Challenges for Nutrition Recommender Systems

Stefanie Mika

Department of Computer Science 8
Friedrich-Alexander-University Erlangen-Nuremberg
Haberstraße 2, D-91058 Erlangen
stefanie.mika@cs.fau.de

Abstract. Obesity or being overweight in general often leads to other more severe diseases. With more than half of the population in the western countries being overweight or even obese, systems are developed to help users make healthier food choices by recommending healthier recipes or food items according to the user's needs or likes and dislikes. But there are many challenges designing such systems. In this paper we try to identify the difficulties encountered and discuss ways to deal with them.

Key words: Recommender System, Nutrition, E-Health

1 Introduction

With more than 50% of the German population either overweight or obese [1] and many health issues caused by obesity [2][3], information about nutrition and a healthy diet is evermore sought-after. There are a lot of forums or newsgroups online with regard to nutrition and diet[1], where the interested user can find support and/or information regarding a healthy diet. The challenge here is to find appropriate and trustworthy information [4]. Several sources of information are available ranging from government websites to recipe databases to online magazines [2]. But here the uninformed user finds an often confusing amount of information as well as conflicting information or information that needs expertise to interpret correctly leaving them at a loss at what to do or how to change their diet. A potential solution to help the user cope with this large amount of data concerning healthy life choices and diet is to develop recommender systems, which can automatically make suitable suggestions based on the user's personal profile.

In this paper we will look at the difficulties for recommender systems in nutrition as well as the challenges to overcome in implementing and designing these systems. Section 2 presents related work in nutrition recommendation, section 3 identifies challenges and problems giving nutritional recommendations.

[1] e.g. diaet.abnehmen-forum.com/ (last visited: 15.08.2011), forum.ernaehrung.de/ (15.08.2011) or www.kilosweg.de/ (15.08.2011)

[2] e.g. the websites of the world health organisation www.who.int/topics/nutrition/en/ or the national health service for England www.nhs.uk/livewell/healthy-eating/Pages/Healthyeating.aspx (both last visited on 15.08.2011)

In section 4 we discuss how to handle to issues of section 3 and section 5 gives an outlook on future work.

2 Related work

Giving good recommendations about nutrition does not only rely on the algorithm used to calculate the recommendations. It also depends on the knowledge about a user's eating habits. Therefore it is advisable to record what foods and the amount of foods a user consumes. The information gathered can not only be used to identify the problems with the user's diet, but also be used in the algorithm to improve the recommendations.

2.1 Recording food consumption

Professional nutritionists utilise a broad selection of tools to establish an individual's nutritional intake. These include food diaries, 24-hour recalls or questionnaires, with food diaries or 24-hour recalls being preferred giving superior measures[5]. With food diaries the user records all amounts of consumed foods and beverages over a period of time preferably at the time they were consumed, usually no more than 3 or 4 days are recorded. For 24-hour recalls the user has to give a report during a meeting with the nutritionist of every food or beverage consumed over the last 24 hours or the last day. Questionnaires are often food frequency questionnaires, where a user notes how often they consumed each food from a list of given foods over a given period of time [6]. All of these methods rely on the ability of the participant to accurately remember what they have eaten or to remember to record everything they consume. As with any kind of diary study the reliabilirty or the accuracy of these techniques are open to severe doubt [7].

More recently systems were developed which calculate the amount of food intake with the help of pictures of meals the user has taken. An example of such a system is FoodLog [8] which identifies food pictures in the user's web archives and calculates the dietary balance of meals.

2.2 Food recommender systems

There are two ways to recommend healthier food choices. The first one is to recommend recipes for healthier meals, the second one is to suggest healthier food items themselves.

Most recipe recommender systems calculate the best choice based on a database of recipes as well as user ratings. Some algorithms can cope with a more extensive amount of recipes, while others need a larger number of user ratings. Many different approaches have been implemented in the past. Some put more emphasis on user ratings, while others take the nutritional needs of the user into more consideration.

There are two types of recipe recommender systems. The first use similarity measures to recommend recipes which are most similar to meals the user likes, while being a healthier option. These are either calculated according to ingredients [9] or user ratings. Freyne et al. [10] compared different algorithms using user ratings either for recipes or single food items confirming that decomposing recipes into ingredients is beneficial for recommendation purposes.

