

Recipe Networks and the Principles of Healthy Food on the Web*

Charalampos Chelmis, Bedirhan Gergin

Computer Science, State University of New York at Albany
1400 Washington Avenue
Albany, NY 12222 USA
cchelmis@albany.edu, bgergin@albany.edu

Abstract

People increasingly use the Internet to make food-related choices, prompting research on food recommendation systems. Recently, works that incorporate nutritional constraints into the recommendation process have been proposed to promote healthier recipes. Ingredient substitution is also used, particularly by people motivated to reduce the intake of a specific nutrient or in order to avoid a particular category of ingredients due to, for instance to allergies. This study takes a complementary approach towards empowering people to make healthier food choices by simplifying the process of identifying plausible recipe substitutions. To achieve this goal, this work constructs a large-scale network of similar recipes, and analyzes this network to reveal interesting properties that have important implications to the development of food recommendation systems.

Introduction

People use the Internet daily to make food-related choices (Li, Mirosa, and Bremer 2020), such as meal planning using ingredients they have readily available (Boulos et al. 2015), modifying their diets by substituting ingredients in their recipes to avoid allergens or in order to adhere to some dietary constraint (e.g., diabetes) (Çelik 2015), or establishing and maintaining healthy and balanced eating habits (Watanabe-Ito, Kishi, and Shimizu 2020).

Unfortunately, recent studies have shown that popular recipe websites exhibit a high prevalence of unhealthy recipes (Trattner and Elsweiler 2017). Furthermore, highly rated and/or popular recipes on such websites have been shown to positively correlate with high fat, sugar, cholesterol, and calorie levels (Chelmis and Gergin 2021). With the goal of recommending healthier recipes, healthy food recommendation systems have recently been proposed (Elsweiler, Hauptmann, and Trattner 2022). Such systems are designed to either filter recipes based on specific health-related properties, create long-term meal plans, or substitute specific ingredients in recipes to obtain healthier recipe alternatives. On the other hand, recipe websites provide lists

*This material is based upon work supported by the National Science Foundation under Grant No. CRII:III-1850097. Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

of common ingredient substitutions¹ for people themselves to make their own choices according to their needs. However, determining “healthiness” has proven to be an unignorable challenge in itself (Howard, Adams, and White 2012), whereas determining if potential ingredient substitutions are “healthier”, “better” with respect to one’s specific dietary needs, or even viable options with respect to overall taste and flavor given the other ingredients in a recipe, has proven to be challenging for both humans and machine alike (Ahn et al. 2011). In summary, prior work focuses mainly on ingredient substitution, with only a limited number of research examining ingredient substitution in conjunction with constraints on nutritional information (Elsweiler, Hauptmann, and Trattner 2022).

Our work takes a complimentary approach towards empowering people to make healthier food choices by simplifying the process of identifying plausible *recipe substitutions* as opposed to ingredient substitutions. To achieve this goal, we construct a network of recipes from RecipeKG (Chelmis and Gergin 2022), a knowledge graph of 77,835 recipes sourced from Allrecipes.com², their popularity and rating, ingredients and corresponding nutritional information, and healthiness information calculated based on two international nutritional standards. We show that the network of recipes has two interesting properties. First, recipes with many common ingredients exhibit similar healthiness scores. Second, given a recipe, a healthier recipe substitution can be obtained from the recipe’s immediate network, hinting towards computationally efficient solutions for future recommendation systems.

The remainder of this paper is structured as follows: Related and prior work is reviewed in Related Work Section. The Dataset Section outlines the dataset used to create the network, while the Recipe Networks Section explains the network construction process. Sections Paradoxes in Recipe Networks, and Results, present and discuss the findings of our analysis of the recipe networks. Finally, conclusion, potential impact and future research directions are discussed in

¹For example, Allrecipes ingredient substitutions guide is available at: <https://www.allrecipes.com/article/common-ingredient-substitutions/>.

²Allrecipes.com is one of the most extensive food-focused social network and recipe website, with the largest food-related traffic volume on the Internet (AELIEVE 2021).

the Conclusion and Potential Impact Section.

Related Work

Healthiness of Web Recipes. Only few works have studied the healthiness of online recipes. Specifically, (Howard, Adams, and White 2012) compared dishes from television chefs and ready meals sold in the supermarket. Even though that study was limited to only 100 recipes for each of the two categories, obtained from only 3 supermarkets and 5 recipe books, the subsequent work presented in (Trattner and Elsweiler 2017) included recipes from Allrecipes.com to the comparison. More recently, (Chelmis and Gergin 2021) conducted a large-scale analysis of the healthiness of Allrecipes.com recipes, corroborating previous findings while more accurately capturing recipes categorization, and automating the process of calculating recipe healthiness using dietary guidelines from the World Health Organization (Consultation 2003), and the United Kingdom Food Standards Agency (FSA 2014).

