

Personalized, Health-Aware Recipe Recommendation: An Ensemble Topic Modeling Based Approach

Mansura A. Khan*
mansura.khan@ucd.ie

Barry Smyth*
barry.smyth@ucd.ie

Ellen Rushe*
ellen.rushe@ucdconnect.ie

David Coyle*
d.coyle@ucd.ie

ABSTRACT

Food choices are personal and complex and have a significant impact on our long-term health and quality of life. By helping users to make informed and satisfying decisions, Recommender Systems (RS) have the potential to support users in making healthier food choices. Intelligent user-modeling is a key challenge in achieving this potential. This paper investigates Ensemble Topic Modelling ($E_{ns}TM$) based **Feature Identification** techniques for efficient user-modeling and recipe recommendation. It builds on findings in $E_{ns}TM$ to propose a reduced data representation format and a smart user-modeling strategy that makes capturing user-preference fast, efficient and interactive. This approach enables personalization, even in a cold-start scenario. We compared three $E_{ns}TM$ based variations through a user study with 48 participants, using a large-scale, real-world corpus of 230,876 recipes, and compare against a conventional Content Based (CB) approach. $E_{ns}TM$ based recommenders performed significantly better than the CB approach. Besides acknowledging multi-domain contents such as taste, demographics and cost, our proposed approach also considers user's nutritional preference and assists them finding recipes under diverse nutritional categories. Furthermore, it provides excellent coverage and enables implicit understanding of user's food practices. Subsequent analysis also exposed correlation between certain features and healthier lifestyle.

CCS CONCEPTS

• **Information systems** → **Recommender systems**.

KEYWORDS

Recommender Systems; Personalization; Topic Modeling; Food Features; Recipe Recommender; Collaborative Filtering

*Insight Centre for Data Analytics, University College Dublin

HealthRecSys'19, September 20, 2019, Copenhagen, Denmark

© 2019 Copyright for the individual papers remains with the authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0). This volume is published and copyrighted by its editors

ACM Reference Format:

Mansura A. Khan, Ellen Rushe, Barry Smyth, and David Coyle. 2019. Personalized, Health-Aware Recipe Recommendation: An Ensemble Topic Modeling Based Approach. In *Proceedings of the 4th International Workshop on Health Recommender Systems co-located with 13th ACM Conference on Recommender Systems (HealthRecSys'19)*. Copenhagen, Denmark, September 20, 2019, 7 pages.

1 INTRODUCTION

Food has a direct, complex and multifaceted relationship with our lifestyle and personality. People have explicit preferences regarding activities around food, such as cooking, plating, grocery and eating-out. Studies showed people are becoming more mindful towards healthier lifestyles and the fact that healthy eating/cooking impacts psychosocial and physical well-being [6] However, finding food-ideas/recipes that acknowledge one's circumstance and preference remains a challenge for many people. Food Recommender Systems (FRS) have the potential to assist users in navigating through the overwhelming amount of online resources on food/recipes and guide them towards healthier choices.

Recommending food is challenging as our choices are defined by many cross-domain factors including demographic and contextual factors, health awareness, social and ethical factors, together with practical considerations such as cost, cooking-time and methods, and the availability of ingredients. In order to develop effective FRS, we must design user-models that capture user data across these diverse factors. Approaches are also required that enable Recommender Systems (RS) to fit user's preference data on a massive information space around food. As Teng et al. note, there are millions of food-items/recipes as different ingredients are grown at different geographical locations and recipes originate from different cultural groups worldwide [27]. In this context coverage and diversity are important constraints, where coverage corresponds to the percentage of items for which a RS is able to generate a prediction [15]. Higher coverage enables the RS to implement varying diversity approaches and draw from more options. Taken together, these challenges necessitate FRS that can (1) identify the attributes/features which are significant for human food-choices, (2) capture user's preference on the identified features, (3) filter a large

information-space, (4) generate recommendations efficiently and finally (5) guide users towards healthier choices.

We explored Ensemble Topic Modelling ($E_{ns}TM$) [7] accompanied by a series of custom text-preprocessing to extract significant food features. The aim was to identify representative or agent contents of diverse domains connected to human food choice. In our study 288 features and their corresponding significance scores were extracted from a corpus of 230,876 recipes. Which later worked as the basis for our intelligent user-modeling approach. As summarized in Table 1, the identified feature set is rich in contents representing multiple domains. The paper describes a foreshortened data representation format based on the extracted features which aims to reduce computational complexity of food recommendation.

