

# SHARE: A Framework for Personalized and Healthy Recipe Recommendations

Konstantinos Zioutos<sup>1</sup>, Haridimos Kondylakis<sup>2</sup> and Kostas Stefanidis<sup>3</sup>

<sup>1</sup>Computer Science Department, UOC

<sup>2</sup>FORTH-ICS

<sup>3</sup>Tampere University

## Abstract

This paper presents a personalized recommendation system that suggests recipes to users based on their health history and similar users' preferences. Specifically, the system utilizes collaborative filtering to determine other users with similar dietary preferences and exploits this information to identify suitable recipes for an individual. The system is able to handle a wide range of health constraints, preferences, and specific diet plans, such as low-carb or vegetarian. We demonstrate the usability of the system through a series of experiments on a large real-world data set of recipes. The results indicate that our system is able to provide highly personalized and accurate recommendations.

## Keywords

Collaborative Filtering, Content-Based, Machine Learning, Personalization, Recommendation Systems

## 1. Introduction

Recommendation systems (RS) nowadays are becoming increasingly important across many industries, using data mining and machine learning to analyze large amounts of data and make personalized recommendations to users [1, 2]. By providing tailored, relevant recommendations they can help users discover new products or services while helping businesses increase sales, engagement, and customer retention [3, 4, 5, 6]. Eminent organizations such as Amazon, Netflix, and YouTube have implemented recommender systems [7, 8, 9] to improve user experience. Tailored recommendations are generated based on viewing histories and preferences of customers in order to deliver more applicable content quickly; significantly reducing browsing time for users.

Recommendation systems have become a fundamental part of our lives [10, 11, 12]. While applications are prevalent in many areas, their implementation for food and recipes is surprisingly limited. Nevertheless, it could be highly beneficial as individuals with special dietary needs or health issues can make informed choices about what recipes to make. Therefore, there exists an immense potential for utilizing recommenders within this field.

On the other hand, the prevalence of chronic health conditions has risen to alarming levels, and their association with morbidity and mortality is well-documented [13, 14]. To reduce the risk of developing or exacerbating these conditions - as could be caused by poor dietary

choices - adopting a healthy well-balanced diet appears crucial. However, finding recipes that meet specific dietary requirements and preferences can be challenging and time-consuming for individuals as they need to search for recipes that fit their needs manually. This is especially true for those who suffer from a particular condition that requires them to follow special diet plans, such as high-protein, gluten-free, or other specific diets.

To address this problem, we develop a hybrid recommendation system that utilizes a combination of collaborative filtering (CF) and content-based filtering (CB) methods to make personalized recommendations. The CF component of the system finds other users with similar preferences and uses this information to recommend recipes that are tailored to the individual user's needs. The CB component of the system analyzes the nutritional information of the recipes to recommend recipes that are suitable for the user's health condition(s). Our approach centers around the idea of recommending healthy personalized recipes uniquely tailored to each user's specific needs. It can accommodate a wide variety of health constraints. This enables us to provide helpful suggestions suitable even for those suffering from chronic diseases who require strict attention when choosing recipes.


Specifically, in this paper, we describe the design and implementation of a python framework called SHARE. Our framework contains a recommender engine, that contains collaborative and hybrid filtering approaches for recipes recommendation. To evaluate the usability of SHARE, we conducted with 40 real users a survey, consisting of 6 questions, each one designed to assess different aspects of the framework's usability. These aspects are the accuracy, personalization, user acceptance, overall coverage, and explainability of the results. The

Published in the Workshop Proceedings of the EDBT/ICDT 2023 Joint Conference (March 28-March 31, 2023), Ioannina, Greece

✉ csd3946@csd.uoc.gr (K. Zioutos); kondylak@ics.forth.gr

(H. Kondylakis); konstantinos.stefanidis@tuni.fi (K. Stefanidis)

© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

survey respondents also indicated their preferred method. The results indicate that the system provides a wide variety of highly relevant personalized recipes to users and excellent justifications

The rest of the paper is as follows: Section II presents related works. Section III presents our framework. Section IV evaluates the SHARE framework, and finally, Section V ends with a conclusion and future work.

## 2. Related Work

There has been a significant amount of previous work on recommendation systems in the field of food and recipe recommendation based on user health.

