

# Towards a Knowledge-aware Food Recommender System Exploiting Holistic User Models

Cataldo Musto  
Università degli Studi di Bari  
Bari, Italy  
cataldo.musto@uniba.it

Alain Starke  
University of Bergen  
Bergen, Norway  
alain.starke@uib.no

Christoph Trattner  
University of Bergen  
Bergen, Norway  
christoph.trattner@uib.no

Giovanni Semeraro  
Università degli Studi di Bari  
Bari, Italy  
giovanni.semeraro@uniba.it

## ABSTRACT

Food recommender systems typically rely on popularity, as well as similarity between recipes to generate personalized suggestions. However, this leaves little room for users to explore new preferences, such as to adopt healthier eating habits.

In this short paper, we present a recommendation strategy based on knowledge about food and users' health-related characteristics to generate personalized recipes suggestions. By focusing on personal factors as a user's BMI and dietary constraints, we exploited a *holistic user model* to re-rank a basic recommendation list of 4,671 recipes, and investigated in a web-based experiment (N=200) to what extent it generated satisfactory food recommendations. We found that some of the information encoded in a users' *holistic user profiles* affected their preferences, thus providing us with interesting findings to continue this line of research.

## KEYWORDS

Food Recommender Systems, User Modeling

### ACM Reference Format:

Cataldo Musto, Christoph Trattner, Alain Starke, and Giovanni Semeraro. 2020. Towards a Knowledge-aware Food Recommender System Exploiting Holistic User Models. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20), July 14–17, 2020, Genoa, Italy*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3340631.3394880>

## 1 INTRODUCTION & BACKGROUND

In recent years, *food recommender systems* have achieved major progress in suggesting what a person could eat [5, 20, 21]. The field has followed in the footsteps of most other RecSys domains by focusing on similarity between items [4, 22]. This has pointed out how recipes can be linked in terms of ingredients, taste, meal type, and other characteristics [24].

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

UMAP '20, July 14–17, 2020, Genoa, Italy

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6861-2/20/07...\$15.00

<https://doi.org/10.1145/3340631.3394880>

However, food recommender research has only a limited understanding of the relevant user characteristics and goals [23]. In particular, developments to design health-aware recommender systems have focused on simple re-ranking models based on recipe features [12, 22], while there has been little research on dietary constraints and health goals in personalized recommendations [13, 21].

One domain-overarching problem is that users tend to stick popular items [3, 10, 18]. As popular food recipes tend to be unhealthy [24], food recommender systems seem to limit users in discovering healthier alternatives, particularly when using a collaborative approach to recommendation [16].

We need a different user model to alleviate this 'popularity problem', which might be possible by focusing on user characteristics and goals [3, 6]. In food recommender systems, recipe factors are typically prioritized over user characteristics [21]. We see this as a missed opportunity, as there might be multiple personal characteristics (e.g., body-mass index) and contextual factors (e.g., mood) that determine how appropriate a recipe is for a user, as well as how likely that user is to try a different, healthier recipe. Moreover, food recommenders often do not consider constraints [16, 21], such as dietary preferences (e.g., vegan users, gluten intolerance).

This work proposes a first step to incorporate more personal factors and constraints into a food recommender system. To do so, we propose a novel, but simple recommendation strategy that combines knowledge about food and users to generate personalized recipe suggestions, which we refer to as a *holistic user model* (HUM). Our recommendation algorithm exploits common-sense knowledge about food requirements (e.g., a user's BMI) to identify the most suitable recipes. To assess its performance, we investigate to what extent users preferred such knowledge-aware recommendations over those generated by a popularity-based algorithm, using a strict baseline to assess the model's performance in a one-shot recommender user study. We posit the following research questions:

- *RQ1*. To what extent can we predict food choices, using a holistic model in a knowledge-aware recommender system?
- *RQ2*. Which user characteristics and goals, as well as recipe features affect a user's food preferences?

## 2 APPROACH

This section describes a simple, but explainable method that considers personal characteristics to provide recipe recommendations.

**Table 1: Aspects that are part of the Holistic User Model.**

User Aspect	Factor
Demographics	Gender/Sex, Age
Affect	Mood
Domain knowledge	Cooking Experience
Behavioral Data	Level of Physical Activity
Health Data	Food Requirements, Amount of Sleep, Mood, Stress level, Weight (BMI)

First, we introduce the *holistic user model* and how it encodes user needs and constraints. Second, we describe the knowledge-aware recommendation strategy that aims to identify the most suitable recipes for a given user, which is later tested in a user study.