We implemented three distinct $E_{ns}TM$ based personalized FRS: a Food Feature based Recommender (FFbR), a Weighted Food Feature based Recommender (WFFbR), and a Food Feature based Collaborative Filtering (FFbCF). To evaluate these approaches we conducted a user study comparing $E_{ns}TM$ based recommenders to a conventional Content Based (CB) approach. Results show that all $E_{ns}TM$ based approaches significantly outperformed CB approach. In contrast to prior work, the $E_{ns}TM$ approach also effectively supported recommendations across diverse social and cultural groups, even in a first recommendation scenario. Finally, the strong adaptation of the concept of dislike across all three methods proved effective in implicitly identifying user's food practice (e.g. vegetarian, halal) and filtering accordingly. Further exploratory analysis exposed previously unknown pattern in user's interactions towards certain features. That is, some features are more popular than others among healthier user-groups. The existing correlation between healthier user-groups and certain food features argue for further research on feature based FRS with healthiness cues.

2 RELATED WORK

Previous research has produced seminal contributions towards FRS, aimed at ensuring user-preference, diversity and nutritional development in diet. Freyne et al. [12, 13] describe an ingredient-based approach where they inferred user's preference on a new recipe as the cumulative sum of his/her preference for each ingredient in that recipe. This formed the basis of their novel user-based K-NN Collaborative Filtering (CF) approach [12], which has been influential and was applied by others including [19, 26]. Subsequently, more advanced methods emerged for tackling different challenges such as, Teng et al. [27] used item-centric CF and applied an ingredient-network to identify similar recipes, where the ingredient-network was generated based on co-occurrence of ingredients within recipes and menus. Kuo et al. [21] proposed a weighted graph based menu planning

approach where ingredients were grouped into subsets and each subset was considered as contents. However, while these approaches are very interesting, they focus purely on ingredients.

Ge et al. [14] proposed a method that leverages tags and latent factors to recommend recipes. Pinxteren et al. adopted a different approach [34] where, first they added custom annotations to each recipe in their corpus, then asked users to rate individual recipes and finally recommended recipes that share annotations with those rated positively by the user. This method was successful in addressing more food-choice factors, but the annotation set was relatively small and specific to their recipe corpus. As they mentioned, this limited their FRS from automatically adopting to new user groups. Further notable work includes: Gu et al. [17] case-based FRS based on user's previous consumption cases; Sobeck et al. [26] hybrid FRS incorporating fuzzy inference with stereotype demographic filtering; and Bianca et al. [8] hybrid model incorporating meta-heuristic and genetic algorithms. Elswiele et al. [10] and Ueta et al. [33] discussed automatic meal planning approach to support balanced nutrition. While effective in constrained contexts, each of these approaches depends on sufficient pre-existing user preference data. They are thus susceptible to failure in cold-start scenarios [8]. Trattner et al. [31] proposing a novel method to recommend recipes to people in a cold-start scenario.

There was also a significant number of interesting research work producing domain specific knowledge to facilitate future research interests. [29] is a seminal work from Trattner et al. on summarizing, "to which extent current recommendation algorithms can adopt healthy recipes recommendation?" and "what resources are out there?". [24, 25, 30, 32] showed how online recipe repositories could be potential sources for knowledge discovery to support personalized and group-based recipe recommendations. [5, 11, 19] looked into patterns in users' online activity around food.

Contributions of This Work. The related works unveil seminal solutions available to address the 5 dominant challenges (as summarized in introduction) in FRS research. Unlike our $E_{ns}TM$ based approach that consider multi-domain food-features, most of these solutions focus on ingredients while generating recommendation. While some of the existing work proposed significant approaches to consider socio-cultural and contextual features, they are often limited to their food-corpus and user-group. Differently from our approach, many existing FRS approaches depend on pre-existing recipe ratings from user. Also, there dose not exist many works which try to reduce the food data format in the aim of enabling the FRS to perform with large recipe corpus (e.g., 230,000+ recipes). Our contributions are summarized as follows:

- a novel method to identify significant multi-domain Food Features from any food-corpus.
- a Food Feature based intelligent user-modeling technique that fosters higher personalization since cold-start scenario.
- fine-grained recommendation algorithms that considers user's preference on multi-domain food features.
- a reduced data representation format that enables FRS to perform faster and at the same time preserves the integrity of the recipe information.
- a substantial user study that showed the recommendation approach achieves the level of user-satisfaction that it thrives for.