One line of research has focused on developing content-based RS, which recommends items based on their features and characteristics. For example, Herlocker et al. [15] present a content-based RS for music, which recommends songs based on the audio features of the songs and the user's previous listening history. Similarly, Freyne et al. [16, 17] propose a CB recipe recommender that recommends recipes based on the ingredients and cooking methods of the recipes.

Another line of research has focused on collaborative filtering methods, which recommend items based on the preferences of similar users. For example, Resnick et al. [18] present a CF method for a movie recommendation, which recommends movies to users based on the ratings of similar users. Likewise, Freyne et al. [19], propose a collaborative filtering method for recipe recommendation, which recommends recipes to users based on the ratings of similar users.

In the health domain, [20, 21] proposes a semantic similarity function that takes into account the patients' medical profiles and shows its superiority over traditional similarity measures in group recommendations. [22] focuses on the notion of fairness, devising an aggregation method for ensuring that if the group recommendation list provides a high relevant document for a patient, then that patient may be tolerant of the existence of documents that are not relevant to him/her. More recently, [23] exploits as well additional properties for producing group recommendations, like the education and health literacy level, and the psycho-emotional status of the group.

As concerns, to recipe recommendations based on dietary preferences or restrictions, Agapito et al. [24] proposed a personalized recipe RS that takes into account users' health profiles and chronic diseases, such as CKD, hypertension, and diabetes. In a similar way, Yang et al. [25] developed a food RS that takes into account users' nutritional expectations, dietary restrictions, and fine-grained food preferences.

Regarding the incorporation of user-generated data, such as recipe ratings, and reviews into the recommen-

ation process, Tian et al. [26] proposed a recipe RS that incorporates user history behavior and user feedback such as the ratings toward recipes, which make the recommendation accounts for user interest and preferences. Similarly, Pessemier et al. [27] present a food recommendation strategy for patients in a care facility that utilizes explicit ratings for menu items, implicit feedback based on the patient's eating behavior and the amount of food that was eaten, and inferred preferences for the ingredients of the menu items.

In addition, there are other approaches such as Hybrid methods, which combine two or more different recommendation techniques, for example, Gaudani and Hetal [28] proposed a hybrid approach that combines CB and CF algorithms to recommend recipes.

SHARE is closely related to these previous studies and builds upon them by combining both content-based and collaborative filtering methods to develop a personalized recipe recommender system that takes into account users' preferences and health history. However, SHARE differs from the above works, because we also combine a knowledge-based component. Also, in content-based filtering where we use this component to analyze the nutritional information of the recipes to recommend recipes that are suitable for the user's chronic disease. Overall according to our research, there is no other hybrid approach that combines collaborative filtering, content-based filtering, and knowledge-based methods where the CB component extracts tags from recipe descriptions including nutritional information and other relevant characteristics.

## 3. The SHARE Framework

In this section, we describe the SHARE framework that develops our recommendation system. We begin by describing the collaborative filtering approach that we used to generate recommendations, secondly, we discuss a personalized recommendation method. After that, we describe the application of a personalized filtering approach. Finally, we discuss the explanations that SHARE provides to users about why they are receiving each recommendation.

### 3.1. User-Based Collaborative Filtering

To generate recommendations for a given user, we apply CF using the user's ratings and the ratings of other similar users.

Before applying the similarity measure, we normalize the ratings of all users by subtracting the mean rating of each user from their ratings. This has the effect of centering the ratings around the mean, with positive ratings indicating ones that are higher than the mean

and negative ratings indicating ratings that are lower than the mean[29].

Normalizing the ratings in this way helps to take into account the differences in the absolute rating scales used by different users. For example, one user may tend to rate all recipes as 5 stars, while another user may rate the same recipes as 1 star. Without normalization, the similarities between these two users would be artificially low due to the differences in their rating scales. Normalization assists in fixing this by adjusting the ratings to a common scale. Normalization is an important step in the CF process because it helps to ensure that the similarities between users are based on their relative preferences rather than their absolute rating scales [29].

