

A Robotic Assistant for Personalised Diet Recommendation

Luca Raggioli¹, Francesco Ciccarelli¹, Silvia Rossi¹, Alessandra Rossi¹

Abstract—Food recommender systems have become valuable tools across various domains, including health-oriented applications that provide personalised dietary advice. Recent studies have shown the potential of integrating recommender systems with assistive robots to promote healthy eating habits, especially among older adults. While transformers and Large Language Models showed advanced reasoning capability for effective recommendation systems, they might have limited knowledge and understanding of the users’ personal preference and requirements. This lack of information can negatively affect their effectiveness and user’s satisfaction. We present a novel transformer-assisted, multi-interface recommendation system for generating food recommendations based on user profiles using a custom dataset including dietary and nutritional information. We conducted a user study with 40 participants for evaluating whether a robot is able to persuade users’ in accepting its food recommendation. Our study found that participants responded positively to the interactions with the robot, showing high satisfaction and trust in the recommendations.

I. INTRODUCTION

Food recommender systems have gained increasing importance in various domains, from cooking inspiration to restaurant recommendations. In particular, health recommender systems have emerged as useful tools to provide personalised nutrition advice to users with applications in different health domains, from diet planning for the older people [1] to chronic disease management [2]. Combining immediately available data (i.e., age, gender, weight and height) and clinical data (i.e., allergies, diets, and intolerances) is indeed an emerging trend, which allows generating recommendations taking into account different types of preferences and dietary restrictions [3]. The potential of integrating assistive robots and recommender systems to support healthy eating habits has shown promising results in supporting older adults with nutritional needs [4]. These systems can help users find alternative food options to promote a healthy eating regimen [5]. Trust and embodiment can be fundamental factors for an effective robot’s recommendation [6]. Preference elicitation features can significantly increase user trust in robot recommender systems, while embodiment type may affect persuasiveness depending on the recommendation approach [7]. In this direction, the use of robots for behavioural change [8] has been demonstrated to effectively influence users’ behaviours through personalized strategies. They can assist in

improving habits such as healthy eating, promoting physical activity, improving psychological well-being, and in general, providing continuous support. In recent years, due to their computational power and memory efficiency, transformers have gained significant attention in recommender systems, showing to be effective across different areas of application [9]. Transformers constitute the basis structure for Large Language Models (LLMs), which have recently emerged as a powerful solution to formalise recommender systems, thanks to their generalization and reasoning capabilities [10]. However, it has been noted that the need for carefully crafted prompt with high level of specificity, and the fact that LLMs are generally relying on textual input may limit the quality of the results, and not fully take into account user behaviours and reactions [11]. For this reason, we aim to integrate a humanoid robot into a transformer-based recommendation system workflow to leverage its interaction capabilities to improve the users’ experience and motivate them to take the recommendations under consideration [12]. We explore techniques, such as inner speech and inner thoughts, for enhancing people’s compliance with a robot’s recommendation. Inner speech, self-talk are thinking out loud strategies that can have strong self-regulating effects that bring people at altering their behaviour, resisting to temptation, changing mood, making a choice, or filtering irrelevant information [13]. Think-aloud strategies have recently been introduced in robotic applications, since they can boost users’ trust, perception of animacy, intelligence and anthropomorphism [14]. In this work, we combine the reasoning power of transformer-based approaches with the interactional advantages provided by a robot. The main contributions consist in the design and implementation of a transformer-assisted multi-interface recommendation system, which was trained with a new food recipes dataset containing dietary needs and nutritional values to tailor food recommendations based on different user profiles. We tested the approach in a user study, using a self-talk strategy to convince the users to accept the recommendations, and showed that our approach satisfies the users’ expectations. To the best of our knowledge, this is among the first studies to combine an embodied agent and a transformer-based recommender system to adapt to the user’s preferences and requirements.

II. RELATED WORK

In recent years, machine learning techniques have been investigated to formalise preference learning and predictive models based on user preferences [15]. A way to collect users’ personal food preferences could be that of recording “food diaries”, with their eating habits together with

This work was supported by the FAIR project (PE0000013) funded by the Italian Ministry for Universities and Research (L. Raggioli), the ADVISOR project funded by European Union - Next Generation EU, Mission 4 Component 1 CUP E53D23016260001, PRIN 2022 PNRR (A. Rossi and S. Rossi).

