

Food Recommendation: Framework, Existing Solutions, and Challenges

Weiying Min , *Member, IEEE*, Shuqiang Jiang , *Senior Member, IEEE*, and Ramesh Jain, *Fellow, IEEE*

Abstract—A growing proportion of the global population is becoming overweight or obese, leading to various diseases (e.g., diabetes, ischemic heart disease and even cancer) due to unhealthy eating patterns, such as increased intake of food with high energy and high fat. Food recommendation is of paramount importance to alleviate this problem. Unfortunately, modern multimedia research has enhanced the performance and experience of multimedia recommendation in many fields such as movies and POI, yet largely lags in the food domain. This article proposes a unified framework for food recommendation, and identifies main issues affecting food recommendation including incorporating various context and domain knowledge, building the personal model, and analyzing unique food characteristics. We then review existing solutions for these issues, and finally elaborate research challenges and future directions in this field. To our knowledge, this is the first survey that targets the study of food recommendation in the multimedia field and offers a collection of research studies and technologies to benefit researchers in this field.

Index Terms—Artificial intelligence, knowledge based systems, image recognition, data mining, health information management.

I. INTRODUCTION

FOOD is always central to the human life. Besides the air we breathe, food is the only physical matter, which humans take into the body. In the early days, humans faced the task of identifying and gathering food for their survival. At present, the dietary choice is becoming vital in satisfying diverse needs, such as basic nutrition, calorie, taste, health and social occasions. According to the International Diabetes Federation, about 415 million people worldwide suffer from diabetes, and the rate

of diabetes incidence is projected to further increase by more than 50% by 2040, becoming one great threat to global health. The dietary factor is one main cause of the dramatic increase in the incidence of obesity and diabetes [1], [2]. The Global Burden of Disease Study also indicates dietary factors as a major contributor to levels of malnutrition, obesity and overweight, and unreasonable diets lead to 11 million avoidable premature deaths per year [3]. Food computing [4] is emerging as a new field to ameliorate these issues. As an important task in food computing, food recommendation intends to find suitable food items for users to meet their personalized needs, and thus plays a critical role in human dietary choice.

A balanced diet is crucial to maintain one's physical health. However, nutrients that need to be ingested vary greatly depending on personal food preference and health conditions. Therefore, how to provide personalized food recommendation according to different personal requirements is very important. The past decade has witnessed the rapid growth of internet services and mobile devices. It has been more convenient for people to access huge amounts of online multimedia food content from various sources, such as forums, social media, recipe-sharing websites and customer review sites. Although this growth allows users to have more choices, it also brings problems for users to select preferred food items from thousands of candidates. Therefore, food recommendation is becoming increasingly essential for serving potentially huge service demand and can help users easily discover a small subset of food items which are enjoyable and suitable for them.

Compared with recommendation in other fields, food recommendation has its own characteristics. For example, food preference learning is an important step towards food recommendation. However, food preference involves various factors, such as taste preference, perceptive difference, cognitive restraint, cultural familiarity and even genetic influence.¹ Therefore, it makes accurate food preference learning more difficult. Furthermore, food recommendation should consider more context information. Besides basic context information captured from familiar mobile devices [5], such as time, location and environmental information (e.g., temperature and PM2.5), various body state-related signals, such as steps taken, heart rate, sleep quality, body acceleration and even affective states can also be captured from new sensing devices, such as watches, wearable fitness trackers and bracelets [6]. These signals can describe users' actual

Manuscript received May 15, 2019; revised August 25, 2019 and November 17, 2019; accepted November 18, 2019. Date of publication December 9, 2019; date of current version September 23, 2020. This work was supported in part by the National Natural Science Foundation of China under Grants 61532018 and 61972378, in part by the Beijing Natural Science Foundation under Grant L182054, and in part by the National Program for Special Support of Eminent Professionals and National Program for Support of Top-Notch Young Professionals, and in part by the State Key Laboratory of Robotics. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Shaoen Wu. (*Corresponding author: Shuqiang Jiang.*)

W. Min is with the Key Laboratory of Intelligent Information Processing, Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100190, China, and also with State Key Laboratory of Robotics, Shenyang Institute of Automation, Chinese Academy of Sciences, Shenyang 110016, China (e-mail: minweiying@ict.ac.cn).

S. Jiang is with the Key Laboratory of Intelligent Information Processing, Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100190, China, and also with University of Chinese Academy of Sciences, Beijing 100049, China (e-mail: sqjiang@ict.ac.cn).

R. Jain is with the Department of Computer Science, University of California, Irvine, CA 92697 USA (e-mail: jain@ics.uci.edu).

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TMM.2019.2958761

¹[Online]. Available: https://www.eurekalert.org/pub_releases/2017-04/eb2-cgi041217.php

body conditions comprehensively in real-time and is of crucial importance for food recommendation. However, they are with different statistical properties, either discrete or continue, and make effective multi-sensor context fusion challenging. As a result, simply borrowing methods from recommendation methods in other fields without considering characteristics of food recommendation probably leads to inferior performance. This negligence will eventually result in a bottleneck in advancing intelligent food recommendation for use in real-world applications.

Existing multimedia research has made great progress in improving the recommendation performance and experience in many fields such as movies and POI, yet largely lags in the food domain. To the best of our knowledge, although relevant works on food recommendation [7]–[9] have received more attention in the multimedia community, there are very few systematic reviews, which provide a unified framework and comprehensive summary of current efforts in food recommendation. Because of huge potentials in food recommendation, the time has come for the multimedia field to give a survey on food recommendation, which can help researchers from relevant communities better understand the strength and weakness of existing methods. In this paper, we give a survey on food recommendation. Particularly, the objective of this article is as follows: to propose a unified framework for food recommendation and identify main issues affecting food recommendation (Section II); to review existing progress for these issues (Section III), and to outline research challenges and future directions in this field (Section IV and Section V). Finally, we conclude the survey in Section VI.

Note that recently there is one survey work on food recommendation [10]. Our survey and [10] both summarize existing methods, challenges and future directions, and are both very complementary to benefit researchers and practitioners working in this field. However, there are three important differences between this work and [10]: (1) To our knowledge, we first systematically define the problem of food recommendation and summarize its three unique aspects. (2) We propose a unified food recommendation framework, and put previous studies together under this unified umbrella. Consequently, our taxonomy on existing food recommendation works is different from [10]. (3) We put more emphasis on multimedia-oriented food recommendation, where how to utilize and fuse multi-modal signals (e.g., images, voice and text) and rich context for food recommendation is highlighted. In contrast, [10] mainly focuses on food recommendation based on recommendation techniques.

II. PROPOSED FRAMEWORK

Definition: Food recommendation aims to provide a list of ranked food items for users to meet their personalized needs. Here, food is a more broad concept, and it includes all food-related items, such as meal, recipes, coffee shops and restaurants. Food recommendation is typically multidisciplinary research, including nutrition, food science, psychology, biology, anthropology, sociology, other branches of natural and social sciences.²

Compared with other types of recommendation, there are mainly three aspects unique for food recommendation. (1) Food recommendation involves different context and domain knowledge. Rich user context (e.g., the heart rate and steps taken) and external environmental context (e.g., physical activity-relevant and health-relevant context) captured from different sensors describe users' actual physical conditions and their surroundings, and thus provide valuable information for exact match between user requirement and food items. For example, food recommendation probably recommends one user food items with much water and protein after exercise captured from sensors. Furthermore, food recommendation is very relevant to health. Therefore, medical knowledge, dietary knowledge and other relevant domain knowledge should also be incorporated into the food recommender system for constraint optimization and computing. (2) From the user aspect, the most significant difference is that food recommendation is very relevant to user's health. Therefore, an ideal food recommendation system should self-adaptively build a trade-off between personalized food preference/interest and personalized nutrition/health requirement. For example, even one diabetic likes sweet food, it is more reasonable for food recommendation to recommend him/her the food with less sugar than before. Besides user's health needs, food recommendation should also consider other complex and various fine-grained user needs, such as allergies and life-style preferences (e.g., the desire to eat only vegan or vegetarian food). Therefore, a personal model should be built to take these factors into consideration. (3) Food has its own characteristics, and many unique factors, such as cooking methods, ingredient combination effects, preparation time, nutritional breakdown, and non-rigid visual appearance should be analyzed to obtain high-level semantic concepts and attributes for food recommendation. For example, nutritional and calorie intake can be calculated by analyzing captured food images to create a quantitative nutrition diary. Therefore, multimodal heterogeneous food analysis is necessary to accurately obtain high-level understanding of the food type and other levels of semantics.

