

Bringing the “healthy” into Food Recommenders

David Elsweiler¹, Morgan Harvey², Bernd Ludwig¹, and Alan Said³

¹ University of Regensburg, Germany
david@elsweiler.co.uk,bernd.ludwig@ur.de

² University of Northumbria, UK
morgan.harvey@northumbria.ac.uk

³ Recorded Future, Sweden
alansaid@acm.org

1 Introduction

Lifestyle-related illnesses such as diabetes and obesity are a major problem in the modern world but can be prevented and sometimes even reversed through good nutrition [3]. Food recommender systems have been touted as a potential means to assist people nourish themselves more healthily e.g. [2, 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 then it seems more likely that the user will accept that change and stick with it. 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 briefly outline two ways in which the recommendation problem can be reformulated to also encompass nutritional aspects and not just user preferences.

2 The “Want to Eat - Should Eat” Trade-off

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. 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 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 , $\text{Max}(r(\hat{i}))$ 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 $\lambda=0.5$ is probably preferable initially as it gives equal weighting to rating and nutrition. Note that all of these estimates are implicitly conditioned on a specific user u .

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

This approach could be implemented 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 calories / fat per gram value for this set, as well as the mean predicted rating. The prediction task would then be to recommend meals with less fat or calories per gram by minimally reducing the predicted rating. The effectiveness of recommender algorithms would be measured using the linear combination above.

2.1 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.

2.2 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.

2.3 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.

3 Building Recipe Plans

A second approach 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:

1) 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, proposed by Rozal et al. [6], estimates an individual's basal metabolic rate (BMR) and daily kilocalorie requirements. Other guidelines are used to calculate the proportion of calories that should be sourced from fats, proteins and carbohydrates etc. 2) 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.

Initial investigations have shown that it is possible to create plans for a large percentage of potential users with diverse calorie needs and eating preferences [1]. We took an algorithmically simple approach to generate plans for each user by taking the top x recommendations for each user profile and splitting the list into breakfasts and main meals. We then performed 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.

We also examined the properties of users for which making plans was more difficult (few plans were possible, or plans had very low ratings). Challenging users had high calorie requirements or extremely non-diverse food tastes. To address such cases, we are developing more complicated algorithms that allow for adjusting portion sizes or different combinations of meals in a day. Further details can be found in [1].

4 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 appealing meal plan. Future users studies are required to establish what actually makes an appealing 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.

5 Summary and Conclusions

In this position statement we have discussed recommender systems 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. David Elsweiler and Morgan Harvey. Towards automatic meal plan recommendations for balanced nutrition. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 313–316. ACM, 2015.
2. J. Freyne and S. Berkovsky. Intelligent food planning: personalized recipe recommendation. In *15th Int. Conf. on Intelligent User Interfaces, IUI '10*, pages 321–324, New York, NY, USA, 2010. ACM.
3. Diabetes Prevention Program Research Group. Reduction in the incidence of type 2 diabetes with lifestyle intervention or metformin. *New England J. of Medicine*, 346(6):393–403, 2002.
4. M. Harvey, M. J. Carman, I. Ruthven, and F. Crestani. Bayesian latent variable models for collaborative item rating prediction. In *Proc. CIKM'11*, pages 699–708. ACM, 2011.
5. Morgan Harvey, Bernd Ludwig, and David Elsweiler. You are what you eat: Learning user tastes for rating prediction. In *String Processing and Information Retrieval*, pages 153–164. Springer, 2013.
6. Allan M Roza and Harry M Shizgal. The harris benedict equation reevaluated: resting energy requirements and the body cell mass. *The American journal of clinical nutrition*, 40(1):168–182, 1984.