial item representation, our model adopts the commonly used pre-trained GloVe word vectors [26] [4] to map each food item name I to a d_a-dimension embedding $E^I = \{e_1^I, \cdots, e_{d_a}^I\}$.

For each item, calorie is a numerical value, which is in a very different order of magnitude than the GloVe and BERT embedding (10^2 vs. 10^{-3}). Thus, we cannot directly input it into the model. Look-up Embedding is a commonly used method in deep learning to convert a numeric value into an embedding. It first randomly initializes a trainable embedding matrix for calorie values, and update it throughout the training process. However, the embedding matrix may not be trained sufficiently, because some calorie values only appear few times during the training process. This is confirmed in our preliminary experiments. Therefore, the following transformation strategy is designed for convert calories to an embedding.

This strategy is inspired by the position encoding in Transformer [44], which employs sinusoidal vectors to encode the position of each term. Similarly, we also want to map a numerical calorie value to an embedding, but our requirements are different:

(1) **Order preserving:** Calories are real-valued scalars, which carries physical meanings. Given two calorie values x_1, x_2 with $x_1 < x_2$, we desire to preserve the order of values in each dimension of the vectors, i.e. $f(x_1)_{(i)} < f(x_2)_{(i)}, \forall i = 1 \cdots n$, where $f(x_1)_{(i)}$ is the i^{th} dimension of the mapped vector for calorie value x_1.

(2) **Non-periodicity:** Moreover the transformation need to be strictly monotonous, i.e. given $x_1 \neq x_2$, each dimension of the mapped vector should have different values $f(x_1)_{(i)} \neq f(x_2)_{(i)}, \forall i = 1 \cdots n$. This is because we do not want to map two different calorie values to the same embedding vector.

Taking the above desired properties into account, we choose hyperbolic tangent (tanh) with some translation and stretching as our transformation function. Therefore, for a numeric calorie value of item i, the transformation function generates $E_i^H = [e_1^H, e_2^H, \cdots, e_{d_b}^H]$ as its encoding. The calculation of the value in j-th dimension e_j^H is presented in Equation 1.

$$e_j^H = 2\tanh\left(\frac{x}{500 + \frac{1000}{d_b - 1} i} - 0.5\right), 0 \leq i < d_b, \tag{1}$$

where d_b is the dimension of the mapped vector, x represents the calorie value ($x \in [0, 1000]$).

This mapping transforms a calorie value into an embedding containing real numbers in the same order of magnitude as item embeddings. In our preliminary experiments, we tested several alternative methods for calorie embedding, and this one worked the best.

Then the joint item representation is created by concatenating the item and calorie embeddings, i.e. for j-th item:

$$C_j = \text{CONCAT}(E_j^I, E_j^H), \tag{2}$$

where $C_j \in \mathbb{R}^{|d_a + d_b| \times 1}$. The representation of the t-th basket B_t with N item is $B_t \in \mathbb{R}^{|d_a + d_b| \times N}$.

4.4 Intra-basket Representation Modeling

Given a sequence of items in a basket $B_t^u = \{C_j^t | j = 1, 2, \cdots, N\}$ of user u, we first apply an item-level bi-directional Long Short-Term Memories (BiLSTM) [16] to capture the interactions between the items within the basket. Then the resulting representation of the last item in the basket is used as the representation V_t^u of the whole basket B_t^u.

4.5 Inter-basket Sequential Modeling

In order to model a user's sequential inter-basket behavior, another basket-level LSTM is applied to the basket representations. Given a sequence of baskets $\{V_1^u, V_2^u, \cdots, V_T^u\}$ for user u, this LSTM will generate another representation for each basket. We take the representation of the last basket as the representation R_T^u for the whole sequence of baskets of user u.

[4] We also tested with BERT, which led to similar performance.

The output scores S^u can be calculated through multiplication of item matrix M and the whole basket sequence representation R^u_T, which is formulated as follows:

$$S^u = MR^u_T + b. \quad (3)$$

We have $S^u_i \in \mathbb{R}^{|I|\times 1}$, i.e., an element of S^u_i represents the interaction score between an item i and a user u. A higher score indicates that the user is more likely to consume the corresponding item.