To identify similar users, we first compute the similarity between each pair of users based on the cosine similarity measure. The cosine similarity measures the similarity between two non-zero vectors of inner product space and is defined as the cosine of the angle between the vectors.<sup>1</sup> The resulting similarity ranges from -1 meaning completely dissimilar, to 1 meaning completely similar, with 0 indicating orthogonality, while in-between values indicate intermediate similarity or dissimilarity. We use it to measure the similarity between the ratings of two users.<sup>2</sup> The formula for the cosine similarity is as follows:

$$\text{sim}(A, B) = \frac{AB}{\|A\| \|B\|} \quad (1)$$

To compute the cosine similarity, we first convert the ratings of each user into a term-frequency representation, which represents the frequency with which each rating appears in an individual's rankings<sup>3</sup>. For example, if someone rated 3 recipes as 5 stars and then 2 more at 4 with 1 left over at 3; their corresponding term-frequency would be [3,2,1].

Once we have the term-frequency representations of the ratings of two users, we compute the dot product of the vectors by multiplying the corresponding elements of the vectors and summing the results. The dot product is then divided by the product of the magnitudes of the vectors to give the cosine similarity. The magnitude of a vector is the square root of the sum of the squares of the elements of the vector<sup>4</sup>. Intuitively, we use the cosine similarity because it can handle sparse data and does not require the ratings to be normally distributed<sup>5</sup>.

To generate recommendations for a given user, we use the ratings of a set of similar users to predict the rating

<sup>1</sup><https://medium.com/@riyasonline/cosine-similarity-is-a-measure-of-similarity-between-two-non-zero-vectors-of-an-inner-product-caa3cd05c10f>

<sup>2</sup>[https://en.wikipedia.org/wiki/Cosine\\_similarity](https://en.wikipedia.org/wiki/Cosine_similarity)

<sup>3</sup><https://blog.marketmuse.com/glossary/term-frequency-definition/>

<sup>4</sup><https://wumbo.net/formulas/magnitude-of-vector/>

<sup>5</sup>[https://en.wikipedia.org/wiki/Cosine\\_similarity](https://en.wikipedia.org/wiki/Cosine_similarity)

that the user would give to each recipe. The formula for the predicted rating of a recipe,  $r_t$ , by user  $u_a$  is as follows:

$$\text{Pred}(u_a, r_t) = \bar{r}_{u_a} + \frac{\sum_{n \in N} \text{sim}(u_a, u_n) \text{rat}(u_n, r_t)}{\sum_{n \in N} \text{sim}(u_a, u_n)} \quad (2)$$

In this formula, the  $N$  ("neighbors") are the users who are most similar to the target user  $u_a$ , as determined by the cosine similarity measure. The predicted rating is a weighted average of the ratings of the neighbors, with the weights being the cosine similarities  $\text{sim}(u_a, u_n)$ , between the users. Finally, the average rating  $\bar{r}_{u_a}$  of the target user  $u_a$  is included in the prediction to take into account the fact that different users may have different overall rating tendencies. Because as we discussed above for "Normalization", one user may tend to give higher ratings to recipes overall, while another user may tend to give lower ratings. Including the average rating of the user in the prediction formula helps to adjust for these differences in rating tendencies so that the predictions are more accurate and reflective of the user's true preferences.

Once the predictions have been computed for all recipes, we can rank them and recommend the top-rated recipes to the user.

### 3.2. Health Personalized Recommendation Method

To take into account the health needs of the individual user, we incorporated personalization techniques that allow the system to consider the user's health history.

To achieve this, we combined collaborative filtering, content-based filtering, and knowledge-based methods. CF relies on the preferences of similar users to recommend items [30], while CB uses the characteristics of the item itself to make recommendations [31]. Knowledge-based recommendation systems use a combination of explicit knowledge about the items and the preferences of the users to make recommendations [32].

For the content-based component, tags were extracted from recipe descriptions including nutritional information and other relevant characteristics of the recipes. These features formed a vector for each recipe, which incorporated the recipe's nutritional information and other relevant attributes.

For the knowledge-based component, we identify the specific nutrients that are suitable for every chronic health condition supported by SHARE. The data was collected

by official statistics<sup>6789</sup>. We used this information to calculate the nutritional profile that is most suitable for the user's health needs, taking into account their specific health condition(s). We then used this nutritional profile to create a target vector, which represents the types of recipes that are suitable for the user based on their specific health needs.

Finally, we used the cosine similarity between the feature and target vectors to identify recipes that are less suited to the nutritional profile of the user and we exclude them from the collaborative filtering we apply after.