3 RECOMMENDER STRATEGIES

To create a recipe data-set, we developed a web-scrapers for geniuskitchen.com [2]. Our final data-set comprises of 230,876 recipes. Each recipe was stored as a plain-text document that included information on ingredients, instructions, servings, cuisine, cooking-time, cooking-approach, cooking equipment, context, taste (e.g. sour or spicy) and nutrition data.

The first aim of our work was to uncover common food-features across the recipe data-set that could then be used to model user-preference and resolve user-to-recipe relationships. One traditional approach to achieving this is to apply TF-IDF [23]. This provides a term (word) frequency matrix that favors intra-document dominance of a word over intra-corpus dominance. However, it does not produce any knowledge about the term beyond the occurrence frequency. Topic Modelling (TM) is an alternative and widely investigated approach, which attempts to discover the underlying thematic structure within a text corpus as derived from co-occurrences of words across the documents [7]. A Topic Model typically consists of k topics, each represented by a ranked list of strongly-associated terms/words. Each topic represents trend or theme of the contents of the document. Belford et al. [7] extended TM in their $E_{ns}TM$. They built on evidence by Topchy et al. [28] that ensemble procedures encourage diversity and improve quality by integrating results across multiple iterations of individual algorithms.

To extract a set of significant features from our recipe corpus, we proceeded with $E_{ns}TM$ [7] based on the generation and integration of the results produced by 100 runs of TM based on non-negative matrix factorization [20]. This produced a Topic-Term Weight Matrix where each column is a topic and each row determines the level of association between {Topic, Term} pair. To achieve a diverse and novel feature set we selected the top 30 topics and top 15 terms within each of these topics. We followed [16] for deciding on the number of topics and number of terms-per-topic. Term number $t=15$ gave the highest stability score [16] for our

recipe corpus. Some terms appeared over multiple topics as they are involved in multiple food-trends.

We consider the value of each {Topic, Term} pair in the Topic-Term Weight Matrix as the significance weight w_i for each term i within the corresponding topic. For terms existing over multiple topics we assigned w_i as the cumulative sum of their weight over all the corresponding topics. This produced a final set of 288 unique terms representing diverse aspects of food, e.g. cooking-approach, ingredient, equipment, serving-techniques, preservation-techniques and context. These 288 terms, summarized in Table 1, are our identified Food Features and their corresponding weight are the proposed Feature Scores¹.

Feature-Type	Features
context	holiday-food, beginner-cook, week-night, inexpensive, 6-people-or-more, potluck
cuisine	italian, hawaiian, tex-mex, chinese, cajun
equipment	saucepan, thermomix, wok, dutch-oven
cooking process	few-steps-recipe, less-than-one-hour, fried, slow-cooked, marinated, 4-hours-or-more
ingredient	poultry, feta, spaghetti, ham, shredded-meat
category	risotto, lasagna, stew, appetizer, pot-roast
nutrition	high-calcium, low-cholesterol, egg-free

Table 1: Summary¹ of the extracted features from ETM

In this work, we adopted a simple recipe-to-feature relationship by representing each recipe as a vector of 288 features, where each feature value corresponds to its TF-IDF within the recipe. The transformation of the recipe corpus into a recipe-to-feature matrix, as shown in figure 1, reduces the bulk overload of food data while still holding enough information to retrieve each recipe.

Recipes	Plaintext				
R_1	Document ₁				
R_2	Document ₂				
.				
R_n	Document _n				
	f_1	f_2	.	.	f_{288}
R_1	0.79	0	.	.	.31
R_2	0	0	.	.	0
.
R_n	0.61	1	.	.	.08

Figure 1: Recipe plain-text to feature vector transformation

In the next step we used the identified food-features to learn user's preference. During their initial interaction with our FRS, users are asked to choose features with a like or dislike. (Note there was no requirement for users to rate all 288 features). To build the user-to-feature matrix the FRS assigns +5 to liked features, -5 to disliked features and 0 to any feature that has not been selected by the corresponding user. Unlike typical RS approaches we assigned an extreme

¹The complete set of 288 features, their corresponding weights and set of food features correlated to healthier lifestyle are available at <https://github.com/MAK273/SupportingFileForHealthRecsys2019>

negative value to disliked features. This was an important design decision and was done with the view to producing insights beyond user's food preferences, by enabling our system to implicitly capture important considerations such as nutritional restrictions or foods which users deliberately avoid.