4.6 The Loss Function for Optimization

For a user u and his/her previous baskets $B^u_{1,t}$, we define the probability of an item i being taken in the next basket B^u_{t+1} by sigmoid function:

$$p(i \in B^u_{t+1}|u, B^u_{1,t}) = \frac{1}{1 + e^{-S^u_i}}. \quad (4)$$

To effectively learn from the training data, we adopt a weighted cross-entropy as the optimization objective at each step of LSTM, which is defined as:

$$L = \sum_{u\in U} \sum_{B^u_t \in B^u} \sum_{I\in I_t} (-m \cdot y_i \log p_i - n \cdot (1 - y_i) \log(1 - p_i)), \quad (5)$$

where p_i is the probability of an item i being consumed in the next basket in our model. If item i is consumed in the the next basket, $y_i = 1$, otherwise, $y_i = 0$. m and n are the weights of positive and negative instances (consumed or not in the next basket). The reason of using different weights is to cope with the fact that there are usually much more negative instances than positive instances in our dataset. After training, given a user's historical records, we can obtain the probability of each item i being taken in the next basket according to Equation 4. We then rank the items according to their probability, and select top K results as the final recommended items to the user.

5 EXPERIMENT

5.1 Dataset and Evaluation Metrics

Dataset. Our experiments are conducted on a large set of real-world anonymized user behavior data for the whole year of 2016 from a mobile health application. The dataset is described in Section 3.2. The pre-processed data are segmented into different sub-datasets according to BMI and intent, as described in Section 3.1. Each week is associated to a weight outcome: increase/maintenance/decrease. The statistics of the dataset are summarized in Table 2. The data are split into training/validation/test sets.

Evaluation method. To evaluate the effectiveness of a food recommendation system for the purpose of weight management, the ideal setting is to test it with real users in real situations. However, this test is difficult to implement now, as updating the app would require much more evaluations before. It also takes a long time to run tests in real situations. As a proxy, we use the collected data for our test. The assumption is that if our recommendation model can make recommendations that correspond better to the foods consumed by the user during a period which produces the intended outcome, then the model is considered to be better. For example, for a test period where the user indeed loses weight, our model that takes into account the BMI and the "lose-weight" intent should recommend foods more similar to those consumed by the user during this period, than the compared baseline model. The same for other intents.

Table 2: Statistics of the dataset. No/Ov/Ob means the BMI groups: Normal weight/Overweight/Obese, while In/De/Ma means weight outcome/intent: increase/decrease/maintenance.

			# Food Records			
Datasets	**# User**	**# days**	**Train**	**Valid**	**Test**	**Avg. Cal**
All	7,457	1,292,251	15,143,793	76,950	73,660	137.6
Normal	2,286	424,814	5,340,033	25,441	24,037	124.1
Overweight	2,576	459,566	5,330,912	26,542	25,110	142.8
Obese	2,595	407,871	4,472,848	24,967	24,513	147.3
No_In	2,079	105,271	1,301,246	24,823	24,685	125.9
No_De	2,191	129,999	1,622,665	26,719	25,898	120.9
No_Ma	2,098	189,544	2,316,590	23,745	23,140	125.4
Ov_In	2,251	102,778	1,165,460	24,603	24,774	146.3
Ov_De	2,504	158,977	1,855,092	28,178	27,401	136.7
Ov_Ma	2,310	197,811	2,209,588	23,996	23,472	146.1
Ob_In	1,970	71,805	764,300	20,431	20,472	155.0
Ob_De	2,514	174,147	1,951,510	26,380	25,900	138.1
Ob_Ma	2,214	161,919	1,670,981	21,076	21,278	154.5

Evaluation metrics. We examine the quality of top K (i.e, K = 1, 5, 10) items recommended by a model. Five evaluation metrics are selected to evaluate the performances of our recommendation systems: Precision@K, R-Precision, Recall@K, and NDCG@K. These are common metrics used in the literature [15, 17, 22, 23, 28, 35–37].

5.2 Implementation Details

Following the previous work [4, 49], we take the last basket of each user as the testing data, the penultimate basket as the validation set to optimize parameters, and the remaining baskets as the training data. The statistics of these splits are shown in Table 2. We implement our models in Python using the library PyTorch. The loss function in Equation 5 is optimized by Adam with a batch size of 128. Due to the disparity of the amount of positive and negative instances, we set m (for positive instances) in Equation (5) 300 times larger than n in our experiments to punish the error of mistaking the positive instances. The learning rate is set to 0.001 and the units in item-level BiLSTM and basket-level LSTM are set to 128 and 256, respectively. The GloVe item embedding size d_a is set to 300, and the calories factor embedding size d_b = 128.

5.3 Model Combinations

We train different models to capture the eating patterns of different groups of users, including:

- a general model for all the users (**General**): the model is trained with all the training data mixed together.
- models for each BMI group (**BMI**): A BMI model is trained with data from each BMI group.
- models for each intent (**Int**): These models are trained on data from periods corresponding to the intent. We expect these models capture some common eating patterns of users with the same intent.