¹Department of Electrical Engineering and Information Technologies, University of Naples Federico II, Naples, Italy {luca.raggioli, silvia.rossi, alessandra.rossi}@unina.it

information such as time of day, mood and health during meals [16]. Such a strategy requires long data collection times, where the user has to record data for several days, weeks or even months before being able to fully use the application. Another approach, that is closer to the aims of our work, is presented by Toledo et al. [17] used a multi-criteria decision analysis tool to filter inappropriate foods and an optimization model to generate a menu that maximizes the user’s preferences based on the nutritional requirements. Yang et al. [18] proposed a system that employs a web-based interface to gather user preferences through visual choices and recommend dishes based on visual similarity. Although our system also presents dishes via a website, it collects explicit user preferences rather than relying on visual similarity, which generate more tailored recommendations.

Transformers are a popular Deep Learning architecture which was first introduced in 2017 [19] and is currently at the base of foundational and generative models such as OpenAI’s ChatGPT¹, Google’s Gemini², and Microsoft’s Copilot³. They have recently become a very pervasive methodology in machine learning strategies and across diverse types of applications. Transformers’ versatility and computational power make them a very suitable methodology also for food and recipe recommendation systems. Some examples include RecipeBowl, which it is a cooking recommendation system that suggests ingredients and recipes based on input ingredients and cooking tags [20]. Another similar example is MeniAI, which combines OCR and a transformer-based Learning to Rank model to rank and suggest dishes, based on users’ nutritional preferences, from the picture of a menu [21]. A solution that is particularly interesting and fitting for the scope of our work is the Behaviour Sequence Transformer [22], which was used to process and leverage sequential signals underlying users’ behaviour sequences to produce e-commerce recommendations.

A. Robots in Food Recommendation

Recent studies have explored the potential of robots in food recommendation scenarios, showing how embodiment can play a crucial role in enhancing persuasiveness and user satisfaction with robotic recommendations [23]. In particular, physical robot has been shown to lead to higher user satisfaction compared to virtual ones [24]. De Carolis et al. [25] investigated the use of a social robot as an interface for providing personalised nutritional information based on automatically recognised demographic traits like gender and age, observing improved persuasive effects, including higher perceived informational quality, motivational strength, and social believability. Social robots have also been shown to effectively deliver psychotherapeutic treatment programs to help people change their eating habits and lose weight [26]. This requires different types of strategies. For instance, verbal and physical cues for humanoid robots were evaluated and tested to persuade children aged 8-9 to eat more

fruits and vegetables [27]. In this direction, health-oriented applications have shown that robot-delivered motivational interventions can be effective at reducing high-calorie snack consumption, facilitating weight loss [8].

III. BST-ASSISTED ROBOTIC FOOD RECOMMENDER

Based on trends in the literature, in this work we aimed to combine strategies that rely on user preferences, intolerances and dietary needs, collected through visual choices and leveraged to generate a food recommendation using a transformer-based approach. One of the main aims of this work is to use a social robot to provide personalised support in dietary choices. Finally, we want to assess whether characterised robot’s behaviours can impact people’s choices.

To this extent, we designed a personalised diet recommendation system as a client-server application in which there is a clear separation between the frontend and backend to allow the adoption of different robotic systems. In our architecture, the server takes care of processing the data to produce the recommendation options for the user, a graphical web interface is used to collect the user preferences and display the recommendations, and the robot is used for guiding the users during their use of the system, and persuading the user to accept the recommendation. We have devised a database to store the feedback provided by the users during the interaction to allow the robot to adapt coherently to the user’s needs across different interactions.

