Recommending Recipes for Balanced Nutrition

David Elsweiler
I:IMSK
University of Regensburg
93053 Regensburg, Germany
david@elsweiler.co.uk

Morgan Harvey
Mathematics and Information Sciences Department
Newcastle upon Tyne, UK
morgan.harvey@northumbria.ac.uk

Alan Said
Recorded Future
Vstra Hamngatan 24
Gteborg, Sweden
alansaid@acm.org

Bernd Ludwig
I:IMSK
University of Regensburg
93053 Regensburg, Germany
bernd.ludwig@ur.de

Abstract

In this short position statement we discuss the role of food recommenders in helping people achieve a healthy nutritional intake. We present two approaches of incorporating nutrition into the recommendation problem, which is typically formulated so that users are simply recommended food items which the system estimates they will rate highly. The first approach involves investigating the trade-off between the recipes the system estimates the user will rate the highest and a set of healthier recipes the system believes the user will still like. The second approach involves recommending plans that conform to a set of nutritional guidelines established for the user. We conclude with a brief discussion of the potential utility and limitations of both approaches.

Introduction

Lifestyle-related illness is a major problem in the modern world. A plethora of statistics reveal that we are indeed eating ourselves sick. The World Health Organisation (WHO) reports that worldwide obesity has nearly doubled since 1980 [9] and predicts that the number of obese adults worldwide will reach 2.3 billion by 2015 [9]. 347 million people worldwide have diabetes and in 2012, an estimated 1.5 million deaths were directly caused by the disease [8]. Moreover, there is a large body of evidence that both obesity and diabetes, as well as other lifestyle-related illnesses, can be prevented and sometimes even reversed through good nutrition [10, 2].

Nutritionalists can be employed to create long-term meal plans and help people to make better choices for their regular intake of food, however this solution is not really practical or economically feasible for everyone. As a result, food recommender systems have been touted as a potential means to assist people nourish themselves more healthily [1, 13]. Recommender Systems are software tools and techniques designed to provide suggestions of items which may be of use to a user [11]. In the context of food this means recipes that the user is likely to enjoy based on ratings given to other recipes in the system [1, 5].

Anecdotally it makes sense to utilise food recommenders as part of a strategy for behavioural change because if you can suggest a change that is less painful, i.e. based on something the user might like, then it is more likely that the user will accept that change and stick with it. In this sense, food recommenders are likely to be effective at predicting which changes will be painful or not as such systems have been shown to estimate with relative accuracy how a given user will rate a given recipe on a 5-point scale [1, 5], despite such rating decisions being highly context dependent [5]. If we are interested in recommending meals to provide a balanced diet, however, such systems have a major limitation: the way they work means they learn user preferences for ingredients and food styles, which, of course, leads to users who like and tend to eat fat- and calorie-laden meals being recommended fat- and calorie-laden meals - an outcome not conducive to improving nutritional habits.

In this position statement we propose two ways in which the recommendation problem can be reformulated to encompass nutritional aspects and not just user preferences. Both of the suggested approaches represent ideas we hope to pursue, meaning they are very much work-in-progress and not yet close to being polished research contributions. Nevertheless, we believe the ideas are useful and will lead to profitable discussion at the workshop. We briefly outline both problem formulations, explain what progress we have made in implementing them and conclude with a brief discussion of how the approaches compare and what we believe they mean in terms of the larger picture of using technology to support people in making healthy eating choices. The approaches presented view the problem of nutrition from the perspective of information retrieval / recommender systems, however, there are many open questions regarding the approach, which need insight from experts from other areas. As a first step we demonstrate how recipe ratings data may be helpful to discriminate those users of a large population with particular need of support.

Targeting the Correct Users

In previous work we were able to demonstrate, based on ratings data for recipes recommended in context, that users can be grouped based on the "healthiness" of the recipes they prefer [5]. We collected 4,472 ratings from 124 users over a period of 9 months and asked participants to also provide us with explanations for their ratings via a short questionnaire.

