

# Advancing Visual Food Attractiveness Predictions for Healthy Food Recommender Systems

Ayoub El Majjodi  
MediaFutures, University of Bergen  
Bergen, Vestland, Norway  
ayoub.majjodi@uib.no

Sohail Ahmed Khan  
MediaFutures, University of Bergen  
Bergen, Vestland, Norway  
sohail.khan@uib.no

Alain D. Starke  
Amsterdam School of Communication  
Research (ASCoR), University of  
Amsterdam  
Amsterdam, Netherlands  
MediaFutures, University of Bergen  
Bergen, Vestland, Norway  
a.d.starke@uva.nl

Mehdi Elahi  
MediaFutures, University of Bergen  
Bergen, Vestland, Norway  
mehdi.elahi@uib.no

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

## ABSTRACT

The visual representation of food on digital platforms affects the foods chosen by users, including in the context of recommender systems. Previous studies show that small changes in visual features can influence human decision-making, regardless of whether the food is healthy. This paper reports on a study aimed at better understanding how users perceive the attractiveness of food recipe images in the digital world. In an online mixed-methods survey ( $N = 192$ ), users provided visual attractiveness ratings of food images on a 7-point scale, along with textual assessments. We found robust correlations between fundamental visual features (e.g., contrast, colorfulness) and perceived image attractiveness. The analysis also revealed that, among other user factors, cooking skills positively affected perceived image attractiveness. Regarding food image dimensions, appearance and perceived healthiness were significantly correlated with user ratings of food image attractiveness.

## KEYWORDS

Food recommender systems; User modeling; Image attractiveness; Health; Personalization; Digital nudges

### ACM Reference Format:

Ayoub El Majjodi, Sohail Ahmed Khan, Alain D. Starke, Mehdi Elahi, and Christoph Trattner. 2024. Advancing Visual Food Attractiveness Predictions for Healthy Food Recommender Systems. In *In Proceedings of the 6th International Workshop on Health Recommender Systems co-located with 18th ACM Conference on Recommender Systems (HealthRecSys'24), Bari, Italy, October 18, 2024, 6 pages.* ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

Visual cues and attractiveness play a crucial role in everyday food choices [23]. Even when only presented with a food image, humans

tend to instantly assess a food's energy density, expected taste and other characteristics [18]. As such, images are one of the key determinants of food preferences [18, 20], tapping into emotional and hedonic processes of an individual [1].

The importance of visual attractiveness also applies to digital choice context, including food recommender systems [5]. Our previous research has shown the capability of recommender systems to influence food behaviors via visual features, including the promotion of either high-fat or low-fat food choices [6], as well as encouraging the search for healthier options [20]. Additionally, our earlier work has established that visual attractiveness significantly contributes to predicting the online popularity of food items [21], and these visual features can also be leveraged to infer cultural backgrounds [25].

What is currently missing is an in-depth examination of image feature modeling. Although previous studies have extracted image features and examined the relation between those features, visual attractiveness and user preferences [6, 20], these models have not been optimized. Moreover, to date, image features have not been related to user characteristics (e.g., demographics, food knowledge), which are also important determinants of food preferences [16].

We present the results of a mixed-method study that explores the determinants of visual attractiveness in digital recipe images more comprehensively. Our approach builds upon previous work by modeling perceived visual attractiveness based on low-level image features [10, 15, 20]. Additionally, we seek to optimize this model by integrating user characteristics that have been employed in knowledge-based food recommender systems to promote healthier recipe choices [4, 11, 19].

Finally, we inquire more qualitatively on user justifications for provided visual attractiveness ratings, asking to motivate their quantitative judgment. We formulate the following research questions:

- **RQ1:** To what extent do the latest deep learning methods predict visual attractiveness compared to state-of-the-art low-level features?

This paper is authored by an employee(s) of the United States Government and is in the public domain. Non-exclusive copying or redistribution is allowed, provided that the article citation is given and the authors and agency are clearly identified as its source.