Taking all these factors into consideration, we propose a unified framework (as shown in Fig. 1), which can jointly utilize rich context and knowledge, user information, and heterogeneous food information for food recommendation. In order to combine these content information with user-food item interaction, a hybrid recommendation strategy is utilized [11] in the food recommendation framework, where joint collaborative learning is adopted to learn both user latent vectors and item latent vectors after embedding and integration on context-dimension, user-dimension and food item-dimension, respectively. For each dimension, the signals from different sources are first changed into the feature representation after embedding, and are then fused into final representation via integration, where different embedding and integration methods can be adopted. Final recommendation is generated from these learned user vectors and item vectors under context constraints. Particularly, food recommendation mainly consists of three components, namely context and knowledge incorporation, personal modal construction and heterogeneous food analysis. Note that the context and knowledge can not only put more constraints on food recommendation

²[Online]. Available: <https://www.wur.nl/en/show/Food-preferences.htm>

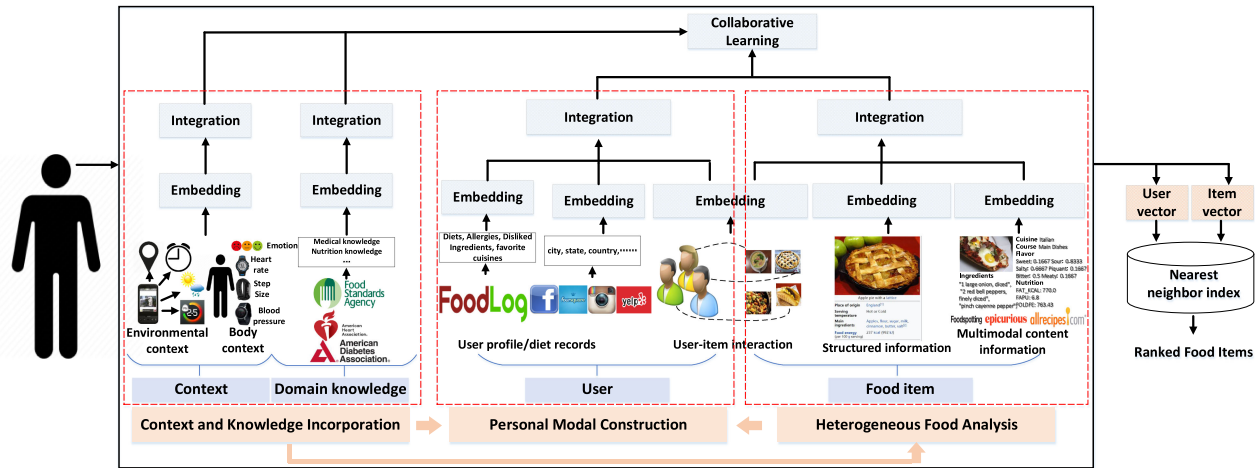


Fig. 1. The general framework of food recommendation, where collaborative learning is adopted to learn both user latent vectors and item latent vectors after embedding and integration on context-dimension, user-dimension and food item-dimension, respectively. For each dimension, the signals from different sources are first changed into the feature representation via embedding, and are then fused into final representation via integration.

independently, but also be used to support the personal model construction and heterogenous food analysis. In addition, heterogenous food analysis is also helpful for building the complete personal model. Next, we will identify detailed requirements of three components, respectively.

A. Context and Knowledge Incorporation

Basic context information (e.g., time and location) can help filter out irrelevant items for recommendation. Compared with other types of recommendation, food recommendation involves more complex, diverse and even dynamic factors. Rich user context and external environmental context information describe users' actual physical conditions and their surroundings, and thus provide valuable information for exact match between user requirement and food items of interest.

Over the last decade, a great variety of wearable electronic devices and ambient sensors have been developed [12]. They can monitor personal body conditions and environmental changes everywhere in real time by connecting users to surrounding machines. The new sensing devices, ranging from social sensors to biosensors [6] can measure physiological parameters to produce various signals, such as step sizes, heart rate, sleep quality, body acceleration, blood pressure and affective states (Fig. 2). In addition, various health and fitness mobile apps such as MyFitnessPal, Endomondo and Fitbit can also help people keep track of what they eat, when they exercise, and how well they sleep. Exploiting these context information will enable more reasonable and accurate food recommendation. For example, if the level of sugar captured from the sensor is high for one user, it is necessary to recommend this user the food with little sugar or low conversion sugar. Food recommender system probably recommends one user the food with much water and protein after exercise, captured from sensors. Besides rich context, health-relevant expert knowledge (e.g., medical knowledge and nutrition knowledge) should also be considered. Food recommendation via expert knowledge is a potential key to

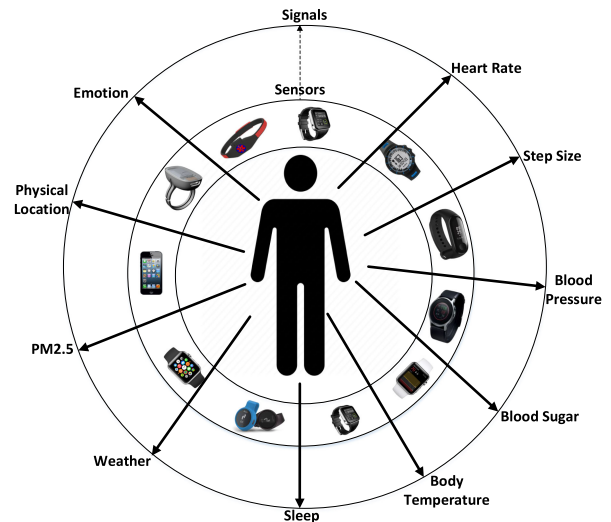


Fig. 2. Various signals from sensor devices. The new sensing devices can measure physiological parameters to produce various signals, such as step sizes, heart rate, sleep quality, body acceleration, blood pressure and even emotion.

unlock healthy diets, and thus gives more precise recommendation results.

There are mainly two issues to solve when applying rich context and knowledge for food recommendation: (1) Current wearable sensing technologies face many challenges, such as inaccuracy and uncertainty of measurement, multi-functional integration and adverse impacts on the environment. Therefore, there are still many body-related signals, which are harder to capture accurately or even not available. More new wearable sensing technologies should be developed to obtain comprehensive human signals relevant to food recommendation. For example, researchers recently designed new sensors to detect muscle motions involved in every chew, and finally can accurately create a time-stamped visual record of the food consumed [13]. (2) Multiple-sensor context predominantly is segregated with the

taxonomy of categorical/discrete and ordinal/continuous features in time and space. For example, some signals such as step count is discrete while others such as electrocardiogram are continuous. In addition, the categorical features probably vary widely in their cardinality. Some are binary while others have many possible values. It is difficult to model user states to accommodate a mixture of discrete and continuous variables in a joint model, not to mention external knowledge. Inappropriate representation and fusion methods for rich context and knowledge are even not as effective as methods without any context. In addition, the variety and number of contexts makes it hard to measure all available data, which in turn introduces new uncertainty levels for food recommendation. Therefore, an effective multi-sensor fusion method should be developed to overcome these problems.

B. Personal Model Construction

Personal model construction involves collecting and fusing relevant information of users to generate a user model for food recommendation. A food recommendation agent cannot function accurately until the personal model has been well constructed. The system needs to know the information from the user as much as possible in order to provide reasonable and precise recommendation.

A user profile is a collection of personal information associated with a specific user, and can be considered as a simple personal model. In contrast, food recommendation considers more complex, dynamic and various constrained factors, such as attitudes and beliefs about food and recipes, personal food preference, lifestyles, hobbies, and even cultural and social aspects. In order to effectively construct the personal model for food recommendation, we should capture and fuse various signals from above-mentioned sensors, websites and social media. People have started recording health-related information using sensors ranging from simple wearable accelerometers that could be classified and recorded in simple activities (such as walking, jogging and climbing stairs) to other measurements (such as body temperature, heart rate, perspiration rate, galvanic skin resistivity, and many other deeper parameters). We can use this to model user's real-time states continuously. In addition, we are leaving digital traces of all kinds of activities online in the social media. As food is central in our life, a significant fraction of online content is about food. A huge amount of information about eating habits is now being recorded digitally and available. Besides shared multimedia food content, we can obtain both basic user information (e.g., the age, gender and residence) and more detailed information (e.g., diets and allergies). For example, we can obtain fine-grained information from Yummly, such as diets (e.g., vegetarian or low fodmap), allergies (e.g., seafood or tree nut), disliked ingredients (e.g., sugar or beef), favorite cuisines (e.g., Chinese and Indian cuisine), taste (e.g., sour, salty, sweet, bitter and meaty) and food restrictions (e.g., vegetarian, vegan, kosher and halal). All these user information can also be taken into consideration for personal model construction via heterogeneous media processing.

Among these different data sources, Foodlog is an important source for personal model construction. Foodlog has recorded users' food intake by taking photos of their meals or writing the diary for food assessment and journaling. Traditional methods of keeping a food journal resort to manually recording meals in as much detail as possible by including the portion size, number of servings and calories, time, location, or even the people around us. Currently, multimodal food journal is developed [14]. The rich and detailed recording information provides us with fine-grained and accurate user preference on food, and finally leading to accurate personal model construction. However, to the best of our knowledge, few works exploits rich structured information from Foodlog to learn personalized food preference for personal model construction.