A. Food Recipes Dataset

As a first step to train the transformer model chosen on data that are consistent with the intended scenario, we created a new food recipes dataset. We followed a similar procedure to Fossemò et al. [28], which contained 101 recipes. In our case, the use of a deep learning approach for the food recommendation imposed the need for a much larger dataset. For this reason, we retrieved 4284 recipes from a famous Italian cooking website⁴ and standardised the recipes’ information, such as the ingredients names (e.g., “*Parmigiano Reggiano DOP*”, “*Grated Parmigiano Reggiano DOP*”, and “*Flaked Parmigiano Reggiano DOP*” are all standardised to “*Parmigiano Reggiano DOP*”). Each recipe in the dataset contains: an ID associated with the food (*FID*), the *recipe name*, the *servings category* (i.e., starters, first course, second course, side dish, main course, leavened, dessert), *dietary needs* (i.e., lactose-free, gluten-free, vegetarian), the list of *ingredients*, the *nutritional values*.

Based on the recipes retrieved, we created different profiles for perspective users of the system, taking into account different serving categories, age groups (children, adolescents, young adults, adults, older adults), intolerances and food constraints, and the number of calories in the dish (see Table I). We produced 280 user profile archetypes for capturing possible participants profiles for the system.

We defined a mapping dataset containing the IDs of possible user profiles (*UID*), the foods IDs (*FID*), dietary

¹<https://chatgpt.com>

²<https://gemini.google.com>

³<https://copilot.microsoft.com>

⁴GialloZafferano <https://www.giallozafferano.it/>

TABLE I: Examples of profiles generated.

UID	Category	Age Group	Dietary Needs	Calories
1	First Course	Children	None	350
2	Second Course	Adolescent	Lactose-Free	800
3	Side Dish	Young Adult	Gluten-Free	450
4	Leavened	Adult	Vegetarian	650
5	Dessert	Older Adult	None	500

needs, nutritional values, serving category, and ingredients. We will refer to this as the Users-Foods set, which contains 29073 instances. We created this set for training the model BST with a large collection of data that are specifically coded to formalise our intended use case. This allows to use the participant’s selected preferences to produce similar recipes recommendations, without retraining the model.

B. Behaviour Sequence Transformer

We used a variation of the Transformer, called Behavior Sequence Transformer [22]. The main idea consists in analysing the history of articles with which the user has interacted to identify possible correlations and create personalised recommendations, increasing the probability that the user will select the proposed article. We have adapted an open-source python implementation⁵ of the model presented by Chen et al. [22]. In order to process more information than the original implementation, we added embedding units to also consider other recipe details that must be considered by the model for its inference. In particular, each instance of the dataset consisted of the UID associated with six FIDs, each followed by 3 binary values to signal the presence of any dietary need, an integer value for the category (1 to 7), 3 binary values for the nutritional values (calories, carbohydrates and proteins), and a 1204 items array representing the ingredients in the dataset, where 1 identifies ingredients that are contained in recipe FID, and all the other are set to 0.

While investigating the choice of optimal parameters for the BST was beyond the aims of this research, we trained and validated the model on the dataset with a K-Fold Cross Validation (K=5). The metrics used to evaluate the quality of the recommendations is the Normalised Discounted Cumulative Gain [29], which is a widely used metric to assess the efficacy of ranking algorithms by also considering the order in which the recommendations are provided. This methodology allows setting the number of the top N highest ranking items in the output to be evaluated. We set this value to 25, as used in the recipes’ generation phase. The average NDCG computed from the K-Fold was 0.716 (std. 0.003). It is important to note that the quality of the recommendations is strongly related to the amount of available choices at the model’s disposal. For instance, the number of recipes that are lactose-free, gluten-free and vegetarian will be lower compared to having no dietary needs, and might lead to recommendations that are not related to the query in input).

⁵<https://github.com/jiwidi/Behavior-Sequence-Transformer-Pytorch/tree/master>

C. Database

The database is structured to store the participants preferences and the food recommendations produced based on them. The database is implemented in SQLite, and consists of three tables: the *Participant* table contains information regarding the human participants; the *Preferences* table contains the recipes they were presented during the profile setting up (whether they liked or disliked the food, and any dietary needs); and the *Recommendation* table stores foods recommended based on the participant’s preferences and whether they were accepted. Between Participant and Preferences tables, and between Preferences and Recommendation tables, we established 1-to-many relationships (i.e., every user can have multiple preferences settings, and for each setting multiple recommendations can be generated).