We implemented three $E_{ns}TM$ based recommendation algorithms: FFbR, WFFbR, and FFbCF. Each uses the recipe-to-feature matrix to transform user's positive and negative scores on features to user's scores on recipes.

- **Food Feature based Recommender (FFbR):** This strategy assigns a preference score P for user u_a on a target recipe r_n based on the cumulative sum of u_a 's rating (dis/like) for all features $f_{i(1,2,\dots,m)}$ present in r_n . Where f_{i,u_a} is u_a 's rating on a feature f_i and m is the total number features consisting r_n .

$$P(u_a, r_n) = \left(\sum_{i=0}^m f_{i,u_a} \right)^{(0,5)} \quad (1)$$

Instead of taking an average, we normalized the cumulative sum to a range $\{0$ to $5\}$ to favor recipes with more liked features over others. FFbR treats all food-features equally, assuming that each feature has an equal impact on user preferences.

- **Weighted Food Feature based Recommender (WFFbR):** With WFFbR we aimed to account for the differing impact of different food features. It scales u_a 's preference on a feature f_b with its corresponding feature score w_b and predicts u_a 's preference on r_n as the cumulative sum of the weighted preferences on all m features within r_n .

$$P(u_a, r_n) = \left(\sum_{i=0}^m f_{i,u_a} \times w_i \right)^{(0,5)} \quad (2)$$

- **Food Feature based Collaborative Filtering (FFbCF):** FFbCF applies the CF proposed by Freyne et al. [12] in order to increase the knowledge on user's preference and predict user's preference score on food-features not been liked or disliked by the user. When user u_a first interacts with it the FFbCF identifies u_a 's nearest neighbors based on similar ratings on overlapping features. We implemented KNN clustering [9] to identify top n nearest neighbours of u_a . For a new feature f_b FFbCF predicted u_a 's preference as,

$$P(f_b, u_a) = \frac{\sum_{i=0}^n f_{b,u_i}}{n} \quad (3)$$

With this more densely populated user-to-feature matrix FFbCF generates $P(u_a, r_n)$ using equation 1.

To compare proposed $E_{ns}TM$ based recommenders we implemented the generic CB [13] approach as our baseline.

- **Content-Based(CB):** CB predicts $P(u_a, r_n)$ based on u_a 's explicit preference on the ingredients $Ing_{i(1,2,\dots,m)}$ comprising r_n . Where m is the total number ingredients in r_n .

$$P(u_a, r_n) = \frac{\sum_{i=0}^m Ing_{i,u_a}}{m} \quad (4)$$

4 EVALUATION

In order to test the $E_{ns}TM$ base FRS strategies, we conducted a user study with 48 users of varying nationalit and ethnicity. The user-group belongs to an age-range of 21 to 65' and comprises of students, professionals and athletes. 45% of our participants identified them as female and 55% as male. Participants were recruited though social media groups within UCD. All participants were entered into a draw for a 50€ gift voucher. Ethics permission for this study was provided by UCD office of research ethics.

A smaller recipe-corpus of 92,539 recipes with valid images was used as the primary recipe data-set. The study compared four approaches: the three $E_{ns}TM$ based FRS strategies and a CB approach. Each approach predicted user's preference on all 92,539 recipes. For each recommendation strategy, the top 2,100 recipes with highest prediction score were divided into 7 equal sized epochs and from each epoch one recipe was randomly selected. This approach was taken to support diversity and allow users to have more options at their disposal.

We developed a website². and hosted it under the university domain. Participants were first required to access the website and indicate their informed consent and then create a user-name and password. They could then log into a secure website that displayed an interactive panel of images representing all 288 features, in the order of their feature weight. They were asked to select at least 20 features which they like and at least 20 features which they dislike. This information was used to create a user profile. Once created, participants could log into their profile and browse the features to update their likes and dislikes. To populate user's profile for the baseline approach participants were asked to elicit the ingredients they like or eats frequently. Each user had to type in at least 20 ingredients. Participants also selected an appointment time for the main experiment.