## 2.1 Basics of Holistic User Modeling


As stated in [8], a holistic user model (HUM) is a representation of the user that is obtained by merging and processing heterogeneous data from social networks, smartphones and personal devices. Based on this paradigm and literature on food choices (e.g., [15]), each user’s needs and preferences can be modeled through eight aspects (also called: ‘facets’): *demographics, interests, affects, knowledge and skills, psychological traits, behaviors, social connections, and health data*. The aim is to learn about each aspect through various behavioral traces, such as posts on social networks (e.g., Twitter), fitness app activities (e.g., Strava), visited locations, and personal data gathered from wearable devices [2]. These data types can be processed using machine learning and natural language processing techniques to infer higher-level characteristics of the person.

In this paper, we assume that a HUM of the target user has been previously built and it is available. For the sake of brevity, please refer to [8] for a complete description of the workflow to build a HUM. However, as a rich source of data is not always available from the onset (e.g., in a cold start situation), we design a representation that facilitates basic needs and food constraints (e.g., a user is vegan), and excludes users’ food preferences and social relations. In doing so, we approximate the user model for our knowledge-aware food recommendation strategy by inquiring on *five user aspects*. Table 1 lists the employed aspects (e.g., demographics), along with their underlying factors (e.g., gender, age). By ignoring a user’s food preferences, we try to avoid simple food recommendations that were based on a similarity calculation, as has been done in earlier studies [21].

## 2.2 Designing a Knowledge-aware Food Recommender System

The workflow carried out by our *food recommender system* is depicted in Figure 1. Our system can be classified as a knowledge-based recommender system [1], as its recommendation strategy exploits common-sense knowledge about food choices to identify the most suitable recipes for a target user (cf. [8, 19]). In a nutshell, the output of the recipe recommendation process is based on three main components: A PROFILER, a FILTER step, and a RANKER step.

The process is initiated by a user who asks for a specific meal type: a *main course*, a *second course*, or a *dessert*. In the first step of the workflow, we build a Holistic User Model of the target user. The to-be-encoded information can be obtained by explicitly asking the

**Figure 1: Workflow of our Knowledge-aware Food Recommender System.**

user or by querying external profiling platforms, such as Myrror (cf. [8]). Once the profile is built, the FILTER generates a preliminary set of candidate recipes by filtering non-compliant recommendations. This step is carried out by analyzing the user’s food restrictions and cooking experience, and subsequently removing recipes from the list of candidates which contain ingredients that a user wishes to avoid (e.g., lactose, meat), or that are too complex to prepare.

After filtering, the set of candidate recommendations is still large and requires further re-ranking, which is done by the RANKER component. Given a user  $u$ , the goal of this component is to assign to each recipe  $r$  a  $score(r, u)$ , to rank all the candidate recipes and to identify the top-1 that matches the user best, in terms of the user’s characteristics and constraints. To do so, we propose two scoring mechanisms. First, we propose a simple *popularity-based* score, which is adopted as a baseline in our user evaluation study:

$$popScore(r) = avgRating(r) * \log_{10}(count(r)) \quad (1)$$

Formula 1 shows that our baseline awards higher scores to popular recipes, where  $count(r)$  is a rating counter stored in our database of 4,761 recipes. However, such a basic scoring scheme does not consider the needs, nor the constraints of the target user. Conversely, we aim to take advantage of the richness of personal information available in a HUM to better tailor our recipe recommendations. Accordingly, we propose a somewhat more sophisticated *holistic* scoring function:

$$holistic(r, u) = popScore(r) * knowledge(r, u) \quad (2)$$

The knowledge-aware part (i.e.,  $knowledge(r, u)$ ) is a modifier that increases or decreases the score of a recipe by using some general knowledge about food choices. Such knowledge is encoded as *rules*, having the form  $factor \rightarrow modifier$ , where *factors* are operationalizations of user aspects, and *modifiers* are recipe characteristics. In short: if the left part of a rule is satisfied, the modifier is applied on the recipe by increasing or decreasing the popularity score.

Although we cannot discuss all the details of our knowledge-based scoring formula due to space reasons, we emphasize that these rules are based on common-sense knowledge about food choices. For example, our knowledge-aware recommender awards a lower score to recipes that are high in calories, if the user has a high BMI. Moreover, we consider insights from recent studies about the link between user factors and food consumption, such as the relation between stress and the amount of salt in recipes<sup>1</sup>. Table 2 reports some of the rules we encoded in our recommender system, listing user aspects and factors, and the corresponding modifiers.

