

A Food Recommender for Patients in a Care Facility

Toon De Pessemer
iMinds-WiCa-Ghent University
G. Crommenlaan 8 box 201
B-9050 Ghent, Belgium
toon.depessemier@ugent.be

Simon Dooms
iMinds-WiCa-Ghent University
G. Crommenlaan 8 box 201
B-9050 Ghent, Belgium
simon.dooms@ugent.be

Luc Martens
iMinds-WiCa-Ghent University
G. Crommenlaan 8 box 201
B-9050 Ghent, Belgium
luc1.martens@ugent.be

ABSTRACT

In this research, a food recommendation strategy for patients in a care facility is proposed. Since many of these patients cannot express their personal preferences, a recommender system can assist the caregivers in the selection of the menu items that match the patients' preferences. Recommendations are generated based on three information sources: explicit ratings for menu items, implicit feedback based on the patient's eating behavior and the amount of food that was eaten, and inferred preferences for the ingredients of the menu items. In addition, monitoring the amount of food that was eaten by each patient can provide insights into the optimal amount of each menu item that has to be served to each patient. Furthermore, monitoring food consumption allows to detect irregularities in the eating behavior of the patient, which can be a symptom of illness.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous

Keywords

recommender system; food; implicit feedback

1. INTRODUCTION

Three times a day, patients in care facilities receive a prepared meal tailored to their personal needs and health condition. For each meal, various menu options are available for the patients to choose. Some of these menu options themselves can be tailored according to the preferences of the patients, such as the filling of a sandwich, the presence of sauce or not, or the accompanying beverage. However, due to illness or age, many patients of the care facility are unable to express their personal preferences for the meal options they receive. E.g., some patients cannot speak anymore or are unintelligible and patients suffering from dementia might not be able to link the name of a dish to the actual food.

Until now, if patients are unable to express their choice, the caregivers make the decision instead of the patients. Since caregivers have many different patients, and a patient can receive meals from different caregivers, caregivers are not able to optimally estimate the food preferences of patients. Therefore, the selection of meal options made by the caregivers can be alternating (e.g., Monday option 1, Tuesday option 2, etc.) or even be random.

The goal of this research was to develop a recommender system that provides food recommendations for patients who are unable to make or express their own choices. By using this recommender, caregivers can offer patients a meal that is more tailored to the patients' personal preferences.

2. RELATED WORK

In the past years, various systems have been developed to assist people in the kitchen or to provide personal suggestions for (healthy) meals. Meals can be suggested by providing the user with the recipe, explaining how to prepare the food. However, the lack of competence and confidence is a common barrier towards the preparation of a (new) dish [8]. Therefore, the cooking competence is an important factor to take into account while recommending recipes. This cooking competence can be estimated by tracking the user's activities in the kitchen by means of sensors in the kitchen utensils [8]. The number of ingredients and activities involved in preparing a meal is an indication for the difficulty of preparing the meal. If users have no experience with preparing difficult meals, the recommender system can take this knowledge into account and avoid complicated recipes. Moreover, this knowledge can be used to assist people to gradually increase their cooking competence [5].

Besides the cooking competence, various other factors, both context and content related, influence the user's preferences for meals and ratings for recipes. The most common reasons for a negative evaluation of a recipe are that the recipe contains a particular disliked ingredient, the combination of ingredients, or the recipe would take too long to prepare and cook. The most common reasons for appreciating a recipe (and providing positive feedback) are the ease and quickness of preparation, the user's own preference for the recipe or an ingredient, and the type of dish or recipe being novel or interesting [4]. So in the context of food recommendations, research has pointed to the highly significant influence of the individual tastes of the user and the content of the recipes in terms of ingredients [4].

Based on these insights, ratings for a meal can be broken down into ratings for its ingredients or component foods, and

ratings for ingredients or component foods can be composed into ratings for meals [4]. Using these inferred ratings for ingredients, a fine-grained food item strategy can take into account the various ingredients of a meal during the recommendation process. Research has shown that more accurate food recommendations can be made when calculations are based on implied user ratings on ingredients rather than on recipes [4].

In some circumstances, the food recommendations are not intended for assisting individuals in the food selection process, but rather as suggestions for groups of people, for example in a family context. This raises the need to uncover the most appropriate group recommendations, taking into account the preferences of all group members in the recommendation process. In the context of recipe recommendations, two strategies to generate group recommendations have been compared: a strategy that aggregates the individual user preferences into a model representing the group preferences and then calculates group recommendations based on this model, and a strategy that calculates recommendations for each individual member of the group and then aggregates the recommendations of the individual members into group recommendations [2].