During the main experiment participants were shown a series of four recommendation lists corresponding to each of our recommendation algorithms. Each list consisted of seven recipes. The order in which the recommendation lists were presented was fully counter-balanced across the 48

²Demo of the website could be found at <https://youtu.be/ujaB0FiqRwk>

participants. Within each list, participants were required to rate each individual recipe on a 5 star rating scale, where 0 and 5 represented "not like at all" and "liked very much" respectively.

RESULTS

Accuracy: The accuracy of the recommendations has been evaluated based on participant ratings of recipes. For each participant, the average rating across the seven-item list generated by each recommendation strategy was calculated. Figure 2 shows the mean score of each algorithm across all users. The pure CB approach was the poorest performer. This was confirmed through statistical analysis. We first conducted a repeated measures analysis of variance that compared the mean ratings of participants across the four algorithms. The result, $F(3,188) = 14.42229, p < 0.001$, indicates a significant difference within the results. Paired sample t-tests were then conducted between the individual algorithms, with a null hypothesis in each case of no difference in the mean ratings. We do not find a significant difference between participants ratings across the E_{nsTM} approaches, indicating that they all performed equally well in terms of accuracy. There was however a significant difference in participants ratings between each of the E_{nsTM} approaches and the CB baseline, with $p < 0.001$ in each case. This suggests that each E_{nsTM} based approach performed significantly better than the baseline CB approach.

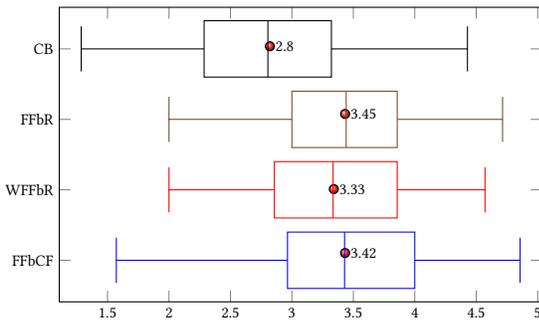


Figure 2: Cumulative preference score from each user

Coverage: Here we consider the coverage achieved by each algorithm across all users, that is, the percentage of recipe-user pairs where the algorithm was able to generate a prediction. Figure 3 details the coverage achieved by each algorithm. The notable outlier is CB, which produced coverage of only 20%. FFbR and WFFbR both had user’s preferences for an average of 51 of our 288 features and both produced a coverage of 91.57%, with predictions for all recipe-user pairs. FFbCF, with a more densely populated user-to-feature matrix,

provided 100% coverage, with predictions for all recipe-user pairs.

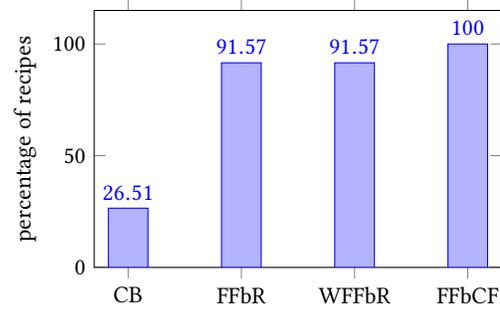


Figure 3: Coverage produced from each RR approach

Implicitly capturing food practices: Another practical aspect of knowledge building for a FRS is an algorithm’s ability to predict important aspects of a user’s food practices from available user information. For example, while both vegetarians and vegans eat vegetables, eggs should only be recommended to vegetarians. Figure 4 shows that the CB baseline performed poorly in this regard. In contrast FFbR, WFFbR identified user’s food practice 100% accurately. Here the feature-to-recipe direct relationship extends the dislike property of the FRS as an effective identifier tool. The reason FFbCF failed to predict food practice for some users is the collaborative effect of their neighbour’s food practice.