- models for different intent/outcome groups within each BMI group (**BMI+Int**): Each model is trained with the specific intent sub-group (e.g. No_In) within a BMI group.

The following model combinations are tested.

(1) **Combining Two Models:** For a user of a BMI group having a weight change intent, we employ the corresponding BMI+Intent model, combined with the BMI or Intent model. The assumption is that the more general **BMI/Int** model can provide general recommendations appropriate for the whole BMI/Intent group, while the **BMI+Int** model can provide more specific recommendations for the specific sub-group. This combination also aims to solve the data sparsity problem: as **BMI+Int** is trained on less data, it is possible that some eating patterns are not well covered in it. The combination with a more general **BMI** or **Int** model could increase the coverage. The score of recommendation of a food item x is determined as follows:

$$S(x) = (1-\alpha) \cdot S_{\text{BMI/Int}}(x) + \alpha \cdot S_{\text{BMI+Int}}(x), \quad \alpha \in [0,1], \quad (6)$$

where α is a parameter controlling the impact of BMI on our recommendation. This parameter is tuned on validation data. When α equals 0, it means we apply the completely middle-model score as the new score. Similarly, it means we only utilize the **BMI+Int** sub-model when α equals 1.

(2) **Combining Three Models:** In this scenario, the recommendation for a user combines the corresponding **BMI+int** model, **BMI** model and the **General** model. The two latter models serve as the background models to complement the specific **BMI+Int** model for the user. The recommendation score is given by Equation 7:

$$\begin{aligned} S(x) = &(1-\alpha-\beta) \cdot S_{\text{General}}(x) + \alpha \cdot S_{\text{BMI}}(x) + \beta \cdot S_{\text{BMI+Int}}(x), \\ &\alpha, \beta \in [0,1], \quad \alpha+\beta \leq 1, \end{aligned} \quad (7)$$

where α and β are two parameters to be tuned with validation data.

5.4 Experimental Results

The performances of NDCG@10 and R-Precision of all models are presented in Table 3 and Table 4, respectively. In the first part of these tables, we show the effectiveness of each individual model. In the middle, we show the effectiveness of 2-model combinations. In the last part, we show the 3-model combination. As we can see, the combination of three models achieves the best results in general (except one case on R_Precision), with an average improvement of 14.2% on NDCG@10 and 15.6% on R-Precision over the **General** model. All these improvements are statistically significant. This result clearly demonstrates that a recommendation model that takes into account the specific BMI and the intent of the user can make better recommendations for the periods that produce the intended outcome. The approach is equally effective across different BMI groups and intents.

We now examine the impact of each factor.

Impact of Intent: We can observe that the performance of Intent-specific model is higher than the general model on both NDCG@10 (average 0.2386 vs 0.2308) and R-Precision (average 0.2412 vs 0.2328). Recall that the intent-specific model is trained on data from the periods of all users producing the intended outcome. The higher effectiveness of the intent models indicates that the eating patterns of different BMI groups of users with the same intent bear some similarity, which is captured by the intent models. Applying the intent models to the test periods with the intended outcome produces better predictions of foods. Taking this observation in the reverse direction, we could expect that the recommendations with the intent model could produce the intended weight outcome.

When the intent models are trained for each specific BMI group (**BMI+Int** models), the effectiveness is further improved. This shows that it is more appropriate to recommend foods based on the BMI of the user. This observation is intuitive: people with different BMI would have different dietary needs, so should have recommendations specific for the BMI group. Although we only incorporated BMI, we expect similar results on other physical characteristics.

Impact of BMI: BMI-specific models produce results (average 0.2532 for NDCG@10) not only better than the general model (average 0.2308 for NDCG@10), but also slight higher than these of Intent model (average 0.2386 for NDCG@10). This suggests that BMI is a more important factor than Intent to consider for food recommendation. In other words, people with similar BMI may have more similar eating patterns. The **BMI** models also produce better results when they are combined with other models, confirming once again the usefulness of model combination.

Impact of model combination: Comparing combination models with individual models, we can see that when more models are combined, better results obtain: 3-model combinations are generally better than 2-model combinations, which in turn are better than individual models. This confirms our earlier intuition that, although **BMI+Int** models can capture the specific eating patterns of a subgroup of user, it may suffer from poor coverage; so, combining it with more general models helps improve their coverage. This is also corroborated by the amount of training data: As we can see in Table 2, each **BMI+Int** model is trained on about 1/3 of the data of **BMI** or **Int** models, which in turn uses about 1/3 of the whole data. Model combination is thus a possible solution to address the data sparsity problem.