<sup>1</sup><https://oklahoman.com/article/feed/687315/did-you-know-salt-reduces-stress>

**Table 2: Selection of rules encoded in our knowledge-aware recommender.**

User Aspect	Factor	Modifier
Demographics	<i>BMI&gt;25</i>	low-calorie recipes ↑ high-calorie recipes ↓
	<i>BMI&lt;25</i>	high-calorie recipes ↑
Mood	<i>bad</i>	recipes with sugar ↑
Behaviors	<i>&gt;30 min/day activity</i>	high-protein recipes ↑↑ high-calorie recipes ↑
Health	<i>stress=yes</i>	high-sodium recipes ↑
Health	<i>sleep=low</i>	high-magnesium recipes ↑
Health	<i>depression=yes</i>	high-fat recipes ↓

### 3 USER EVALUATION

To examine the effectiveness of our holistic user model, we compared it to our popularity-based baseline in a web-based experiment on food recipes. Users were presented three different pairs of recipes (i.e. the top-1 recipes of our holistic and popular recommendation models), and asked to choose the recipe they preferred the most. Recipes were sampled from a database of 4,671 recipes, which is available online: <https://tinyurl.com/recipes-uniba>. The recipes were obtained from the popular Italian food community platform GialloZafferano<sup>2</sup>, and translated to English. The recipes contain information about their name, category, preparation difficulty, as well as their ingredients, macro-nutrients, calories, rating count, and average website rating. Moreover, they also include several binary tags, such as *vegetarian*, *vegan*, *lactose-free*, and *low-nickel*.

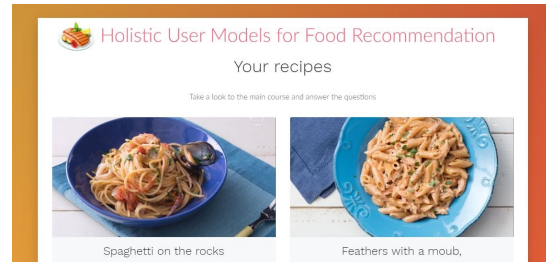
#### 3.1 Participants

We invited a total of 200 participants on Amazon MTurk to participate in a choice study on food recipes. To ensure high quality feedback, we only sampled users who had a HIT acceptance rate of 99% and had successfully completed at least 500 HITS in the past. Participants were compensated with 0.5 USD for a HIT, which took them on average five minutes to complete. Eventually, we had to remove nine participants due to missing stimulus data.

#### 3.2 Experimental Procedure & Measures

Users interacted with a Web Application<sup>3</sup>. First, to build a user model in line with Table 1, we asked users to disclose personal information. We inquired on the user’s gender, age, BMI (5-point scale), recipe website usage (4-point scale), cooking experience (5-point scale), and mood (i.e., ‘good’, ‘neutral’, or ‘bad’). In addition, we also asked users about their typical sleep length, feelings of stress or depression (yes/no), dietary goals (lose or gain weight, or none), and dietary constraints (e.g., *vegan*, *low-nickel*).

Subsequently, we ran our food recommendation pipeline. Each user was presented three pairs of recipes, where each pair represented a different part of a meal: main courses (i.e. mostly pasta dishes), second courses (i.e. mostly meat-based dishes), and desserts. Each pair consisted of one recipe that was generated using a simple popularity-based ranking, and another recipe that was obtained by applying our holistic user model scoring. A screenshot of the

**Figure 2: Screenshot of the user study recommending the user two recipe choices (holistic vs popular).**

interface is provided in Figure 2, which shows the representation of the two alternatives. For each pair of recipes, we asked users which of the two recipe they preferred the most, or whether they preferred none of them. In addition, we also inquired on their underlying motivations for choosing either recipe (if any), presenting four propositions about the chosen recipe on 5-point Likert scales: “It seems savory and tastier”, “It helps me to eat more healthily”, “It would help me to lose/gain weight”, and “It seems easier to prepare”.

### 4 RESULTS

To address *RQ1*, we compared preferences for popular and holistic recommendations per course (i.e., pasta meal, meat meal, dessert), as well as all courses combined in two separate logistic regression analyses. To address *RQ2*, we examined whether choices were determined by specific recipe factors, or user factors and motivations.