Besides complex food preference, personal nutritional and health factors should also be emphasized in building the personal model. For example, nutritional fitness offers an objective way to prioritize recommended foods for each physical and dietary condition [15]. That is, food recommendation should consider not only what he likes, but also what is nutritionally appropriate. When we incorporate nutrition information into food recommendation, it should be also personalized. Each user's nutrition is complex and uncertain on many levels [16]. This contrasts with most public nutrition and health advice, which is generic, non-specific healthy eating advice e.g., eat at least five portions of fruit and vegetables daily. Therefore, how to accurately build a personal model is an important issue, and should be deeply explored.

C. Heterogeneous Food Analysis

Food analysis is one basic component of food recommendation, as it provides prerequisites for obtaining a high-level understanding of the type (e.g., food category and ingredients), the volume of consumed food, nutrition and calorie intake by the user. Particularly, heterogeneous food analysis can be considered from the following three aspects.

- 1) Effective representation of single modality. Take food images as an example, with the prosperity of mobile services, more and more users tend to take photos they eat instead of text recording. Food images has become more prevalent. Visual analysis provides not only ingredient and calorie information, but also visual perceptual information. It has proved that visual information is helpful for food preference learning [17]. However, it is very difficult for fine-grained visual food representation because of their indistinctive spatial configuration and various types of geometrical variants [18]. Different from general object recognition, many types of food do not exhibit distinctive spatial configuration. They are typically non-rigid, and the structure information can not be easily exploited. Therefore, existing recognition approaches on general objects probably do not perform well for lack of effective visual food representation.
- 2) The multimodality of food data. The food actually involves multi-modalities. We see food, feel its texture, smell its

odors, taste its flavors and even hear its sounds when chewing. For example, when we bite into an apple, not only do we taste it, but also hear it crunch, see its red skin, and feel the coolness of its core [19]. Visual information of a food product such as color and texture can exert an influence on the acceptance of food products [20]. Furthermore, with the explosive growth of multimedia food content and service, food items consist of different modalities or structures, such as food images, cooking videos, ingredient lists, cooking instructions and various attributes. We should fuse different types of modalities for effective food representation. Despite these advances, multimodal food fusion still faces some challenges. For example, signals might not be aligned, such as continuous food image signals and discrete ingredient signals. Each modality might exhibit different types and different levels of noise at different points.

- 3) The heterogeneity of food data. Food data has also different structures, and thus contains richer semantic information. For example, in food-oriented Knowledge Graph (KG), entities usually include food items and relevant attributes (e.g., ingredients and taste), and the relationship between entities (e.g., belong to “Kung Pow Chicken belongs to Chinese cuisine,” and include “Chow mein includes noodles, soy sauce and vegetables ingredients”). KG can benefit the recommendation in many aspects, such as improving the precision and diversity of recommended items. Because the dimension of knowledge graph is higher and the semantic relationship is richer, the processing of KG is more complex and difficult.

III. EXISTING SOLUTIONS

Taking the above-mentioned issues into consideration, this section first surveys existing solutions from the following three aspects, corresponding to each component of the proposed food recommendation framework: 1) incorporating context and knowledge for food recommendation, 2) personal model construction for food recommendation, 3) heterogeneous food analysis for food recommendation. In addition, considering the importance of nutrition and health in food recommendation, we therefore surveys existing solutions on 4) nutrition and health oriented food recommendation.

A. Incorporating Context and Knowledge for Food Recommendation

Food recommendation is highly dependent on current context. Any such context would need to be detected either by direct user input or sensing environmental variables via various sensors, such as GPS and accelerometers to heart-rate monitoring and blood oxygen level. The location context has been widely used for food recommendation [8], [21], [22]. For example, Yanjie *et al.* [21] integrated the geographic proximity into one generative probabilistic model to capture the geographic influence for restaurant recommendation. Another important dimension is time, and it is very relevant to food recommendation because of food popularity [23]. Nag *et al.* [8] further fused more types

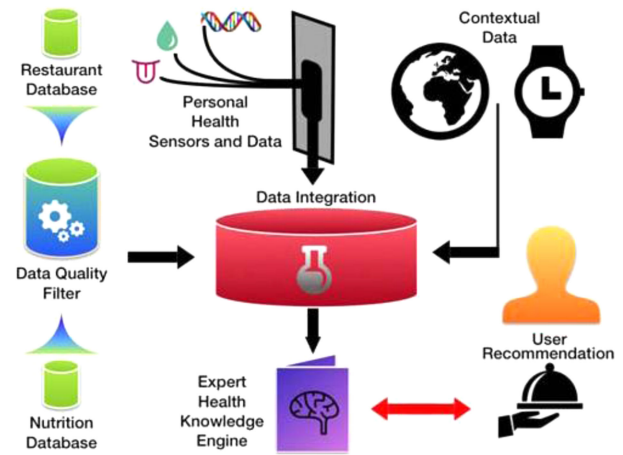


Fig. 3. Live context aware personalized dietitian system [8].

of context data via one pre-defined equation, such as barometer and pedometer output to estimate user’s daily nutritional requirements for local dish recommendation. The emerging sensor technologies such as social sensors and biosensors [6], [24] provided richer sensing context, especially health-relevant one. These context information put more constraints on recommendation to improve the accuracy of recommendation system.

Besides rich context from various sensors, efforts in modeling expert knowledge are made. For example, nutrition facts can be readily available for all major restaurant chains. For packaged items, algorithms that use this information are most promising for immediate consumer use and health impact. USDA Food Composition Databases,³ North American derived Nutrient Rich Foods Index 6.3 (NRF) [25] and British FSA [26] have been built based more heavily on available nutrition facts. However, they have not been established to capture expert knowledge of dietitians.

As representative work, Nag *et al.* [8] proposed a live personalized nutrition recommendation system. As shown in Fig. 3, a decision support system is made by fusing timely, contextually aware, personalized data to find local restaurant dishes to satisfy user’s needs. Multi-modal contextual data including GPS location, barometer, and pedometer output is fused through one pre-defined equation for calculating daily values. The calculated daily values are then used to rank the meals for local meal recommendation.

Discussion: In this section, we identify different types of context and their combination for food recommendation. The common context such as GPS and time information has been extensively used for food recommendation. One simple and effective method is to directly use these context to filter out irrelevant food items as the constraints, such as [8], [22]. Besides common spatio-temporal context, we can obtain richer context information, such as the physical state and health state captured from various sensors. In this case, a joint model is necessary to accommodate a mixture of both discrete and continuous context variables, such as [21]. However, the noise generated from these sensors makes effective context modeling more difficult. As for

³[Online]. Available: <https://ndb.nal.usda.gov/ndb/search/list>

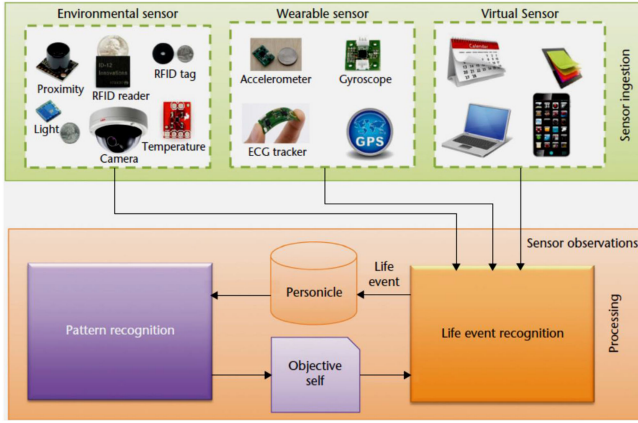


Fig. 4. High-level architecture of personal model construction system [27].

the domain knowledge, one reasonable method is to construct the food knowledge graph for food recommendation. However, there are few food knowledge graphs available. As a result, food knowledge has not been fully exploited by existing methods.

B. Personal Model Construction for Food Recommendation

The important step in making effective food recommendation is to build a personal model. Few works focus on personal model construction. The earlier work [27] proposed a high-level architecture of an objective self system for quantified self, namely personal model construction. As shown in Fig. 4, the proposed architecture consists of three main components data ingestion, life event recognition and pattern recognition. Data ingestion uses various sensors and preprocessing modules to extract appropriate attributes. Life event recognition predicted the most appropriate life activity from a set of predefined event classes based on pattern recognition. Later, Nag *et al.* [8] further utilized the objective self [27] for lifestyle recommendation.

In building the personal model, learning user's food preference is vital as it provides personalized information for effective food targeting and suggestions. Users' food preference learning is mainly conducted via various types of interactions, including survey based methods, web activity (e.g., ratings, browsing history and implicit feedbacks) based methods, online learning methods, food log based methods and dialogue based methods.

A typical on-boarding survey method should ask a number of multi-choice questions about general food preferences. Such methods are generally adopted by many commercial systems, such as Zipongo,⁴ Shopwell⁵ and PlateJoy.⁶ For example, a daily meal planner app PlateJoy elicits preferences for healthy goals and dietary restrictions with some questions, such as "Are there any ingredients you prefer to avoid? avocado, eggplant, eggs, seafood.....". Yummly also asks users to select allergies (e.g., seafood or tree nut), disliked ingredients (e.g., sugar or beef) and favorite cuisines (e.g., Chinese and Indian cuisine). Survey based preference elicitation methods are generally coarse-grained, and it cannot comprehensively capture user's preference.