Healthiness Paradox. The friendship paradox states that on average, most people have fewer friends than their own friends (Feld 1991). The paradox has been empirically demonstrated for both online and offline social networks (Hodas, Kooti, and Lerman 2013; Eom and Jo 2014). This work postulates that a similar paradox occurs in online recipe websites, much similar to the popularity paradox in online social networks. This new paradox has a surprising twist: recipes in the immediate neighborhood of any given recipe in the recipe network are healthier on average. Thus, given a recipe, a healthier recipe substitution may be easily obtained from the recipe network. This finding may have a profound implication on how food recommendation systems are designed, which can in turn affect the well-being of millions of individuals.

Healthy Food Recommendation Systems. Incorporating healthiness into food recommendation systems has only recently gained attention (Elsweiler, Hauptmann, and Trattner 2022), with simple calorie-based filters (Ge, Ricci, and Massimo 2015) being used to propose “healthy” recipes to users, and long-term meal plans being offered to compensate intake for healthiness (Elsweiler and Harvey 2015). On the other hand, ingredient networks (e.g., frequently co-occurring, or functionally equivalent ingredients derived from user-generated) have been used for recipe modifications and recommendations (Teng, Lin, and Adamic 2012), with ingredients consumed in similar contexts having been used to make healthier recipe substitutions (Achananuparp and Weber 2016). More recently, (Tang, Zheng, and Lai 2019) proposed to predict users’ ratings while substituting ingredients and calculating a modified recipe’s healthiness to improve recommendations.

Similar to this work, (Elsweiler, Trattner, and Harvey 2017) explored the possibility of recommending suitable replacements for unhealthy online recipes. To find similar recipes with healthier nutritional properties, that work proposed to compute the cosine similarity between all pairs of recipes, retain those pairs above a threshold, and replace

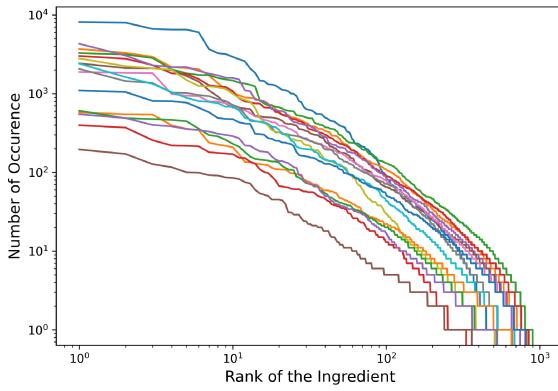
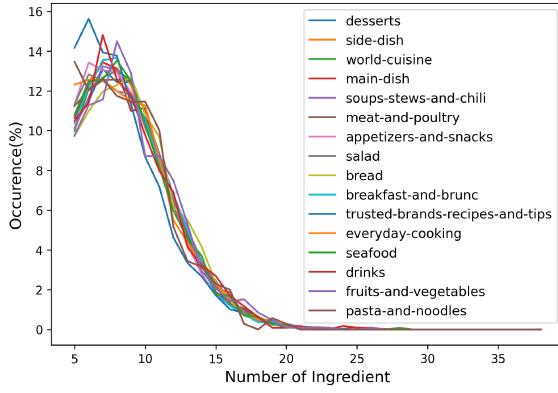
recipes with similar, healthy and comparably, or better rated recipes. Unlike our analysis, which considers all possible recipe pairs in the full collection of RecipeKG, that study examined only small subsets of recipe pools of various sizes drawn randomly from the collection.

Dataset

This study is based on a publicly available dataset of 77,835 recipes published on the main site of Allrecipes.com between the years 1997 and 2021 (Chelmis and Gergin 2022). The dataset, collectively referred to as *RecipeKG*, includes for every recipe: (i) its name, (ii) the average rating and number of reviews it has received, (iii) its category, (iv) the year of its publication, (v) the recommended number of servings, (vi) nutritional information (vii) the complete list of ingredients and their corresponding quantities, as well as (viii) two health score derived by internationally recognized standards for measuring the healthiness of meals, namely the “Dietary Guidelines for Americans” by the United States Department of Agriculture (USDA) (Committee 2020), and the “Guide to creating a front of pack (FoP) nutrition label” by the United Kingdom Food Standards Agency (FSA) (FSA 2014). For reference, we summarize below how these two health scores are derived:

- **USDA Score.** The USDA score is computed over the 7 most important macronutrients (i.e., carbohydrates, protein, fat, saturated fat, sugar, sodium, and fiber) that should be considered in a daily meal plan. Specifically, the limit of content in grams by the given percentage of energy for each macronutrient is derived for a given daily calorific intake (e.g., 2,000). Each recipe is then awarded a point for every macronutrient for which its requirement is met (e.g., the level of calories is below the threshold determined for a given age group), for a total score between 0 (totally unhealthy) and 7 (very healthy).
- **FSA Score.** To derive a “traffic light labeling” score (i.e., front-of-pack three color-coded nutrition label (Jones et al. 2019) used to make nutrition information noticeable and easily understood to consumers) that indicates whether food has high, medium, or low amounts of fat, saturated fat, sugars, and salt, an integer value is assigned for each nutrient. Specifically, 2 for green (low), 1 for amber (medium), and 0 for red (high). By summing the scores for each macronutrient, a score between 0 (very unhealthy) and 8 (very healthy) is derived for each recipe.