If a user, at any point during the study, justified a rating based on a health-related aspect they were grouped together. Similarly, users who never used health or healthiness as a justification were placed in a second group. Clear differences were observed between these two groups: the "health-aware" users demonstrated a strong linear trend showing that the more fat per gram or calories per gram a recipe contained, the lower the recipe was likely to be rated, while for the other group no such trend was present in the data. In fact, there was a slight trend in the opposite direction. For those users, recipes higher in fat and calories were likely to be rated higher, although the correlation was not quite as strong.

In a second study we have been analysing ratings data collected from the well-known recipe website www.allrecipes.com. In all 3.1 million ratings of 52,049 recipes from 24,719 users were collected for a time period covering nearly 14 years. Combining this data with statistics published by the County Health Rankings [1], we were again able to show strong trends in the data. For example, the recipes rated by users in the 10 most obese counties were significantly higher in terms of fat and energy content than those rated by users in the bottom 10 counties. Moreover, when one examines the users who appear in the top 25% of the population in terms both fat and energy content of the recipes they rated, these users were located in counties with higher rates of obesity. The difference in county obesity ranking between these users and the rest was highly significant.

The analyses summarised above indicate that recipe ratings data can be used to make estimates of nutritional habits and to target users who rate the recipes with the highest quantities of fat and calories, but how should we target and assist these users? The following two sections outline two potential ways of formulating the problem.

The "Want to Eat - Should Eat" Tradeoff

As the summarised results above suggest, providing the user with the recipes he is most likely to eat is perhaps not the best thing to do if we want to improve nutritional habits. This does not mean, however, that learning what a user likes is not useful. As an example, imagine we were to learn that a user prefers fatty dishes, but especially likes tomatoes. In this case it is perhaps sensible to recommend that user slightly less fatty dishes whilst giving preference to those that contain tomatoes. Similarly, if we can learn that a user values recipes that are quick and easy to prepare, perhaps we can exploit this in recommendations of less fatty meals. This means that from ratings we can determine sets of nutritionally positive and negative characteristics and employ these when recommending recipes in future.

One potential way to formulate the problem is to understand the trade-off between giving the user our best prediction of what he wants and giving him something which is healthy or at least healthier than what he is currently choosing. This could be investigated by measuring the cost incurred in terms of the rating and the benefit achieved in terms of the reduction in energy / fat content. We could operationalise this as one metric consisting of a normalised, weighted linear combination of the two scores as shown in the equation below. Here i is a given recipe, $r\hat{(i)}$ is the estimated rating for recipe i, Max($r\hat{(i)}$) is the maximum estimated rating over all recipes. $n(i)$ is the nutritional "error" incurred when recommending this recipe (relative to some ideal set of nutritional values). λ is a free parameter that we can set to suit our priorities, although λ=0.5 is probably preferable as it gives equal weighting to rating and nutrition. Note that all of these estimates are implictly conditioned on a specific user u.

$$Score(i) = \lambda \frac{\hat{r}(i)}{Max(r\hat{(i)})} + (1-\lambda) - 1 \times \frac{n(i)}{Max(n\hat{(i)})}$$

[1] http://www.countyhealthrankings.org

This approach could be operationalised in the following way. In a first step, the best state of the art prediction algorithm available would be used to estimate the top recipes for each user (i.e. recipes with predicted probability above a certain percentile). This set of recipes would be treated as a gold standard i.e. we assume no error. The next step would involve calculating the cals / fat per gram value for this set, as well as the mean predicted rating. The prediction task would then be as the follows: We want to understand how we can recommend meals with less fat or calories per gram by minimally reducing the predicted rating. The effectiveness of recommendation algorithms would be measured using the linear combination above.

Potential algorithm idea 1
Partition the recipe collection based on nutritional (fat and energy) content, i.e. create a sub-set of low-fat, low-calorie recipes to base predictions on. A simple approach would be to train a recommender on the full set of ratings, but only make predictions on the partitions with recipes with lower fat and energy content. We can try various partitions to see how this influences trade-offs.

Potential algorithm idea 2
Rank recipes in "healthier" partitions based on the similarity to those in the gold-standard set. Similarity could be measured with various distance metrics.