Impact of combination factors With the same 2-model combination (Equation 6), we examine the impact of the combination factor α in Figure 4: how the relative weight between **BMI** and **BMI+Int** model impact the effectiveness. The left-most points correspond to the utilization of **BMI** models only, while the right most points to the **BMI+Int** models. We can see that the best setting is to balance both models (with slightly higher weight for **BMI+Int** for the Overweight group). This observation confirms that in the combination, all combined models contribute in the recommendations. The model combination strategy used in this work is effective.

6 CONCLUSION

This paper explored a novel approach to food recommendation for weight management. The model takes into account the food nutrition information (calories) and user's BMI information. To recommend foods for a user to achieve a weight goal, the approach follows the good examples in real data that produced the intended outcome. This approach is different from the traditional food recommendation systems which either follow the user's habits or impose dietary constraints. By following good examples, our recommendation model can guide the user smoothly toward the intended goal

Table 3: Performance comparison of different models based on NDCG@10. The best performance based on each sub-dataset is boldfaced. All improvements over the baseline (i.e., General model) are statistically significant with $p < 0.01$ in t-test.

Model	No_In	No_De	No_Ma	Ov_In	Ov_De	Ov_Ma	Ob_In	Ob_De	Ob_Ma	Average
General	0.2412	0.2470	0.2399	0.2251	0.2243	0.2294	0.2133	0.2257	0.2201	0.2308
Intent	0.2625	0.2494	0.2486	0.2402	0.2347	0.2305	0.2199	0.2326	0.2222	0.2386
BMI	0.255	0.2758	0.266	0.2403	0.2494	0.2495	0.2206	0.2503	0.2455	0.2532
BMI+Intent	0.2694	0.2736	0.2609	0.2449	0.2578	0.2464	0.2435	0.2433	0.2354	0.2535
Int+(BMI+Int)	0.2764	0.2778	0.2664	0.2526	0.2603	0.2506	0.2515	0.2488	0.2382	0.2584
Improvement	14.6%	12.5%	11.0%	12.2%	16.0%	9.2%	17.9%	10.2%	8.2%	12.0%
BMI+(BMI+Int)	0.2779	0.2843	0.2730	0.2517	0.2601	0.2518	0.2508	0.2574	0.2476	0.2626
Improvement	15.2%	15.1%	13.8%	11.8%	16.0%	9.8%	17.6%	14.0%	12.5%	13.8%
Gen+BMI+(BMI+Int)	**0.2793**	**0.2851**	**0.2731**	**0.2541**	**0.2609**	**0.2546**	**0.2529**	**0.2574**	**0.2476**	**0.2635**
Improvement	15.8%	15.4%	13.8%	12.9%	16.3%	11.0%	18.6%	14.0%	17.0%	14.2%

Table 4: Performance comparison of different models based on R-Precision. The best performance on each sub-dataset is boldfaced. All improvements over the baseline (i.e., General) are statistically significant with $p < 0.01$ in t-test.

Model	No_In	No_De	No_Ma	Ov_In	Ov_De	Ov_Ma	Ob_In	Ob_De	Ob_Ma	Average
General	0.2399	0.2412	0.2419	0.2298	0.2310	0.2327	0.2141	0.2314	0.2208	0.2328
Intent	0.2624	0.2450	0.2423	0.2434	0.2444	0.2351	0.2274	0.2395	0.2296	0.2412
BMI	0.2586	0.2738	0.2691	0.2472	0.2593	0.2620	0.2278	0.2570	0.2591	0.2602
BMI+Intent	0.2636	0.2691	0.2625	0.2505	0.2680	0.2541	0.2484	0.2566	0.2508	0.2592
Int+(BMI+Int)	0.2733	0.2752	0.2648	**0.2592**	0.2694	0.2597	0.2530	0.2585	0.2522	0.2635
Improvement	13.9%	14.1%	9.5%	12.8%	16.6%	11.6%	18.2%	11.7%	14.2%	13.2%
BMI+(BMI+Int)	0.2775	**0.2809**	0.2740	0.2578	0.2695	0.2643	**0.2542**	0.2643	**0.2639**	0.2687
Improvement	15.7%	16.5%	13.3%	12.2%	16.7%	13.6%	18.7%	14.2%	19.5%	15.4%
Gen+BMI+(BMI+Int)	**0.2787**	**0.2809**	**0.2742**	0.2588	**0.2698**	**0.2650**	**0.2542**	**0.2646**	**0.2639**	**0.2691**
Improvement	16.2%	16.5%	13.4%	12.6%	16.8%	13.9%	18.7%	14.3%	19.5%	15.6%