The categories that each recipe is associated with are organized in a taxonomy of subcategories (e.g., “Burger Recipes” is a subcategory of “Main Dishes”), allowing for the analysis of recipes by category. Fig. 1 shows the distribution of the number of ingredients per recipe, across sixteen main categories each of which has at least fifty associated recipes. The average number of ingredients used in a recipe, regardless of its category, is approximately eight, with recipes with very large number of ingredients, being rare. In contrast, the frequency of use of specific ingredients across recipes varies significantly, even if an approximately invariant distribution is observed across categories (Fig. 2).



Out of a total of 6,906 ingredients, the most frequent ingredient, salt, occurs 30,205 times, while onion and egg follow in the ranking with 23,370 and 21,643 occurrences, respectively. It is also worth noting that butter, white sugar, all-purpose flour, and olive oil are among the top ten most used ingredients. However, some of these ingredients (e.g., butter and white sugar) are not recommended to be regularly consumed (or in high quantities) due to their negative impacts on health (Freeman et al. 2018; Consultation 2003).

Recipe Networks

Here we introduce a network–based approach to study the healthiness of recipes sourced from the Allrecipes.com website. Specifically, we begin by querying the RecipeKG dataset of 13 million triples, using Apache Jena Fuseki SPARQL server. Then a Recipe–Ingredient matrix (77,309x6,906) is created by setting the corresponding cell to 1, for every ingredient in a recipe. This matrix effectively represents a bipartite graph comprising two types of nodes, namely ingredients, and the recipes using them, as shown in Fig. 3. Before proceeding, we removed 10,969 recipes with less than five ingredients (%15 of total number

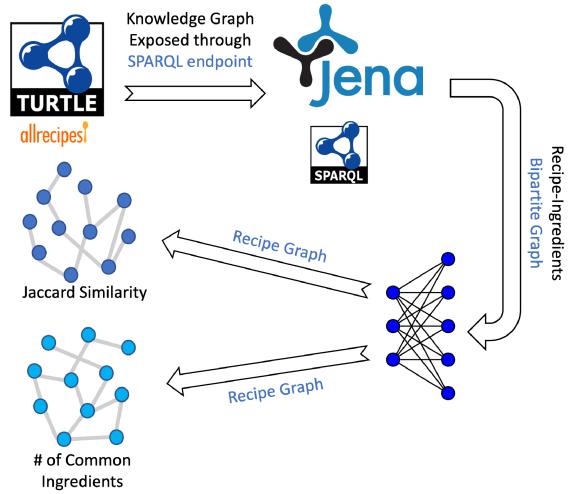
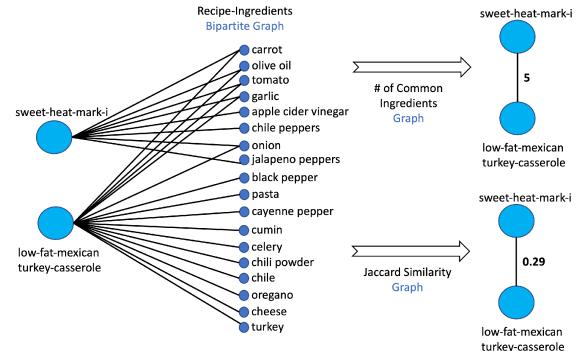


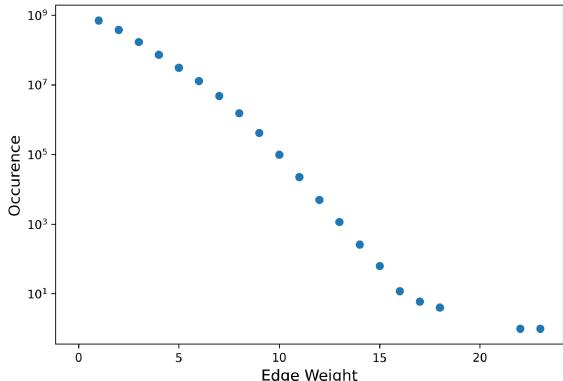
Figure 3: Overview of the pipeline from data collection to network analysis.