Once the CF process is finished, the system will have produced a list of recommended recipes that are customized to meet the individual's health requirements and personal preferences.

### 3.3. Personalized Filtering

We apply a personalized filtering technique in both methods to improve the accuracy and relevance of recommendations in the two methods we discussed above.

SHARE offers personalized filtering, allowing users to customize the recommendations based on their tastes and dietary needs. Through personalized filtering, users are able to narrow down their recommendations based on tags extracted from recipe descriptions (e.g. vegan or gluten-free) and desired nutritional values like calories, protein, and saturated fat.

This elevated level of customization helps ensure that our recommended recipes accurately meet individual preferences, resulting in highly relevant and useful results for maximized satisfaction.

### 3.4. Explainability

One of the challenges of using machine learning methods for recommendation systems is the lack of explainability of the results. This can make it difficult for users to understand why a particular recommendation was made.

To address these issues, SHARE includes an explainability component that provides users with a clear and concise explanation of the reasons behind each recommendation.

The explanations are generated using a combination of natural language and domain-specific knowledge that includes recipe tags, nutritional properties as well as information about the preferences and dietary restrictions of the users. The system is designed to process this

<sup>6</sup><https://www.greenfacts.org/en/diet-nutrition/index.html/>

<sup>7</sup><https://www.hopkinsmedicine.org/health/conditions-and-diseases/cancer/cancer-diet-foods-to-add-and-avoid-during-cancer-treatment/>

<sup>8</sup><https://stanfordhealthcare.org/medical-clinics/cancer-nutrition-services.html/>

<sup>9</sup><https://www.eatingwell.com/our-food-nutrition-philosophy/>

information and use it to generate human-readable explanations for the recommended recipes. To do this, SHARE uses a database for storing and retrieving information about recipes, users, and chronic diseases, allowing the system to generate explanations that are personalized to the user based on their preferences and specific health needs.

For example, if a recommendation is made to a user with obesity, the explanation might highlight the fact that the recommended recipe is low in saturated fat and high in fiber, and that these characteristics are beneficial for managing obesity.

The justifications are provided by the system through tables like Table I below.

The explainability component of our recommender system is designed to provide users with a greater understanding of the reasoning behind the recommendations and to enable them to make more informed decisions about which recipes to try.

## 4. Experiments

### 4.1. Dataset / Set-Up

The dataset used for our experiments was obtained from food.com<sup>10</sup> containing real-world recipe ratings collected from numerous users, forming an ideal source for our RS. The dataset consisted of user IDs and associated recipe IDs with corresponding rating scores on a scale of 1 to 5 that reflected the level of satisfaction expressed by each individual user towards the respective recipe in question. All entries were arranged into a CSV file wherein the first column corresponded to the ID signifying specific individuals, followed by a column representing the recipe identifiers, and the last one corresponding to the given rating.

By utilizing the health history of each user, we are able to make more personalized recommendations tailored to their health needs. To do so, we gather data from official statistics<sup>111213141516</sup> and we assigned a collection of Boolean attributes outlining whether the users of the dataset have any chronic conditions or not. This allows us to generate suggestions that take into account people's medical history making it easier for them to manage

<sup>10</sup><https://www.kaggle.com/datasets/shuyangli94/food-com-recipes-and-user-interactions>

<sup>11</sup><https://www.who.int/news/item/04-03-2022-world-obesity-day-2022-accelerating-action-to-stop-obesity>

<sup>12</sup><https://diabetesresearch.org/diabetes-statistics/>

<sup>13</sup><https://www.cdc.gov/heartdisease/facts.htm>

<sup>14</sup><https://www.wcrf.org/cancer-trends/worldwide-cancer-data/>

<sup>15</sup><https://www.who.int/news-room/fact-sheets/detail/oral-health>

<sup>16</sup><https://www.bonehealthandosteoporosis.org/wp-content/uploads/2015/12/Osteoporosis-Fast-Facts.pdf>

**Table 1**

System's justification example

Recommended Recipes	Health History	The reason
emeril's essence	Obese	low-saturated-fat,high calcium
pumkin biscuits	Obese	high calcium
microwave cornbread	Obese	high calcium
pressed Cuban sandwich	Obese	high calcium
turkish cornbread	Obese	high calcium

chronic illnesses when proceeding with any decisions relating to food.