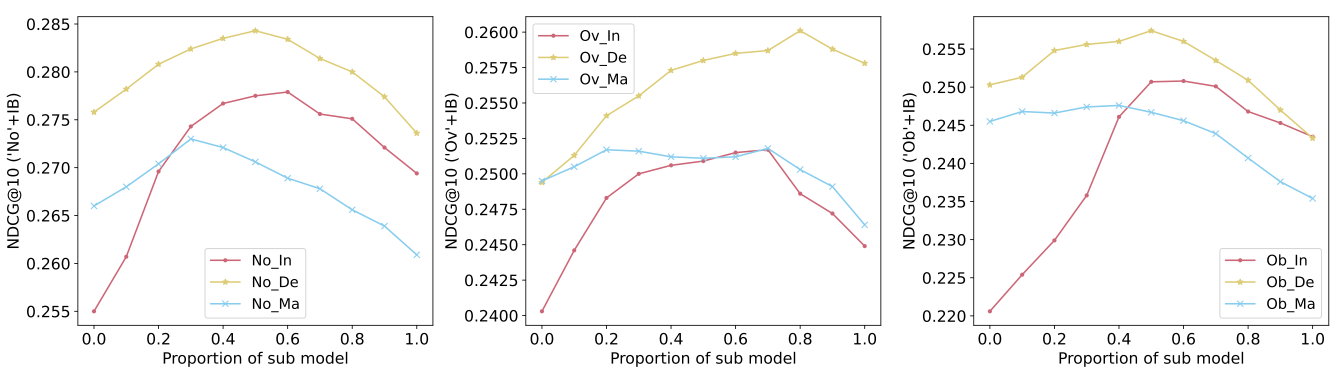

Figure 4: Combining BMI+Intent and BMI models with different parameters: effectiveness on different BMI groups. The left endpoint represents BMI-specific models while the right endpoint the BMI+Int models.

while trying to keep user's preferences at the same time. The principle of "following good examples" could make the recommendations less disruptive from user's habits, thereby increasing the chance of their adoption by users.

Notice that the proposed approach is complementary to the existing ones. It would be possible to combine our approach with that using dietary constraints for a more strict control on diet when necessary.

In this first exploration, we have not considered all the nutrition- and health-related factors such as food energy density, categories of food, etc. These factors can be explored in the future.

REFERENCES

[1] J. Aberg. 2006. Dealing With malnutrition: A meal planning system for elderly. In *AAAI Spring Symposium: Argumentation for Consumers of Healthcare.*

[2] Damoon Ashtary-Larky, Reza Bagheri, Amir Abbasnezhad, Grant M Tinsley, Meysam Alipour, and Alexei Wong. 2020. Effects of gradual weight loss v. rapid weight loss on body composition and RMR: a systematic review and meta-analysis. *British journal of nutrition* 124, 11 (2020), 1121–1132.

[3] Fatemeh Azizi-Soleiman and Leila Azadbakht. 2014. Weight loss maintenance: A review on dietary related strategies. *Journal of research in medical sciences : the official journal of Isfahan University of Medical Sciences* 19 (03 2014), 268–275.

[4] Ting Bai, Jian-Yun Nie, Wayne Xin Zhao, Yutao Zhu, Pan Du, and Ji-Rong Wen. 2018. An Attribute-aware Neural Attentive Model for Next Basket Recommendation. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, Ann Arbor, MI, USA, July 08-12, 2018*, Kevyn Collins-Thompson, Qiaozhu Mei, Brian D. Davison, Yiqun Liu, and Emine Yilmaz (Eds.). ACM, 1201–1204. https://doi.org/10.1145/3209978.3210129

[5] Ruth Brown, Karissa Canning, Michael Fung, Dishay Jiandani, Michael Riddell, Alison Macpherson, and Jennifer Kuk. 2015. Calorie Estimation in Adults Differing in Body Weight Class and Weight Loss Status. *Medicine and science in sports and exercise* 48 (10 2015). https://doi.org/10.1249/MSS.0000000000000796

[6] Jingjing Chen, Chong-Wah Ngo, and Tat-Seng Chua. 2017. Cross-modal Recipe Retrieval with Rich Food Attributes. In *Proceedings of the 2017 ACM on Multimedia Conference, MM 2017, Mountain View, CA, USA, October 23-27, 2017*, Qiong Liu, Rainer Lienhart, Haohong Wang, Sheng-Wei "Kuan-Ta" Chen, Susanne Boll, Yi-Ping Phoebe Chen, Gerald Friedland, Jia Li, and Shuicheng Yan (Eds.). ACM, 1771–1779. https://doi.org/10.1145/3123266.3123428

[7] Zafra Cooper and Christopher G Fairburn. 2001. A new cognitive behavioural approach to the treatment of obesity. *Behaviour research and therapy* 39, 5 (2001), 499–511.