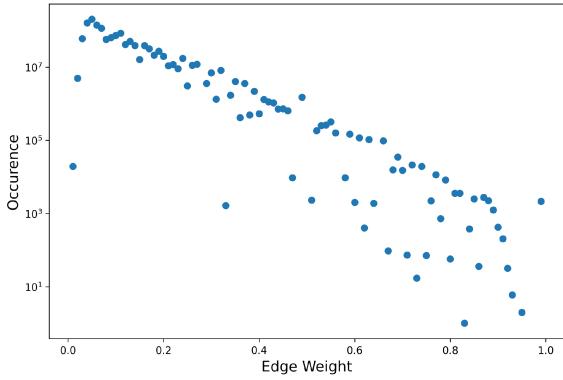
of recipes) and 5,890 ingredients that occurred less than 25 times (%85 of the total number of distinct ingredients). We performed this preprocessing step to ensure that the computed similarity of recipes was estimated based on “enough” ingredients, and that therefore our analysis is meaningful. In the end, our analysis is based on a set of 66,340 recipes using 1,016 unique ingredients. For illustration purposes, Fig. 4 shows a bipartite graph of two recipes and their corresponding ingredients. Since any two recipes may have some common ingredients, we obtain two networks, namely

- \mathcal{N}_i : the network created based on number of common ingredients, and
- \mathcal{N}_J : the network created using Jaccard index,

by projecting the bipartite graph into the recipe space. Specifically, we derived network \mathcal{N}_i by retaining recipe pairs with more than 5 common ingredients. In the resulting network comprises 66,302 nodes and 123,513,257 edges. Similarly, we derived network \mathcal{N}_J by computing the Jaccard similarity between each pair of recipes. The



(a)



(b)

Figure 5: Edge weights and their occurrences in networks (a) \mathcal{N}_i and (b) \mathcal{N}_j .

Jaccard similarity of recipes R_A and R_B is computed as $Jaccard(R_A, R_B) = \frac{|I_A \cap I_B|}{|I_A \cup I_B|}$, where I_A and I_B denote the set of ingredients for each of the two recipes accordingly. We selected Jaccard similarity for two reasons. First, its is easy to compute, even for a large dataset. Second, because of the high dimensionality of our data, other metrics such as cosine similarity, would not produce meaningful similarity scores. For our analysis, we focus on the subgraph of recipes with at least 25% similarity. This network comprises 66, 314 nodes, and 85, 892, 575 edges.

Both networks are too dense for visualization. We therefore illustrate the projection using an example in Fig. 4 (top and bottom right accordingly). Figs. 5a and 5b show, for each network, the edge weights and their occurrences. In both cases, the links between recipes inform us about the ingredient content matching between recipes.

Paradoxes in Recipe Networks

According to the friendship paradox, on average, the neighbors of any node u are better connected than u (i.e., have higher degree) (Feld 1991). By generalizing the paradox to arbitrary node characteristics (Eom and Jo 2014), the paradox has been shown to “hold” for other individual characteristics including happiness (Bollen et al. 2017).

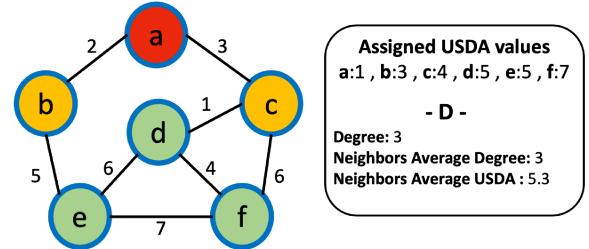


Figure 6: Toy network and calculations for node d . Edge weights represent the number of common ingredients between recipes, and node colors indicate the healthiness of recipes according to USDA Score.

The two recipe networks allow us to explore if similar paradoxes hold for the following recipe characteristics: (i) healthiness (both USDA and FSA scores), (ii) popularity, and (iii) average rating. Specifically, considering a recipe r in either of the two networks \mathcal{N}_i and \mathcal{N}_j , and one of its characteristics (e.g., healthiness) x_r , a paradox for that characteristic can be observed if the following condition is satisfied:

$$x_r < \frac{\sum_{j \in \Gamma_r} x_j}{k_r}, \quad (1)$$

where Γ_r and $k_r = |\Gamma_r|$ are the set of neighbors of recipe r , and its degree, accordingly.

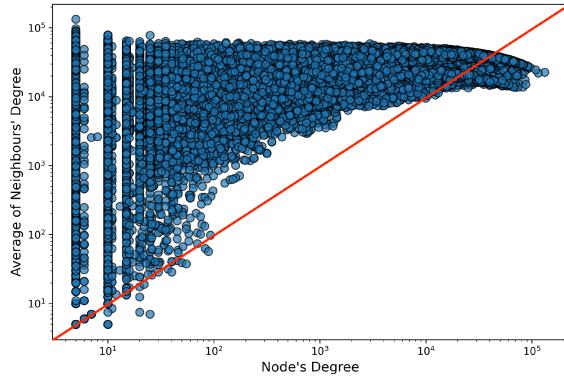
To empirically validate the paradoxes for each characteristic, we begin by computing Eq. 1 for each recipe in the network, and then average over all distinct values of that property as $\frac{1}{|\mathcal{R}_v|} \sum_{z \in \mathcal{R}_v} \frac{\sum_{j \in \Gamma_z} x_j}{k_z}$, where \mathcal{R}_v denotes the set of recipes with property value v for characteristic x . Finally, we compute the ratio of the average of the neighbors’ value of a characteristic to a recipe’ own value.