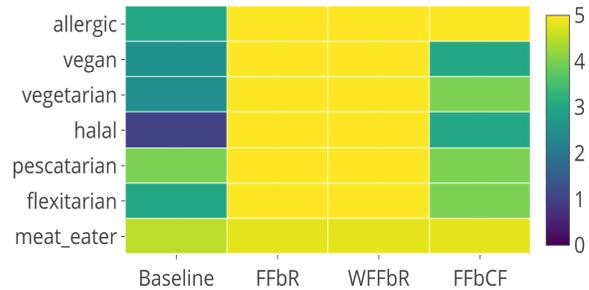


Figure 4: How successfully each method identified user’s food practice

Correlation between lifestyle and food-features: Further analysis on the data-set collected from the user study exposed interesting associations between users’ lifestyle and their feature-preference. Users were categorized under different health-groups based on three different healthiness measures: activity_level, BMI and average food_healthScore. User’s activity_level was a self reported assessment by user. BMI was calculated from users’ height and weight following [4]. User’s average food_healthScore was defined as the average FSA health-score [18] of all recipes user liked (rated 4 or

more). Table 2 summarizes the category labels corresponding to each healthiness measure and the guideline associated with each categorization criteria.

The activity_level and food_healthScore based categorization showed agreement on the healthiness of user's lifestyle preference. Figure 5 illustrates the spread of the 48 participants over different activity based categories. It also illustrates the percentage of each food_healthScore based categories within each activity based categories. The proportion of LessHealthy user-group decreased with the increase in activity level. The BMI based categorization was not predictive of either of activity_level and food_healthScore based categorization.

Scale	Guideline	User Group
Activity level	FAO: activity level, energy intake [3]	sedentary, lightly_active, moderately_active, extra_active
BMI	WHO: BMI [4]	underweight, normal_weight, pre_obesity, obesityclass_1
Food choices	FSA: nutrient intake guideline [1]	less_healthy, moderately_healthy, very_healthy

Table 2: user-groups based on different health variable.

The aim of the categorization was to investigate, if there is any pattern in the interactions between certain health-group and any food features. Finding the correlation between these two variables allows us to assess whether healthier users tend to like or dislike a particular feature. A natural approach for such analysis is the application of machine learning classification algorithms to access the predictive capabilities of these features, although due the small sample size (48 users) and the high degree of imbalance in the class size across all three scales, a simple correlation analysis is used in favour of these methods in this instance.

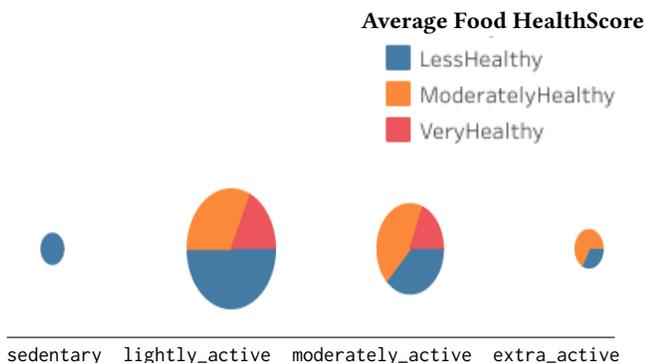


Figure 5: Average Food HealthScore of users in different activity level category

Results expressed interesting associations between health-groups and features. Given that the group/category-level

associated with activity_level and food_healthScore are ordinal in nature, we conducted a Spearman rank correlation analysis [22] to find the degree of association between preference (positive/negative) for features and health-groups. Table 3 shows the strongest significant features with $p < 0.05$ for a sample of 48 users. The coherence between user's personality factors, food_choice and activity_level, negotiates for the features popular among the healthier user-group to be leverages as initial recommendations for new users who are looking for inspiration on healthier food-ideas/recipes.

Average Food HealthScore		Activity Level	
Feature	r	Feature	r
peanut-butter	0.447989	wing	0.441152
granola	0.365171	tuna	0.430467
lentil	0.360767	tilapia	0.363502
indian	0.356347	salmon	0.359852
cauliflower	0.352353	hawaiian	0.346401
low-cholesterol	0.350818	canadian	0.322470
maple	0.321131	smoothy	0.314174
vegetable	0.307459	chicken-thighs-legs	0.314059
wheat	0.303326	halibut	0.310990
carrot	0.303052	main-dish	0.303345

