User Nutrition Modelling and Recommendation - Balancing Simplicity and Complexity

Hanna Schäfer*
Technical University of Munich

Mehdi Elahi
Free University of Bozen-Bolzano

David Elsweiler
University of Regensburg

Georg Groh
Technical University of Munich

Morgan Harvey
Northumbria University

Bernd Ludwig
University of Regensburg

Francesco Ricci
Free University of Bozen-Bolzano

Alan Said
University of Skövde

ABSTRACT

In order to use and model nutritional knowledge in a food recommender system, uncertainties regarding the users nutritional state and thus the personal health value of food items, as well as conflicting nutritional theories need to be quantified, qualified and subsumed into falsifiable models. In this paper, we reflect on different error sources with respect to nutrition and consider how such issues can be tackled in future systems. We discuss the integration of general nutritional theories into information systems as well as user specific nutritional measures and different approaches to evaluating the utility of a given nutritional model.

KEYWORDS

Nutrition Recommender Systems; Health Informatics; Patient Modeling; Nutrition Modeling; Online Nutrition Interventions;

1 INTRODUCTION

Food Recommender Systems have been touted as a useful tool in combating global diet-related epidemics, such as diabetes and obesity [2, 7, 15, 21, 25]. Developing such systems is, however, for many reasons particularly challenging including the need to select and use multiple nutritional and preference models. The contrasting goals of providing users with what they want and what they should eat has been discussed [2], however, the additional complexity of modelling divergent nutritional opinions has yet to be addressed.

In user modelling, one of the core problems is a lack of reliable input (e.g. what the user eats) or contextual data (e.g. why, how and when something is being eaten). To model the nutritional behavior of a user, we must often deal with missing and inaccurate data, incomplete and/or contradictory models, as well as missing output parameters. Here we suggest that such error sources can be diminished by understanding why nutritional models are so complex and how we can take them one step closer to becoming both clear and algorithmically-processable. This effort has led to several approaches which will be discussed in the following sections.

*Email: hanna.schaefer@tum.de

UMAP'17 Adjunct, July 09-12, 2017, Bratislava, Slovakia

ACM ISBN 978-1-4503-5067-9/17/07…$15.00
https://doi.org/10.1145/3099023.3099108

2 STATE OF THE ART IN NUTRITION MODELLING

To understand nutritional models, we should first understand the different components of nutritional data. A nutritional model consists of three interacting core components. The first model, the *user model*, includes continuous or approximated intake tracking of content and volume, food preferences, availability, social context and many other contextual factors. Second, the *food items* themselves must be modelled on a nutritional level to understand their nutrient content or, in case of complex items such as recipes, their subcomponents, including their culinary functions and cultural roles. Finally, the combination of user models and the food item models must be mapped onto an array of target values e.g. with a rule system or *algorithmic model*. The most complex among these targets is the health utility defined by, for example, long or short term risks and benefits. Others targets may include taste, cost, sustainability, or availability. Furthermore, in the course of analyzing and learning these complex models, we need to define ways to *evaluate* the error propagation, continuity properties, or stability of the resulting models in a reproducible way.

2.1 Nutritional User Information

The user model is prone to many sources of error and uncertainty, originating from measurement error, misinterpretation, misperception and a lack of knowledge about the user. Many of these errors relate to a lack of compliance from the user, reflecting the level of effort involved in providing the data required to establish an appropriate model and keep it up to date (e.g. by wearing or using sensors, keeping a food diary or describing goals). Below we illustrate sources of error relating to different aspects of the user's nutritional information and means to limit these. We discuss intake tracking, portion size estimation, the use of sensors, user preferences and user context.

Intake Tracking. The most frequently used tools in measuring nutritional intake are food frequency questionnaires (FFQs), consumption recalls and food records. While FFQs and recalls retrospectively collect data about the last month/week/day, food records are prospectively collected at the time of consumption [23]. Both approaches have drawbacks: retrospective methods suffer from inaccurate memories and generalization; prospective methods rely on significant day-to-day user effort. Both also suffer from active or subconscious under- or over-reporting, estimation or perception errors and inadequate communication tools [3].