Fig. 6 illustrates the calculation of the paradox for a given recipe (i.e., d) in a toy recipe network of six recipes. Specifically, node d has degree 3, which in this case is equal to the average degree of its neighbors. Its USDA score is 5, which is less than the average USDA score of its neighbors (i.e., 5.3). Thus, the healthiness paradox holds for recipe d .

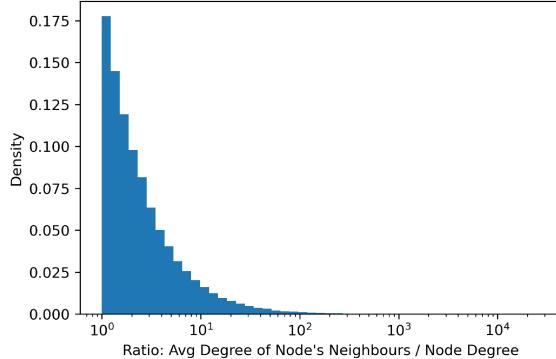
Results

Figs. 7a and 8a show the average degree of a recipe’s neighbors (y -axis) as a function of a recipe’s own degree for both networks. The red line shows, in both cases, equality of the left and right terms in Eq. 1. Figs. 7b and 8b show the probability distribution of the ratio of the average neighbors’ connectivity to a recipe’s own degree, for each of the two networks accordingly. Evidently, the ratio is greater than 1 for 90.2% of the recipes in \mathcal{N}_i and 84.8% in \mathcal{N}_j , respectively. This indicates that given any recipe, one is expected to find similar recipes (i.e., neighboring recipes) that are, on average, more connected.

Figs. 9 and 10 plot the average USDA score of recipes in the neighborhood of a given recipe, versus the USDA score of that recipe. Specifically, Fig. 9a gives the average of the neighbors of recipes with that specific USDA



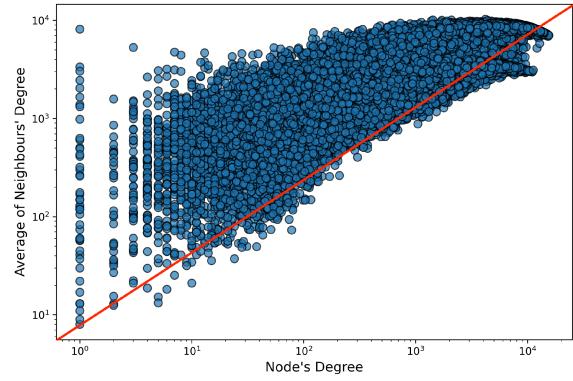
(a)



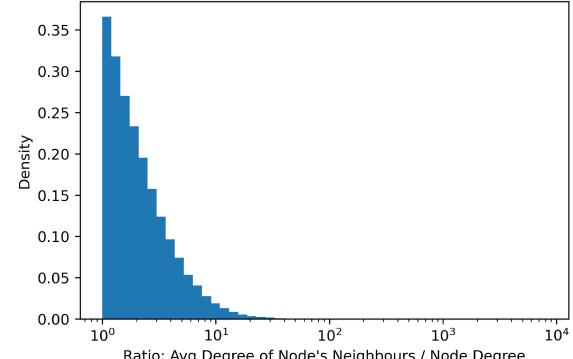
(b)

Figure 7: (a) Degree of a recipe and average of its neighbours' degree in \mathcal{N}_i . (b) Ratio frequency table in log binning for Average Degree of Nodes' Neighbours over Node Degree in \mathcal{N}_i .

score. In the case of network \mathcal{N}_i , recipes with USDA score of 1 (unhealthy) have on average, neighbors whose average USDA score is 2.5 (healthier), whereas recipes with a USDA score of 6 (healthy) have on average, neighbors whose average USDA score is 4 (less healthy). Note that the case of USDA score 0 is an outlier, since only one recipe in the network is included with that USDA score. A similar trend can be observed for FSA scores. For instance, recipes with an FSA score of 1 (unhealthy) have on average, neighbors whose average FSA score is 2.5 (healthier), whereas recipes with an FSA score of 8 (healthy) have on average, neighbors whose average FSA score is 4.8 (less healthy). Similar results can be observed for network \mathcal{N}_J . The probability distribution of the ratio of recipes' neighbor's average USDA value over their own USDA scores for \mathcal{N}_J is shown in Fig. 9b. The distribution peaks at 1 for all three groups (i.e., unhealthy and healthy recipes, and overall), meaning that similar recipes in terms of healthiness tend to cluster together. More importantly, much healthier recipes can be found in the neighborhood of unhealthy recipes (USDA score ≤ 2), whereas equally healthy (or slightly less healthy) recipes can be found in the neighborhood of healthy recipes (USDA score ≥ 5).



