

Recommending Better Food Choices for Clinically Obese Users

Bernd Ludwig
Selina Meyer
Julian Dietz
Gregor Donabauer
Andreas Pfaffelhuber
David Elswailer

Bernd.Ludwig@sprachlit.uni-regensburg.de
University of Regensburg Chair for Information Science, University Regensburg
Regensburg, Germany

ABSTRACT

In this demo, we present an interactive recommender system that suggests recipes to participants of a weight loss programme. Nutritional constraints imposed by the programme serve as initial information to tackle the crucial cold start problem.

CCS CONCEPTS

• **Applied computing** → **Health care information systems.**

KEYWORDS

food recsys, nutrition, obesity, user interface

ACM Reference Format:

Bernd Ludwig, Selina Meyer, Julian Dietz, Gregor Donabauer, Andreas Pfaffelhuber, and David Elswailer. 2021. Recommending Better Food Choices for Clinically Obese Users. In *Adjunct Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '21 Adjunct)*, June 21–25, 2021, Utrecht, Netherlands. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3450614.3461677>

1 INTRODUCTION

An abundance of food and inability for people to make good choices have led to personalised food recommender systems being proposed as a useful tool in the fight against obesity [3–7, 13, 14]. One approach that has been shown empirically to hold promise is the idea of post-filtering recommendations to push healthier choices [14]. Nevertheless, like many ideas in the literature, this has only been tested using simulated analyses with a collection of online recipes. Here, we present a system, designed for a specific group of **real** users, who suffer from obesity and are under medical supervision as part of the OPTIFAST52 programme. Programme participants will use the recommendation engine in practice, during and after treatment. We explain the challenges for such a system as well as our solutions.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

UMAP '21 Adjunct, June 21–25, 2021, Utrecht, Netherlands

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8367-7/21/06.

<https://doi.org/10.1145/3450614.3461677>

2 THE OPTIFAST52 PROGRAMME

OPTIFAST52 is a multidisciplinary weight loss programme primarily directed at those suffering from clinical obesity. 12 weeks of meal replacement are followed by transition and stabilisation phases, where home cooked meals are gradually introduced and increased to enable participants to maintain their new weight. Participants learn and stabilise healthy behaviour and are continually supported in managing their long-term weight goals by a team of experts from different disciplines as well as support groups [1, 2, 8]. After the meal replacement phase, patients are introduced to a point system (OPTIFAST score) consisting of green (carbohydrates), yellow (meat or dairy) and red (fat) points to help nutritional choices. The system aligns with recommendations from the German Nutrition Association (DGE) [11], focusing on vegetables and fruit which are low in fat and carbohydrates, while allowing smaller amounts of calorie dense foods. The score allocated per day is initially very low but increases over the course of the transition phase. Although slow, this transition still causes participants to regain some of the weight they have lost during phase one [15]. The approach has nevertheless been shown to lead to significant weight loss and increased quality of life, as well as a decrease in comorbidities (e.g. hypertension and metabolic syndrome) [1, 2]. Large-scale, longitudinal studies have revealed, however, that programme adherence is low (58%). Dropouts often occur at the transition phase where patients naturally started to regain weight, which affects motivation. Patients unable to complete the programme typically have lower scores on autonomy orientation and internal health locus of control than those who finish [15].

3 SUPPORTING OPTIFAST USERS

In collaboration with the medical supervision team at the OPTIFAST obesity centre at our university hospital we developed a web application to mitigate high dropout rates and boost motivation by empowering patient food choice. The design of the described features were motivated by an interview study conducted with patients who took part in the OPTIFAST course in 2019 [9]. The conclusions from the interviews were supported by the nutrition experts in three different German OPTIFAST centres who have received feedback from course participants for more than ten years. The current version of the system simplifies activities that are integral to the treatment, but difficult and tedious for patients. For example, participants are required to **document the food they**