Table 3: Top¹ 10 positively correlated features to healthier user-groups

5 CONCLUSIONS AND FUTURE WORK

This work presents an initial evaluation of $E_{ns}TM$ based FRS. Results show that $E_{ns}TM$ based approaches performs significantly better than a conventional CB approach. It provides a universal feature extraction approach that can generate a set of significant food-features from any recipe/ menu/ food corpus. The features have the added advantage of being human understandable and allowed us to directly model user preferences. $E_{ns}TM$ based feature identification resolves the limitation of user-group dependency and is capable of making food recommendations for users from diverse nationality, ethnicity and culture. It allows for the generation of recommendations without the need for existing user ratings on recipes, helping to address the cold start problem. By working with a reduced feature set, $E_{ns}TM$ also enables computationally efficient recommendation. Furthermore the the subset of nutritional features within our food features supports the proposed $E_{ns}TM$ approaches to personalize the Reclist according user's nutritional preference.

While there was no significant difference between the three $E_{ns}TM$ based approaches in terms of users' recipe ratings, the use of $E_{ns}TM$ in combination with CF provided best coverage, predicting user preferences across 100% of our recipe corpus. However, the CF based approach performed

more poorly in terms of implicit understanding of users' food practices. In future work we aim to focus on applying the $E_{ns}TM$ based recommenders to support diet/menu planning by incorporating health-aware filtering strategies, with the view to providing long-term, guided and healthier food choices. The positive and negative popularity of features among certain health-groups also inspired us to investigate food feature in comparison with healthiness clues for user modeling and recipe recommendation.