HealthRecSys '24, October 18th, 2024, Bari, Italy  
© 2024 Copyright held by the owner/author(s).  
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

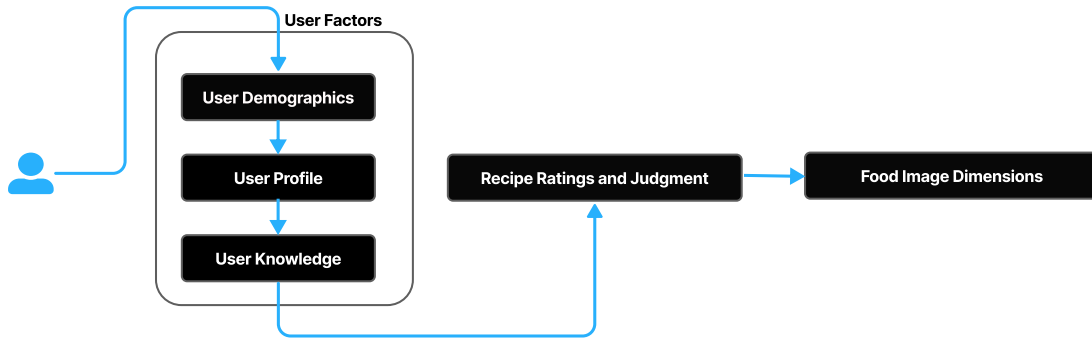


Figure 1: Steps of the user flow designed for the online survey.

- **RQ2:** To what extent do user characteristics, including demographics, food knowledge, and eating goals, predict food image attractiveness?
- **RQ3:** What dimensions determine the attractiveness of food image?

### 1.1 Contributions

Compared to our extensive previous work in the field of visual attractiveness and food recommender systems [6, 19, 21, 25], this study offers novel insights into several key aspects:

- Previous work mostly relied on low-level image attractiveness features, while this study shows how new deep-learning models compare to these old features.
- This work, compared to any before, also shows as to what extent demographic features play a role in predicting visual food attractiveness. To our knowledge, no other work has shown this before.
- Finally, this study tries to go beyond traditional quantitative black box approaches and reveals why images are rated less or more attractive.

## 2 STUDY DESIGN

To perform our study, we employed a dataset sourced from the well-known recipes website AllRecipes.com, with the addition of new recipe photos [4, 20]. The dataset comprised various recipe features, including image URL, ingredients, amount of fats and sugar, and instructions and ingredients. To generate a diverse set of images, we randomly selected 200 recipes with relatively from the dataset of 58,000. As most images in this dataset were relatively unattractive [20], we used the recipe’s title in search engines and image websites (e.g., Unsplash) to look for more attractive images for 100 of these recipes. To validate this process, three computational food researchers, including a co-author, voted on which of the two photos was the most attractive to ensure a diverse set of recipe images in terms of expected attractiveness.

The study involved a survey design, as depicted in Figure 1. Participants first provided demographic information, as well as responded to items that measured their subjective food knowledge (4 items) and cooking skills (6 items), using 5-point Likert scales based on earlier work [7, 8, 13]. We also used questions from earlier work on a knowledge-based food recommender [4], to inquire

on other user characteristics, including recipe website usage and home cooking frequency, cooking experience and dietary goals. Afterwards, users were invited to rate the visual attractiveness of 12 semi-randomly selected recipe images, on 7-point attractiveness scales. In addition, to address [RQ3], they were asked to write at least one sentence about why they had given this rating. Finally, to support our examination of [RQ3], we used 5-point Likert scales on food image dimensions [26], to ask to what extent a recipe’s appearance, expected taste, healthiness, and familiarity affected their attractiveness ratings.

We employed the Prolific crowdsourcing platform to recruit 192 users (65% male;  $Age = 33.54$ ) to participate in our study. The study took approximately 11 min to complete and participants were reimbursed with GBP 1.65<sup>1</sup>.

## 3 RESULTS

To address our research questions, we primarily employed linear regression models. This helped to understand the principal impacts of image attributes and user characteristics on image attractiveness, the latter derived from user ratings. For our thematic analysis, the images were split into attractive and unattractive based on the mid-point of the rating scale (4) ( $M = 4.33$ ,  $SD = 1.80$ ). Details of used materials and conducted analyses can be accessed through the following URL [3].

### 3.1 RQ1: Predicting Visual Attractiveness

We first modeled perceived visual attractiveness based on the underlying image features. We extracted diverse low-level visual features using the OpenIMAJ Java Framework (cf. [21]). Subsequently, we conducted a linear regression analysis to predict attractiveness based on these extracted visual features. The results are outlined in Table (1.A), revealing that several image features significantly affected the attractiveness of a recipe image:  $F(8, 2100) = 32.66$ ,  $p < 0.001$ . Specifically, Colourfulness, Brightness, Naturalness, and Entropy demonstrated a positive association with image attractiveness. In contrast, Saturation, Sharpness, and RgbContrast negatively

<sup>1</sup>Our study complied with the ethical guidelines of the Research Council of Norway and the guidelines of University of Bergen for scientific research. It was judged to pass without further extensive review.