The specific use case of food recommendations for patients in a care facility has some direct consequences on the factors that have to be taken into account in the recommendation process. E.g., since meals are prepared by professional cooks, the cooking competences of the end-users (i.e., the patients) are irrelevant. The personal preferences of the patients are an important influencing factor as well as diet or medical restrictions such as diabetes, lactose intolerant, allergies, etc. Since food choices can be made on an individual basis, no group recommendation strategy has to be used.

3. RECOMMENDATION STRATEGY

In order to enable personal recommendations for menu items in the context of a care facility, a framework was developed to store data about the preferences, activities, and behavior of the patients regarding the meals. This framework stores explicit personal ratings of the patient for the menu items (if available), the ingredients of each menu item, and for each meal, details about the amounts of food that each patient received from the caregivers, the amounts of food that were eaten by each patient, and the choices regarding menu items that each patient (or the caregiver) made.

Explicit ratings can be used to specify the patient's favorite meals or his/her aversion for menu items or ingredients. These explicit preferences can be specified by the patient (in an early stage of the disease), the caregivers, or even by family members who visit the patient.

The amount of food that was eaten by the patient is considered as implicit feedback for the item. Menu items that are completely eaten receive positive feedback, whereas menu items that are not eaten receive negative feedback, thereby relying on the assumption that patients eat more if they like the food. This feedback is acquired by means of a mobile application, available on the smartphones and tablets used by caregivers to register the activities and behavior of the patients as well as special situations.

All this information is used by a hybrid recommendation algorithm, which combines content-based [7] and collaborative filtering techniques [3] in order to generate personal recommendations for the patients. For each menu choice,

(i.e., for each food option that patients have), the recommendation algorithm calculates a *preference score* that estimates how much the patient will like the food option. Subsequently, the menu item(s) with the highest preference score will be recommended to the patient (if no restrictions apply). During the calculation of this preference score for a menu item, the following three different cases are considered in the recommendation algorithm.

1. An explicit rating for the item is available.
2. No explicit rating is available, but implicit feedback for the item is available
3. No explicit rating and no implicit feedback are available. The ingredients of the menu item are the only available information.

3.1 Prediction Based on Explicit Ratings

For the food recommendations, the explicit ratings are considered as an important information source for estimating the patient's preferences. So if the patient has specified an explicit rating for a menu item, this will strongly influence the patient's preference score for this item.

In order to take into account the possible changes in the preferences of patients, the implicit feedback for the item is also considered in the calculation process in addition to the explicit ratings, which might not be updated frequently. Moreover, the implicit feedback, which reflects the eating behavior of the patient, may refine the preferences inferred from the explicit ratings.

As a result, the *Explicit Prediction* of the patient's preference of an item consists of a weighted average of the patient's explicit rating and a prediction of the patient's preference based on his/her implicit feedback for the item, which is referred to as *Implicit Prediction*.

$$\begin{aligned} \text{ExplicitPrediction} = & \alpha \cdot \text{ExplicitRating} \\ & + (1 - \alpha) \cdot \text{ImplicitPrediction} \end{aligned} \quad (1)$$

In this formula, α is a weighting factor that has a value ranging from 0 to 1 if an explicit rating for the menu item is available, and that is 0 if no explicit rating is available. The explicit prediction is used as input for the *AlreadyKnown* recommendation algorithm of the Duine recommender framework [7]. This recommendation algorithm generates a preference score for the menu item, thereby taking into account the fact that the user has already provided feedback for that item.

3.2 Prediction Based on Implicit Feedback

If the patient (or someone who knows the patient's preferences) has not specified an explicit rating for the menu item, the patient's preference for the menu item is estimated based on his/her eating behavior (enthusiastically or reluctantly) and the amount of food that the patient has eaten. The eating behavior and the amount of food (i.e., the implicit feedback for the item) is registered by the caregivers using their mobile device as part of the monitoring process of the patients.

However, this implicit feedback is not always a perfect predictor of the patient's preferences, since the patient's eating behavior and the amount of food that was eaten can be influenced by other factors besides the patient's preference for

the item. Unpredictable factors, such as the current appetite of the patient or illness, have also a significant impact on the eating behavior and the amount of food that was eaten. Consequently, the implicit feedback for an item can deviate from the true preferences of the patient.

Therefore, a prediction of the patient’s preference for the item, based on the patient’s preferences for the ingredients of the item, is added to the calculation process. So, the *Implicit Prediction* is calculated by means of a weighted average of the implicit feedback for the item, and a prediction based on the ingredients of the item, i.e., *Ingredient Prediction*.

$$\begin{aligned} \text{ImplicitPrediction} = & \beta \cdot \text{ImplicitFeedback} \\ & + (1 - \beta) \cdot \text{IngredientPrediction} \end{aligned} \quad (2)$$

In this formula, β is a weighting factor that has a value ranging from 0 to 1 if implicit feedback for the menu item is available, and that is 0 if no implicit feedback is available. The implicit prediction is used as input for the *Already-Known* recommendation algorithm, in order to calculate a preference score, which estimates how much the patient will like the menu item.