Estimation of Portion Sizes. While users are often aware of the nature of what they have eaten, estimation of portion sizes is a key error source in dietary tracking. Errors result from a limited knowledge of the processing of the food and a lack of ability to estimate weight or size. A further source of error is the use of non-standardized measures for cooking recipes e.g. one carrot. The variances are reported to be around 50% for the estimation of food intake and 20% for nutrients [18]. Even with modern sensing techniques such as image-based volume estimation or 3D modelling of food, the error is still around 20% [27].

Learning from Sensors. One approach to tackling such inaccuracies in tracking is to build new and more advanced ways of sensing. For example, to sense food with images and predict type and volume [14] or to sense explicit information given e.g. by barcodes. Other potential uses of such technology include smart vessels and other utensils that measure food intake [10]. Intake can also be measured post-consumption by means of biological sensors, which can, for example, continuously measure blood sugar [5].

User Preferences. Eliciting accurate user preferences involves overcoming a number of well-documented challenges. 1) The cold start problem, where, upon initial use, almost no information on user preferences is available [22]. 2) individual food preferences are highly context-dependent and can change over time [7]. 3) Despite some components of the preference model being easy to attain, such as no meat for a vegan, user preferences when choosing between similar or complex food items will be more challenging [24]. 4) The preferences of a user are not equally informative of her taste, and hence, it is important to carefully design a proper preference elicitation strategy for identifying and obtaining a small set of preferences which still conveys the most valuable information.

User Context. The preferences of a user are highly dependent on the current context [8]. For example a family provider might have a different preference cooking for the family than when going for dinner in a restaurant. Also available time can place constraints on the decision process. Any such dependencies would need to be detected either by direct user input or by sensing the environmental variables, such as location, time or short term social context. The variety and number of contexts makes it hard to measure all available data, which in turn introduces new uncertainty levels for the intent and preference of the user.

2.2 Nutritional Food Information

It is extremely challenging to objectively obtain the nutritional properties of consumed food. There are multiple available datasets, which define the nutritional composition of food items, either measured or approximated. Naturally, these datasets contain inconsistencies since the individually-measured food items may be nutritionally divergent. Sensors can help to achieve this kind of individual precision, e.g., using bar-code scanners for standardized items or chemical sensors that measure the composition of a food item. Several databases are available that allow real world food items to be associated with their nutritional and other properties.

Nutritional Databases. Depending on the level of detail required, the nutritional information may be accessible for different types of food items. For example, in the EU every food product must be labeled with its macro-nutritional content (e.g. sugar). Such data is often widely available but quite coarse. More detailed information can be found in nutritional databases such as the BLS, NEVO or USDA [1], where, on the other hand, concrete (e.g. branded) product data is often not represented at all.

Product Databases. In addition to the nutritional content, other data on a food item may be unavailable in a general database, such as the price or availability. Including such parameters into the nutritional model thus requires a mapping between *general food items*, e.g. apple, *nutritional food items*, e.g. raw, Gala apple and the specific *food item product*, e.g. 500g package of organic Gala apples from Spain.

Recipe Databases. From the perspective of a user, food is rarely consumed as a single ingredient. Most interaction with food is based on recipes. To understand the internal structure of recipes and the role of its subcomponents it is necessary to find similar recipes or even replacements for subcomponents within these recipes. Additionally, most recipes are not formally linked to any of the abovementioned information sources. This results in a need to match each individual ingredient not only with the appropriate entry in the database [16], but also with the processed state of the ingredient (e.g. whether it is cooked or raw).

Taste and Sensory Properties. To build taste models of food items, different approaches might be implemented. In theory, a sensory effect could be related to the molecular setup of these items. While a known configuration of molecules create a reproducible sensory impression in certain users, not all users might react in the same way. Furthermore, not all such senso-chemical relations are known, since the analysis of food items is very time consuming. A much easier approximation would be a user-based tagging of recipes with certain labels, e.g. "soup", "hot", "vegan", etc. in order to find a correlation between the labels and user preferences.

2.3 Modelling the Utility of Food for Users

The most complex utility model is the relation between user, item and a measure of healthiness. There are various guidelines, e.g. the DGE [9], that inform consumers, on a rule basis, how to eat correctly. These guidelines have several drawbacks from a modelling perspective. Firstly, they are not always conclusive in their evaluation metric and secondly, depending on the organization, the reference values might have differences and contradictions, due to the missing conclusive validation of these values.