**Table 1: Linear regression models predicting visual attractiveness ratings for recipe images: (A) with low-level image visual features, (B) with deep learning-based visual features. \*\*\* $p < 0.001$ , \*\* $p < 0.01$ , \* $p < 0.05$ .**

(A)		(B)		
Low-level Image Features		Image features Extractor		
	$\beta$ (S.E)	VGG16	ResNet	CLIP
Colourfulness	6.725 (1.521)***			
Brightness	2.136 (0.155)***			
Naturalness	1.925 (0.530)***			
Entropy	1.026 (0.154)**			
Saturation	-3.976 (1.020)***			
Sharpness	-1.182 (1.187)*			
RGBContrast	-1.782 (3.808)			
Contrast	7.401 (11.101)			
Constant	-6.884 (1.243)***			
$R^2$	0.110***	0.351***	0.349***	0.357***
RMSE	1.753	1.500	1.491	1.501

affected image attractiveness. In line with [20], these results suggested that users perceived colorful, bright, and naturalistic food images as more attractive.

Going beyond low-level visual image features, we used deep learning architecture models. Our toolkit included established models, such as VGG16 [17] and ResNet [9], along with the well-known transformer [22] architecture for visual feature extraction, CLIP<sup>2</sup> [14]. Table (1.B) outlines the performance of these different models, outperforming our regression model in terms of  $R^2$  and RMSE. This aligns with previous research where deep learning embeddings also outperformed low-level visual features within the context of food [2, 21].

### 3.2 RQ2: User characteristics and Image Attractiveness

We further examined whether user factors affected the perceived visual attractiveness of images. Accordingly, we divided user characteristics into different categories: User demographics, User profile, which represented the backbone of a food knowledge-based recommender system, and User knowledge, which measures the user's food knowledge and cooking skills. A confirmatory factor analysis, reported in Table 2, showed that both subjective food knowledge and cooking skills adhered to internal consistency guidelines ( $\alpha > 0.70$ ) while they also met the guidelines for convergent validity ( $AVE > 0.5$ ).

Table (3.A) presents the outcomes of the linear regression model aimed at forecasting the attractiveness of image recipes:  $F(9, 2090) = 3.60$ . Among the various user factors examined, only two significantly affected recipe attractiveness: cooking skills ( $\beta = 0.34$ , p-value= 0.00021) and recipe website usage ( $\beta = 0.18$ , p-value= 0.020). However, none of the other user aspects affected user ratings for a given image recipe. Additionally, we also analyzed a combined

model of image features and user factors, but this led to results similar to the separate models reported in Tables (1 and 3.A). This suggested that low-level visual features had a more significant impact on food image attractiveness than user features, largely in line with preliminary findings in previous research [20, 26].

### 3.3 RQ3: Justifications for Visual Attractiveness

To assess the influence of different food image dimensions on user ratings for food images, we modeled visual attractiveness based on the reported importance of food image dimensions. Table 3 outlines the results of the regression model:  $F(4, 21) = 2.41$ .

Two factors significantly impacted attractiveness. First, appearance significantly impacted user ratings ( $\beta = 0.12$ ,  $p = 0.03$ ). Second, the expected healthiness from the images also demonstrated a significant impact ( $\beta = 0.07$ ,  $p = 0.03$ ). However, perceived taste and familiarity did not impact user ratings.

To gain insights into the reasons behind the visual attractiveness ratings, we collected qualitative justifications from users. Natural Language Processing (NLP) techniques, such as punctuation removal, repeated character elimination, and stop word filtering, were applied to analyze 2,019 user justifications for both attractive and unattractive images. From these responses, we generated two word clouds to highlight the most prominent terms. Figure 2 presents the most frequent responses associated with both attractive and unattractive images. These findings are discussed in relation to the themes of 'appearance' and 'health' (cf. Table (3.B)).