D. Webpage Interface

The web interface allows the following (see Figure 1): 1) Login: the user logs in to use the system’s previously acquired knowledge about their preferences; 2) Registration: participant inputs new username and password, gender, age range, and any dietary needs; 3) Food categories: the user chooses what kind of course they want and any calories limit. If a combination course-calories limit can not be satisfied, the robot apologises and asks to amend the option; 4) Preferences choice: based on the food categories and calories limits, the systems retrieves recipes from the database, which can be selected by pressing on the “Like” and “Don’t Like” buttons; 5) Recommendation: for each category selected, a recommendation is generated by the BST model based on the dietary needs and accepted using a “Like” button.

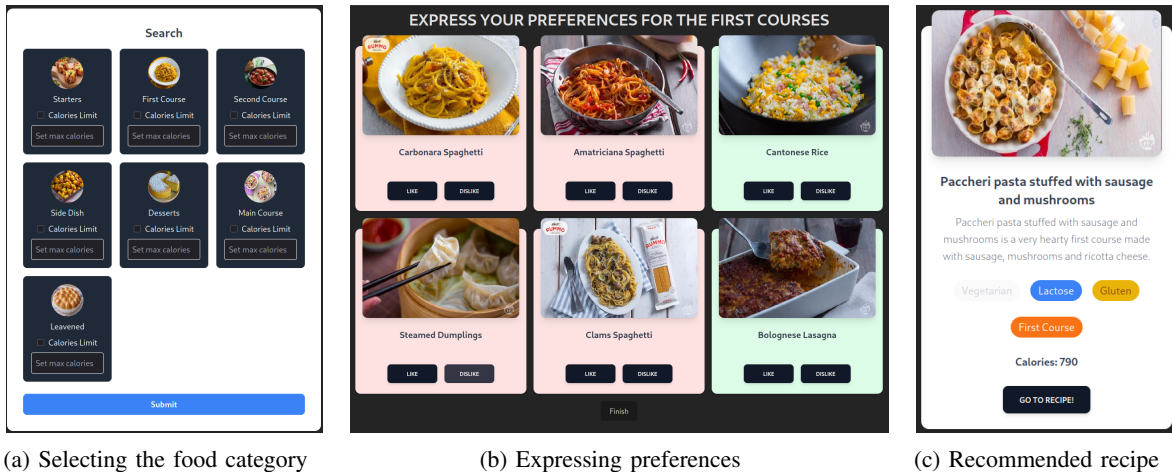
When the users are faced with the recommendation view, they can accept or decline the recipe provided.

IV. EXPERIMENTAL APPROACH

In this work, we aim at investigating whether our recommended system is capable of providing satisfying and acceptable recommendations, and a social robot is able to increase the user’s perception and the acceptance rate of the recommended food options. The robot chosen is the humanoid robotic head Furhat⁶. In the experimental setting designed to deploy our model, a participant sits in front of a computer with the robot placed next to it.

The experimental procedure is articulated in three main steps: registration and login, creation of a preferences model, recommendation and user feedback. The interaction with the web interface and the robot lasts approximately from 3 to 5 minutes. Participants were told to imagine having Furhat as their robotic companion that helps them in their daily food planning based on their dietary requirements. The interaction is started by the robot, which explains to the users that they can log in or register as a new user. Once the user is registered and logged in, the user has to choose food categories and any calories limits from the web interface. The robot guides the participant in the process. Participants are

⁶<https://www.furhatrobotics.com>



(a) Selecting the food category

(b) Expressing preferences

(c) Recommended recipe

Fig. 1: Example of the main three views, where the participants choose the food category, with any calories limits, their food preferences, and received the food recommendation.

free to select multiple food categories, and for each chosen they may decide to specify a limit of calories. Participants are then showed the database recipes that are consistent with their profile and the selected food options. After the users have selected (liked or disliked) the recipes, these are used for a one-shot training of the BST model. Finally, for every food category selected by the user, participants are shown with recipe recommendations produced by the BST. In this phase (see Figure 1c), the robot asks the participants whether they are happy to accept the recommended dish. In case of a rejected recommendation, the robot can have one of two possible behaviours, in which it either provides a new recommendation or tries to persuade the participants to select the last recommendation generated by the BST. In the end, the server prompts the robot to thank them and say goodbye.