⁴[Online]. Available: <https://meetzipongo.com/>

⁵[Online]. Available: <http://www.shopwell.com>

⁶[Online]. Available: <https://www.platejoy.com/>

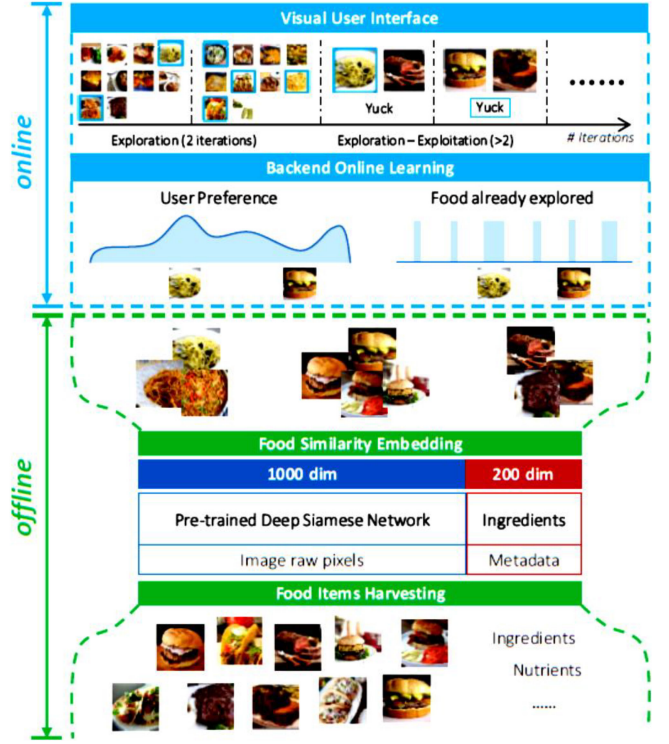


Fig. 5. PlateClick system pipeline [30].

Current popular methods resort to hybrids of historical records and item ratings for food preference learning [28], [29]. For example, Mouzhi *et al.* [29] extended matrix factorization by including additional parameters used for modelling the dependencies between assigned tags and ratings. For matrix factorization, let $\mathbf{p}_u \in R^k$ and $\mathbf{q}_m \in R^k$ denote the vector of user u and recipe m . In matrix factorization, ratings are estimated by computing the dot product of vectors $\hat{r}_{um} = \mathbf{p}_u^T \mathbf{q}_m$. After introducing the tags assigned to the user and recipe into the model, ratings are now estimated as follows:

$$\hat{r}_{um} = \left(\mathbf{p}_u + \frac{1}{|T_u|} \sum_{t \in T_u} \mathbf{x}_t \right)^T \left(\mathbf{q}_m + \frac{1}{|T_m|} \sum_{s \in T_m} \mathbf{y}_s \right) \quad (1)$$

where \mathbf{x}_t , \mathbf{y}_s denote the feature vector from user's and recipe's tags, respectively. T_u and T_m are the set of tags assigned by the user u and recipe m , respectively.

However, such methods suffer from the scarcity of user feedback. In contrast, Some works [17], [30], [31] resorts to using visual content to learn user's food preference. For example, Longqi *et al.* [17], [30] learned users' fine-grained food preference with only visual content. They proposed a food recommendation system PlateClick (Fig. 5), which consists of two stages, namely offline visual food similarity embedding and online food preference learning. In the first stage, a deep Siamese network is trained for visual food similarity embedding from pairwise food image comparisons. This network is to learn a low dimensional feature embedding that pulls similar food items together and pushes dissimilar food items far away. Contrastive loss is

selected as the loss function and can be expressed as:

$$L = \left\{ \sum_{(i,j)} [yL_o(\mathbf{x}_i, \mathbf{x}_j) + (1 - y)L_s(\mathbf{x}_i, \mathbf{x}_j)] \right\} \quad (2)$$

where $L_o(\mathbf{x}_i, \mathbf{x}_j)$ and $L_s(\mathbf{x}_i, \mathbf{x}_j)$ denote the contrastive constraint for ordered image pairs and the similar constraint for un-ordered image pairs, respectively. When incorporating deep feature learning into Eqn. (2), they are denoted as:

$$L_o(\mathbf{x}_i, \mathbf{x}_j) = \frac{1}{2}(\max(0, 1 - (f(\mathbf{x}_i) - f(\mathbf{x}_j)))^2$$

$$L_s(\mathbf{x}_i, \mathbf{x}_j) = \frac{1}{2}(f(\mathbf{x}_i) - f(\mathbf{x}_j))^2 \quad (3)$$

$f(\mathbf{x}_i)$ is one low-dimensional feature embedding for \mathbf{x}_i . $y \in \{0, 1\}$ indicates whether the input pair of food items $\mathbf{x}_i, \mathbf{x}_j$ are similar or not ($y = 0$ for similar, $y = 1$ for dissimilar).

In the second stage, a novel online learning framework is explored for learning users' preferences based on a small number of user-food image interactions. The core of online learning lies in the preference propagation in locally connected graphs. Because of the subjectivity and uncertainty of visual perception from users, such methods can not accurately learn food preference.

Recently, there are also some new methods for food preference learning. For example, Jie Zeng *et al.* [32] proposed a dialogue system to elicit user's food preference via human-robot interaction. To guarantee the conversation about food, the taste/texture expressions about dishes/ingredients are extracted from Twitter, and are registered as the default knowledge base of the dialogue system.

Food journal/diary/log records users' food intake and daily diet by taking photos of their meals or writing the diary in detail by including the portion size, number of servings and calories, time, location, or even the people around us [20]. This detailed description is effective, and thus we can learn users' food preference based on these historical records. For example, Nutrino,⁷ a personal meal recommender, asks users to log their daily food consumption and learn users' fine-grained food preferences. As is typical of systems relying on user-generated data, food journaling suffers from the cold-start problem, where recommendations cannot be made or are subject to low accuracy when the user has not yet generated a sufficient amount of data. This is because traditional methods of keeping a food journal are manually recorded and it is very easy to forget or procrastinate logging food entries. In order to solve this problem, some methods have been proposed via image recognition to recognize food items. For example, Foodlog [14] has contributed to a record of users' food intake simply by taking photos of their meals. Recently, multimodal Foodlog [33] is proposed to automatically recognize the starting moment of eating, and then prompt the user to begin a voice command food journaling method. With the popularity of mobile devices and advanced sensing technologies, learning user's food preference via exploring food logs will be one effective way to accurately learn fine-grained user's food

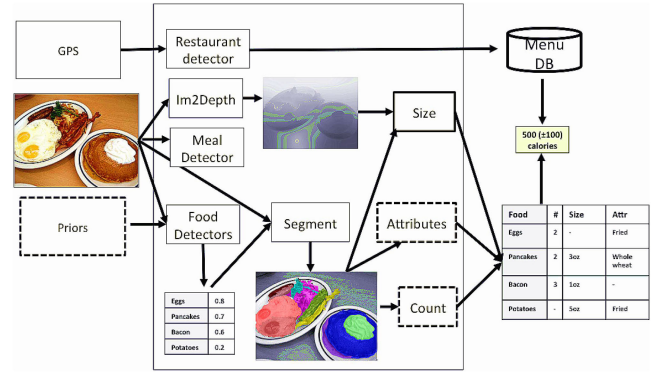


Fig. 6. Illustration of the overall system [34].

preference. Foodlog contains richer and more comprehensive dietary records, exploring Foodlog for accurate personal model construction will be a promising direction. However, because of heterogeneous multimodal content from Foodlog, it is not easy to build the personal model based on the Foodlog.

Discussion: Accurate personal model construction for food recommendation is still challenging because of its complexity of modeling and the diversity of factors to consider. As one important factor for one personal model, food preference learning has been widely exploited for food recommendation via various ways, such as survey based methods, historical records based methods and dialog-based methods. Some technologies such as matrix factorization and deep Siamese network are utilized. However, the former suffers from the scarcity of user data while the latter needs to construct a lot of item pairs. On the other hand, Foodlog comprehensively records users' food intake, daily diet and other user profiles. Therefore, Foodlog-oriented food preference learning is expected to be very promising in the future.

C. Heterogeneous Food Analysis for Food Recommendation

For single modality-oriented representation, the basic analysis method is visual food recognition with or without context information [7], [35]–[39]. We can obtain a high-level understanding of the type and the amount of food consumed by the user via food recognition. It has been used in Foodlog [37], food preference elicitation [17], [40] and dietary tracking [41] for further food recommendation. As representative work, Myers *et al.* [34] presented a system which can recognize the content of the meal from one image, and then predicted its nutritional content, such as calories. It used a segmentation based approach to localize the meal region of the food photo, and then applied CNN-based multilabel classifiers to label these segmented regions. Once the system segmented the foods, it can estimate their volume. There are also other works for food volume estimation via the 3D model based on multiple food images [42].