**3.3.1 Appearance-based justifications.** Figure 3 shows a few examples of user textual justifications. Several participants, including user ( $U_a$ ), expressed the term 'crispy' in their assessments of attractive images, mainly referring to appearance. The word 'simple' is frequently used by users, such as user ( $U_b$ ), to convey the simplicity of recipe content. In contrast, 'mess' was more commonly associated with judgments of unattractive food images, indicating

<sup>2</sup>Contrastive Language-Image Pre-training (CLIP).

**Table 2: Results of the principal component factor analysis across different subjective food knowledge and cooking skills. Items were measured on 5-point Likert scales. Cronbach's Alpha is denoted by  $\alpha$ , AVE is the average variance explained. Items in grey and without loading were omitted.**

Aspect	Item	Loading
Subjective Food Knowledge $\alpha = 0.866$ AVE = 0.858	Compared with an average person, I know a lot about healthy eating.	0.777
	I think I know enough about healthy eating to feel pretty confident when choosing a recipe.	0.885
	I know a lot about how to evaluate the healthiness of a recipe.	0.773
	I do not feel very knowledgeable about healthy eating.	0.932
Cooking skills $\alpha = 0.783$ AVE = 0.591	I can confidently cook recipes with basic ingredients.	0.751
	I can confidently follow all the steps of simple recipes.	
	I can confidently taste new foods.	0.737
	I can confidently cook new foods and try new recipes.	0.869
	I enjoy cooking food.	0.655
	I am satisfied with my cooking skills.	0.816

**Table 3: Linear regression models predicting user rating for recipe image attractiveness: (A): with user factors, (B): with food image dimensions. \*\*\* $p < 0.001$ , \*\* $p < 0.01$ , \* $p < 0.05$ .**

(A)		(B)	
User Factors	$\beta$ (S.E)	Food Image Dimension	$\beta$ (S.E)
<b>User Demographic</b>			
Age	-0.047 (0.116)	Appearance	0.129 (0.061)*
Education	-0.424 (0.320)	Healthiness	0.077 (0.035)*
Gender	-0.077 (0.088)	Taste	-0.005 (0.050)
<b>User Profile</b>		Familiarity	0.0231 (0.038)
Recipe Website Usage	0.201 (0.086)*	Constant	3.487 (0.365)***
Home Cooking	-0.009 (0.078)	R <sup>2</sup>	0.011***
Cooking Experience	-0.052 (0.079)	RMSE	1.855
Eating Goals	0.019 (0.063)		
<b>User Knowledge</b>			
Subjective Food Knowledge	-0.213 (0.138)		
Cooking Skills	0.315 (0.086)***		
Constant	4.001 (0.570)***		
R <sup>2</sup>	0.015***		
RMSE	1.845		

their unappealing appearance. Moreover, the repeated use of the term 'fat' suggested that fatty foods were generally perceived as unattractive, as in judgments by users ( $U_c - d$ ).

**3.3.2 Healthiness-based justifications.** Judgments related to health frequently appeared in connection with the food's appearance, such as by user ( $U_e$ ) in Figure 4. The term 'restaurant' was employed in various user judgments, often associated with presentation and healthiness, as described by the user ( $U_f$ ). Conversely, the concept

of unhealthiness was linked to fatty foods and messy representation, as evident in the judgments of users ( $U_{g-h}$ ) in Figure 4.

## 4 CONCLUSION & FUTURE WORK

This work has explored different aspects of the relationship between the user and food images. Through an online user study, we have found that various visual features can predict the attractiveness of a given image (i.e. colorfulness, brightness, naturalness). This prediction accuracy could be slightly improved using image features extracted using deep learning techniques (RQ1). In line with earlier

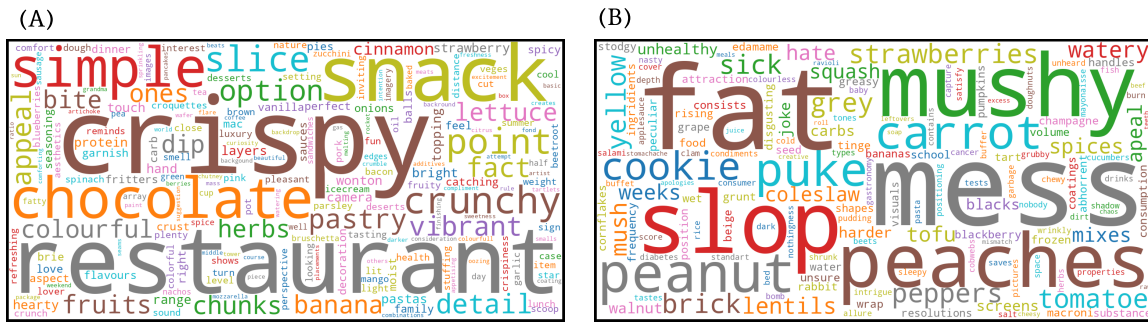


Figure 2: Word cloud for terms in the user judgment: (A) : judgments for attractive images, (B): judgments for unattractive images.