#### 4.1 Preferences per Meal Type

Overall, a sequence of two-sample t-tests showed that users were more likely to choose popular recipes for the main course (54.0% (popular) vs 32.5% (holistic)):  $t(190) = -3.27, p < 0.01$ , as well as for dessert recipes (54.5% vs 36.1%):  $t(190) = -2.70, p < 0.01$ . In contrast, users preferred holistic recipes (49.2%) over popular recipes (33.0%) for the second course:  $t(190) = 2.51, p = 0.013$ . These mixed results suggested that the holistic user model did not entirely outperform the popular baseline, and that differences might boil down to specific user factors and motivations.

To examine *RQ2*, we used three different logistic regression models (excluding ties), to predict which recipe type was chosen in each course, based on a set of user factors and motivations. Table 3 describes the performance of each model, pointing out that the ‘2nd course’ model had a low predictability and was not significant, while the predictive value of the ‘Main Course’ model ( $R^2 = .24$ ) was higher than that of the ‘Dessert’ model ( $R^2 = .14$ ), as it also included more significant personal predictors.

Nonetheless, the predictive value of the user factors was small. Although the knowledge-aware recommender used factors such as BMI and cooking experience for its HUM, most of them did not steer user preferences towards the holistic recipes. Although no user factors affected dessert recipe choices, we did find that male participants were more likely to choose second course holistic recipes ( $p < 0.05$ ). Since second course recipes were typically meat or fish-based, this was consistent with previous findings that men prefer meat-based dishes over sweet dishes [11]. With respect to the main course, users who reported to be in a good mood preferred

<sup>2</sup><https://www.giallozafferano.it/ricette-cat/>

<sup>3</sup>It can be found at: <http://90.147.102.243:8080/foodrecsys/>

**Table 3: Logistic regression models predicting choices for holistic recipes (compared to popular). Reported are regression coefficients. \*\*\*  $p < 0.001$ , \*\*  $p < 0.01$ , \*  $p < 0.05$ .**

Holistic	Main Course	2nd Course	Dessert
<i>Personal factors</i>			
BMI	.44	.078	.22
Website use	.19	.32	.082
Cooking freq.	-.49	.024	-.50
Mood	<b>.88*</b>	-.36	.52
Gender (male)	.13	<b>.96*</b>	-.53
Age	.013	-.031	-.011
<i>Choice Motivation</i>			
Taste-related	<b>-.82**</b>	-.11	<b>-.55**</b>
Health-related	<b>.98***</b>	.024	.30
Weight-related	.28	.044	-.068
Ease to prepare	<b>-.71***</b>	-.045	<b>-.62***</b>
Weight-gain goal	<b>-3.04**</b>	-.93	.18
Weight-loss goal	<b>-1.15*</b>	-.87	<b>-.99*</b>
Constant	2.03	.53	<b>4.43***</b>
<i>Adjusted R<sup>2</sup></i>	<b>.24***</b>	.062	<b>.15***</b>

the holistic recipe. This is a notable finding, as mood is a common factor in the music recommender domain [14, 25], but not in food.

A user’s reported motivation for choosing either recipe provided more insight. Table 3 shows that users chose popular recipes due to their taste and ease to prepare, while users with health-related goals were more likely to choose the holistic main course. In contrast, users with either weight-gain or weight-loss goals (not part of the holistic model) were more likely to choose popular main courses and desserts (weight-loss only). This suggested that weight-related goals were distinct from general healthy eating habits, and that the holistic user model supported the latter.

## 4.2 Recipe Factors vs User Factors

We also investigated to what extent different recipe characteristics<sup>4</sup> could predict user choices (RQ2), and how they compare to user characteristics across all meal types. In Table 4, we report the results of two multilevel logistic regression models, for either holistic or popular recipes. Although our database comprised more recipe characteristics than reported in Table 4 (e.g., sugar, fat, etc.), we excluded those that had high cross-correlations (e.g., calories and carbs:  $r > 0.8$ ). We also excluded personal factors that had no predictive value in Table 3.

We found that both models have little predictive quality, as their  $R^2$  values were small: 0.021 for holistic recipes, 0.036 for popular recipes. Although Table 4 shows that personal factors did not play a role in user decision-making, different recipe factors affected preferences for either recipe type. For one, Table 4 suggests a ‘health divide’ between holistic and popular recipes, as low-carb and high-protein *holistic* recipes were more likely to be chosen than high-carb, low-protein recipes, which was in line with the health-related motivations to choose a holistic recipe. In contrast, *popular* recipes were more likely to be chosen if they contained more carbs and more saturated fat, which supported findings in [22, 24].