## 4.2. Experimental Methodology

In the experiment, four methods were utilized including, Collaborative Filtering(CF) discussed in subsection III/A, Health Collaborative Filtering(HCF) discussed in subsection III/B, Personalized Collaborative Filtering(PCF) which combines the CF method and Personalized filtering discussed in subsection III/C, Health Personalized Collaborative Filtering(HPCF) which combines the HCF method and Personalized filtering discussed in subsection III/C.

The components of each method are shown in Table III.

**Table 2**

System's methods' components

Method	Component
CF	Collaborative Filtering
PCF	CF,Personalized Filtering
HCF	CF,CB,KB
PHCF	CF,CB,KB,Personalized Filtering

To evaluate the usability of our four methods, we conducted a survey with 40 real users with 40% being between 18-24 years old, 40% between 25-49 years old, and 20% being 50-59 years old. The participants had a diverse range of educational and professional backgrounds, including 40% with a computer science background, 30% with a healthcare background, 10% with a cooking background, and 20% with a general background.

The survey consisted of 6 questions, which were designed to assess different aspects of the methods' usability. Each question focuses on a specific aspect of usability. Specifically:

- The first question assessed the accuracy of the results in each method separately. This question asked participants to rate on a 5-point rating scale, their satisfaction with the provided recommendations based on users' past behavior. This can provide valuable insights into how well the RS is meeting the needs of the users.
- The second question assessed the personalization of the results in each method separately. This

question asked participants to rate on a 5-point rating scale whether the recommendations provided were helpful and relevant for the specific health problem they are facing. This question can help us understand whether the system is effective at providing useful and relevant recommendations to users who are seeking recipes suitable to specific health problems.

- The third question assessed the user acceptance of the results in each method separately. This question asked participants to state how many of the suggested recipes they find appealing. This question can help us understand whether the system is effective at providing recommendations that the users are interested in or that meet their needs.
- The fourth question asked participants to state which method they preferred. This question can help us understand which method is the most famous among users.
- The fifth question assessed the overall coverage of the results. This question asked participants to rate on a 5-point rating scale, their satisfaction with the variety of provided recommendations. By assessing the coverage of the results, we can get an indication of how well the recommender is able to suggest a wide variety of recipes to users.
- The sixth question assessed the overall explainability of the system. This question asked participants to rate on a 5-point rating scale the justification provided for the recommended recipes is sufficient for them to make a decision and whether the popularity of the recipes influences their decision-making process. These questions can help us understand how users make decisions based on the recommendations provided, and whether the system is providing sufficient information to support those decisions.

The questions presented to the real users are shown in Table III.

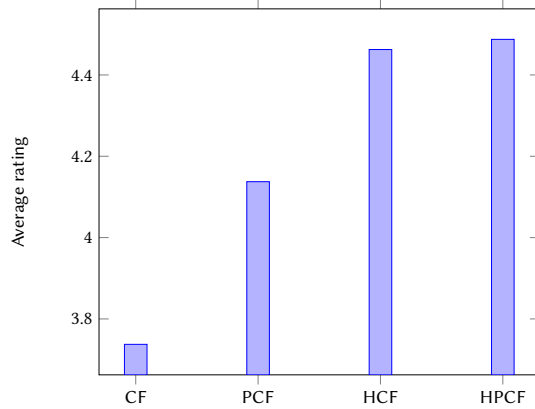
## 4.3. Results

The results of the survey presented in the previous section are shown in the sequel.

**Table 3**  
Survey Questions

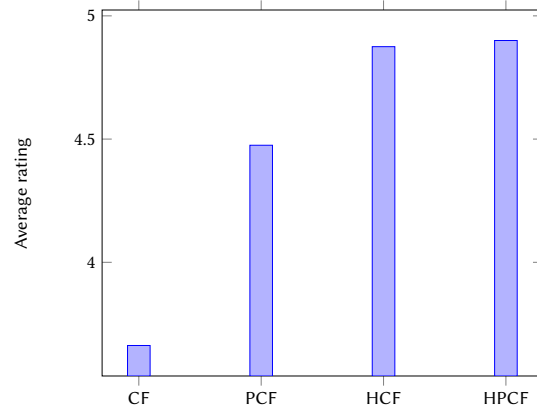
No.	Question
1	Do the results reflect your personal preferences?
2	Are the results helpful for the health problem you are facing?
3	How many of the suggested recipes do you find appealing?
4	Which method do you prefer?
5	Is the variety of proposals satisfactory?
6	Is the justification sufficient to choose a recommended recipe?