[1] http://www.eurofir.org/food-information/food-composition-databases-2/

Nutrition guidelines. Nutritional guidelines are available from different sources and on different evaluation levels [11]. The highest level used by one-size-fits-all guidelines, e.g. the DGE [9], are *food groups* organized, e.g., in food pyramids. Here 5-7 food categories are connected with a guideline rule such as "5 portions of fruit and vegetables per day". This approach does not offer the potential to model intake continuously or to infer to what extent a user's health state deviates from the ideal at a given time point. The next level is to evaluate the *energy balance*, where the necessary calorie intake is determined e.g. by the basal metabolic rate multiplied with an physical activity factor [3]. For *macro-nutrients*, the energy intake is further separated into different energy sources, e.g. fat or carbohydrate. Healthy nutrition is then defined as a balance between the different energy sources, such as 7-20% of energy from protein [13]. A simple model would be to combine total energy and macro-nutrient based energy intake guidelines, since such data is widely available in food item databases. Additionally, these macro-nutrient balances are largely invariant to processing such as cooking. On the other hand, such a model could not be used to account for many deficiencies in important micro-nutrients such as *vitamins* or *minerals and traces*. While these are important for many body functions, some also have toxic effects when consumed excessively [3]. When modelling these nutrients, *dietary reference intake standards*, e.g. [20], are an important source of information to build a health model. They represent the intake as a Gaussian distribution between the lower deficiency border and the upper toxicity border. Unfortunately, depending on the source, these distribution and their border values can differ.

Personalization. Another approach is to avoid the uncertainties in the literature by personalizing the healthiness functions with *phenotype* and even *genotype* information [17]. This approach is hindered by an as-yet immature understanding of the relation between genetics and nutrition [19] and the high cost of analyzing genetic measures. Additionally, some ethical concerns were raised regarding the usage of genetic information in this context [4].

Uncertainty modelling. Two types of uncertainty need to be modelled: 1) uncertainty in user- or food-specific input data and 2) uncertainty from a lack of scientific knowledge or a lack of adequate modeling of existing knowledge. The former can be handled by building target functions that can deal with probabilistic input parameters. The latter can be solved by, for example, employing fuzzy goal functions [1, 26]. Improvements could be achieved with advanced machine learning techniques and probabilistic models.

User Intent. A user's intent in a specific situation is based on the priorities of a variety of goals. One obvious goal is the desire to eat tasty food, in the sense of food matching one's own taste preference. Others include the desire to eat healthily or in a sustainable manner. Some goals may require more context, for example, to optimize the time requirements for shopping and preparing the food. The array of such weighted goal functions create the user's current intent. Since the goal functions might be contradictory, they need to be weighted by the user's current, context-depended priorities, in order to find a compromising optimal solution.

2.4 Quality Assessment of Nutritional Models

Finally, many of the contradicting or inconclusive nutritional guidelines are due to missing validation parameters. Traditionally, correlation with diseases such as cancer, heart attacks and similar NCDs are referenced. These are highly dependent on other influence factors and require long-term studies with large sample sets. A more objective measure is the age of death of study participants, which is easy to measure but hard to correlate to nutritional information only. Modelling the nutritional benefits on a day-to-day level generates new requirements for the evaluation of such models. Instead of searching for a true health model, e.g. guaranteed maximal cancer prevention, the digital nutritional models should prove to be stable against uncertain input, flexible against inaccurate nutritional correlations and minimize potential harm to the user. Thus to evaluate such a digital meta model, stability, error propagation, continuity and similar metrics need to be investigated.

3 VISIONARY SCENARIOS

It is clear from the summary above that a complete and sound nutritional model for recommender systems is some way off. We conjecture, however, that combining the insights offered by existing models in intelligent ways, can lead to the limitations being circumvented to a degree where the user will not suffer from them. Below we discuss pragmatic but intelligent solutions that aim at reaching this balance, such as contextual constraints, uncertainty modelling, learning of good defaults, suggestive recommendations and integrating the user in-the-loop. Each has been used in one or another version before, but needs to be developed further to close the gap with the conjectured complex models discussed above.

3.1 Context as a Constraining Parameter

One option to improve the accuracy of the system's nutritional advice is to put constraints on the set of options being optimized. For example, if the menu options in a restaurant or a canteen are known to the system, the task could be defined as finding the best combination(s) of the available items. Compared to finding the perfect advice, this task can be solved with much higher certainty. The same kind of context can be generated in other situations, such as shopping for groceries. Here, a given grocery list could provide the necessary context, to find the best replacement for certain critical items, instead of completely generating the list. In the future this approach should be extended by generating the missing contextual constraints from sensor data. One example would be integrating an intelligent kitchen that automatically limits the choices to fit the available tools and ingredients. Another direction would be to apply bar-codes to restaurant and canteen menus, so their data is always to hand.