Besides visual information, there are other types of content information available, such as ingredients and various food attributes for food recommendation. For example, Peter *et al.* [43] incorporated ingredient information as linear constraints to extend the matrix factorization model for recipe recommendation. Devis *et al.* [44] selected relevant recipes via the comparisons among features used to annotate both users' profiles and recipes.

⁷[Online]. Available: <http://nutrino.co/>

Min *et al.* [45] considered one recipe as one document and one ingredient as one word, and proposed a probabilistic topic model for recipe representation. Other recipe attributes are also incorporated into the model to discover cuisine-course specific topic distribution to enable both cuisine and ingredient oriented recipe recommendation. Lin *et al.* [46] fused features from different types of information, such as ingredients, dietary, preparation, courses, cuisines and occasion into unified representation, and then used a content-driven matrix factorization approach to model the latent dimension of recipes, users, and features, respectively.

Combination of different food modalities leads to the multimodality of food, which particularly lies in the following two aspects: (1) One important factor determining our food choice is how we perceive food from its certain characteristics. We perceive food through a number of simultaneous sensory streams—we see food objects, hear its sounds when chewing, feel its texture, smell its odors and taste its flavors. Therefore, food perception actually involves multi-modalities. Food perception has an important affect on food preference learning. Accurate food perception should also learn multimodal sensory information [47]. (2) We are living in the age of the social web, and leave digital traces of all types of food-relevant activities online. For example, online recipe websites often have rich modalities and metadata about recipes. Each food item in Yummly consists of the visual food photo, textual content (e.g., name and ingredients) and attributes (e.g., cuisine and course). With the advent of Twitter, Facebook and Instagram, it is a common practice to upload food relevant text, share food photos. The multimodal information contains rich, detailed information about what people eat, how much, when, where, and even with whom. Such multimodal food information can provide a valuable signal for diet profiling and food preference elicitation.

Multimodal fusion methods [45], [48], [49] are needed to effectively combine different modalities. For example, Nag *et al.* [9] combined food images and GPS context for Foodlog. They [50] then fused multiple user source data streams along with the domain knowledge for cross-modal health state estimation. Recently, Markus *et al.* [51] used different kinds of features from different modalities, including a recipe's title, ingredient list and cooking directions, popularity indicators (e.g., the number of ratings) and visual features to estimate the healthiness of recipes for recipe recommendation. In addition, Chu *et al.* [52] combined text information, metadata and visual features for restaurant attributes and user preference representation. Two common recommendation approaches, i.e., content-based filtering and collaborative filtering are then integrated for restaurant recommendation. Recently, the capability of Deep Learning (DL) in processing heterogeneous data [53], [54] brings more opportunities in recommending more accurate and diverse items. For example, Min *et al.* [53] utilized the deep belief network to jointly model visual information, textual content (e.g., ingredients), and attributes (e.g., cuisine and course) to solve recipe-oriented problems. Salvador *et al.* [54] developed a multi-modal deep neural model which jointly learns to embed images and recipes in a common space for shared multi-modal representation.

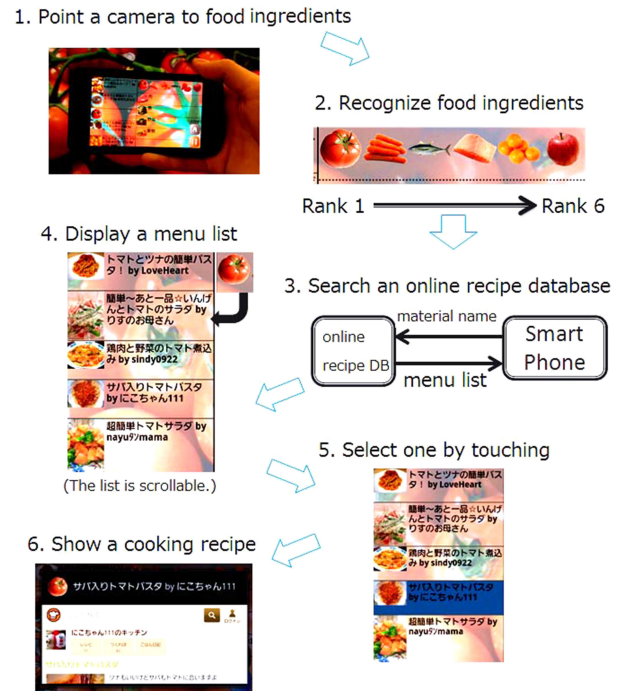


Fig. 7. Illustration of the overall recipe recommendation system [55].

As one representative work, Fig. 7 shows the flow of recipe recommendation system from the smart phone [7], [55]. The user first points a smartphone camera toward food ingredients at a grocery store or at a kitchen to recognize food ingredients, and then searches online cooking recipe databases with the name of the recognized food ingredient as a search keyword, and retrieve a menu list. Finally, the user selects one menu from the menu list and displays the corresponding cooking recipe including a list on necessary ingredients and seasonings and a cooking procedure on the pop-up window.

Compared with multimodal data, KG can organize heterogeneous food data into one effective semantic structure to enable food recommendation. It has been widely used for news recommendation [56] and movie recommendation [57]. To our knowledge, there is only one work [58], which constructed the food KG for semantic recipe retrieval.

Discussion. Rich food information is generally heterogeneous and complex, and is roughly divided into two types: multi-modal data and KG. Multimodal food analysis is becoming one fundamental stage for food recommendation. Deep learning methods have been one popular method for multimodal food analysis. However, some traditional models, such as probabilistic graphical models are still being used, especially for tasks with limited training samples. Despite these advances, multimodal food analysis still faces a number of difficulties. For example, some modalities such as smell and taste are very important for food recommendation, but is harder to quantify, not to mention the following multimodal fusion. In addition, there are few KG-oriented food recommendation methods because of the lack of a large-scale food KG. Constructing the food KG to enable food recommendation will help improve the recommendation performance and interpretability.

D. Nutrition and Health Oriented Food Recommendation

Many people are facing the problem of making healthier food decisions to reduce the risk of chronic diseases such as obesity and diabetes, which are very relevant to what we eat. Therefore, food recommendation not only caters user's food preference but should be also able to take user's health into account. Correspondingly, how to build the model to balance these two components becomes the core problem for health and nutrient-based food recommendation.

Most methods have tried to incorporate healthiness into the recommendation process by substituting ingredients [40], [59], incorporating calorie counts [60], generating food plans [61], and incorporating nutritional facts [9], [22], [62]. For example, Ge *et al.* [60] simply calculated the weighted mean between the preference component and health component, where the weights are manually adjusted by the user. Nag *et al.* [9] utilized food recognition method to obtain food-relevant information, and then obtain nutritional information for recommendation. In addition to health-aware food recommendation for general public, there are some works on recommendation for toddlers. For example, Yiu-Kai *et al.* [62] proposed a personalized toddler recipe recommendation system which incorporates the standard nutrition guideline published by the US government with users' food preference for toddlers. A hybrid recommendation approach, which incorporates content information directly as a linear constraint to provide additional insights about the contents themselves is adopted. Ribeiro *et al.* [63] considers more factors including nutrition, food preferences and the budget for meal recommender system for older adults.

Discussion. One important difference between food recommendation and other types of recommendation is that users' health information should be considered. Current methods mainly consider the nutrition information via visual food analysis or external nutrition table look-up for nutrition-oriented recommendation. However, user's real-time health-relevant state information is neglected. With the development of various sensor technologies, such information can be obtained from various portable devices, such as watches and bracelets. In addition, existing methods balance user's food preference and user's health via simple fusion, such as weighted summation [60], which are probably not accurate. More effective non-linear fusion methods between these two factors should be explored.

IV. RESEARCH CHALLENGES

Food recommendation recently has received more attention for its potential applications in human health. Thus, it is important to discuss existing challenges that form major obstacles to current progress. This section presents key unresolved issues.

A. Various Sensor Signal Fusion

Compared with recommendation in other fields, besides common spatio-temporal context, we should understand and model people states in food recommendation based on multiple signals, such as the affective state, physical state and health state captured from various sensors. The context information is with different types of distribution. Therefore, a joint model is necessary to

model the user state to accommodate a mixture of both discrete and continuous context variables. In addition, although we can resort sensors to obtain various context, there are still many signals, which are harder or inconvenient to obtain. The measure inaccuracy, uncertainty and unreliability from these sensors also exist. Therefore, further innovations on sensors are also needed for materials and devices, with the aim of more accurate sensing, the capability of detecting more context and real-time monitoring. In addition, how to filter out noise and conduct effective context fusion from these sensors is worth studying in the future. When the semantics from different modalities are not consistent, the conflict disambiguation and reasonable decision mechanism are further needed, but ignored in existing solutions.