3.3 Prediction Based on Ingredients

For new menu items that never have been served to the patient, explicit ratings or implicit feedback may not be available. Therefore, the preference of the patient for a new menu item is estimated based on the ingredients of the menu item and the patient’s preferences for these ingredients.

As an alternative for a reasoning based on the ingredients, a standard collaborative filtering approach could be used to predict the patient’s preferences for menu items. Nevertheless, research has shown that the menu breakdown into ingredients with reasoning on the ingredients provides more accurate recommendations than a collaborative filtering algorithm reasoning on the menu items [4].

First, the menu item is decomposed into high-level ingredients. E.g., the menu item “macaroni” is split into “cheese”, “ham”, and “pasta”. Next, the patient’s preferences for these ingredients are inferred. Similar to the approach of Freyne and Berkovsky [4], ratings for menu items are transferred to the ingredients of these menu items. In other words, the ingredients (e.g., cheese, ham, and pasta) get the same rating as the item that contains the ingredients (e.g., macaroni). Hereby, cooking processes and combinations effects are ignored, and all ingredients are considered to be equally weighted in a menu item.

For an ingredient that is new for the patient, i.e., the patient has not yet (implicitly or explicitly) evaluated a menu item containing that ingredient, ingredient preferences cannot be inferred from the item preferences of the patient. If the patient’s historical data provide no clue about his/her preferences for some ingredients, preferences have to be inferred by using the *collaborative filtering* implementation of the Duine framework [7].

Through collaborative filtering, personal preferences for new ingredients are inferred from the preferences of neighboring patients, i.e., patients with similar preferences. Research has shown that more accurate neighbor determination and food prediction occurs when neighbors are based on implied user ratings for ingredients rather than on ratings for menu items [4]. So in our developed recommender system,

neighboring patients are determined by calculating the Pearson similarity between the patients’ inferred preferences for the ingredients of the menu items.

To summarize, the patient’s preferences for ingredients are inferred from his/her ratings and implicit feedback for menu items that contain these ingredients. If the patient has never (implicitly or explicitly) evaluated a menu item containing a specific ingredient, his/her preference for this ingredient is calculated using collaborating filtering. The *Ingredient Prediction* for an item is calculated by averaging the patient’s preferences for the n ingredients of the item.

$$\text{IngredientPrediction} = \frac{1}{n} \sum_{i=1}^n \text{IngredientPreference}_i \quad (3)$$

However, (the taste of) some ingredients can be predominant in a menu, so that if these ingredient are disliked, the entire menu will be disliked, irrespective of other likable ingredients in the same menu item. If more information about the (relative importance of) ingredients becomes available, the compensatory approach of the average function can be replaced by a more informed aggregation method such as a weighted average in which the weights reflect the importance of the ingredients in the menu. Finally, the patient’s preference for a menu item (i.e., the preference score) is estimated by the *content-based* recommender of the Duine framework [7], which reasons on the user’s preferences for the ingredients of the item.

3.4 Rule-Based Restrictions

The main goal of the recommender system is to generate accurate recommendations by suggesting menu items that match the patient’s (inferred) preferences. However being accurate is not enough for a recommender system [6]; also other characteristics like diversity are essential for a valuable list of suggestions. Especially for a food recommender, diversity is an important quality aspect. Without diversity, patients risk to get recommended their favorite dish every day. Therefore, caregivers can specify how often a menu item can be recommended to a patient. Last but not least, diet or medical restrictions can be specified for specific menu items or in terms of the ingredients of the menu items. The restrictions of each patient are taken into account during the recommendation process, thereby also considering the previous meals. E.g., some patients can only eat a limited amount of sugar each day, so the choice of the breakfast influences the options for the lunch. Menu items or ingredients not fulfilling these personal requirements are eliminated from the patient’s recommendations.

4. ANALYZING THE EATING BEHAVIOR

4.1 Predicting the Optimal Amount of Food

Besides assisting the patients in the food selection process by providing personal recommendations, the system analyzes the gathered data in order to provide insights into the eating habits of the patients. This data analysis is performed by the monitoring engine of the system, which operates in parallel with the recommendation engine. To facilitate the usage of the system, the monitoring engine is self-learning without the need for configuration of rules.

By analyzing the amount of food eaten by every patient over a long period of time, the monitoring engine can make a prediction about the amount of food that will be eaten by a patient in the future. For each menu item that is served to a patient, the system can make such a prediction. Predicting the amounts of food that patients will eat, can be useful for determining the size of the food portions as well as for food ordering and stock management of the care facility.