3.2 Modelling Uncertainties

Another way to work with the different uncertainties occurring in a nutritional user model is to communicate both results of these uncertainties and their reasons. For example, if the uncertainty is due to missing information on the user's part, simple information about this could incentivize the user to provide more input for better results. If, on the other hand, the uncertainties are not connected to a solvable reason, the user should still be made aware of them in order to empower an informed decision. Similar to having different weather forecasts, a range of nutritional evaluations and their probabilities could be displayed to the user [6]. In the future one step would be to build a nutritional system that can explicitly model and describe these uncertainties. Additionally, the effect of such a system needs to be carefully evaluated. While having the additional information empowers the user's decision, the knowledge of uncertainty might also lessen his/her trust in the system.

3.3 Learning Default Models

Since many of the uncertainties are based on missing input, e.g. when using FFQs or recalls [3], there is a need for accurate default behavior models. Recently the retrieval and analysis of online recipe platforms has been shown to provide a lot of insight into the cooking behavior depending e.g. on culture or season [12]. In the future these defaults could be further improved by using machine learning techniques on different dietary study datasets or commercial user intake data sets such as MyFitnessPal[2]. Even learning common shopping patterns from supermarket datasets or payback customer information could increase the insight into user behavior.

3.4 Suggestive Recommendations

In some cases finding a good nutritional option for the current context is not possible. If, for example, the user is at a dinner with friends, he will not be cooking vegetables at home. In such cases the system should be able to give preventive suggestions for future situations, e.g. you don't have any veggies here, so you can eat your steak, but please go shopping for some tomatoes tomorrow. The advantage of such suggestive recommendations to assume an ideal context for the future time scope. To realize such suggestive recommendations, the systems needs to have a deep insight into the user's everyday contexts and a predictive system for the likely future contexts, e.g. balancing the recommended intake over a week.

3.5 Keeping the User-In-The-Loop

In the end, improving the accuracy and uncertainty tolerance of the model is not enough to create meaningful system output for every situation. Too many context parameters are unknown to the system. To solve this conflict, systems should integrate the user into the decision process and interactively update the output set based on the user's input. For example if a user plans to cook for friends, the system might suggest a steak night. Then the user can specify that his friends are vegetarian. In that case, the system would have to update its suggestion to the new parameters.

4 CONCLUSION FOR FUTURE WORK

In conclusion, while the modelling of a user's nutrition is complex and uncertain on many levels, some measures can help to deliver approximations that are reliable enough for most information systems building on then. By working together with domain experts and learning from each other's methodologies and experiences we believe future nutritional models may even inherently include the given uncertainties as a feature of precision instead of an obstacle.