Potential algorithm idea 3
Modern recommender algorithms estimate ratings based on a number of biases, which tailor suggestions to individual users based on their preferences for a number of factors [4]. A more complicated model in our case may consider incorporating a number of user biases based on, for example, the preparation time or the complexity of the recipe (#number of ingredients / length of description etc.), both of which have been shown to influence the decisions of different users to different degrees [5].

We are currently setting up experiments to test these algorithms using the datasets described earlier.

Building Recipe Plans

A second approach to incorporating nutrition into the food recommendation problem is to use recommendations as a basis to algorithmically derive balanced meal plans. This means that rather than simply recommending individual meals, the task is to recommend complete meal combinations that meet nutritional guidelines for the user. We approach this problem in two stages:

In a first step, we calculate the nutritional requirements of the user based on their personal profile (gender, height, weight, level of physical activity etc.) using an updated version of the HarrisBenedict equation (also called the HarrisBenedict principle) [3], proposed by Rozal et al. [12]. This method used to estimate an individual's basal metabolic rate (BMR) and daily kilocalorie requirements. The estimated BMR value is multiplied by a number corresponding to the individual's activity level with the resulting number being the recommended daily kilocalorie intake to maintain current body weight. Nutritionalists have additional recommendations with respect to where these calories should be sourced: 45 to 65% of calories eaten should come from carbohydrates, 20% to 35% should come from fat and 10 to 35% of calories eaten should be proteins [7]. In determining the nutritional requirements for our plans, we assume that 20% of the required energy will come from drinks and between-meal snacks (fruit, confections, etc.). This is slightly lower than typical values, but nevertheless a principled and sensible target [6].

The second step is to combine recipes we believe users will like (the gold-standard set for each user as described above) in such a way that they correspond to these nutritional requirements.

We have made inital investigations regarding the feasibility of creating daily plans for a representative sample of 100 user personas (50 male, 50 female) drawn from a large sample (n=9,338) of the US population. The sample was obtained from the National Health and Nutrition Examination Survey (NHANES) project web site and reports, where for each participant, general demographic information as well as detailed information about their weights, heights, BMIs and ages at the time of the survey [7]. Activity levels and weight goals were assigned to each subject at random assuming unconditional uniform distributions for both.

Additionally we take 64 users taste profiles for which we have 10 or more ratings. These were established via our food recommendation system http://www.quizine.me. This system is an active platform for our research in this area and has provided several thousand food ratings over the last few years including the explained rating data described above. Combining each persona with each profile gives us a total of 6400 test users with a broad range of demographics and tastes and represents a rich platform to test the feasiblity of our planning approach.

We take an algorithmically simple approach to generate plans for each user (persona-profile combination). Plans are created by taking the top x recommendations for each user profile and splitting the list into breakfasts and main meals. We then perform a full search on every combination of those recipes (breakfast, main meal, main meal) to determine whether the combination meets the target nutritional requirements as establish above within a small error bound.

Our inital results show that even using this simplistic approach it is possible to derive plans for the majority of users. While it is not the case that we are able to generate plans for all combinations of persona and user profile, if we consider the top 100 recommendations, it is possible to generate plans for 4025/6400 cases (63%). It was possible to generate at least 1 plan for 57 out of the 64 user profiles. Similar to the previous approach there is a trade-off to be made between number of healty plans and the predicted ratings. Naturally, the lower the value of x, the fewer plans can be generated. Even using the top 100 recommendations there were several users for whom it was difficult or impossible to generate plans using this approach. We are currently analysing these users in more detail to establish which factors make generating plans difficult. The hope is this will allow us to derive better planning algorithms in the future.

Discussion

In this short position paper we have presented two ways of incorporating healthy nutrition into the food recommendation problem. We see the approaches as having differing potential utility in varying real-life use cases. The first approach could be utilised in the context of a food portal when users are viewing one particular recipe, in such cases the system could make healthier suggestions, perhaps in a sidebar with the header "users who enjoyed this meal also liked ..." or "lower fat alternatives that you might also like to try are ..."