A. Experimental Procedure

We designed an in between-subject study, where participants were randomly assigned to only one of two experiment conditions (*Neutral* and *Persuasive*). The robot’s behaviours varied in the two experimental conditions only in the recommendation phase when the participants reject the food recommended by the robot: 1) *Neutral*: the robot apologises because the participant did not like the recommendation, and provides a new recommendation. The behaviour is repeated until the user accepts the recommendation; 2) *Inner Thoughts*: the robot tries to persuade the participant to reconsider the recipe, by reasoning out loud on the appropriateness of the recommendation for the participant. For example, the robot says “*It is such a shame they do not like this recipe, everybody usually love it!*” or “*You should really try this recipe, it is very healthy and tasty!*”.

After the interaction is concluded, the participants are asked to fill out the Perceived Social Intelligence (PSI) questionnaire [30] to explore more in depth different features of the robot’s social intelligence. The scales cover three main areas of perceived social intelligence in robots: Cognition, Emotions, and Behaviours. Since the task devised is less con-

cerned with the Emotions components, we investigated only the following dimensions: Recognizes Human Behaviours (**RB**); Recognizes Human Cognitions (**RC**); Adapts to Human Behaviours (**AB**); Adapts to Human Cognitions (**AC**); Identifies Humans (**IH**); Social Competence (**SOC**); Friendly (**FRD**); Helpful (**HLP**); Trustworthy (**TRU**).

At the end of the experimental session, participants were also asked to rate the recommendations, if they would use a system like this on a daily basis, and to justify their answers with an open-ended question. Participants rated each set of questions using 5-point Likert scales ranging from 1 (strongly disagree / not all) to 5 (strongly agree / very much).

B. Participants

We recruited 40 adult participants (10 females, 29 males, and 1 wrote “other” in the gender field), aged between 19 and 48 years old (*mean*: 26.7, *std*: 5.9). The gender imbalance was due to the fact that the recruitment was mostly carried out among students and staff at the Engineering Department of the University, in addition to a science faire. Participants stated to be Italians, 16 of which had never had any type of encounter with robots, 9 had already seen a robot in digital media, 12 had already interacted with a robot, 12 had already participated in a study with robots, and 5 work with robots. Before starting the interaction with the system, participants were informed of the task, and signed a written consent. Participants were evenly distributed over the two robot’s behavioural conditions (*Neutral* and *Inner Thoughts*).

V. RESULTS

In this work, we evaluate participants’ perception of the recommendations provided and the robot.

The acceptance rate of the first recommendation was very high in both the *Neutral* and the *Inner Thoughts* conditions (80% in each). However, we observed that participants did not always accept the first suggestion of the robot. In the *Neutral* case, participants accepted the second recommendation in 15% of the cases, and the third one in the remaining

TABLE II: Satisfaction rates for all the participants and divided by conditions. We report the overall rates, and those observed in the case of acceptance and rejection of the recommendations.

	Overall	Accepted	Rejected
All	3.875	4.083	2
Neutral	4.1	4.1	n/a
Inner Thoughts	3.65	4.063	2

5%. In the *Inner Thoughts* condition, instead, only the 20% did not accept the first recommendation, and the “persuasive” robot was not able to convince them to accept it.