(a)



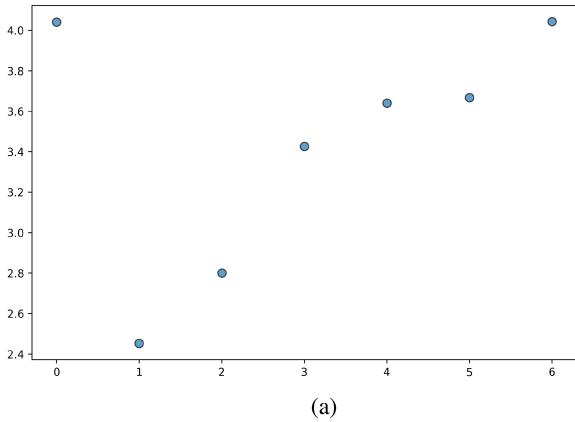
(b)

Figure 8: (a) Degree of a recipe and average of its neighbours' degree in \mathcal{N}_J . (b) Ratio frequency table in log binning for Average Degree of Nodes' Neighbours over Node Degree in \mathcal{N}_J .

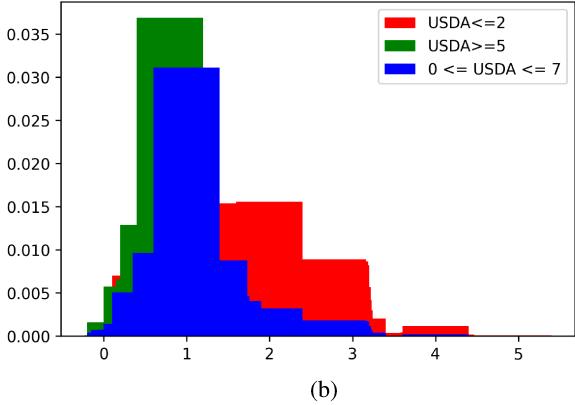
On the other hand, the probability distribution of the ratio of recipes' neighbor's average USDA value over their own USDA scores for \mathcal{N}_J exhibits two peaks at 1 and 3 as shown in Fig 10b for both unhealthy recipes (i.e., USDA score ≤ 2) and overall. In contrast, healthy recipes have just one peak on ratio 1. In other words, similar recipes to healthy recipes are healthy, whereas similar recipes to unhealthy recipes are either equally unhealthy, or have a high chance of being 3 times healthier.

Finally, Fig. 11 shows the recipes' neighbor's average rating (y -axis) as compared to a recipe's own rating (x -axis) for both networks. Unhealthy (USDA score ≤ 2) and healthy (USDA score ≥ 5) recipes are shown in red and green, accordingly. The rating paradox is more profound for recipes with low rating, regardless of their health score. Specifically, the rating of recipes neighboring recipes with a rating of 1 is 4 times higher, on average. On the other hand, it becomes increasingly difficult for recipes neighboring highly rated recipes to have higher rating, on average.

This result has an important implication for recommendation systems. Given an unhealthy recipe, one can find a similar recipe, which is at least as healthy, but additionally has a rating that is up to 4 times larger than that of the input



(a)



(b)

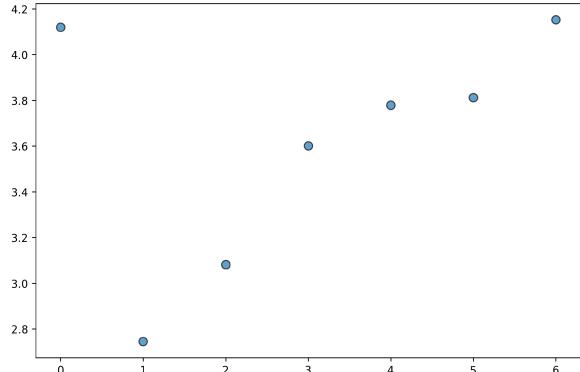
Figure 9: (a) USDA score and average neighbors’ USDA score for the recipes with that score in \mathcal{N}_i . (b) Probability distributions for the ratio of recipe neighbours’ average USDA score over recipe’s own USDA for \mathcal{N}_i .

recipe. Similarly, given a healthy recipe, one can find a similar recipe, which is at least as healthy, and at the same time has similar, if not better, rating.