While user likes and dislikes are also important for the second type of recipe recommender systems, these are more concerned with the user's nutritional needs. Here recipes are recommended which fulfill most of the user's nutritional needs, indentified beforehand by health care providers [11].

The second type of food recommender systems does not recommend whole meals, but rather suggests which foods to replace with which healthier option. Mankoff et al. [12] designed a system which analyses shopping receipts and then recommends healthier food choices supplying the nutrients missing in the user's diet. Unlike recipe recommender systems, food recommender systems do not take user likes and dislikes into account.

3 Challenges for nutrition recommender systems

There are challenges to overcome in building the system itself as well as challenges with regard to the user and their interaction with the system. Considering the dependence of the system on the user, we first identify the challenges in connection with the user. Then we determine the problems concerning the algorithms used in recommender systems.

3.1 Challenges with regard to the user

One of the main challenges of food recommender systems is getting enough information about the user. The data needed to give user-tailored recommendations is:

- nutritional needs of the user
- user ratings of specific foods or recipes
- information about the user's last meals

For the first the WHO provides general guidelines[3]. In the case of individuals with special medical requirements or elderly individuals with carers, more specific guidelines can be provided by medical or care staff. For this kind of information no direct interaction with the users themselves is needed. User ratings and data about an individual user's former meals can only be procured by the user's continuing cooperation. Here are the main challenges. First:

- How to get valid information about the user's meals?
- How do you convince a user to continue providing this data?

[3] www.who.int/topics/nutrition/en/ (15.08.2011)

It is well-known that the means to record nutritional intake are not without fault. People tend to forget what they have eaten, if it is not written down right after consumption [5][6]. Many also tend to underestimate or under-report what they have consumed [13]. The systems invented to battle these issues, like FoodLog, have not matured enough to be able to give accurate nutritional information about the meals consumed, even though they are able to calculate a balance of the different kinds of food in a meal (i.e. amount of meat, starchy foods, vegetables). But even these systems encounter the problem of users ceasing to use them [8].

And second:

- How do you get enough user ratings?
- How do you keep getting user ratings?

Considering the purpose of a recommender system, which is that users heed the advice they are given, the system needs to pay attention to user likes and dislikes. Systems, which do not gather user ratings to avoid the cold-start problem, end up at the problem described above. They need information about what the user has eaten to recommend similar meals [9]. Therefore you need to collect enough user ratings for the algorithms to work properly and give good recommendations, while keeping user effort to a minimum, as it is unrealistic to expect the user to rate every food or recipe [10]. The problem of persuading the user to keep rating dishes is the same as the one of getting the user to keep reporting the food they consumed above.

The problems stated up to now deal with user input. But for a recommder system to work the user needs to actually follow the instructions given (i.e. use the recipes recommended). Therefore the challenge with regard to the system output is:

- How do you encourage user compliance?

Non-compliance (the user not following instructions or only incorrectly following them) is a huge obstacle in health care regardless the specific therapy [14]. The same holds for e-Health systems like nutrition recommender systems, which can be seen as a different form of nutritional therapy.

The challenges in designing a recommender system are therefore to get the user to use the system, keep using it and follow the recommendations given.

3.2 Challenges with regard to the algorithm

Any algorithm to calculate user-tailored nutritional recommendations needs several information:

- user information (likes, dislikes, food consumed and nutritional needs)
- a database of recipes and their nutritional information
- a set of constraints and/or rules

User information is needed to tailor recommendations to a specific user. The challenges in getting this information have already been stated above. The questions not mentioned above are:

- How to deal with the cold-start problem?
- How to deal with missing or inaccurate user data?

When a system is first used it has only the information provided beforehand. Most systems need some time to gather information about the user, while they are used. The information learned can then be used to improve the recommendations given. But even during this learning-period the system should be able to make suitable suggestions. This is called the cold-start problem and has to be addressed when designing the system.

As already mentioned, user information is often inaccurate and sometimes even missing. The system therefore needs to be able to either supplement the missing data or correct false data in one way or another.

There are two questions to be answered with regard to the database of recipes:

- How many recipes?
- How to get accurate nutritional information of the recipes?

The first one is tricky in a way as the system needs enough recipes to accommodate the likes and dislikes of many users, as well as to be able to vary the recommended recipes, while still giving recommendations in a timely manner. There is no use in a system which keeps repeating the same recommendations. Just as well user satisfaction decreases the longer a system needs to respond, decreasing the probability of continued use of the system [15].