[8] Cem Ekmekcioglu and Yvan Touitou. 2011. Chronobiological aspects of food intake and metabolism and their relevance on energy balance and weight regulation. *Obesity reviews : an official journal of the International Association for the Study of Obesity* 12 (01 2011), 14–25. https://doi.org/10.1111/j.1467-789X.2010.00716.x

[9] M. Elahi, M. Ge, F. Ricci, I. Fernandez-Tobıas, and D. Berkovsky, S. andMassimo. 2015. Interaction design in a mobile food recommender system. In *IntRS@recsys, CEUR-WS.org, CEUR workshop proceedings*, Vol. 1438. 49–52.

[10] David Elsweiler and Morgan Harvey. 2015. Towards Automatic Meal Plan Recommendations for Balanced Nutrition. In *Proceedings of the 9th ACM Conference on Recommender Systems, RecSys 2015, Vienna, Austria, September 16-20, 2015*, Hannes Werthner, Markus Zanker, Jennifer Golbeck, and Giovanni Semeraro (Eds.). ACM, 313–316. https://dl.acm.org/citation.cfm?id=2799665

[11] Mouzhi Ge, Francesco Ricci, and David Massimo. 2015. Health-aware Food Recommender System. In *Proceedings of the 9th ACM Conference on Recommender Systems, RecSys 2015, Vienna, Austria, September 16-20, 2015*, Hannes Werthner, Markus Zanker, Jennifer Golbeck, and Giovanni Semeraro (Eds.). ACM, 333–334. https://dl.acm.org/citation.cfm?id=2796554

[12] Gijs Geleijnse, Peggy Nachtigall, Pim van Kaam, and Luciënne Wijgergangs. 2011. A personalized recipe advice system to promote healthful choices. In *Proceedings of the 16th International Conference on Intelligent User Interfaces, IUI 2011, Palo Alto, CA, USA, February 13-16, 2011*, Pearl Pu, Michael J. Pazzani, Elisabeth André, and Doug Riecken (Eds.). ACM, 437–438. https://doi.org/10.1145/1943403.1943487

[13] Drishti P Ghelani, Lisa J Moran, Cameron Johnson, Aya Mousa, and Negar Naderpoor. 2020. Mobile apps for weight management: A review of the latest evidence to inform practice. *Frontiers in Endocrinology* 11 (2020), 412.

[14] Kelly L. Haws, Peggy J. Liu, Joseph P. Redden, and Heidi J. Silver. 2017. Exploring the relationship between varieties of variety and weight loss: When more variety can help people lose weight. *Journal of Marketing Research* 54, 4 (Aug. 2017), 619–635. https://doi.org/10.1509/jmr.15.0206 Publisher Copyright: © 2017, American Marketing Association. Copyright: Copyright 2018 Elsevier B.V., All rights reserved.

[15] Ruining He, Wang-Cheng Kang, and Julian J. McAuley. 2017. Translation-based Recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems, RecSys 2017, Como, Italy, August 27-31, 2017*, Paolo Cremonesi, Francesco Ricci, Shlomo Berkovsky, and Alexander Tuzhilin (Eds.). ACM, 161–169. https://doi.org/10.1145/3109859.3109882

[16] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Comput.* 9, 8 (1997), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

[17] Haoji Hu, Xiangnan He, Jinyang Gao, and Zhi-Li Zhang. 2020. Modeling Personalized Item Frequency Information for Next-basket Recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, Jimmy Huang, Yi Chang, Xueqi Cheng, Jaap Kamps, Vanessa Murdock, Ji-Rong Wen, and Yiqun Liu (Eds.). ACM, 1071–1080. https://doi.org/10.1145/3397271.3401066

[18] Rosemary Johnson. 1990. Restructuring: An emerging theory on the process of losing weight. *Journal of advanced nursing* 15, 11 (1990), 1289–1296.