B. Personal Model Construction

Few works focus on personal model construction for food recommendation. Simple person model can be defined by user's context and some inherent health parameters such as weight, height and activity steps. In order to achieve comprehensive and accurate personal model construction, according to the objective self [27], we should complete it through data ingestion, life event recognition, and pattern recognition. However, each component is hard to achieve and the challenges derive from three-fold: (1) data ingestion uses various sensors (e.g., smartphones and various wearable devices) and preprocessing modules to extract appropriate attributes from raw sensor measurements. However, the information obtained from different sensors varies in many aspects. Methods to convert data to information and the reliability of information could be entirely different for different sensors. Furthermore, as new sensing technologies emerge and are now becoming omnipresent in daily lives, more types of sensor data will be generated. It is difficult to effectively fuse heterogeneous multimedia information from different sensors via a unified model. (2) the proliferation of recipe-sharing websites (e.g., Yummly, Meishijie, foodspotting and Allrecipes) has resulted in huge user-uploaded food data collections. These recipes are associated with rich modality and attribute information. Such recipe data with rich types can be exploited to answer various food related questions in food recommendation. Besides recipe-sharing websites, the social media, such as Twitter, Foursquare, Flickr, and Instagram also provide large-scale food data uploaded by users. An increasing amount of user-shared food-related data presents researchers with more opportunities for personal model construction. However, it also presents researchers with challenges, such as much noise and the sheer size of user-shared food data. (3) accurate personal model construction benefits from fast development of many technologies, such as natural-language processing, machine learning and computer vision methods. For example, activity recognition techniques are necessary for life event recognition, and is still one hot problem in the computer vision. Lower recognition performance leads to inaccurate personal model construction, which conversely affects the experience of food recommendation.

C. Visual Food Analysis

Visual food analysis can obtain a high-level understanding of the type (e.g., the food category and ingredients), the amount of

food consumed by the user and even the calorie, and thus is very essential for food recommendation. This category can broadly be divided into different types, such as food category recognition [18], [64], [65], food ingredient recognition [7], cooking instruction recognition [54] and food quantity estimation [34], [42]. However, accurate visual food analysis is very challenging. Food images have their own distinctive properties. A large number of food dishes have deformable food appearance and thus lacks rigid structures. Therefore, it is hard for us to achieve satisfactory analysis results based on existing visual analysis methods. In addition, visual food analysis involves fine-grained visual feature representation learning. However, we can not simply use existing fine-grained feature learning methods for food analysis. One important reason is that existing fine-grained feature learning methods generally assume that there are fixed semantic patterns in each image and the task is to discover these patterns. However, the concepts of common semantic patterns do not exist in many food images. Therefore, we should design a new visual feature learning paradigm particularly for food.

D. Multi-Modality Food Analysis

Multi-modality is commonly the case in food recommendation. For example, each recipe item in Yummly consists of the food image, textual content (e.g., ingredients and cooking instructions), attributes (e.g., cuisine and taste), calories and nutrition facts. Users adopt both captured food images and text information for diet recording in the Foodlog. How to effectively integrate various types of information into the recommendation algorithm according to the characteristics of the specific recommendation scene has been a challenging problem, and how to enable effective feature extraction and fusion from various types of multimedia content information is also one core problem in this research field, not to mention the tight coupling between multi-modal fusion and collaborative filtering.

V. FUTURE DIRECTIONS

Food recommendation is still in its infancy in the multimedia community, and much more needs to be done to promote the development in this field. In this section, we discuss some open research directions and new research perspectives of food recommendation.

A. Food Preference Learning

Compared with user's preference on other objects, user's food preference has its own special features. Factors that guide food preference are various, and it is a result of a combination of biological, psychological, social, cultural and historical influences. Benefiting from the rapidly growing body of publicly available food data, it is convenient to utilize these data to learn user's food preference. However, they mainly focused on modeling one or several aspects from these data and neglected many other factors, which are probably hard to obtain or model. Borrowing methodologies from other areas such as psychology and neuroscience probably provides new perspectives for accurate food preference learning. In addition, most existing food recommender systems

assumed that user's food preference is static. However, users' food preferences are dynamic in reality [66]. A reasonable solution is to adopt deep reinforcement learning to model temporal dynamics of food preferences for food recommendation.

B. Large-Scale Benchmark Dataset Construction

Like MovieLens-1 M (<https://grouplens.org/datasets/movielens/1m/>) for movie recommendation, a large-scale dataset with rich user-food interaction and multimodal content information (e.g., food images, ingredients and user comments for one dish) is also a critical resource for developing advanced food recommendation algorithms, as well as for providing critical training and benchmark data for such algorithms. Although some datasets for food recommendation [67] are available, there are few public large-scale datasets, especially with multimodal information (e.g., food images and other rich attributes) for food recommendation. Therefore, collecting and releasing a large-scale multimodal benchmark dataset for the food recommendation task is in urgent need.

C. Large-Scale Food KG Construction

A complete KG can provide a deeper and longer range of associations between items. It has been proved effectively in recommendation system for many fields, such as movies and books. Unfortunately, there is no food-oriented KG available. Therefore, we should construct food-oriented KG. In order to obtain comprehensive food profile, we should crawl and parse metadata from various media sites with structured and semi-structured data. In many websites, metadata associated with content is poorly structured without a well defined ontology, especially in various social media. As a result, it is very hard to explore. Towards the food KG, we should define nodes and edges, and establish the standard for the food domain. A good starting point is to utilize food items from Wikipedia⁸ to build the food KG. In addition, we can even associate visual information and other modality information with this KG to build multimodal food knowledge graph.

D. Tight Combination Between DL and KG

Like other types of recommendation, food recommendation systems work on a lot of heterogeneous multi-modal information. Current systems mostly leverage heterogeneous information to improve the recommendation performance, while a lot of research efforts are needed regarding how to jointly leverage heterogeneous information for food recommendation. These include a wide range of research tasks such as multi-modal alignment on multiple different modalities and transfer learning over heterogeneous information sources for food recommendation. Recently, recommendation has undergone a fundamental paradigm shift towards using deep learning as a general-purpose solution. One key advantage of adopting DL is that arbitrary continuous and categorical features can be easily added to the model. Some works [17] have utilized deep learning for food

⁸[Online]. Available: https://en.wikipedia.org/wiki/Lists_of_foods

recommendation. We expect the wide use of deep learning to enhance the performance of food recommendation.

On the other hand, incorporating KG into recommender systems has attracted increasing attention in recent years. Compared with multimodal information, by exploring the relationships within a KG, the connectivity between users and items can be discovered and provides rich and complementary information to user-item interactions. However, to our knowledge, there is no work on food recommendation. The probable reason is that there is no constructed food-specific KG. Constructing the KG in the food domain requires us carefully define nodes and types of relationships, which is not easy to complete. Even we have available food KG, how to effectively combine KG embedding and collaborative filtering in the unified framework is still one hot topic and needs further exploration. Finally, effectively combining DL and KG jointly for food recommendation in a two-wheeled driven way will be needed.

E. Explainable Food Recommendation

Explainable recommendation refers to personalized recommendation algorithms that address the problem of why-they not only provide users with the recommendations, but also provide explanations to make the user aware of why such items are recommended. The same is true for food recommendation. Because food recommendation considers more complex and various factors, in this way, explainable food recommendation helps to find out main factors affecting recommended results and thus improve user satisfaction of recommendation systems. Incorporating KG into food recommendation is a promising direction for explainable food recommendation. In addition, the research community has been leveraging deep learning techniques for food recommendation. Current approaches focus on designing deep models to improve the performance. However, the research of leveraging deep models for explainable food recommendation is still in its initial stage, and there is much more to be explored in the future.

VI. CONCLUSION

Food recommendation is a promising and important research direction for its importance to quality of life for people and potential applications in human health. It provides people with ranked food items using rich context and knowledge, personal model constructed dynamically using Foodlog, and heterogeneous food analysis to understand nutrition and taste characteristics. However, there are very few systematic reviews. We believe that a critical and exhaustive review presented here may encourage researchers to shape this area. We provide a comprehensive and in-depth summary of current efforts as well as detail open problems in this area. This paper contains a comprehensive overview of food recommendation by proposing and using a unified food recommendation framework. Main issues affecting food recommendation are identified. The existing solutions are also introduced.

The area of food recommendation is still in its infancy with many challenges and open questions. Many of the challenges cannot be addressed using techniques from only one discipline. We believe that multidisciplinary research that combines

nutrition, food science, psychology, biology, anthropology and sociology will lead to more powerful approaches and technologies to handle food recommendation. Thus, considering huge potentials in human health and great commercial applications, we will look forward to seeing the surge in this research field in the foreseeable future.