The results of the first question, focusing on the accuracy of the results, are presented in Fig. 1. As shown, methods HCF and HPCF have the most accurate results with HPCF results being slightly more accurate than the results of HCF. The PCF method has a moderate rating compared to the above. On the other hand, the CF method has the lowest rating among all.



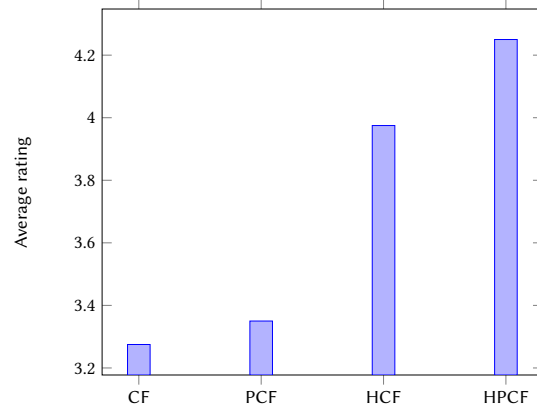
**Figure 1:** Results on the accuracy aspect.

The results for the second question, focusing on the personalization of the results, are shown in Fig. 2. According to the figure HCF and HPCF have more personalized results with HPCF results being slightly more personalized than the results of HCF. The PCF method has a moderate rating compared to the above. On the other hand, the CF method has the lowest rating among all.



**Figure 2:** Results on the personalization aspect.

The results of the third question, which asked about user acceptance of the results, as shown in Fig. 3, indicate that HPCF has the most appealing results. The HCF method has a moderate rating compared to HPCF. Finally, the other two methods have a considerably lower score than the above methods.



**Figure 3:** Results on user acceptance.

The results of the fourth question, which asked about the most favored method, as shown in Fig. 4, indicate that HPCF is clearly the most famous method among users. The other three methods have much lower results,

with CF having an overall score equal to zero and PCF also close to it.

Finally, the results of questions 5 and 6, which asked about the overall coverage and justification of the results of all methods, are shown in Table IV. The results indicate that the system provides a wide variety of recipes to users and excellent justifications.

Overall, the results of the survey indicate that the



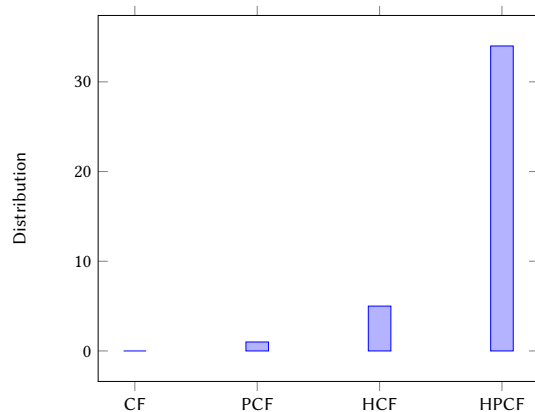


Figure 4: Results on the preferred method.

Table 4

Average Users' Rating	Questions	
	Question 5	Question 6
	4.6	5

Health Personalized Collaborative Filtering method is the most accurate, personalized, and well-received method among users with a wide variety of well-justified results.

## 5. Conclusion and Future Work

In this paper, we presented a personalized recommendation system that recommends recipes to users based on their health history and the preferences of similar users. Overall, the SHARE framework combines user tastes and nutritional information about the recipes in order to provide recommendations for recipes that meet the user's preferences and specific health needs. We also offer personalized filtering for the users of the system. Finally, we evaluate its usability through a series of experiments on a large real-world data set of recipes. Our experiments demonstrate the system's ability to provide highly relevant personalized recommendations.