<sup>4</sup>Recipe characteristics have previously been investigated in [4, 24], but not simultaneously with various user factors as is done in this work.

**Table 4: Multilevel logistic regression models predicting choices for any recipe, clustered at the user level, per model (holistic vs popular). \*\*\*  $p < 0.001$ , \*\*  $p < 0.01$ , \*  $p < 0.05$ .**

	Holistic Recipes	Popular Recipes
<i>Recipe factors</i>		
Carbohydrates	<b>-.0094*</b>	<b>.011*</b>
Protein	<b>.013*</b>	-.014
Saturated fat	.023	<b>.051***</b>
Fiber	.011	-.031
<i>Personal factors</i>		
BMI	.065	-.086
Mood	.084	-.085
Gender (male)	.30	-.32
Constant	<b>-.57*</b>	-.096
<i>Pseudo R<sup>2</sup></i>	<b>.021*</b>	<b>.036***</b>

## 5 CONCLUSIONS & FUTURE WORK

We have investigated to what extent we could predict food choices using a *holistic user model* (RQ1). The results are mixed, as recipes based on a popularity ranking are chosen more often among main course and dessert recipes, while recipes taken from our holistic user model are preferred among second course recipes. This suggests that the set of aspects and factors in our HUM requires further improvement to meet user preferences. However, we wish to emphasize that we have used a strict popular baseline, which has shown to be hard to ‘beat’ [24], particularly for hedonic dishes as desserts. Moreover, it seems that user models are more effective when focused on specific meal types, as most of their predictive value is lost when aggregated across all recipes.

With regard to which specific user characteristics, goals, and recipe features affect user preferences (RQ2), the results are also mixed. Most of the user factors encoded in our holistic model have not led to preference differences, except for mood and gender in specific meal types. However, the underlying motivations of user choices, such as taste, health, and ease to prepare, do signal that a holistic user model can appeal to users who wish to pursue healthy food choices. A knowledge-aware recommender should also capture such motivations, as they seem to support behavioral change [3, 9].

To follow up on this study’s results, it should be addressed how weight-loss goals can be discerned from health-related goals. For example, users might not be interested in reducing their calorie and fat intake, but to consume specific nutrients [13]. We propose to investigate: *how can a holistic user model support healthy maintenance behaviors?* Moreover, there is currently a limited understanding on how personalized recipe suggestions can support behavioral change in the longer-term. To this end, psychological theories as Prochaska’s transtheoretical model of behavioral change [9] and explanation strategies [7] might help, but have yet to be applied in personalized contexts. Finally, food recommender research should also consider other behavioral goals that underlie our food consumption, such as sustainability [16, 17].

## ACKNOWLEDGMENTS

We thank Anthony Pellicani and Angelo Sparapano for the implementation of the food recommender system. This paper was supported by a grant from the Niels Stensen Fellowship.