A. Overall Satisfaction Rate

The qualitative analysis of the overall satisfaction revealed that in the case of rejection of the recommendation, the participants were all driven by their dislike towards the recipe, rather than by any discomfort or annoyance towards the robot. In such cases, however, the strategy was to be considered ineffective in swaying the users’ intentions. Regarding, instead, the cases in which the participants accepted the recommendations (in both the *Neutral* and the *Inner Thoughts* conditions), they reported that the suggestions were in line with their preferences. In the former, in particular, the average satisfaction rates by number of recommendations required to accept the suggestions were respectively 4.31 for the first one, 3 for the second, and 4 for the third recommendation. The quantitative satisfaction rates are shown in Table II. We can see how, in the case of acceptance, the overall satisfaction is almost the same across the two conditions. Results also showed that, in the case of acceptance, the overall satisfaction rates (4.083) is almost the same across the two conditions (4.1 for the *Neutral* and 4.063 for the *Inner Thoughts* conditions, respectively). In case of rejections of the recommendations, the overall satisfaction rates are 2 (2 for the *Inner Thoughts*, and n/a for the *Neutral* conditions, respectively). Finally, 87.5% of the participants said they would enjoy using such a system on a daily basis to have new ideas to differentiate their diet, while 5% reported they would not like to integrate such a system in their lives.

B. User’s Perception

We assessed users’ perception by analysing the average aggregate value for each of the scales of the PSI questionnaire (see Figure 2). Using a MANOVA to compare the two strategies, *Neutral* and *Inner Thoughts*, we did not observe any statistically significant difference between the two strategies, as the results show in Table III.

C. Discussion

The analysis of the acceptance rate shows that the recommendations of the robot were generally consistent with participants preferences (with 80% of first recommendations being accepted), and that the user profiles archetypes created successfully allowed to deploy the BST model in real time to provide recommendations coherently. In the case of the



Fig. 2: Average aggregates of the scales.

TABLE III: Comparison of PSI subscales between the *Neutral* and *Inner Thoughts* robot’s behaviours, using the univariate MANOVA, with mean values.

Scales	F	df	p	mean Neutral	mean Inner Thoughts
RB	4.022	1	0.052	2.99	2.40
RC	0.573	1	0.454	3.01	3.19
AB	2.783	1	0.103	3.24	2.69
AC	0.028	1	0.866	3.13	3.09
HI	<0.01	1	1.000	3.35	3.35
SOC	0.145	1	0.705	2.33	2.44
FRD	0.003	1	0.959	2.80	2.81
HLP	0.296	1	0.590	3.84	3.98
TRU	1.460	1	0.234	3.99	3.70

Neutral condition, 3 participants accepted the second suggestion, arguing for the rejections based on their disliking for a specific ingredient in the recipes which was not correctly reflected by the recommendations (e.g., the cauliflower). The participant, who accepted the third recommendation, stated that they didn’t know the meal suggested by the system. In the *Inner Thoughts* condition, the robot tries to subtly convince the participants to reconsider the rejected recommendations, but they reported that they did not like either the recipe, or some of the ingredients. The acceptance rate results are reflected in the satisfaction rates, whereas the participants who accepted the recommendations rated the system above 4 in average. It is interesting to note that, in the *Neutral* condition, while the participants who accepted the second recommendation rated the system lower than those who accepted the first one (satisfaction rate 3 and 4.31 respectively), the participant who accepted at the third recommendation rated the system with 4. This might indicate that this user appreciated the system’s capacity to differentiate and come up with a recipe that pleases them.

Participants generally perceived the robot as helpful and trustworthy, with 3.91 and 3.84 respectively. Participants also positively believed that the robot was able to adapt its behaviour appropriately based upon people’s thoughts and beliefs, to detect human presence, their thoughts and beliefs. This suggests that the robot was perceived as attentive to the users’ thoughts and beliefs, managing to provide recommendations that were in line with their expectations. However, the robot was generally perceived as less friendly and adaptive in regard to the participants’ behaviours. This is

to be expected, as the robot's behaviours were mostly aimed at presenting the options effectively, rather than adapting to the users' demeanours. This is even more clear when observing the lowest rates for social competence (2.38 - SOC subscale), which shows that the participants did not find the robot socially engaging. If we consider the results by condition, we do not observe any significant differences. This is likely due to the fact that only four participants rejected the food recommendation of the robot, and therefore experienced the robot's *Inner Thoughts* behaviours. The majority of the participants in this condition accepted immediately the first recommendation proposed by the BST. Participants (87.5%) positively perceived the system, and would use it daily.