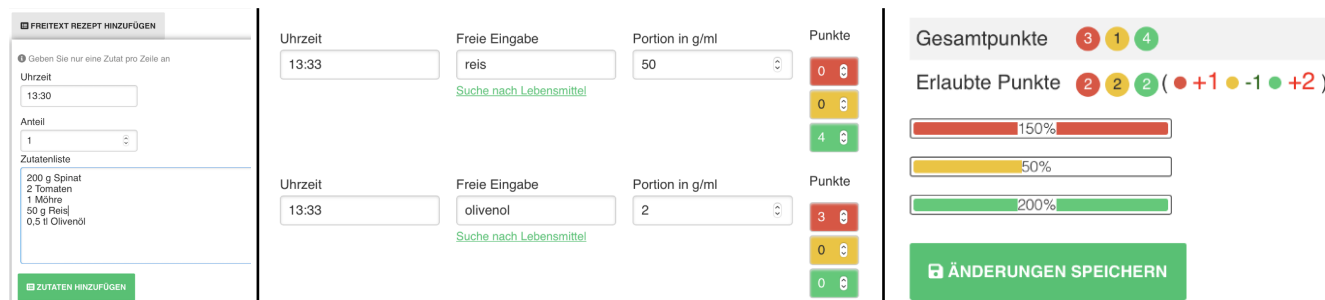


Figure 1: The screenshots on the left shows a self-edited recipe, in the middle the calculated nutritional value and OPTIFAST score for two of the ingredients are displayed, while on the right the complete recipe is scored. In this case, the recipe is beyond the threshold as the red and green points are too high.

consume. Our system provides a food diary interface, which simplifies this process. Moreover, every food choice during the programme necessitates that patients calculate the score associated with a food-item, which means analysing (and often modifying) dishes. This is a cognitively challenging and time-consuming task that patients struggle with. The result is that they rely on the few sample recipes included with OPTIFAST instructions, which are not adapted to participants' individual food preferences and quickly become repetitive. We provide a NutriCalc interface, based on our past work [10], that **automatically estimates the nutritional content and OPTIFAST score for foods and recipes.**¹ This allows patients to experiment and learn about the nutritional value of foods and the impact they have on their diets. A visual overview of these interfaces and the recommendation interface is provided in supplementary material.

We use the output of NutriCalc for a context-aware recommender system, which aims to broaden patient food choices by recommending personalised recipes from the web. Our current database comprises the 10,000 most popular recipes out of 400,000 on the German recipe portal <https://www.kochbar.de>. Our main goal is to implement a ready-to-use system with basic components to collect a corpus of recipes rated by OPTIFAST participants in each phase. The data will help to train a recommender that can take the particular situation of the weight loss programme into account.

The recommender engine is implemented via librec (<https://pypi.org/project/librec/>). In order to build a first baseline system, we calculate recommendations using the SVDpp algorithm for context-aware recommendations. In the UI, each recommendation is labelled with a marker indicating whether the ingredients' calculated OPTIFAST scores are within the threshold for the user's current OPTIFAST phase (see figure 3). In this way, **recommendations are personalised by recipe ratings and nutritional validity given current dietary requirements.** Other constraints, such as allergic reactions to ingredients, are not addressed in the current version of the tool, but will be in the future when we have collected a sufficient amount of naturalistic data.

The current implementation achieves nutritional validity by post filtering results of SVDpp. This is intended as a baseline that can be



Figure 2: Users are recommended recipes based on their preferences. Recipe choices are presented as a carousel. Recipes can be post-filtered depending on the amount of OPTIFAST points available.

used actively in OPTIFAST course and delivers data valuable for improvements. Our tool currently provides four ways for participants to communicate their preferences:

1. Via the food diary function.
2. They can provide explicit feedback for recommended recipes: In our UI, recommendations are presented sequentially in a carousel. From the user interactions (acceptance or rejection of recommendations) we elicit preferences for ingredients (see figure 2).
3. via self-edited recipes (see figure 1).
4. They can enter ingredients as a means to search or filter recipes.

Independently of how preferences are communicated to our tool, from the retrieved data it estimates a probability distribution for an ingredient to be a component of a recipe rated highly or even be cooked by a user. A similar distribution (the probability of an ingredient to be needed for a recipe) is calculated for each recipe in our data set. On this basis, recipes that are similar to the preferences of a user (fourth way above) are identified (currently using the cosine heuristics). Using such an ordering for recipes with respect to known user preferences, we apply Bayesian Personalized

¹The NutriCalc component converts an ingredient's name and given quantity into a vector of nutritional values. It is one of the few systems for German language foods and is available for other researchers via an API <http://smarthome.ur.de/naehrwertrechner/>