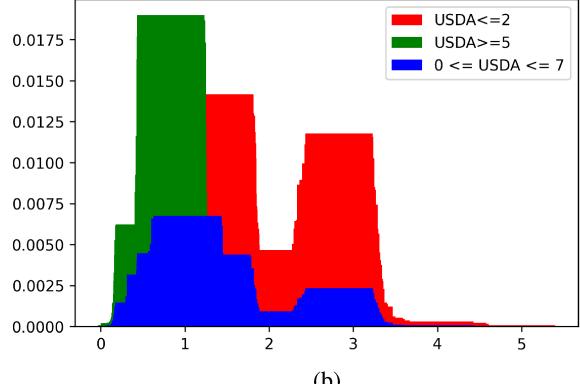
Potential Impact and Ethical Considerations

This work presents the first large-scale study of online recipes’ healthiness using a network-based approach. To generate the two recipe networks we leverage a dataset of recipes sourced from Allrecipes.com, the potential limitations of which are discussed in detail in (Chelmis and Gergin 2022). To the best of our knowledge, the authors of (Chelmis and Gergin 2022) complied with the terms of use of Allrecipes.com while crawling the dataset. The dataset itself is licensed under the Apache License 2.0. It is worth mentioning that the main findings of this study are strongly restricted by this recipe collection. Therefore, generalizations should be avoided, unless the findings are reproduced with recipes sourced from other websites (e.g., [epicurious.com](#) or [simplyrecipes.com](#)) or datasets, such as the Recipe1M dataset (Marin et al. 2019).

Finally, this study neither implements nor empirically evaluates recipe substitution as a method for healthier food recommendations. However, by demonstrating the existence of a healthiness and rating paradox, our study paves the way



(a)



(b)

Figure 10: (a) USDA score and average neighbors’ USDA score for the recipes with that score in \mathcal{N}_J . (b) Probability distributions for the ratio of recipe neighbours’ average USDA score over recipe’s own USDA for \mathcal{N}_J .

for future work in this area by suggesting that given a recipe, similar, yet healthier recipes can be found using the network of recipes. A thorough evaluation of algorithmic solutions to recipe substitutions is necessary. From a user perspective, there is much to be learnt regarding user perception of recipe substitution compared with ingredient substitution, and how either approach relates to their health, dietary restrictions, or other constraints.

Conclusion

Recipe websites and apps are frequently used by people to make food-related choices. However, prior research has shown that highly rated and well-known recipes on these platforms tend to have high levels of fat, sugar, cholesterol, and calories. To address this issue researchers have been focusing on developing recommendation systems for healthier food choices by emphasizing on ingredient substitution. This work took a complementary approach by creating networks of “similar” recipes, that can be used for recipe substitution. Analysis of the recipe networks resulted in the discovery of interesting properties that have important implications to the development of food recommendation systems. Informed by these findings, we plan to investigate the ability of a recommendation system to generate healthier recipe

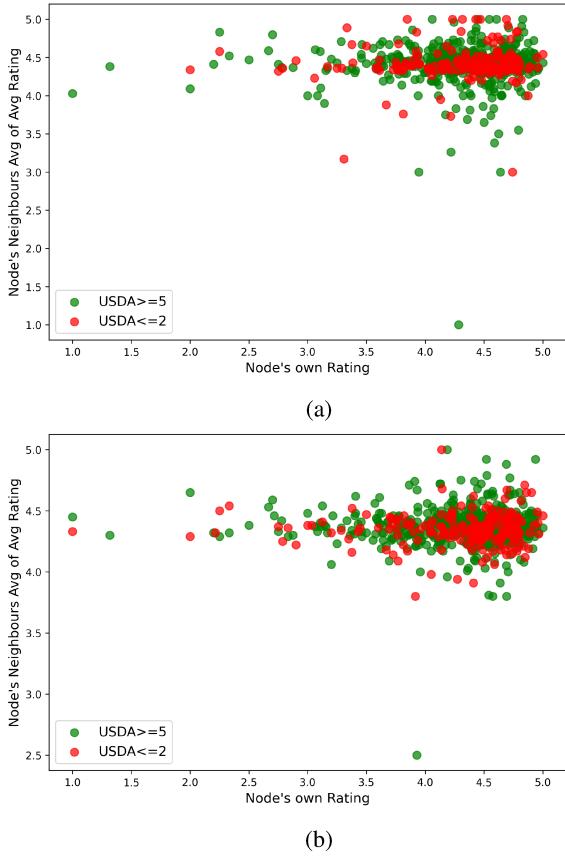


Figure 11: Rating average and average neighbors' rating score for the recipes with that rating based on USDA Range (a) \mathcal{N}_i , and (b) \mathcal{N}_j

substitutions in future work. We additionally plan to analyze recipes sourced from other websites or datasets, so as to confirm (or disprove) the generality of our findings.

References

Achananuparp, P.; and Weber, I. 2016. Extracting Food Substitutes From Food Diary via Distributional Similarity. *ArXiv*, abs/1607.08807.

AELIEVE. 2021. Best Recipe Websites. <https://aelieve.com/rankings/websites/category/food-beverage/best-recipe-websites/>. Accessed: 2022-02-14.

Ahn, Y.-Y.; Ahnert, S.; Bagrow, J.; and Barabasi, A.-L. 2011. Flavor network and the principles of food pairing. *Scientific Reports*, 1.

Bollen, J.; Gonçalves, B.; van de Leemput, I.; and Ruan, G. 2017. The happiness paradox: your friends are happier than you. *EPJ Data Science*, 6(1): 4.