VI. CONCLUSIONS

This work presented a multi-modal approach for food recommendations based on user preferences and facilitated by a humanoid robot. Our results show that the system manages to effectively provide recommendations that were deemed acceptable by the users, with a high satisfaction ratings by the participants. The robot was also perceived as helpful and trustworthy. Two robot's behaviours were investigated, *Neutral* and *Inner Thoughts*, but due to the low number of participants who experienced the latter, no significant difference was observed between the two. The low number of cases exposed to a persuasive robotic behaviour is a main limitation, which will be addressed in future extensions of this study. Future works will also target the deployment of this system in a more structured setting and for a long term dietary plan implementation. Future iterations of this work will also investigate the use of LLMs, and will also be compared with different persuasive techniques to allow a more thorough evaluation of the results.

REFERENCES

- [1] V. Espín, M. V. Hurtado, and M. Noguera, "Nutrition for elder care: a nutritional semantic recommender system for the elderly," *Expert Systems*, vol. 33, no. 2, pp. 201–210, 2016.
- [2] G. Agapito, M. Simeoni, B. Calabrese, I. Caré, T. Lamprinouidi, P. H. Guzzi, A. Pujia, G. Fuiano, and M. Cannataro, "Dietos: A dietary recommender system for chronic diseases monitoring and management," *Computer Methods and Programs in Biomedicine*, vol. 153, pp. 93–104, 2018.
- [3] S. Sandri and A. Molinari, "Preference learning in food recommendation: the "myfood" case study," *2023 3rd International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME)*, pp. 1–6, 2023.
- [4] S. Lukasik, S. Tobis, K. Wiczorowska-Tobis, and A. Suwalska, "Could robots help older people with age-related nutritional problems? opinions of potential users," *International Journal of Environmental Research and Public Health*, vol. 15, no. 11, 2018.
- [5] C. Trattner and D. Elsweller, "Food recommender systems: Important contributions, challenges and future research directions," 2017.
- [6] S. Rossi, M. Staffa, and A. Tamburro, "Socially assistive robot for providing recommendations: Comparing a humanoid robot with a mobile application," *International Journal of Social Robotics*, vol. 10, no. 2, pp. 265–278, 2018.
- [7] S. Herse, J. Vitale, M. Tonkin, D. Ebrahimian, S. Ojha, B. Johnston, W. Judge, and M. Williams, "Do you trust me, blindly? factors influencing trust towards a robot recommender system," in *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 7–14, 2018.
- [8] N. L. Robinson, J. Connolly, L. Hides, and D. J. Kavanagh, "Social robots as treatment agents: Pilot randomized controlled trial to deliver a behavior change intervention," *Internet Interventions*, vol. 21, p. 100320, 2020.
- [9] Y. Tay, M. Dehghani, D. Bahri, and D. Metzler, "Efficient transformers: A survey," *ACM Comput. Surv.*, vol. 55, Dec. 2022.
- [10] Z. Zhao, W. Fan, J. Li, Y. Liu, X. Mei, Y. Wang, Z. Wen, F. Wang, X. Zhao, J. Tang, and Q. Li, "Recommender systems in the era of large language models (llms)," *IEEE Transactions on Knowledge and Data Engineering*, vol. 36, no. 11, pp. 6889–6907, 2024.
- [11] J. Zhang, Z. Wang, W. Liu, X. Liu, and Q. Zheng, "A unified approach to designing sequence-based personalized food recommendation systems: tackling dynamic user behaviors," *International Journal of Machine Learning and Cybernetics*, vol. 14, no. 9, pp. 2903–2912, 2023.
- [12] V. Chidambaram, Y.-H. Chiang, and B. Mutlu, "Designing persuasive robots: how robots might persuade people using vocal and nonverbal cues," in *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction, HRI '12*, (New York, NY, USA), p. 293–300, Association for Computing Machinery, 2012.