The second approach requires more discipline from the user in terms of adhering to constructed plans. If plans are derived in such a way that they consist of recipes that the user actually likes and are within their abilities to easily prepare, then such plans may be a useful means to support dieting i.e. a deliberate attempt to eat in a way that will result in healthy weight loss.

Although we believe the approaches have potential future utility, there are a number of open issues with both. In particular, we wish to highlight that our planning algorithm currently does not account for a number of potentially important factors such as the suitability of combining meals, ingredients, cooking time or food styles. There is no guarantee that just because a user would like three separate meals individually, that the combination of these three would make an appealling meal plan. Future users studies are required to establish what actually makes an apealling plan. Equally, or perhaps even more important, is that just because the meal plans meet high-level nutritional guidelines in terms of fat, calories and protein, it does not automatically follow that the plan is healthy or balanced. We would very much like to work together with nutritional experts on this issue.

Summary and Conclusions

In this position statement we have discussed recommender systems - software tools commonly studied in the information retrieval and recommender systems communities - in the context of healthy nutrition. Although recommender systems have previously been proposed as useful tools for helping people achieve a balanced diet, past work has focused purely on estimating what dishes people will like. Here we have outlined two ways in which the recommendation problem can be reformulated to incorporate aspects of healthy nutrition and demonstrated how these approaches may be implemented. Additionally, we showed that the recipe ratings data supplied by users of a recommender system can be used to highlight users who may benefit most from technical assistance.

References

[1] Freyne, J., and Berkovsky, S. Intelligent food planning: personalized recipe recommendation. In *15th Int. Conf. on Intelligent User Interfaces*, IUI '10, ACM (New York, NY, USA, 2010), 321–324.

[2] Group, D. P. P. R. Reduction in the incidence of type 2 diabetes with lifestyle intervention or metformin. *New England J. of Medicine 346*, 6 (2002), 393–403.

[3] Harris, J. A., and Benedict, F. G. A biometric study of basal metabolism in man. *Publication no 279, Carnegie Institution of Washington* (1919).

[4] Harvey, M., Carman, M. J., Ruthven, I., and Crestani, F. Bayesian latent variable models for collaborative item rating prediction. In *Proc. CIKM'11*, ACM (2011), 699–708.

[5] Harvey, M., Ludwig, B., and Elsweiler, D. You are what you eat: Learning user tastes for rating prediction. In *String Processing and Information Retrieval*, Springer (2013), 153–164.

[6] Livescience. 25 http://www.livescience.com/14769-snacking-calories-increase.html, Last Accessed, Dec,2014, August 2014.

[7] NHANES. National health and nutrition examination survey report. http://www.cdc.gov/nchs/nhanes.htm, Last Accessed, Dec,2014, August 2014.

[8] Organisation, W. H. Diabetes fact sheet no. 311. http://www.who.int/mediacentre/factsheets/fs312/en/, Last Accessed, Dec,2014, November 2014.

[9] Organisation, W. H. Obesity and overweight fact sheet no. 311. http://www.who.int/mediacentre/factsheets/fs311/en/, Last Accessed, Dec,2014, August 2014.

[10] Ornish, D., Brown, S., Billings, J., Scherwitz, L., Armstrong, W., Ports, T., McLanahan, S., Kirkeeide, R., Gould, K., and Brand, R. Can lifestyle changes reverse coronary heart disease?: The lifestyle heart trial. *The Lancet 336*, 8708 (1990), 129 – 133.

[11] Ricci, F., Rokach, L., Shapira, B., and Kantor, P., Eds. *Rec. Systems Handbook*. Springer, 2011.

[12] Roza, A. M., and Shizgal, H. M. The harris benedict equation reevaluated: resting energy requirements and the body cell mass. *The American journal of clinical nutrition 40*, 1 (1984), 168–182.

[13] Suksom, N., Buranarach, M., Thein, Y. M., Supnithi, T., and Netisopakul, P. A knowledge–based framework for development of personalized food recommender system. In *The Fifth International Conference on Knowledge, Information and Creativity Support Systems* (2010), 25–27.