Figure 3: Example of images used in the study, associated with users’ textual judgment related to the appearance. (U<sub>a–b</sub>) were textual justifications for attractive images, (U<sub>c–d</sub>) were textual justifications for unattractive images.

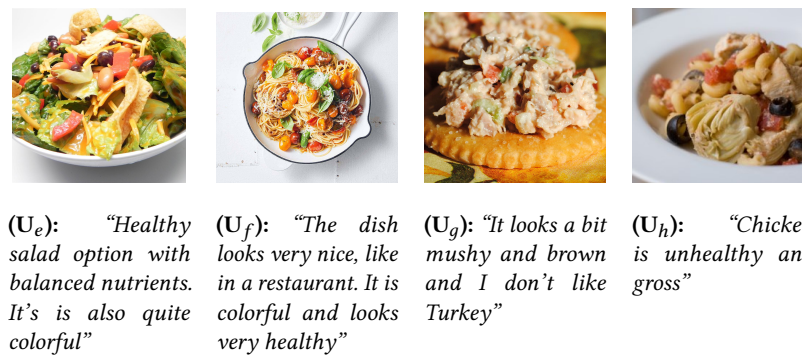


Figure 4: Example images used in the study, associated with users’ textual justifications related to the healthiness. (U<sub>e–f</sub>) were for attractive images, (U<sub>g–h</sub>) for unattractive images.

work [15, 20, 26], this suggests that the visual attractiveness of food images can be enhanced by increasing their colorfulness, brightness, and naturalness, while decreasing other features, such as saturating and sharpness. Obviously, there may be tradeoffs between these features when altering them.

Regarding user characteristics, none of the user demographics are related to food image attractiveness. In contrast, using online recipe websites and cooking skills are positively associated with the attractiveness of food images (RQ2). More novel is our contribution on the user justifications, for which we have found image appearance and perceived healthiness to be important dimensions

of visual attractiveness ratings (RQ3). It seems that attractiveness are related to the expect taste or hedonic food goals (e.g., 'crispy'), while unattractive images focused on poor presentation and disliked ingredients.

Our study offers valuable insights into techniques for image attractiveness selection for various goals and domains. In particular, these techniques can be leveraged to persuade or nudge users towards specific eating goals, such as health [20, 24]. We believe that leveraging the visual appeal of attractive images can address this issue. Our future studies will focus on designing image selection pipelines for the application of food recommender systems tailored to guide people toward healthy food choices without compromising the benefits of personalization. We aim to analyze and categorize the collected textual judgment through thematic analysis to build word dictionaries related to image dimensions. These dictionaries can then be used to train learning models, enabling the evaluation of food image attractiveness based on user textual inputs.

## ACKNOWLEDGMENTS

This work was supported by industry partners and the Research Council of Norway with funding to MediaFutures: Research Centre for Responsible Media Technology and Innovation, through the centers for Research-based Innovation scheme, project number 309339.

The authors acknowledge the use of ChatGPT [12] for checking and correcting the grammar of this article. No new content was generated this way; only existing text was checked and, if needed, corrected.