The second one is even trickier, as different ways of cooking items results in different nutritional values. Even more difficult is to ensure, that the data gathered from nutritional tables is accurate, as comparing different tables of nutritional values of foods sometimes returns varying values for the same items.

The challenge with regard to the set of rules implemented in the system is:

- How to balance the quality of recommendations and system performance?

It seems logical that the more constraints and rules are considered the better and more user-tailored the resulting recommendations are. But when implementing these rules the possibility of contradicting rules has to be taken into account. To many constraints may also lead to performance problems, particularly with a larger number of recipes in the database, as every recipe has to be checked against every constraint. Another issue raised by a large amount of constraints is the problem of contradicting rules, preventing the algorithm to find a solution.

As the user might not be living alone another question has to be considered:

- How to give one recommendation for more than one user at the same time?

Different users have other nutritional needs as well as varying likes and dislikes. A system able to recommend recipes for more than one user (like a family) has to be able to include and accommodate the requirements of different users.

4 Dealing with the challenges

After identifying the main challenges of recommender systems, solutions for them have to be implemented and integrated into the system. Some of these problems have a seemingly more straight-forward solution. Others are obviously more interrelated. Challenges of the first group include problems like:

- the uncertainty of nutritional information of recipes or foods
- the missing or incorrect data from food recording measurements

To deal with uncertainty regarding the nutritional information of recipes or foods the help of nutrition experts is necessary. Effectively only an expert can truly determine which data is accurate and can be used. If the help of an expert is not available, an extensive comparison of multiple sources is essential to properly combine the available data.

Different models have been developed on how to deal with missing or incorrect data in food record measurements [5]. The question is which one to use, how to integrate it into the system and how this will influence the performance of the recommender system. The answer depends solely on where the emphasis of the system lies. Is it acceptable if dishes are suggested which are unfamiliar to the user and which may be too complicated for the user to prepare (as the system was unable to correctly guess or learn the user's cooking skills)? If so, missing values might not be so serious. But with regard to one of the reasons of continued recording of food intake (i.e. verification of user compliance) missing values are more severe. In such a case the system will most likely have to assume that the user actually complied, unless it receives input from other sensors (like weight-scales), which can be interpreted to show the user's compliance or non-compliance.

It can be seen that even these problems with seemingly more straight-forward solutions are dependent on other problems. So it is of no surprise that some of the other challenges identified are even more closely interrelated. Here the solution to one problem aggravates the other issue(s). Therefore there will have to be a trade-off between them. There are two groups of interrelated problems.

The first is:

- number of recipes in the database
- amount of constraints or rules used
- system performance

To increase variety in recommendations as well as to be able to heed the likes and dislikes of many different users a large amount of recipes is needed. The number of rules or constraints required rises with the complexity of the demands on the system, like recommendations for more than one meal, more than one day or more than one user. Enhancing the attractiveness of a system by providing more possible uses therefore needs a larger number of recipes as well as constraints. But this effects the system's performance negatively. More time is needed to calculate the recommendation as every recipe will have to be considered by the constraint solver.

The second group of strongly connected problems is:

- the cold-start problem
- the amount of user effort needed
- the (continued) use of the system by the user
- the information available about the user

There are two ways to solve the cold-start problem. The first is to increase the number of user ratings collected before the system is first used. This highly increases the user effort in advance. The second is to use information about the user's previous meals to calculate similarity measures to recommend new recipes [9]. This increases the user effort not only in advance but also during the use of the system as well as the need for a continued use of the system. A higher user effort however lessens the desire to use the system. Which in turn leads to a decreasing amount of information on the user available to the system. But users actually often cease to use even systems with a moderate or low amount of user effort required. It should therefore be ensured that the user can see the benefits of a continued use of the system, to increase their willingsness to invest as much time and effort into the system as needed. Giving the user constant feedback about their progress [16] as well as personalised information they can relate to their own situation [17] is a way of achieveing this goal.