Boulos, M. N. K.; Yassine, A.; Shirmohammadi, S.; Namahtoot, C. S.; and Brückner, M. 2015. Towards an "Internet of Food": Food Ontologies for the Internet of Things. *Future Internet*, 7: 372–392.

Chelmis, C.; and Gergin, B. 2021. A Knowledge Graph for Semantic-Driven Healthiness Evaluation of Recipes. *Semantic Web Journal*.

Chelmis, C.; and Gergin, B. 2022. A Knowledge Graph for Semantic-Driven Healthiness Evaluation of Online Recipes.

Committee, U. S. D. G. A. 2020. *Dietary Guidelines for Americans, 2020-2025*, volume 9th Edition. U.S. Department of Agriculture and U.S. Department of Health and Human Services.

Consultation, J. W. E. 2003. Diet, nutrition and the prevention of chronic diseases. *World Health Organ Tech Rep Ser*, 916(i-viii): 1–149.

Elsweiler, D.; and Harvey, M. 2015. Towards Automatic Meal Plan Recommendations for Balanced Nutrition. In *Proceedings of the 9th ACM Conference on Recommender Systems*, RecSys '15, 313–316. New York, NY, USA: Association for Computing Machinery. ISBN 9781450336925.

Elsweiler, D.; Hauptmann, H.; and Trattner, C. 2022. *Food Recommender Food recommenderSystems*, 871–925. New York, NY: Springer US.

Elsweiler, D.; Trattner, C.; and Harvey, M. 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. New York, NY, USA: Association for Computing Machinery. ISBN 9781450350228.

Eom, Y.-H.; and Jo, H.-H. 2014. Generalized friendship paradox in complex networks: The case of scientific collaboration. *Scientific reports*, 4(1): 1–6.

Feld, S. L. 1991. Why your friends have more friends than you do. *American journal of sociology*, 96(6): 1464–1477.

Freeman, C. R.; Zehra, A.; Ramirez, V.; Wiers, C. E.; Volkow, N. D.; and Wang, G.-J. 2018. Impact of sugar on the body, brain, and behavior. *FBL*, 23(12): 2255–2266.

FSA, U. 2014. Guide to creating a front of pack (FoP) nutrition label for pre-packed products sold through retail outlets. *Food Standards Agency*.

Ge, M.; Ricci, F.; and Massimo, D. 2015. Health-Aware Food Recommender System. In *Proceedings of the 9th ACM Conference on Recommender Systems*, RecSys '15, 333–334. New York, NY, USA: Association for Computing Machinery. ISBN 9781450336925.

Hodas, N.; Kooti, F.; and Lerman, K. 2013. Friendship paradox redux: Your friends are more interesting than you. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 7, 225–233.

Howard, S. J.; Adams, J.; and White, M. 2012. Nutritional content of supermarket ready meals and recipes by television chefs in the United Kingdom: cross sectional study. *The BMJ*, 345.

Jones, A.; Neal, B.; Reeve, B.; Mhurchu, C. N.; and Thow, A. M. 2019. Front-of-pack nutrition labelling to promote healthier diets: current practice and opportunities to strengthen regulation worldwide. *BMJ Global Health*, 4(6): e001882.

Li, C.; Mirosa, M.; and Bremer, P. 2020. Review of Online Food Delivery Platforms and their Impacts on Sustainability. *Sustainability*, 12(14).

Marin, J.; Biswas, A.; Ofli, F.; Hynes, N.; Salvador, A.; Aytar, Y.; Weber, I.; and Torralba, A. 2019. Recipe1m+: A dataset for learning cross-modal embeddings for cooking recipes and food images. *IEEE transactions on pattern analysis and machine intelligence*, 43(1): 187–203.

Tang, Y. S.; Zheng, A. H.; and Lai, N. 2019. Healthy Recipe Recommendation using Nutrition and Ratings Models. <https://web.stanford.edu/class/cs224w/project/26410425.pdf>. Accessed: 2022-12-10.

Teng, C.-Y.; Lin, Y.-R.; and Adamic, L. A. 2012. Recipe Recommendation Using Ingredient Networks. In *Proceedings of the 4th Annual ACM Web Science Conference, WebSci '12*, 298–307. New York, NY, USA: Association for Computing Machinery. ISBN 9781450312288.

Trattner, C.; and Elsweiler, D. 2017. Investigating the healthiness of internet-sourced recipes: implications for meal planning and recommender systems. In *Proceedings of the 26th international conference on world wide web*, 489–498.

Watanabe-Ito, M.; Kishi, E.; and Shimizu, Y. 2020. Promoting Healthy Eating Habits for College Students Through Creating Dietary Diaries via a Smartphone App and Social Media Interaction: Online Survey Study. *JMIR Mhealth Uhealth*, 8(3): e17613.

Çelik, D. 2015. FoodWiki: Ontology-Driven Mobile Safe Food Consumption System. *The Scientific World Journal*, 2015.