## REFERENCES

- [1] Romain Cadario and Pierre Chandon. 2020. Which healthy eating nudges work best? A meta-analysis of field experiments. *Marketing Science* 39, 3 (2020), 465–486.
- [2] Jing-jing Chen, Chong-Wah Ngo, and Tat-Seng Chua. 2017. Cross-modal recipe retrieval with rich food attributes. In *Proceedings of the 25th ACM international conference on Multimedia*. 1771–1779.
- [3] Ayoub El Majjodi, Sohail Ahmed Khan, Alain D. Starke, Mehdi Elahi, and Christoph Trattner. 2024. Examining the Visual Attractiveness of Digital Recipe Images: Material. <https://github.com/ayoubGL/Health-RecSys-2024>
- [4] Ayoub El Majjodi, Alain D Starke, Mehdi Elahi, Christoph Trattner, et al. 2023. The Interplay between Food Knowledge, Nudges, and Preference Elicitation Methods Determines the Evaluation of a Recipe Recommender System. In *Proceedings of the 10th Joint Workshop on Interfaces and Human Decision Making for Recommender Systems (IntRS 2023)*. 1–18.
- [5] David Elswailer, Hanna Hauptmann, and Christoph Trattner. 2022. *Food Recommender Food recommenderSystems*. Springer US, New York, NY, 871–925. [https://doi.org/10.1007/978-1-0716-2197-4\\_23](https://doi.org/10.1007/978-1-0716-2197-4_23)
- [6] 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*. 575–584.
- [7] Leisa Reinecke Flynn and Ronald E Goldsmith. 1999. A short, reliable measure of subjective knowledge. *Journal of business research* 46, 1 (1999), 57–66.
- [8] Nike Frans. 2017. *Development of cooking skills questionnaire for EFNEP participants in Kansas*. Ph. D. Dissertation. Kansas State University.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [10] Aditya Khosla, Atish Das Sarma, and Raffay Hamid. 2014. What makes an image popular?. In *Proceedings of the 23rd international conference on World wide web*. 867–876.
- [11] Cataldo Musto, Christoph Trattner, Alain Starke, and Giovanni Semeraro. 2020. Towards a knowledge-aware food recommender system exploiting holistic user models. In *Proceedings of the 28th ACM conference on user modeling, adaptation and personalization*. ACM, New York, NY, USA, 333–337.
- [12] OpenAI. 2024. ChatGPT. <https://openai.com/chatgpt/> Accessed: 2024-09-16.
- [13] Zuzanna Pieniak, Joris Aertsens, and Wim Verbeke. 2010. Subjective and objective knowledge as determinants of organic vegetables consumption. *Food quality and preference* 21, 6 (2010), 581–588.
- [14] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PMLR, 8748–8763.
- [15] Jose San Pedro and Stefan Siersdorfer. 2009. Ranking and classifying attractiveness of photos in folksonomies. In *Proceedings of the 18th international conference on World wide web*. 771–780.
- [16] Benjamin Scheibehenne, Linda Miesler, and Peter M Todd. 2007. Fast and frugal food choices: Uncovering individual decision heuristics. *Appetite* 49, 3 (2007), 578–589.
- [17] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [18] Charles Spence, Kosuke Motoki, and Olivia Petit. 2022. Factors influencing the visual deliciousness/eye-appeal of food. *Food Quality and Preference* 102 (2022), 104672.
- [19] Alain D Starke, Cataldo Musto, Amon Rapp, Giovanni Semeraro, and Christoph Trattner. 2023. "Tell Me Why": using natural language justifications in a recipe recommender system to support healthier food choices. *User Modeling and User-Adapted Interaction* (2023), 1–34.
- [20] Alain D Starke, Martijn C Willemsen, and Christoph Trattner. 2021. Nudging healthy choices in food search through visual attractiveness. *Frontiers in Artificial Intelligence* 4 (2021), 20.
- [21] Christoph Trattner, Dominik Moesslang, and David Elswailer. 2018. On the predictability of the popularity of online recipes. *EPJ Data Science* 7, 1 (2018), 1–39.
- [22] Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Neural Information Processing Systems*. <https://api.semanticscholar.org/CorpusID:13756489>
- [23] Iris Vermeir and Gudrun Roose. 2020. Visual design cues impacting food choice: A review and future research agenda. *Foods* 9, 10 (2020), 1495.
- [24] Longqi Yang, Cheng-Kang Hsieh, Hongjian Yang, John P Pollak, Nicola Dell, Serge Belongie, Curtis Cole, and Deborah Estrin. 2017. Yum-me: a personalized nutrient-based meal recommender system. *ACM Transactions on Information Systems (TOIS)* 36, 1 (2017), 1–31.
- [25] Qing Zhang, David Elswailer, and Christoph Trattner. 2020. Visual cultural biases in food classification. *Foods* 9, 6 (2020), 823.
- [26] Qing Zhang, David Elswailer, and Christoph Trattner. 2023. Understanding and predicting cross-cultural food preferences with online recipe images. *Information Processing & Management* 60, 5 (2023), 103443.