[19] Anna-Maria Keränen, Markku Savolainen, Annakaisa Reponen, Mona-Lisa Kujari, Sari Lindeman, Risto Bloigu, and Jaana Laitinen. 2009. The effect of eating behavior on weight loss and maintenance during a lifestyle intervention. *Preventive medicine* 49 (05 2009), 32–8. https://doi.org/10.1016/j.ypmed.2009.04.011

[20] Yehuda Koren and Robert M. Bell. 2011. Advances in Collaborative Filtering. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor (Eds.). Springer, 145–186. https://doi.org/10.1007/978-0-387-85820-3_5

[21] Jessica RL Lieffers, Jose F Arocha, Kelly Grindrod, and Rhona M Hanning. 2018. Experiences and perceptions of adults accessing publicly available nutrition behavior-change mobile apps for weight management. *Journal of the Academy of Nutrition and Dietetics* 118, 2 (2018), 229–239.

[22] Qiao Liu, Refuoe Mokhosi, Yifu Zeng, and Haibin Zhang. 2018. STAMP: Short-term attention/memory priority model for session-based recommendation. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1831–1839. https://doi.org/10.1145/3219819.3219950

[23] Yujie Lu, Shengyu Zhang, Yingxuan Huang, Luyao Wang, Xinyao Yu, Zhou Zhao, and Fei Wu. 2021. Future-Aware Diverse Trends Framework for Recommendation. In *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*, Jure Leskovec, Marko Grobelnik, Marc Najork, Jie Tang, and Leila Zia (Eds.). ACM / IW3C2, 2992–3001. https://doi.org/10.1145/3442381.3449791

[24] Charoula K Nikolaou and Michael EJ Lean. 2017. Mobile applications for obesity and weight management: current market characteristics. *International Journal of Obesity* 41, 1 (2017), 200–202.

[25] Marcia Otto, Cheryl Anderson, Jennifer Dearborn, Erin Ferranti, Dariush Mozaffarian, Goutham Rao, Judith Wylie-Rosett, and Alice Lichtenstein. 2018. Dietary Diversity: Implications for Obesity Prevention in Adult Populations: A Science Advisory From the American Heart Association. *Circulation* 138 (09 2018). https://doi.org/10.1161/CIR.0000000000000595

[26] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, Alessandro Moschitti, Bo Pang, and Walter Daelemans (Eds.). ACL, 1532–1543. https://doi.org/10.3115/v1/d14-1162

[27] Michael G Perri, Arthur M Nezu, Eugene T Patti, and Karen L McCann. 1989. Effect of length of treatment on weight loss. *Journal of consulting and clinical psychology* 57, 3 (1989), 450.

[28] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing personalized Markov chains for next-basket recommendation. In *Proceedings of the 19th International Conference on World Wide Web, WWW 2010, Raleigh, North Carolina, USA, April 26-30, 2010*, Michael Rappa, Paul Jones, Juliana Freire, and Soumen Chakrabarti (Eds.). ACM, 811–820. https://doi.org/10.1145/1772690.1772773

[29] A. Robertson. 2004. Food and health in Europe: a new basis for action. *Academic Search Complete, WHO Regional Office for Europe* (2004).

[30] Barbara J Rolls, Adam Drewnowski, and Jenny H Ledikwe. 2005. Changing the energy density of the diet as a strategy for weight management. *Journal of the American Dietetic Association* 105, 5 (2005), 98–103.

[31] Kathryn Ross, Peihua Qiu, Lu You, and Rena Wing. 2019. Week-to-week predictors of weight loss and regain. *Health Psychology* 38 (09 2019). https://doi.org/10.1037/hea0000798

[32] Badrul Munir Sarwar, George Karypis, Joseph A. Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the Tenth International World Wide Web Conference, WWW 10, Hong Kong, China, May 1-5, 2001*, Vincent Y. Shen, Nobuo Saito, Michael R. Lyu, and Mary Ellen Zurko (Eds.). ACM, 285–295. https://doi.org/10.1145/371920.372071

[33] Julie T Schaefer and Amy B Magnuson. 2014. A review of interventions that promote eating by internal cues. *Journal of the Academy of Nutrition and Dietetics* 114, 5 (2014), 734–760.

[34] Sabrina Schlesinger, Manuela Neuenschwander, Carolina Schwedhelm, Georg Hoffmann, Angela Bechthold, Heiner Boeing, and Lukas Schwingshackl. 2019. Food groups and risk of overweight, obesity, and weight gain: a systematic review and dose-response meta-analysis of prospective studies. *Advances in Nutrition* 10, 2 (2019), 205–218.