There are several directions in which the work presented in this paper could be extended in the future. One possible extension is to incorporate additional types of user data, such as age, gender, allergies, exercise habits, and physical activity levels, to make more informed recommendations. Another possible future direction is enriching the RS by including more factors, such as cultural background or social connections. This could allow the system to suggest recipes that are more likely to be well-received by the user's friends and family. Finally, we believe it is worth trying to expand the system by consid-

ering other factors beyond the user's health history, such as the user's location, the season, the availability, and the cost of ingredients, to make more contextually relevant recommendations. Overall, there are many exciting possibilities for improving the performance and usability of the recipe recommender presented in this paper.

## References

- [1] H. Kondylakis, L. Koumakis, M. Psaraki, G. Troullinou, M. Chatzimina, E. Kazantzaki, K. Marias, M. Tsiknakis, Semantically-enabled personal medical information recommender, in: S. Villata, J. Z. Pan, M. Dragoni (Eds.), Proceedings of the ISWC 2015 Posters & Demonstrations Track, volume 1486 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2015. URL: [http://ceur-ws.org/Vol-1486/paper\\_49.pdf](http://ceur-ws.org/Vol-1486/paper_49.pdf).
- [2] L. K. M. C. E. K. K. M. M. T. Galatia Iatraki, Haridimos Kondylakis, Personal health information recommender: implementing a tool for the empowerment of cancer patients, *Ecancermedalscience* 12 (2018) 851. doi:10.3332/ecancer.2018.851.
- [3] B. Pathak, R. Garfinkel, R. Gopal, R. Venkatesan, F. Yin, Empirical analysis of the impact of recommender systems on sales, *J. of Management Information Systems* 27 (2010) 159–188. doi:10.2307/29780174.
- [4] A. Das, M. Datar, A. Garg, S. Rajaram, Google news personalization: Scalable online collaborative filtering, 2007, pp. 271–280. doi:10.1145/1242572.1242610.
- [5] M. Dias, D. Locher, M. Li, W. El-Deredy, P. Lisboa, The value of personalised recommender systems to e-business, 2008, pp. 291–294. doi:10.1145/1454008.1454054.
- [6] K. Stefanidis, E. Ntoutsis, M. Petropoulos, K. Nørnvåg, H. Kriegel, A framework for modeling, computing and presenting time-aware recommendations, *Trans. Large Scale Data Knowl. Centered Syst.* 10 (2013) 146–172.
- [7] X. Amatriain, J. Basilico, Recommender systems in industry: A netflix case study, 2015, pp. 385–418. doi:10.1007/978-1-4899-7637-11.
- [8] B. Smith, G. Linden, Two decades of recommender systems at amazon.com, *IEEE Internet Computing* 21 (2017) 12–18. doi:10.1109/MIC.2017.72.
- [9] J. Davidson, B. Liebald, J. Liu, P. Nandy, T. Vleet, U. Gargi, S. Gupta, Y. He, M. Lambert, B. Livingston, D. Sampath, The youtube video recommendation system, 2010, pp. 293–296. doi:10.1145/1864708.1864770.
- [10] E. Ntoutsis, K. Stefanidis, K. Nørnvåg, H. Kriegel, Fast group recommendations by applying user clustering, in: *Conceptual Modeling - 31st International*