The optimal amount of a specific menu item for a patient is predicted by using a standard regression analysis based on the knowledge available in the system (independent variables): the menu item and its ingredients, the amounts of the menu item that were eaten by the patient in the past, and the general appetite of the patient, which is measured by the amounts of other menu items that were eaten by the patient during the past meals. To reflect the recent changes in the eating behavior of patients, the most-recently acquired data, i.e., data about the meals of the past week, are analyzed separately from the formerly acquired data.

4.2 Detecting Irregularities

Another task of the monitoring engine is to detect irregularities in the eating behavior of the patients. The system can detect abrupt and unexpected increases or decreases in the amount of food eaten by a patient. A sudden decrease can for example be an indication of illness. Although the monitoring engine passes no judgment about the cause of such a change in eating behavior, detecting these changes can be helpful for the caregivers in order to monitor the health condition of the patients. Data regarding the eating behavior of a patient is analyzed using regression analysis to calculate the variations in the amount of food eaten by the patient, and to determine if these variations are statistically significant. Exceptional small (or large) amounts of food or drinks consumed by the patient are detected using the outlier detection mechanisms of the ELKI framework [1]. Statistically significant trends or exceptional cases in the eating behavior are automatically reported to the caregivers via their mobile device, so that they can intervene if the health of the patient might be in danger.

However, a health problem is not the only possible cause of a variation in the amount of food that patients eat. Significant variations might be caused by external conditions such as the temperature in the dining room. E.g., if the temperature rises a few degrees in the summer, one might expect an increased consumption of water by the patients in the care facility. Information about the environment is retrieved at regular time intervals by means of sensors installed in the care facility. This way, the system can calculate the correlation between the increase or decrease of the amount of consumed food and the environmental conditions such as the temperature or humidity in the room. Moreover, the environmental conditions are used as independent variables in the regression analysis to explain trends. The results of the regression analysis and the correlation values are reported to the caregivers, together with a visual presentation of the variations in the food consumption, to enable them to make a more informed decision.

5. CONCLUSIONS

This paper proposes a hybrid recommendation algorithm for suggesting menu items to patients in a care facility. In the context of food recommendations, suggesting already-

consumed items again is desirable if the user liked the item in the past. As a result, feedback for the potential recommendation might already be available as an additional information source during the calculation process. In contrast, for many traditional application domains, such as recommenders for online shops or movies, users generally prefer not to buy or watch the same item twice. Future work will focus on the evaluation of the proposed recommender system. Therefore, the recommender will be integrated into the central data system of two care facilities in order to evaluate the food suggestions and the analysis of eating behavior. Subsequently, a qualitative evaluation consisting of interviews with patients and caregivers will provide insight into the usefulness of the system.

6. ACKNOWLEDGMENTS

This work was supported by the iMinds aToM (A Touch of Memory) project. This is a project co-funded by iMinds (Interdisciplinary institute for Technology) a research institute founded by the Flemish Government.

7. REFERENCES

- [1] ELKI: Environment for Developing KDD-Applications Supported by Index-Structures, 2012. Online available at <http://elki.dbs.ifi.lmu.de/>.
- [2] S. Berkovsky and J. Freyne. Group-based recipe recommendations: analysis of data aggregation strategies. In *Proceedings of the fourth ACM conference on Recommender systems*, RecSys '10, pages 111–118, New York, NY, USA, 2010. ACM.
- [3] J. S. Breese, D. Heckerman, and C. Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, UAI'98, pages 43–52, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc.
- [4] J. Freyne and S. Berkovsky. Intelligent food planning: personalized recipe recommendation. In *Proceedings of the 15th international conference on Intelligent user interfaces*, IUI '10, pages 321–324, New York, NY, USA, 2010. ACM.
- [5] M. Harvey, B. Ludwig, and D. Elswiler. Learning user tastes: a first step to generating healthy meal plans? In *First International Workshop on Recommendation Technologies for Lifestyle Change (LIFESTYLE 2012)*, page 18, 2012.
- [6] S. M. McNee, J. Riedl, and J. A. Konstan. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *CHI '06 extended abstracts on Human factors in computing systems*, CHI EA '06, pages 1097–1101, New York, NY, USA, 2006. ACM.
- [7] Telematica Instituut / Novay. Duine Framework, 2009. Online available at <http://duineframework.org/>.
- [8] J. Wagner, G. Geleijnse, and A. van Halteren. Guidance and support for healthy food preparation in an augmented kitchen. In *Proceedings of the 2011 Workshop on Context-awareness in Retrieval and Recommendation*, CaRR '11, pages 47–50, New York, NY, USA, 2011. ACM.