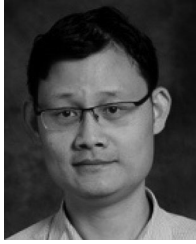
REFERENCES

- [1] A. Tirosh *et al.*, "The short-chain fatty acid propionate increases glucagon and FABP4 production, impairing insulin action in mice and humans," *Sci. Translational Med.*, vol. 11, no. 489, p. eaav0120, 2019.
- [2] I. D. Federation, *IDF Diabetes Atlas*, International Diabetes Federation, Brussels, Belgium, 2015.
- [3] G. R. F. Collaborators, "Global, regional, and national incidence, prevalence, and years lived with disability for 354 diseases and injuries for 195 countries and territories, 1990-2017: A systematic analysis for the global burden of disease study 2017," *Lancet*, vol. 392, no. 10159, pp. 1789-1858, 2018.
- [4] W. Min, S. Jiang, L. Liu, Y. Rui, and R. Jain, "A survey on food computing," *ACM Comput. Surv.*, vol. 52, no. 5, 2019, Art. no. 92.
- [5] W. Min *et al.*, "A survey on context-aware mobile visual recognition," *Multimedia Syst.*, vol. 23, no. 6, pp. 647-665, 2017.
- [6] S. Boll, J. Meyer, and N. E. O'Connor, "Health media: From multimedia signals to personal health insights," *IEEE MultiMedia*, vol. 25, no. 1, pp. 51-60, Jan.-Mar. 2018.
- [7] T. Maruyama, Y. Kawano, and K. Yanai, "Real-time mobile recipe recommendation system using food ingredient recognition," in *Proc. ACM Int. Workshop Interactive Multimedia Mobile Portable Devices*, 2012, pp. 27-34.
- [8] N. Nag, V. Pandey, and R. Jain, "Live personalized nutrition recommendation engine," in *Proc. 2nd Int. Workshop Multimedia Personal Health Health Care*, 2017, pp. 61-68.
- [9] N. Nag, V. Pandey, and R. Jain, "Health multimedia: Lifestyle recommendations based on diverse observations," in *Proc. ACM Int. Conf. Multimedia Retrieval*, 2017, pp. 99-106.
- [10] C. Trattner and D. Elswiler, "Food recommender systems: Important contributions, challenges and future research directions," 2017, *arXiv:1711.02760*.
- [11] R. Burke, "Hybrid recommender systems: Survey and experiments," *User Model. User Adapted Interaction*, vol. 12, no. 4, pp. 331-370, 2002.
- [12] Z. Lou, Li, L. Wang, and G. Shen, "Recent progress of self-powered sensing systems for wearable electronics," *Small*, vol. 13, no. 45, 2017, Art. no. 1701791.
- [13] E. Strickland, "3 sensors to track every bite and gulp [news]," *IEEE Spectr.*, vol. 55, no. 7, pp. 9-10, Jul. 2018.
- [14] K. Aizawa and M. Ogawa, "FoodLog: Multimedia tool for healthcare applications," *IEEE MultiMedia*, vol. 22, no. 2, pp. 4-8, Apr./Jun. 2015.
- [15] S. Kim, M. F. Fenech, and P. J. Kim, "Nutritionally recommended food for semi- to strict vegetarian diets based on large-scale nutrient composition data," *Sci. Rep.*, vol. 8, no. 1, 2018, Art. no. 4344.
- [16] M. Elahi *et al.*, "User nutrition modelling and recommendation: Balancing simplicity and complexity," in *Proc. Adjunct Publication Conf. User Model., Adaptation Personalization*, 2017, pp. 93-96.
- [17] L. Yang *et al.*, "Yum-Me: A personalized nutrient-based meal recommender system," *ACM Trans. Inf. Syst.*, vol. 36, no. 1, pp. 7-1-7-31, 2017.
- [18] L. Bossard, M. Guillaumin, and L. Van Gool, "Food-101-mining discriminative components with random forests," in *Proc. Eur. Conf. Comput. Vision*, 2014, pp. 446-461.
- [19] L. Smith and M. Gasser, "The development of embodied cognition: Six lessons from babies," *Artif. Life*, vol. 11, no. 1/2, pp. 13-29, 2014.
- [20] F. Cordeiro, E. Bales, E. Cherry, and J. Fogarty, "Rethinking the mobile food journal: Exploring opportunities for lightweight photo-based capture," in *Proc. ACM Conf. Human Factors Comput. Syst.*, 2015, pp. 3207-3216.
- [21] Y. Fu, B. Liu, G. Yong, Z. Yao, and X. Hui, "User preference learning with multiple information fusion for restaurant recommendation," in *Proc. SIAM Int. Conf. Data Mining*, 2014, pp. 470-478.
- [22] N. Nitish *et al.*, "Pocket dietitian: Automated healthy dish recommendations by location," in *Proc. Int. Conf. Image Anal. Process.*, 2017, pp. 444-452.
- [23] S. Sanjo and M. Katsurai, "Recipe popularity prediction with deep visual-semantic fusion," in *Proc. ACM Conf. Inf. Knowl. Manage.*, 2017, pp. 2279-2282.