[2]https://www.myfitnesspal.com/

REFERENCES

[1] S.B. Dias, S.J. Hadjileontiadou, J.A. Diniz, J. Barroso, and L.J. Hadjileontiadis. 2016. On Modeling the Quality of Nutrition for Healthy Ageing Using Fuzzy Cognitive Maps. In *International Conference on Universal Access in Human-Computer Interaction.*
[2] D. Elsweiler and M. Harvey. 2015. Towards automatic meal plan recommendations for balanced nutrition. In *Proceedings of the 9th ACM Conference on Recommender Systems.*
[3] M.J. Gibney, H.H. Vorster, and F.J. Kok. 2002. *Introduction to human nutrition.*
[4] U. Görman, J.C. Mathers, K.A. Grimaldi, J. Ahlgren, and K. Nordström. 2013. Do we know enough? A scientific and ethical analysis of the basis for genetic-based personalized nutrition. *Genes & nutrition* (2013).
[5] D.A. Gough, L.S. Kumosa, T.L. Routh, J.T. Lin, and J.Y. Lucisano. 2010. Function of an implanted tissue glucose sensor for more than 1 year in animals. *Science Translational Medicine* (2010).
[6] M. Greis, E. Avci, A. Schmidt, and T. Machulla. 2017. Increasing Users' Confidence in Uncertain Data by Aggregating Data from Multiple Sources. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems 2017.*
[7] M. Harvey, B. Ludwig, and D. Elsweiler. 2012. Learning user tastes: a first step to generating healthy meal plans. In *First international workshop on recommendation technologies for lifestyle change (lifestyle 2012).*
[8] H. Jung and K. Chung. 2016. Knowledge-based dietary nutrition recommendation for obese management. *Information Technology and Management* (2016).
[9] A. Jungvogel, I. Wendt, K. Schäbethal, E. Leschik-Bonnet, and H Oberritter. 2013. Überarbeitet: Die 10 regeln der dge. *Ernährungs Umschau* (2013).
[10] A. Kadomura, C.-Y. Li, K. Tsukada, H.-H. Chu, and I. Siio. 2014. Persuasive technology to improve eating behavior using a sensor-embedded fork. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing.*
[11] A.K. Kant. 1996. Indexes of Overall Diet Quality: A Review. *Journal of the American Dietetic Association* (1996).
[12] T. Kusmierczyk, C. Trattner, and K. Nørvåg. 2016. Understanding and Predicting Online Food Recipe Production Patterns. In *Proceedings of the 27th ACM Conference on Hypertext and Social Media.*
[13] E. Lee, J. Choi, A. Ahn, E. Oh, H. Kweon, and D. Cho. 2015. Acceptable macronutrient distribution ranges and hypertension. *Clinical and Experimental Hypertension* (2015).
[14] A. Meyers, N. Johnston, V. Rathod, A. Korattikara, A. Gorban, N. Silberman, S. Guadarrama, G. Papandreou, J. Huang, and K.P. Murphy. 2015. Im2Calories: towards an automated mobile vision food diary. In *Proceedings of the IEEE International Conference on Computer Vision.*
[15] S. Mika. 2011. Challenges for nutrition recommender systems. In *Proceedings of the 2nd Workshop on Context Aware Intel. Assistance, Berlin, Germany.*
[16] M. Müller, M. Harvey, D. Elsweiler, and S. Mika. 2012. Ingredient matching to determine the nutritional properties of internet-sourced recipes. In *2012 6th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth) and Workshops.*
[17] D.M. Mutch, W. Wahli, and G. Williamson. 2005. Nutrigenomics and nutrigenetics: the emerging faces of nutrition. *The FASEB journal* (2005).
[18] M. Nelson and S.A. Bingham. 1997. 6. Assessment of food consumption. *Design concepts in nutritional epidemiology* (1997).
[19] K. Nordström, N. Juth, S. Kjellström, F.L.B. Meijboom, U. Görman, and Food4Me Project. 2013. Values at stake: autonomy, responsibility, and trustworthiness in relation to genetic testing and personalized nutrition advice. *Genes & nutrition* (2013).
[20] J.J. Otten, J.P. Hellwig, L.D. Meyers, et al. 2006. *Dietary reference intakes: the essential guide to nutrient requirements.*
[21] H. Schäfer, G. Groh, J.H. Schlichter, S. Kolossa, H. Daniel, R. Hecktor, and T. Greupner. 2015. Personalized Food Recommendation.. In *DMRS.*
[22] A.I. Schein, A. Popescul, L.H. Ungar, and D.M. Pennock. 2002. Methods and metrics for cold-start recommendations. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval.*
[23] A Straßburg. 2010. Ernährungserhebungen - Methoden und Instrumente. *Ernährungs Umschau* (2010).
[24] C.-Y. Teng, Y.-R. Lin, and L.A. Adamic. 2012. Recipe recommendation using ingredient networks. In *Proceedings of the 4th Annual ACM Web Science Conference.*
[25] N. Terzimehić, N. Leipold, H. Schäfer, M. Madenach, M. Böhm, G. Groh, and K. Gedrich. 2016. Can an Automated Personalized Nutrition Assistance System Successfully Change Nutrition Behavior?-Study Design. (2016).
[26] B. Wirsam, A. Hahn, E.O. Uthus, and C. Leitzmann. 1997. Fuzzy sets and fuzzy decision making in nutrition. *European journal of clinical nutrition* (1997).
[27] Z. Zhang, Y. Yang, Y. Yue, J.D. Fernstrom, W. Jia, and M. Sun. 2011. Food volume estimation from a single image using virtual reality technology. In *Bioengineering Conference (NEBEC), 2011 IEEE 37th Annual Northeast.*