- Conference ER 2012, Florence, Italy, October 15-18, 2012. Proceedings, volume 7532 of *Lecture Notes in Computer Science*, Springer, 2012, pp. 126–140.
- [11] E. Ntoutsi, K. Stefanidis, K. Rausch, H. Kriegel, "strength lies in differences": Diversifying friends for recommendations through subspace clustering, in: Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, CIKM 2014, Shanghai, China, November 3-7, 2014, ACM, 2014, pp. 729–738.
- [12] M. Stratigi, E. Pitoura, K. Stefanidis, SQUIRREL: A framework for sequential group recommendations through reinforcement learning, *Inf. Syst.* 112 (2023) 102128.
- [13] W. Raghupathi, V. Raghupathi, An empirical study of chronic diseases in the united states: A visual analytics approach, *International Journal of Environmental Research and Public Health* 15 (2018) 431. doi:10.3390/ijerph15030431.
- [14] S. Fanelli, S. Jonnalagadda, J. Pisegna, O. Kelly, J. Krok-Schoen, C. Taylor, Poorer diet quality observed among us adults with a greater number of clinical chronic disease risk factors, *Journal of primary care community health* 11 (2020) 2150132720945898. doi:10.1177/2150132720945898.
- [15] J. Herlocker, J. Konstan, L. Terveen, J. C. Lui, T. Riedl, Evaluating collaborative filtering recommender systems, *ACM Transactions on Information Systems* 22 (2004) 5–53. doi:10.1145/963770.963772.
- [16] J. Freyne, S. Berkovsky, Intelligent food planning: personalized recipe recommendation, 2010, pp. 321–324. doi:10.1145/1719970.1720021.
- [17] J. Freyne, S. Berkovsky, Recommending food: Reasoning on recipes and ingredients, volume 6075, 2010, pp. 381–386. doi:10.1007/978-3-642-13470-8\_36.
- [18] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, J. Riedl, Grouplens: An open architecture for collaborative filtering of netnews (1994). doi:10.1145/192844.192905.
- [19] J. Freyne, S. Berkovsky, G. Smith, Recipe recommendation: Accuracy and reasoning, volume 6787, 2011, pp. 99–110. doi:10.1007/978-3-642-22362-4\_9.
- [20] M. Stratigi, H. Kondylakis, K. Stefanidis, Fairness in group recommendations in the health domain, in: 33rd IEEE International Conference on Data Engineering, ICDE 2017, San Diego, CA, USA, April 19-22, 2017, IEEE Computer Society, 2017, pp. 1481–1488.
- [21] M. Stratigi, H. Kondylakis, K. Stefanidis, The fairgrecs dataset: A dataset for producing health-related recommendations, in: Proceedings of the First International Workshop on Semantic Web Technologies for Health Data Management, SWH@ISWC 2018, Monterey, California, USA, October 9, 2018, volume 2164 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2018.
- [22] M. Stratigi, H. Kondylakis, K. Stefanidis, Fairgrecs: Fair group recommendations by exploiting personal health information, in: Database and Expert Systems Applications - 29th International Conference, DEXA 2018, Regensburg, Germany, September 3-6, 2018, Proceedings, Part II, volume 11030 of *Lecture Notes in Computer Science*, Springer, 2018, pp. 147–155.
- [23] M. Stratigi, H. Kondylakis, K. Stefanidis, Multidimensional group recommendations in the health domain, *Algorithms* 13 (2020) 54.
- [24] G. Agapito, M. Simeoni, B. Calabrese, I. Carè, T. Lamprinouidi, P. Guzzi, A. Pujia, G. Fuiano, M. Cannataro, Dietos: A dietary recommender system for chronic diseases monitoring and management, *Computer Methods and Programs in Biomedicine* 153 (2017). doi:10.1016/j.cmpb.2017.10.014.
- [25] L. Yang, A. Hsieh, H. Yang, J. Pollak, N. Dell, S. Belongie, C. Cole, D. Estrin, Yum-me: A personalized nutrient-based meal recommender system, *ACM Transactions on Information Systems* 36 (2017) 1–31. doi:10.1145/3072614.
- [26] Y. Tian, C. Zhang, R. Metoyer, N. Chawla, Recipe recommendation with hierarchical graph attention network, *Frontiers in Big Data* 4 (2022). doi:10.3389/fdata.2021.778417.
- [27] T. De Pessemier, S. Dooms, L. Martens, A food recommender for patients in a care facility, 2013, pp. 209–212. doi:10.1145/2507157.2507198.
- [28] H. Gaudani, Personalized recipe recommendation system using hybrid approach, *International Journal of Innovative Research in Computer and Communication Engineering* 5 (2016) 192–197. doi:10.17148/IJARCC.2016.5642.
- [29] A. N. Nikolakopoulos, X. Ning, C. Desrosiers, G. Karypis, Trust Your Neighbors: A Comprehensive Survey of Neighborhood-Based Methods for Recommender Systems, Springer US, New York, NY, 2022, pp. 39–89. URL: [https://doi.org/10.1007/978-1-0716-2197-4\\_2](https://doi.org/10.1007/978-1-0716-2197-4_2). doi:10.1007/978-1-0716-2197-4\_2.
- [30] B. Schafer, B. J. D. Frankowski, Dan, Herlocker, Jon, Shilad, S. Sen, Collaborative filtering recommender systems, 2007.
- [31] R. Meteren, Using content-based filtering for recommendation (2000).
- [32] B. Towle, C. Quinn, Knowledge based recommender systems using explicit user models (2000).