- [24] J. Meyer and S. Boll, "Digital health devices for everyone!" *IEEE Pervasive Comput.*, vol. 13, no. 2, pp. 10–13, Apr.–Jun. 2014.
- [25] I. Fulgoni, L. Victor, D. R. Keast, and A. Drewnowski, "Development and validation of the nutrient-rich foods index: A tool to measure nutritional quality of foods," *J. Nutrition*, vol. 139, no. 8, pp. 1549–1554, 2009.
- [26] C. Julia *et al.*, "Validation of the FSA nutrient profiling system dietary index in french adults—findings from suvimax study," *Eur. J. Nutrition*, vol. 55, no. 5, pp. 1901–1910, 2016.
- [27] R. Jain and L. Jalali, "Objective self," *IEEE MultiMedia*, vol. 21, no. 4, pp. 100–110, Oct.–Dec. 2014.
- [28] J. Freyne and S. Berkovsky, "Intelligent food planning: Personalized recipe recommendation," in *Proc. Int. Conf. Intell. User Interfaces*, 2010, pp. 321–324.
- [29] M. Ge, M. Elahi, I. Fernández-Tobías, F. Ricci, and D. Massimo, "Using tags and latent factors in a food recommender system," in *Proc. 5th Int. Conf. Digit. Health*, 2015, pp. 105–112.
- [30] L. Yang *et al.*, "Plateclick: Bootstrapping food preferences through an adaptive visual interface," in *Proc. 24th ACM Int. Conf. Inf. Knowl. Manage.*, 2015, pp. 183–192.
- [31] X. Gao *et al.*, "Hierarchical attention network for visually-aware food recommendation," *IEEE Trans. Multimedia*, p. 1, doi: 10.1109/TMM.2019.2945180.
- [32] J. Zeng, Y. I. Nakano, T. Morita, I. Kobayashi, and T. Yamaguchi, "Eliciting user food preferences in terms of taste and texture in spoken dialogue systems," in *Proc. 3rd Int. Workshop Multisensory Approaches Human Food Interaction*, 2018, Art. no. 6.
- [33] H. Oh, J. Nguyen, S. Soundararajan, and R. Jain, "Multimodal food journaling," in *Proc. 3rd Int. Workshop Multimedia Personal Health Health Care*, 2018, pp. 39–47.
- [34] A. Meyers *et al.*, "Im2Calories: Towards an automated mobile vision food diary," in *Proc. IEEE Int. Conf. Comput. Vision*, 2015, pp. 1233–1241.
- [35] H. Wang, W. Min, X. Li, and S. Jiang, "Where and what to eat: Simultaneous restaurant and dish recognition from food image," in *Proc. Pacific Rim Conf. Multimedia*, 2016, pp. 520–528.
- [36] S. Horiguchi, S. Amano, M. Ogawa, and K. Aizawa, "Personalized classifier for food image recognition," *IEEE Trans. Multimedia*, vol. 20, no. 10, pp. 2836–2848, Oct. 2018.
- [37] K. Aizawa, Y. Maruyama, H. Li, and C. Morikawa, "Food balance estimation by using personal dietary tendencies in a multimedia food log," *IEEE Trans. Multimedia*, vol. 15, no. 8, pp. 2176–2185, Dec. 2013.
- [38] R. Xu *et al.*, "Geolocalized modeling for dish recognition," *IEEE Trans. Multimedia*, vol. 17, no. 8, pp. 1187–1199, Aug. 2015.
- [39] J. Chen and C.-W. Ngo, "Deep-based ingredient recognition for cooking recipe retrieval," in *Proc. ACM Multimedia Conf.*, 2016, pp. 32–41.
- [40] M. Harvey, M. Harvey, and M. Harvey, "Exploiting food choice biases for healthier recipe recommendation," in *Proc. Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2017, pp. 575–584.
- [41] Z.-Y. Ming *et al.*, "Food photo recognition for dietary tracking: System and experiment," in *Proc. Int. Conf. MultiMedia Model.*, 2018, pp. 129–141.
- [42] J. Dehais, M. Anthimopoulos, S. Shevchik, and S. Mougiakakou, "Two-view 3D reconstruction for food volume estimation," *IEEE Trans. Multimedia*, vol. 19, no. 5, pp. 1090–1099, May 2017.
- [43] P. Forbes and M. Zhu, "Content-boosted matrix factorization for recommender systems: Experiments with recipe recommendation," in *Proc. 5th ACM Conf. Recommender Syst.*, 2011, pp. 261–264.
- [44] D. Bianchini, V. De Antonellis, and M. Melchiori, "A web-based application for semantic-driven food recommendation with reference prescriptions," in *Web Information Systems Engineering*. Berlin, Germany: Springer, 2015, pp. 32–46.
- [45] W. Min *et al.*, "You are what you eat: Exploring rich recipe information for cross-region food analysis," *IEEE Trans. Multimedia*, vol. 20, no. 4, pp. 950–964, Apr. 2018.
- [46] C.-J. Lin, T.-T. Kuo, and S.-D. Lin, "A content-based matrix factorization model for recipe recommendation," in *Advances in Knowledge Discovery and Data Mining*. Berlin, Germany: Springer, 2014, pp. 560–571.
- [47] J. V. Verhagen and L. Engelen, "The neurocognitive bases of human multimodal food perception: sensory integration," *Neurosci. Biobehav. Rev.*, vol. 30, no. 5, pp. 613–50, 2006.
- [48] W. Min, S. Jiang, S. Wang, J. Sang, and S. Mei, "A delicious recipe analysis framework for exploring multi-modal recipes with various attributes," in *Proc. ACM Multimedia Conf.*, 2017, pp. 402–410.
- [49] J. Chen, C. Ngo, and T. Chua, "Cross-modal recipe retrieval with rich food attributes," in *Proc. ACM Multimedia Conf.*, 2017, pp. 1771–1779.
- [50] N. Nag *et al.*, "Cross-modal health state estimation," in *Proc. 26th ACM Int. Conf. Multimedia*, 2018, pp. 1993–2002.
- [51] M. Rokicki, C. Trattner, and E. Herder, "The impact of recipe features, social cues and demographics on estimating the healthiness of online recipes," in *Proc. 12th Int. Conf. Web Soc. Media*, 2018, pp. 310–319.
- [52] W. T. Chu and Y. L. Tsai, "A hybrid recommendation system considering visual information for predicting favorite restaurants," *World Wide Web-Internet Web Inf. Syst.*, vol. 20, no. 6, pp. 1313–1331, 2017.
- [53] W. Min *et al.*, "Being a supercook: Joint food attributes and multi-modal content modeling for recipe retrieval and exploration," *IEEE Trans. Multimedia*, vol. 19, no. 5, pp. 1100–1113, May 2017.
- [54] A. Salvador *et al.*, "Learning cross-modal embeddings for cooking recipes and food images," in *Proc. Comput. Vision Pattern Recognit.*, 2017, pp. 3020–3028.
- [55] Y. Kawano, T. Sato, T. Maruyama, and K. Yanai, "mirurecipe: A mobile cooking recipe recommendation system with food ingredient recognition," in *Proc. IEEE Int. Conf. Multimedia Expo. Workshops*, 2013, pp. 1–2.
- [56] H. Wang, F. Zhang, X. Xie, and M. Guo, "DKN: Deep knowledge-aware network for news recommendation," in *Proc. World Wide Web Conf.*, 2018, pp. 1835–1844.
- [57] F. Zhang, N. J. Yuan, D. Lian, X. Xie, and W. Ma, "Collaborative knowledge base embedding for recommender systems," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2016, pp. 353–362.
- [58] U. Zulaika, A. Gutiérrez, and D. López-de Ipiña, "Enhancing profile and context aware relevant food search through knowledge graphs," in *Proc. Multidisciplinary Digit. Publishing Inst.*, vol. 2, no. 19, p. 1218, 2018.
- [59] C.-Y. Teng, Y.-R. Lin, and L. A. Adamic, "Recipe recommendation using ingredient networks," in *Proc. ACM Web Sci. Conf.*, 2012, pp. 298–307.
- [60] M. Ge, F. Ricci, and D. Massimo, "Health-aware food recommender system," in *Proc. Conf. Recommender Syst.*, 2015, pp. 333–334.
- [61] D. Elsweiler and M. Harvey, "Towards automatic meal plan recommendations for balanced nutrition," in *Proc. 9th ACM Conf. Recommender Syst.*, 2015, pp. 313–316.
- [62] Y.-K. Ng and M. Jin, "Personalized recipe recommendations for toddlers based on nutrient intake and food preferences," in *Proc. 9th Int. Conf. Manage. Digit. EcoSyst.*, 2017, pp. 243–250.
- [63] D. Ribeiro, J. Ribeiro, M. J. M. Vasconcelos, E. F. Vieira, and A. C. de Barros, "Souschef: Improved meal recommender system for Portuguese older adults," in *Information and Communication Technologies for Ageing Well and e-Health*. Berlin, Germany: Springer, 2018, pp. 107–126.
- [64] W. Min, L. Liu, Z. Luo, and S. Jiang, "Ingredient-guided cascaded multi-attention network for food recognition," in *Proc. 27th ACM Int. Conf. Multimedia*, 2019, pp. 1331–1339.
- [65] S. Jiang, W. Min, L. Liu, and Z. Luo, "Multi-scale multi-view deep feature aggregation for food recognition," *IEEE Trans. Image Process.*, vol. 29, no. 1, pp. 265–276, Jul. 2019.
- [66] J. Delarue and E. Loescher, "Dynamics of food preferences: a case study with chewing gums," *Food Quality Preference*, vol. 15, no. 7, pp. 771–779, 2004.
- [67] C.-J. Lin, T.-T. Kuo, and S.-D. Lin, "A content-based matrix factorization model for recipe recommendation," in *Advances in Knowledge Discovery and Data Mining*. Berlin, Germany: Springer, 2014, pp. 560–571.

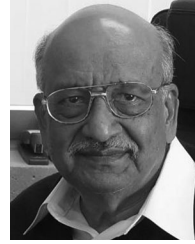


Wei Qing Min received the B.E. degree from Shandong Normal University, Jinan, China, in 2008, the M.E. degree from Wuhan University, Wuhan, China, in 2010, the Ph.D. degree from the National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, Beijing, China, in 2015. He is currently an Associate Professor with the Key Laboratory of Intelligent Information Processing, Institute of Computing Technology, Chinese Academy of Sciences. His current research interests include multimedia content analysis, understanding and applications, food computing and geo-multimedia computing. He has authored and coauthored more than 30 peer-reviewed papers in relevant journals and conferences, including *ACM Computing Surveys*, the *IEEE TRANSACTIONS ON IMAGE PROCESSING*, *IEEE TRANSACTIONS ON MULTIMEDIA*, *ACM TOMM*, *IEEE Multimedia Magazine*, *ACM Multimedia*, *International Joint Conferences on Artificial Intelligence*, etc. He is the Reviewer of more than ten well-known international journals including the *IEEE TRANSACTIONS ON MULTIMEDIA*, *IEEE TRANSACTIONS ON CYBERNETICS*, *IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY*, *IEEE TRANSACTIONS ON NEURAL NETWORK AND LEARNING SYSTEM*, *ACM TOMM*, etc. He is the recipient of 2016 *ACM TOMM* Nicolas D. Georganas Best Paper Award and the 2017 *IEEE Multimedia Magazine* Best Paper Award.



Shuqiang Jiang (SM'08) is a Professor with the Institute of Computing Technology, Chinese Academy of Sciences (CAS), Beijing, China and a Professor with the University of CAS. He is also with the Key Laboratory of Intelligent Information Processing, CAS. His research interests include multimedia processing and semantic understanding, pattern recognition, and computer vision. He has authored or coauthored more than 150 papers on the related research topics. He was supported by the New-Star Program of Science and Technology of Beijing Metropolis in 2008, NSFC

Excellent Young Scientists Fund in 2013, and Young Top-Notch Talent of Ten Thousand Talent Program in 2014. He won the Lu Jiaxi Young Talent Award from Chinese Academy of Sciences in 2012, and the CCF Award of Science and Technology in 2012. He is the Senior Member of CCF, a member of ACM, an Associate Editor for the *IEEE Multimedia*, Multimedia Tools and Applications. He is the Vice Chair of IEEE CASS Beijing Chapter and Vice Chair of ACM SIGMM China chapter. He is the General Chair of ICIMCS 2015 as well as Program Chair of ACM Multimedia Asia2019 and PCM2017. He has also served as a TPC member for more than 20 well-known conferences, including ACM Multimedia, CVPR, ICCV, IJCAI, AAAI, ICME, ICIP, and PCM.



Ramesh Jain is an Entrepreneur, Researcher, and Educator. He is currently a Donald Bren Professor in Information and Computer Sciences with the University of California (UCI), Irvine, CA, USA. His research interests cover control systems, computer vision, artificial intelligence, and multimedia computing. He is addressing health issues using cybernetic principles building on the progress in sensors, mobile, processing, artificial intelligence, computer vision, and storage technologies. He is Founding Director with the Institute for Future Health, UCI. AAAS, ACM, IEEE, AAAI, IAPR, and SPIE have elected him as Fellow of their societies. He co-founded several companies, managed them in initial stages, and then turned them over to professional management. He enjoys new challenges and likes to use technology to solve them. He is participating in addressing the biggest challenge for us all: how to live long in good health.

IEEE, AAAI, IAPR, and SPIE have elected him as Fellow of their societies. He co-founded several companies, managed them in initial stages, and then turned them over to professional management. He enjoys new challenges and likes to use technology to solve them. He is participating in addressing the biggest challenge for us all: how to live long in good health.