- [13] F. Racy and A. Morin, "Relationships between self-talk, inner speech, mind wandering, mindfulness, self-concept clarity, and self-regulation in university students," *Behavioral Sciences*, vol. 14, no. 1, 2024.
- [14] A. Pipitone, A. Geraci, A. D'Amico, V. Seidita, and A. Chella, "Robot's inner speech effects on human trust and anthropomorphism," *International Journal of Social Robotics*, vol. 16, no. 6, pp. 1333–1345, 2024.
- [15] M. De Gemmis, L. Iaquinta, P. Lops, C. Musto, F. Narducci, and G. Semeraro, *Learning Preference Models in Recommender Systems*, pp. 387–407. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011.
- [16] A. A. Metwally, A. K. Leong, A. Desai, A. Nagarjuna, D. Perelman, and M. Snyder, "Learning personal food preferences via food logs embedding," 2021.
- [17] R. Yera Toledo, A. A. Alzahrani, and L. Martínez, "A food recommender system considering nutritional information and user preferences," *IEEE Access*, vol. 7, pp. 96695–96711, 2019.
- [18] L. Yang, Y. Cui, F. Zhang, J. P. Pollak, S. Belongie, and D. Estrin, "Plateclick: Bootstrapping food preferences through an adaptive visual interface," in *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management, CIKM '15*, (New York, NY, USA), p. 183–192, Association for Computing Machinery, 2015.
- [19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems* (I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds.), vol. 30, Curran Associates, Inc., 2017.
- [20] M. Gim, D. Park, M. Spranger, K. Maruyama, and J. Kang, "Recipebowl: A cooking recommender for ingredients and recipes using set transformer," *IEEE Access*, vol. 9, pp. 143623–143633, 2021.
- [21] X. Ju, F. P. W. Lo, J. Qiu, P. Shi, J. Peng, and B. Lo, "Menuai: Restaurant food recommendation system via a transformer-based deep learning model," 2022.
- [22] Q. Chen, H. Zhao, W. Li, P. Huang, and W. Ou, "Behavior sequence transformer for e-commerce recommendation in alibaba," 2019.
- [23] S. Herse, J. Vitale, D. Ebrahimian, M. Tonkin, S. Ojha, S. Sidra, B. Johnston, S. Phillips, S. L. K. C. Gudi, J. Clark, W. Judge, and M.-A. Williams, "Bon appetit! robot persuasion for food recommendation," in *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction, HRI '18*, (New York, NY, USA), p. 125–126, Association for Computing Machinery, 2018.
- [24] K. Sakai, Y. Nakamura, Y. Yoshikawa, and H. Ishiguro, "Effect of robot embodiment on satisfaction with recommendations in shopping malls," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 366–372, 2022.
- [25] B. N. De Carolis, F. D'Errico, N. Macchiarulo, *et al.*, "'keep the user in mind!' persuasive effects of social robot as personalized nutritional coach.," in *PSYCHOBOT*, 2019.
- [26] N. L. Robinson and D. J. Kavanagh, "A social robot to deliver a psychotherapeutic treatment: Qualitative responses by participants in a randomized controlled trial and future design recommendations," *International Journal of Human-Computer Studies*, vol. 155, p. 102700, 2021.
- [27] I. Baroni, M. Nalin, M. Coti Zelati, E. Oleari, and A. Sanna, "Designing motivational robot: How robots might motivate children to eat fruits and vegetables," in *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, pp. 796–801, 2014.

- [28] D. Fossemò, F. Mignosi, L. Raggioli, M. Spezialetti, F. A. D'Asaro, *et al.*, "Using inductive logic programming to globally approximate neural networks for preference learning: challenges and preliminary results," in *CEUR WORKSHOP PROCEEDINGS*, pp. 67–83, 2022.
- [29] O. Jeunen, I. Potapov, and A. Ustimenko, "On (normalised) discounted cumulative gain as an off-policy evaluation metric for top-n recommendation," in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '24*, (New York, NY, USA), p. 1222–1233, Association for Computing Machinery, 2024.
- [30] K. A. Barchard, L. Lapping-Carr, R. S. Westfall, A. Fink-Armold, S. B. Banisetty, and D. Feil-Seifer, "Measuring the perceived social intelligence of robots," *J. Hum.-Robot Interact.*, vol. 9, Sept. 2020.