## REFERENCES

- [1] Robin Burke. 2000. Knowledge-based recommender systems. *Encyclopedia of library and information systems* 69, Supplement 32 (2000), 175–186.
- [2] Marco De Gemmis, Leo Iaquina, Pasquale Lops, Cataldo Musto, Fedelucio Narducci, and Giovanni Semeraro. 2010. Learning preference models in recommender systems. In *Preference Learning*. Springer, 387–407.
- [3] Michael D Ekstrand and Martijn C Willemsen. 2016. Behaviorism is not enough: better recommendations through listening to users. In *Proceedings of the 10th ACM Conference on Recommender Systems*. ACM, 221–224.
- [4] David Elswailer, Christoph Trattner, and Morgan Harvey. 2017. Exploiting food choice biases for healthier recipe recommendation. In *Proc. of SIGIR '17*.
- [5] Jill Freyne and Shlomo Berkovsky. 2010. Intelligent food planning: personalized recipe recommendation. In *Proceedings of the 15th international conference on Intelligent user interfaces*. 321–324.
- [6] Yu Liang. 2019. Recommender system for developing new preferences and goals. In *Proceedings of the 13th ACM Conference on Recommender Systems*. ACM, 611–615.
- [7] Cataldo Musto, Fedelucio Narducci, Pasquale Lops, Marco de Gemmis, and Giovanni Semeraro. 2019. Linked open data-based explanations for transparent recommender systems. *International Journal of Human-Computer Studies* 121 (2019), 93–107.
- [8] Cataldo Musto, Giovanni Semeraro, Cosimo Lovascio, Marco de Gemmis, and Pasquale Lops. 2018. A framework for holistic user modeling merging heterogeneous digital footprints. In *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization*. 97–101.
- [9] James O Prochaska. 2008. Decision making in the transtheoretical model of behavior change. *Medical decision making* 28, 6 (2008), 845–849.
- [10] Al Mamunur Rashid, Istvan Albert, Dan Cosley, Shyong K Lam, Sean M McNee, Joseph A Konstan, and John Riedl. 2002. Getting to know you: learning new user preferences in recommender systems. In *Proceedings of the 7th international conference on Intelligent user interfaces*. 127–134.
- [11] Markus Rokicki, Eelco Herder, Tomasz Kuśmierczyk, and Christoph Trattner. 2016. Plate and prejudice: Gender differences in online cooking. In *Proceedings of the 2016 conference on user modeling adaptation and personalization*. 207–215.
- [12] Hanna Schäfer, Santiago Hors-Fraile, Raghav Pavan Karumur, André Calero Valdez, Alan Said, Helma Torkamaan, Tom Ulmer, and Christoph Trattner. 2017. Towards health (aware) recommender systems. In *Proceedings of the 2017 international conference on digital health*. ACM, 157–161.
- [13] Hanna Schäfer and Martijn C Willemsen. 2019. Rasch-based tailored goals for nutrition assistance systems. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*. ACM, 18–29.
- [14] Markus Schedl, Peter Knees, Brian McFee, Dmitry Bogdanov, and Marius Kaminakas. 2015. Music recommender systems. In *Recommender systems handbook*. Springer, 453–492.
- [15] Benjamin Scheibehenne, Linda Miesler, and Peter M Todd. 2007. Fast and frugal food choices: Uncovering individual decision heuristics. *Appetite* 49, 3 (2007), 578–589.
- [16] Alain Starke. 2019. RecSys Challenges in achieving sustainable eating habits. In *HealthRecSys'19: Proceedings of the 4th Workshop on Health Recommender Systems*. ACM, 29–30.
- [17] Alain Starke, Martijn Willemsen, and Chris Snijders. 2017. Effective user interface designs to increase energy-efficient behavior in a Rasch-based energy recommender system. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*. 65–73.
- [18] Alain D. Starke, Martijn C. Willemsen, and Chris Snijders. 2020. With a Little Help from My Peers: Depicting Social Norms in a Recommender Interface to Promote Energy Conservation. In *Proceedings of the 25th International Conference on Intelligent User Interfaces (Cagliari, Italy) (IUI '20)*. Association for Computing Machinery, New York, NY, USA, 568–578. <https://doi.org/10.1145/3377325.3377518>
- [19] Napat Suksom, Marut Buranarach, Ye Myat Thein, Thepchai Supnithi, and Ponrudee Netisopakul. 2010. A knowledge-based framework for development of personalized food recommender system. In *Proc. of the 5th Int. Conf. on Knowledge, Information and Creativity Support Systems*.
- [20] Thi Ngoc Trang Tran, Müslüm Atas, Alexander Felfernig, and Martin Stettinger. 2018. An overview of recommender systems in the healthy food domain. *Journal of Intelligent Information Systems* 50, 3 (2018), 501–526.
- [21] Christoph Trattner and David Elswailer. 2017. Food recommender systems: important contributions, challenges and future research directions. *arXiv preprint arXiv:1711.02760* (2017).
- [22] Christoph Trattner and David Elswailer. 2017. Investigating the healthiness of internet-sourced recipes: implications for meal planning and recommender systems. In *Proc. of WWW '17*. 489–498.
- [23] Christoph Trattner and Dietmar Jannach. 2019. Learning to recommend similar items from human judgments. *User Modeling and User-Adapted Interaction* (2019), 1–49.
- [24] Christoph Trattner, Dominik Moesslang, and David Elswailer. 2018. On the predictability of the popularity of online recipes. *EPJ Data Science* 7, 1 (2018).
- [25] Alexandra Uitdenbogerd and R Schyndel. 2002. A review of factors affecting music recommender success. In *ISMIR 2002, 3rd International Conference on Music Information Retrieval*. IRCAM-Centre Pompidou, 204–208.