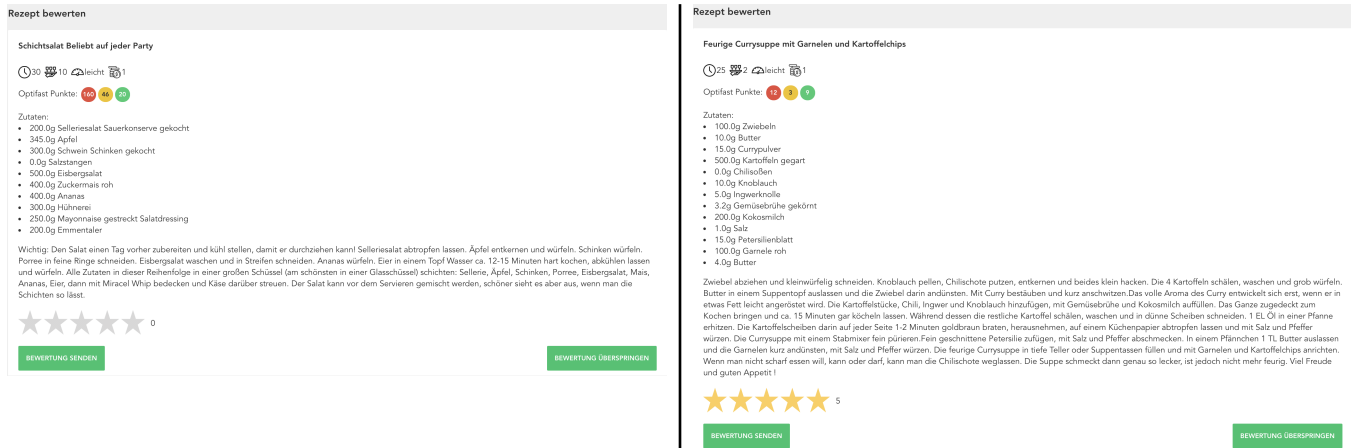


Figure 3: In the food recommender component, users are first prompted to rate recipes between zero and five stars based on their preferences. These ratings, together with the amount of available OPTIFAST points, are taken into account for recipe recommendations.

Ranking [12] for computing recommendations for individual users based on their current nutrition behaviour.

Our recommender has a special version of the cold start problem to tackle: First, each time a new course starts, a group of new users, without a model of eating preferences is created. Second, many of the recipes are not suitable for early phases of OPTIFAST with top level dropout rates as their constraints on acceptable scores are really tight and therefore dramatically limit the recipes that can be recommended without modifying their ingredients somehow.

In order to illustrate the challenging nature of recommendation problem, we analysed the Kochbar.de corpus subset. For each recipe, we calculate the OPTIFAST score. As the ingredient lists for the recipe are user generated content, they are not standardised in terms of names for ingredients and metric units. As a consequence, NutriCalc cannot calculate scores for each lists entry. Therefore, we filter out those recipes for which the score can be calculated for at least 95% of the ingredients. In table 1 we list how many recipes are applicable for the various phases during an OPTIFAST course and the map@5 and ndcg@5 scores computed using librec’s SVDpp algorithm. We note that in particular in the first transition phase (week 13) only few choices exist for OPTIFAST participants. This is nevertheless far more recipes than are available via the programme itself. The data show that predicting rankings based on the available recipe ratings is far from accurate, but largely improves as the number of available recipes grows. For the whole data set, map@5 is 0.0988, and ndcg@5 is 0.1185. Similar results are reported in [14] for the problem of food recommendation.

4 CONCLUSION AND OUTLOOK

Our system is in an early evaluation stage. It is inspired by requirements expressed by former participants, covers all features of the current paper-based OPTIFAST documentation digitally, and provides additional features that are intended to avoid the high dropout rates in the early transition phase of the programme. The system architecture is modular and allows for additions of new features which implement different or additional motivational strategies.

Phase	# recipes	map@5	ndcg@5
transition 1	299	0.0084	0.0107
transition 2	482	0.0053	0.0069
transition 3	593	0.0069	0.091
transition 4	727	0.0175	0.0203
stabilisation 1	877	0.01	0.0126
stabilisation 2	1010	0.01	0.012
stabilisation 3	1117	0.01	0.012
stabilisation 4	1240	0.008	0.01
stabilisation 5	1346	0.0124	0.015

Table 1: # of recipes among <https://www.kochbar.de>’s 10,000 most popular recipes for specific OPTIFAST phases

The current system implements two strategies: (1) recommending phase-specific recipes for an extended range of food choices, and (2) monitoring and visualising nutrition behaviour as a way of feedback that eventually leads to behaviour change.

Our next steps in further developing the framework will be the integration of a chat function in order to ease the logging and to provide natural language explanations for the OPTIFAST score calculated for a recipe. In this way, we aim to gain a better understanding of barriers for behaviour change. We are interested in validating the results of our initial interview study [9], defining strategies to overcome well understood barriers, and integrating them into our tool. By doing so, we can achieve a maximum degree of personalization in applying policies for individual OPTIFAST participants to teach them how to effectively and permanently change their nutrition behaviour.