[35] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3-7, 2019*, Wenwu Zhu, Dacheng Tao, Xueqi Cheng, Peng Cui, Elke A. Rundensteiner, David Carmel, Qi He, and Jeffrey Xu Yu (Eds.). ACM, 1441–1450. https://doi.org/10.1145/3357384.3357895

[36] Yong Kiam Tan, Xinxing Xu, and Yong Liu. 2016. Improved Recurrent Neural Networks for Session-based Recommendations. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems, DLRS@RecSys 2016, Boston, MA, USA, September 15, 2016*, Alexandros Karatzoglou, Balázs Hidasi, Domonkos Tikk, Oren Sar Shalom, Haggai Roitman, Bracha Shapira, and Lior Rokach (Eds.). ACM, 17–22. https://doi.org/10.1145/2988450.2988452

[37] Jiaxi Tang and Ke Wang. 2018. Personalized top-N sequential recommendation via convolutional sequence embedding. *WSDM 2018 - Proceedings of the 11th ACM International Conference on Web Search and Data Mining* 2018-Febua (2018), 565–573. https://doi.org/10.1145/3159652.3159656

[38] ChunYuen Teng, Yu-Ru Lin, and Lada A. Adamic. 2012. Recipe recommendation using ingredient networks. In *Web Science 2012, WebSci '12, Evanston, IL, USA - June 22 - 24, 2012*, Noshir S. Contractor, Brian Uzzi, Michael W. Macy, and Wolfgang Nejdl (Eds.). ACM, 298–307. https://doi.org/10.1145/2380718.2380757

[39] T.N. Trang Tran, M. Atas, and A. et al. Felfernig. 2018. An overview of recommender systems in the healthy food domain. *J Intell Inf Syst* 50 (2018), 501—-526.

[40] Mayumi Ueda, Syungo Asanuma, Yusuke Miyawaki, and Shinsuke Nakajima. 2014. Recipe recommendation method by considering the users preference and ingredient quantity of target recipe. In *Proceedings of the international multiconference of engineers and computer scientists*, Vol. 1. 12–14.

[41] Erica van Herpen and Rik Pieters. 2002. The Variety of An Assortment : An Extension to the Attribute-Based Approach. *Marketing Science* 21 (06 2002), 331–341.

[42] Y. Van Pinxteren, G. Geleijnse, and P. Kamsteeg. 2011. Deriving a recipe similarity measure for recommending healthful meals. In *Proceedings of the 16th international conference on intelligent user interfaces*. 105––114.

[43] Corneel Vandelanotte, Andre M Müller, Camille E Short, Melanie Hingle, Nicole Nathan, Susan L Williams, Michael L Lopez, Sanjoti Parekh, and Carol A Maher. 2016. Past, present, and future of eHealth and mHealth research to improve physical activity and dietary behaviors. *Journal of nutrition education and behavior* 48, 3 (2016), 219–228.

[44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.). 5998–6008. https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html

[45] Neeltje Vogels, Kristel Diepvens, and Margriet Westerterp-Plantenga. 2006. Predictors of Long-term Weight Maintenance. *Obesity research* 13 (01 2006), 2162–8. https://doi.org/10.1038/oby.2005.268

[46] Pengfei Wang, Jiafeng Guo, Yanyan Lan, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2015. Learning Hierarchical Representation Model for Next Basket Recommendation. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, Santiago, Chile, August 9-13, 2015*, Ricardo Baeza-Yates, Mounia Lalmas, Alistair Moffat, and Berthier A. Ribeiro-Neto (Eds.). ACM, 403–412. https://doi.org/10.1145/2766462.2767694

[47] Edward Weiss, Deborah Galuska, Laura Kettel Khan, and Mary Serdula. 2006. Weight-Control Practices among U.S. Adults. *American journal of preventive medicine* 31 (08 2006), 18–24. https://doi.org/10.1016/j.amepre.2006.03.016

[48] Haochao Ying, Fuzhen Zhuang, Fuzheng Zhang, Yanchi Liu, Guandong Xu, Xing Xie, Hui Xiong, and Jian Wu. 2018. Sequential Recommender System based on Hierarchical Attention Networks. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization, 3926–3932. https://doi.org/10.24963/ijcai.2018/546

[49] Feng Yu, Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. 2016. A Dynamic Recurrent Model for Next Basket Recommendation. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval, SIGIR 2016, Pisa, Italy, July 17-21, 2016*, Raffaele Perego, Fabrizio Sebastiani, Javed A. Aslam, Ian Ruthven, and Justin Zobel (Eds.). ACM, 729–732. https://doi.org/10.1145/2911451.2914683