Problems without a seemingly easy solution, but not as interrelated, are:

- How to ensure user compliance
- How to determine user compliance

Multifactorial approaches to ensure compliance have been suggested [14]. In the context of nutrition recommender systems these consist of providing the user with more relevant information, assessing and (if necessary) improving nutrition-related attitudes. And should a user need the help of their relatives or a care team, it is also advisable to enlist their help and support as well. User compliance can either be determined directly or indirectly. For the direct approach users record their food intake, which can then be compared with the recommendations given. This again leads to an increased user effort. On the other hand user compliance might be deduced from other health values, like the developement of the user's weight, blood sugar or activity levels, depending on the user information available by different sensors.

It can be said that there is no overall perfect solution for recommender systems. It has to be analysed thoroughly what the requirements for the recommender system to be developed are. Then the best possible solution for this specific recommender system can be decided.

5 Future work

Our future work is to develop such a recommender system for elderly people living at home in Germany. We are now analysing the problems indentified with regard to this specific population. We want to evaluate different algorithms with regard to the issues stated here and in the end implement a solution to the challenges proposed.

References

1. Statistisches Bundesamt: Mikrozensus - Fragen zur Gesundheit - Körpermaße der Bevölkerung (2009)
2. Bramlage, P., Böcking, W., Kirch, W.: Übergewicht und kardiovaskuläre Folgeerkrankungen. In Kirch, W., Badura, B., eds.: Prävention. Springer-Verlag, Berlin/Heidelberg (2006) 429–442
3. Slawik, M., Beuschlein, F.: Adipositas – Klinische Bedeutung und Behandlungsstrategien. Der Radiologe **51** (2011) 346–351 10.1007/s00117-010-2084-8.
4. Savolainen, R.: Dietary blogs as sites of informational and emotional support (December 2010)
5. Keogh, R.H., White, I.R.: Allowing for never and episodic consumers when correcting for error in food record measurements of dietary intake. Biostatistics (March 2011)
6. Thompson, F.E., Byers, T.: Dietary assessment resource manual. The Journal of Nutrition **124**(11 Suppl) (1994) 2245s–2317s
7. Elsweiler, D., Ruthven, I., Jones, C.: Towards memory supporting personal information management tools. Journal of American Society of Information Sciece and Technology **58**(7) (2007) 924–946
8. Aizawa, K., de Silva, G., Ogawa, M., Sato, Y.: Food log by snapping and processing images. In: Virtual Systems and Multimedia (VSMM), 2010 16th International Conference on. (oct. 2010) 71 –74
9. van Pinxteren, Y., Geleijnse, G., Kamsteeg, P.: Deriving a recipe similarity measure for recommending healthful meals. In: Proceedings of the 16th international conference on Intelligent user interfaces. IUI '11, New York, NY, USA, ACM (2011) 105–114
10. Freyne, J., Berkovsky, S.: Intelligent food planning: personalized recipe recommendation. In: Proceedings of the 15th international conference on Intelligent user interfaces. IUI '10, New York, NY, USA, ACM (2010) 321–324
11. Åberg, J.: Dealing with malnutrition : A meal planning system for elderly. AAAI Spring Symposium on Argumentation for Consumers of Health Care,2006 (2006)
12. Mankoff, J., Hsieh, G., Hung, H., Lee, S., Nitao, E.: Using Low-Cost Sensing to Support Nutritional Awareness. In Borriello, G., Holmquist, L., eds.: UbiComp 2002: Ubiquitous Computing. Volume 2498 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg, Berlin, Heidelberg (September 2002) 371–378
13. Cook, A., Pryer, J., Shetty, P.: The problem of accuracy in dietary surveys. Analysis of the over 65 UK National Diet and Nutrition Survey. Journal of Epidemiology and Community Health **54**(8) (August 2000) 611–616
14. Becker, M.H., Maiman, L.A.: Strategies for enhancing patient compliance. Journal of Community Health **6**(2) (December 1980) 113–135
15. Hoxmeier, J.A., D, P., Manager, C.D.: System response time and user satisfaction: An experimental study of browser-based applications. In: Proceedings of the Association of Information Systems Americas Conference. (2000) 10–13
16. DiClemente, C.C., S.Marinilli, A., Singh, M., Bellino, L.E.: The role of feedback in the process of health behavior change. American Journal Health Behavior **25**(3) (2001) 217–227
17. Campbell, M.K., DeVellis, B.M., Strecher, V.J., Ammerman, A.S., DeVellis, R.F., Sandler, R.S.: Improving dietary behavior: The effectiveness of tailored messages in primary care settings. American Journal of Public Health **84**(5) (1994) 783–787