REFERENCES

- [1] Jamy D. Ard, Kristina H. Lewis, Amy Rothberg, Anthony Auriemma, Sally L. Coburn, Sarah S. Cohen, Judy Loper, Laura Matarese, Walter J. Pories, and Seletha Periman. 2019. Effectiveness of a Total Meal Replacement Program (OPTIFAST Program) on Weight Loss: Results from the OPTIWIN Study. *Obesity* 27, 1 (2019), 22–29. <https://doi.org/10.1002/oby.22303>
- [2] S. C. Bischoff, A. Damms-Machado, C. Betz, S. Herpertz, T. Legenbauer, T. Löw, J. G. Wechsler, G. Bischoff, A. Austel, and T. Ellrott. 2012. Multicenter evaluation

- of an interdisciplinary 52-week weight loss program for obesity with regard to body weight, comorbidities and quality of life - A prospective study. *Int. J. Obes.* 36, 4 (2012), 614–624. <https://doi.org/10.1038/ijo.2011.107>
- [3] David Elswiler and Morgan Harvey. 2015. Towards automatic meal plan recommendations for balanced nutrition. In *Proceedings of the 9th ACM Conference on Recommender Systems*. 313–316.
- [4] David Elswiler, Christoph Trattner, and Morgan Harvey. 2017. Exploiting food choice biases for healthier recipe recommendation. In *Proceedings of the 40th international acm sigir conference on research and development in information retrieval*. 575–584.
- [5] Jill Freyne and Shlomo Berkovsky. 2010. Recommending food: Reasoning on recipes and ingredients. In *International Conference on User Modeling, Adaptation, and Personalization*. Springer, 381–386.
- [6] Mouzhi Ge, Francesco Ricci, and David Massimo. 2015. Health-aware food recommender system. In *Proceedings of the 9th ACM Conference on Recommender Systems*. 333–334.
- [7] Morgan Harvey, Bernd Ludwig, and David Elswiler. 2013. You are what you eat: Learning user tastes for rating prediction. In *International symposium on string processing and information retrieval*. Springer, 153–164.
- [8] H. Hauner, A. Moss, A. Berg, S. C. Bischoff, M. Colombo-Benkmann, T. Ellrott, C. Heintze, U. Kanthak, D. Kunze, N. Stefan, M. Teufel, M. Wabitsch, and A. Wirth. 2014. *Interdisziplinäre Leitlinie der Qualität S3 zur „Prävention und Therapie der Adipositas“*. Technical Report 04. Deutsche Adipositas-Gesellschaft (DAG) e.V., Martinsried. 60–63 pages. <https://doi.org/10.1055/s-0037-1618857>
- [9] Fabian Hoffmann. 2021. *Wie kann die Optifast-App der Universität Regensburg verbessert werden, um die Akzeptanz bei den Teilnehmern zu erhöhen?* Master's thesis. University of Regensburg, Regensburg.
- [10] Manuel Müller, Morgan Harvey, David Elswiler, and Stefanie 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*. IEEE, 73–80.
- [11] H. Oberitter, K. Schäbenthal, A. von Ruesten, and H. Boeing. 2013. The DGE-Nutrition Circle—representation and fundamentals of the food-based recommendations of the German Nutrition Society. *Ernährungs Umschau international* 2, 60 (2013), 24–29.
- [12] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence* (Montreal, Quebec, Canada) (UAI '09). AUAI Press, Arlington, Virginia, USA, 452–461.
- [13] Alain Starke, Martijn Willemsen, and Christoph Trattner. 2020. Nudging Healthy Choices in Food Search Through Visual Attractiveness. *Frontiers in Artificial Intelligence* 20, XX (2020).
- [14] Christoph Trattner and David Elswiler. 2017. Investigating the healthiness of internet-sourced recipes: implications for meal planning and recommender systems. In *Proceedings of the 26th international conference on world wide web*. 489–498.
- [15] Geoffrey C. Williams, Virginia M. Grow, Zachary R. Freedman, Richard M. Ryan, and Edward L. Deci. 1996. Motivational Predictors of Weight Loss and Weight-Loss Maintenance. *J. Pers. Soc. Psychol.* 70, 1 (1996), 115–126. <https://doi.org/10.1037/0022-3514.70.1.115>