REFERENCES

- [1] [n. d.]. FSA Nutrient and Food Guidelines. <https://www.ptdirect.com/training-design/nutrition/national-nutrition-guidelines-united-kingdom> Accessed : March 2018.
- [2] [n. d.]. Geniuskitchen. <http://www.geniuskitchen.com> Accessed : March 2018.
- [3] 2009. FAO energy requirement guideline. <http://www.fao.org/3/y5686e/y5686e07.htm> Accessed :March 2018.
- [4] 2009. WHO : Body mass index. <http://www.euro.who.int/en/health-topics/disease-prevention/nutrition/a-healthy-lifestyle/body-mass-index-bmi> Accessed :March 2018.
- [5] 2019. Investigating and predicting online food recipe upload behavior. *Information Processing and Management* 56, 3 (2019), 654 – 673.
- [6] Carole A Bisogni, Margaret Jastran, Marc Seligson, and Alyssa Thompson. 2012. How People Interpret Healthy Eating: Contributions of Qualitative Research. *Journal of nutrition education and behavior* 44 (07 2012), 282–301.
- [7] Mark Belford, Brian MacNamee, and Derek Greene. 2016. Ensemble Topic Modeling via Matrix Factorization.
- [8] Jesús Bobadilla, Fernando Ortega, Antonio Hernando, and Jesús Bernal. 2012. A Collaborative Filtering Approach to Mitigate the New User Cold Start Problem. *Know.-Based Syst.* 26 (Feb. 2012), 225–238.
- [9] T. Cover and P. Hart. 2006. Nearest Neighbor Pattern Classification. *IEEE Trans. Inf. Theor.* 13, 1 (Sept. 2006), 21–27.
- [10] David Elswailer and Morgan Harvey. 2015. Towards Automatic Meal Plan Recommendations for Balanced Nutrition. In *Proceedings of the 9th ACM Conference on Recommender Systems (RecSys '15)*. 313–316.
- [11] David Elswailer, 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 (SIGIR '17)*. 575–584.
- [12] Jill Freyne and Shlomo Berkovsky. 2010. Intelligent Food Planning: Personalized Recipe Recommendation. In *Proceedings of the 15th International Conference on Intelligent User Interfaces (IUI '10)*. 321–324.
- [13] Jill Freyne and Shlomo Berkovsky. 2010. Recommending Food: Reasoning on Recipes and Ingredients. In *Proceedings of the 18th International Conference on User Modeling, Adaptation, and Personalization (UMAP'10)*. 381–386.
- [14] Mouzhi Ge, Mehdi Elahi, Ignacio Fernáandez-Tobías, Francesco Ricci, and David Massimo. 2015. Using Tags and Latent Factors in a Food Recommender System. In *Proceedings of the 5th International Conference on Digital Health 2015 (DH '15)*. 105–112.
- [15] Mouzhi Ge, Francesco Ricci, and David Massimo. 2015. Health-aware Food Recommender System. In *Proceedings of the 9th ACM Conference on Recommender Systems (RecSys '15)*. 333–334.
- [16] Derek Greene, Derek O'Callaghan, and Pádraig Cunningham. 2014. How Many Topics? Stability Analysis for Topic Models. In *Machine Learning and Knowledge Discovery in Databases*. Springer.
- [17] Hanshen Gu and Dong Wang. 2009. A Content-aware Fridge Based on RFID in Smart Home for Home-healthcare. In *Proceedings of the 11th International Conference on Advanced Communication Technology - Volume 2 (ICACT'09)*. 987–990.
- [18] Morgan Harvey, Bernd Ludwig, and David Elswailer. [n. d.]. Learning user tastes: a first step to generating healthy meal plans?
- [19] Morgan Harvey, Bernd Ludwig, and David Elswailer. 2013. You Are What You Eat: Learning User Tastes for Rating Prediction. In *Proceedings of the 20th International Symposium on String Processing and Information Retrieval - Volume 8214 (SPIRE 2013)*. 153–164.
- [20] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8 (Aug. 2009), 30–37.
- [21] Fang-Fei Kuo, Cheng-Te Li, Man-Kwan Shan, and Suh-Yin Lee. 2012. Intelligent Menu Planning: Recommending Set of Recipes by Ingredients. In *Proceedings of the ACM Multimedia 2012 Workshop on Multimedia for Cooking and Eating Activities (CEA '12)*. 1–6.
- [22] Mavuto Mukaka. 2012. Statistics Corner: A guide to appropriate use of Correlation coefficient in medical research. *Malawi medical journal : the journal of Medical Association of Malawi* 24 (09 2012), 69–71.
- [23] Juan Ramos. 2003. Using TF-IDF to determine word relevance in document queries. (01 2003).
- [24] 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 (UMAP '16)*. 207–215.
- [25] Markus Rokicki, Christoph Trattner, and Eelco Herder. 2018. The Impact of Recipe Features, Social Cues and Demographics on Estimating the Healthiness of Online Recipes. In *ICWSM*.
- [26] Janusz Sobocki, E. Babiak, and M. Slanina. 2006. Application of Hybrid Recommendation in Web-based Cooking Assistant. In *Proceedings of the 10th International Conference on Knowledge-Based Intelligent Information and Engineering Systems - Volume Part III (KES'06)*. 797–804.
- [27] Chun-Yuen Teng, Yu-Ru Lin, and Lada A. Adamic. 2012. Recipe Recommendation Using Ingredient Networks. In *Proceedings of the 4th Annual ACM Web Science Conference (WebSci '12)*. 298–307.
- [28] Alexander Topchy, Anil K. Jain, and William F. Punch. 2005. Clustering ensembles: models of consensus and weak partitions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27, 12 (2005), 1866–1881.
- [29] Christoph Trattner and David Elswailer. 2017. Food Recommender Systems: Important Contributions, Challenges and Future Research Directions. *CoRR* abs/1711.02760 (2017). arXiv:1711.02760 <http://arxiv.org/abs/1711.02760>
- [30] Christoph Trattner and David Elswailer. 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 (WWW '17)*. 489–498.
- [31] Christoph Trattner, Dominik Moesslang, and David Elswailer. 2018. On the predictability of the popularity of online recipes. *EPJ Data Science* 7, 1 (05 Jul 2018), 20.
- [32] Christoph Trattner, Markus Rokicki, and Eelco Herder. 2017. On the Relations Between Cooking Interests, Hobbies and Nutritional Values of Online Recipes: Implications for Health-Aware Recipe Recommender Systems. In *Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization (UMAP '17)*. 59–64.
- [33] Tsuguya Ueta, Masashi Iwakami, and Takayuki Ito. 2011. A Recipe Recommendation System Based on Automatic Nutrition Information Extraction. In *Proceedings of the 5th International Conference on Knowledge Science, Engineering and Management (KSEM'11)*. 79–90.
- [34] Youri van Pinxteren, Gijs Geleijnse, and Paul Kamsteeg. 2011. Deriving a Recipe Similarity Measure for Recommending Healthful Meals. In *Proceedings of the 16th International Conference on Intelligent User Interfaces (IUI '11)*. 105–114.